

Large sample analysis of the median heuristic

Damien Garreau^{*†}

e-mail: damien.garreau@tuebingen.mpg.de

and

Wittawat Jitkrittum[†]

e-mail: wittawat.jitkrittum@tuebingen.mpg.de

and

Motonobu Kanagawa[†]

e-mail: motonobu.kanagawa@tuebingen.mpg.de

[†]*Max Planck Institute for Intelligent Systems*

Abstract: In kernel methods, the median heuristic has been widely used as a way of setting the bandwidth of RBF kernels. While its empirical performances make it a safe choice under many circumstances, there is little theoretical understanding of why this is the case. Our aim in this paper is to advance our understanding of the median heuristic by focusing on the setting of kernel two-sample test. We collect new findings that may be of interest for both theoreticians and practitioners. In theory, we provide a convergence analysis that shows the asymptotic normality of the bandwidth chosen by the median heuristic in the setting of kernel two-sample test. Systematic empirical investigations are also conducted in simple settings, comparing the performances based on the bandwidths chosen by the median heuristic and those by the maximization of test power.

1. Introduction

Kernel methods form an important class of algorithms in machine learning and statistics (Schölkopf and Smola, 2002; Muandet et al., 2017). They make use of rich feature spaces that depend only on the kernel chosen by the user. Given a positive semi-definite kernel k and observations x_1, \dots, x_n , the first step of most kernel-based methods is to compute the Gram matrix $K = (k(x_i, x_j))_{1 \leq i, j \leq n}$. Thanks to the celebrated *kernel trick*, all ensuing computations need only the knowledge of K .

In this paper, we are especially interested in data lying in a metric space (\mathcal{X}, δ) . When this is the case, commonly used kernels are *radial basis function (RBF) kernels* of the form

$$k(x, y) = f(\delta(x, y) / \nu), \quad x, y \in \mathcal{X}, \quad (1)$$

^{*}Corresponding author

[†]Max-Planck-Ring 4, 72 076 Tübingen, Germany

where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function and ν is a positive parameter called the *bandwidth*. In many applications, the space \mathcal{X} is \mathbb{R}^d and δ is derived from the Euclidean norm $\|x\| = \sqrt{\sum_i x_i^2}$, that is, $\delta(x, y) = \|x - y\|$. Numerous kernels used in practice belong to this class of kernels. For instance, $f(x) = \exp(-x^2)$ corresponds to the Gaussian kernel (Aizerman et al., 1964), arguably the most popular positive definite kernel used in applications (see, for instance, Vert et al., 2004). The function $f(x) = \exp(-x)$ yields the exponential kernel—also called Laplace or Laplacian kernel—whereas more exotic f give rise to less common kernels such as the rational quadratic kernel, the wave kernel or the Matérn kernel (see Genton, 2001, and references therein).

It is well-known that the performance of kernel methods depends highly on the kernel choice. In practice, this choice often reduces to the calibration of the bandwidth ν , which may even be more important than the choice of f (Schölkopf and Smola, 2002, Section 4.4.5). Since the Gram matrix depends only on the $\|x_i - x_j\|/\nu$ in this case, it is reasonable to pick ν in the same order as the family of all pairwise distances $(\|x_i - x_j\|)_{1 \leq i, j \leq n}$. As an example, suppose that we settled for the Gaussian kernel. Then when $\nu \rightarrow 0$, the Gram matrix K is the identity matrix, and when $\nu \rightarrow \infty$, the components of K are all equal to 1. All relevant information about the data is lost in both these extreme cases. This is a general phenomenon, even though the values taken by K in the degenerate cases depend on the function f . Hence a reasonable middle-ground for choosing ν is to pick a value “in the middle range” of the $(\|x_i - x_j\|)_{1 \leq i, j \leq n}$, that is, an empirical quantile, which is often set to be the median. This strategy is called the *median heuristic*; see Section 2.2 for a precise definition.

The median heuristic has been extensively used in practice.¹ In fact, in unsupervised learning where no principled way is available for bandwidth selection, the median heuristic may be one of the first choice. This is the case for, among many others, kernel PCA (Schölkopf et al., 1998), kernel CCA (Bach and Jordan, 2002) and kernel two-sample test (Gretton et al., 2012a). It has also been used in supervised learning, *e.g.*, kernel SVM (Boser et al., 1992) or kernel ridge regression (Hoerl and Kennard, 1970). There, one can choose ν from prescribed candidate values by performing cross-validation, and it is common to use the median heuristic to set the scale of these candidate values, *e.g.*, ν may be chosen from $(2^a H_n)_{a=-M}^M$, where H_n denotes the value given by the median heuristic and M is some positive integer. In fact, the empirical median is the default bandwidth choice in the kernel SVM implementation of the `kernlab` R package (Karatzoglou et al., 2004).

Despite its popularity, there is very little theoretical understanding of the median heuristic. To the best of our knowledge, the only work in this direction is contained in Reddi et al. (2015). They observe that the median of all the pairwise distances has to be close to the mean pairwise distance, $\mathbb{E} \|X_i - X_j\|$.

¹ As noted in Flaxman et al. (2016), the origin of the median heuristic is unclear and does not appear in the monograph of Schölkopf and Smola (2002), while it has become the main reference for this heuristic. The earliest appearance of the median heuristic that we know of is in Sriperumbudur et al. (2009, Section 5). Gretton et al. (2012a) refers to Takeuchi et al. (2006) and Schölkopf et al. (1997) for similar heuristics.

Based on this observation, they obtain the asymptotic of the median heuristic when the dimension of the data goes to infinity, using the asymptotic of the mean pairwise distance. This argument can be made rigorous by observing that, given a random variable X with a second order moment, $|\mathbb{E}X - \text{med}(X)| \leq \sqrt{\text{Var}(X)}$ holds (Mallows, 1991). Hence the observation of Reddi et al. (2015) is correct, up to a variance term. We will see in Section 3 that our results make this insight more precise.

The aim of this article is to advance our understanding of the median heuristic, both theoretically and empirically. To this end, we focus on the setting of the kernel two-sample test (Gretton et al., 2012a), which has been used in a wide range of applications including transfer learning (Long et al., 2015) and generative adversarial learning (Li et al., 2017); see Muandet et al. (2017) for a recent extensive survey. As kernel two-sample testing is an unsupervised problem, numerous authors report the use of the median heuristic in their experiments (Sriperumbudur et al., 2009; Arlot et al., 2012; Reddi et al., 2015; Muandet et al., 2016; Zhang et al., 2017; Jitkrittum et al., 2017a; Sutherland et al., 2017). On the other hand, there is a more principled approach for bandwidth (or kernel) selection based on the maximization of the test power (Gretton et al., 2012a; Jitkrittum et al., 2017a; Sutherland et al., 2017). Apart from theoretical analysis, we will conduct empirical comparisons between these approaches, and discuss when the median heuristic works and when it fails in simple examples; see Section 4.

This paper is organized as follows. Our setting is made explicit in Section 2 and we show in the same section how it is relevant for this application. In Section 3, we state our main result: the median heuristic is asymptotically normal when the number of observations goes to $+\infty$. In particular, the median heuristic converges towards the theoretical median of a target distribution that we describe in terms of the samples' distributions. This result is obtained by the mean of an auxiliary proposition that we think has an interest of its own, that is, a central limit theorem for a certain class of U -statistics that we state in the same section. Finally, we use this result in Section 4 to investigate the quality of the median heuristic as a bandwidth choice in simple settings. While we provide sketches of proofs, the complete proofs are deferred to the Appendix.

2. Setting

Given any random variable Z , the notation Z' will stand for an independent copy of Z and we write $Z \sim \pi$ if Z follows the distribution of π . Unless specified in subscript, the expected value is taken with respect to all the random variables that appear in the expression. For instance, $\mathbb{E}[h(X, Y)]$ stands for $\mathbb{E}_{X, Y}[h(X, Y)]$. We also denote by $L^2(P)$ the space of real-valued measurable functions f such that $\mathbb{E}[f(X)^2] < +\infty$ where $X \sim P$, with P a probability distribution. We use $\xrightarrow{\mathcal{L}}$ for convergence in law, and $\xrightarrow{\mathbb{P}}$ for convergence in probability.

In the following, we suppose that we are given a *triangular array* of indepen-

dent \mathbb{R}^d -valued random variables. Namely, for each sample size n , we suppose that we observe an *entirely* new sample $X_{n,1}, \dots, X_{n,n}$. We see later how this asymptotic setting is adapted to the methods we consider. Let X (resp. Y) be a \mathbb{R}^d -valued random variable following the law P (resp. Q). Our main assumption on the distribution of the $X_{n,i}$ is that our observations are split in two contiguous segments such that, on the left segment they follow P , and on the right segment they follow Q . That is,

Assumption 2.1 (Split sample). There is a fixed $\alpha \in (0, 1)$ such that, for any $i \leq \alpha n$, $X_{n,i} \sim P$ and $X_{n,i} \sim Q$ if $i > \alpha n$.

We will assume from now on that αn is an integer. Everything that follows can be readily adapted by replacing αn with $\lfloor \alpha n \rfloor$ when it is needed.

2.1. Connections with kernel two-sample test

Let us briefly recall the modus operandi of the kernel two-sample test with our notation. The goal of two-sample test is to decide whether $P = Q$ or $P \neq Q$ given observations $x_1, \dots, x_{\alpha n}$ and $x_{\alpha n+1}, \dots, x_{\alpha n}$. Gretton et al. (2007) have proposed a kernel method for two-sample testing, that relies on the mean embeddings of P and Q in the reproducing kernel Hilbert space \mathcal{H} associated with k . Let us call μ_P and μ_Q these embeddings; then a good measure of proximity between the distributions P and Q is the so called *maximum mean discrepancy* (MMD), which is defined as $\|\mu_P - \mu_Q\|_{\mathcal{H}}$. It is also possible to write the (squared) MMD as

$$\text{MMD}^2(P, Q) = \mathbb{E}[k(x, x')] - 2\mathbb{E}[k(x, y)] + \mathbb{E}[k(y, y')],$$

where $x, x' \sim P$ and $y, y' \sim Q$ are independent. Setting $M = \alpha n$ and $N = (1 - \alpha)n$, Gretton et al. (2007) provide an unbiased estimate of this quantity,

$$\begin{aligned} \text{MMD}_u^2(P, Q) &= \frac{1}{M(M-1)} \sum_{\substack{i,j=1 \\ j \neq i}}^M k(x_i, x_j) \\ &+ \frac{1}{N(N-1)} \sum_{\substack{i,j=M+1 \\ i \neq j}}^{M+N} k(x_i, x_j) - \frac{2}{MN} \sum_{\substack{i=1 \\ j=M+1}}^{M,N} k(x_i, x_j). \end{aligned}$$

In the following, we refer to MMD_u^2 as the *quadratic-time* statistic, since its computational cost is quadratic in the number of samples. In the case $\alpha = 1/2$, that is $M = N$, Gretton et al. (2012a) also proposed a linear-time estimate,

$$\begin{aligned} \text{MMD}_\ell^2(P, Q) &= \frac{2}{M} \sum_{i=1}^{M/2} k(x_{2i-1}, x_{2i}) + k(x_{M+2i-1}, x_{M+2i}) \\ &- k(x_{2i-1}, x_{M+2i}) - k(x_{2i}, x_{M+2i}). \end{aligned}$$

Letting n grow to infinity corresponds to let both M and N grow to infinity with the ratio M/N constant equal to $\alpha/(1-\alpha)$.

Remark 2.1. The two-sample test problem is very close to that of *change-point detection* when there is a single change-point, the main difference being that α is generally unknown in the change-point problem. Recently, some *kernel* change-point detection methods have been proposed (Harchaoui and Cappé, 2007; Arlot et al., 2012). These methods aim to detect a change-point by minimizing a *kernelized* least-squares criterion, and the median heuristic is also used in this case to set the bandwidth (Arlot et al., 2012; Garreau and Arlot, 2016). The setting described in this section is also appropriate to study the median heuristic in this case.

2.2. Median heuristic

Suppose that we use a kernel that has the form (1) for a fixed f with bandwidth $\nu > 0$. We define the median heuristic as the choice of bandwidth $\nu = \sqrt{H_n}/2$, where $H_n > 0$ is

$$H_n = \text{Med}\{\|X_{n,i} - X_{n,j}\|^2 \mid 1 \leq i < j \leq n\}, \quad (2)$$

where Med is the empirical median. More precisely, H_n is obtained by first ordering the $\|X_{n,i} - X_{n,j}\|^2$ in increasing order, and then setting H_n to be the central element if $n(n-1)/2$ is odd, or the mean of the two most central elements if $n(n-1)/2$ is even. Note that some authors choose simply $\nu = \sqrt{H_n}$.

In order to investigate the asymptotic properties of H_n , rather than using Eq. (2), we are going to define H_n via the empirical cumulative distribution function of the $\|X_{n,i} - X_{n,j}\|^2$, that is,

$$\forall t \in \mathbb{R}, \quad \hat{F}_n(t) = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} \mathbb{1}_{\|X_{n,i} - X_{n,j}\|^2 \leq t}. \quad (3)$$

For any $p \in (0, 1)$, we define the generalized inverse of \hat{F}_n by $\hat{F}_n^{-1}(p) = \inf\{t \in \mathbb{R} \mid \hat{F}_n(t) \geq p\}$. Then H_n may be written as

$$H_n = \hat{F}_n^{-1}\left(\frac{1}{2}\right). \quad (4)$$

Note that other empirical quantiles as $\hat{F}_n^{-1}(p)$ for $p \in (0, 1)$ can be used in practice (e.g., $p = 0.1$ and $p = 0.9$). Though we are mainly concerned with H_n , we will see that our main result still holds for arbitrary p .

3. Main results

3.1. Empirical cumulative distribution function of the pairwise distances

For any $1 \leq i, j \leq n$, set $\delta_{ij}^2 := \|X_{n,i} - X_{n,j}\|^2$. Under Assumption 2.1, there are only three possibilities for δ_{ij}^2 . Namely, for any fixed, distinct indices i and j




FIG 1. Histogram of the $\|X_{n,i} - X_{n,j}\|^2$ with $n = 400$ for Gaussian distributions in dimension $d = 100$ and proportion $\alpha = .25$. Left panel: change in the mean, $X \sim \mathcal{N}(0, \mathbf{I}_d)$ and $Y \sim \mathcal{N}(10^3 \mathbf{1}, \mathbf{I}_d)$. Right panel: change in the variance, $X \sim \mathcal{N}(0, \mathbf{I}_d)$ and $Y \sim \mathcal{N}(0, 2\mathbf{I}_d)$. The error bars correspond to the standard deviation over 10 repetitions of the experiment.

such that $i < j$,

- (i) if $i \leq \alpha n$ and $j \leq \alpha n$, then δ_{ij}^2 has the distribution of $\|X - X'\|^2$;
- (ii) if $i > \alpha n$ and $j > \alpha n$, then δ_{ij}^2 has the distribution of $\|Y - Y'\|^2$;
- (iii) if $i \leq \alpha n$ and $j > \alpha n$, then δ_{ij}^2 has the distribution of $\|X - Y'\|^2$.

In the following, we set random variables $T_{XX} \sim \|X - X'\|^2$, $T_{YY} \sim \|Y - Y'\|^2$, and $T_{XY} \sim \|X - Y'\|^2$. There are $\alpha n(\alpha n - 1)/2$ occurrences of case (i), so case (i) occurs with proportion α^2 as $n \rightarrow +\infty$. Similarly, case (ii) occurs with proportion $(1 - \alpha)^2$ and case (iii) with proportion $2\alpha(1 - \alpha)$.

Define a mixture distribution $T \sim T_{XX}$, $T \sim T_{YY}$ and $T \sim T_{XY}$ with weights α^2 , $(1 - \alpha)^2$ and $2\alpha(1 - \alpha)$ respectively. Thereafter, we will call T the *target* random variable and denote by F its cumulative distribution function. Intuitively, when $n \rightarrow \infty$, T_n should behave like a n -sample of the target T , provided that the dependency between the $\|X_{n,i} - X_{n,j}\|^2$ that have elements in common is not too strong. Indeed, a specialization of a result stated in the next paragraph shows that

$$\forall t \in \mathbb{R}, \quad \widehat{F}_n(t) \xrightarrow{\mathbb{P}} F(t). \quad (5)$$

We believe that Eq. (5) is already a step in the comprehension of the median heuristic, since we are now able to think about H_n “approximately” as the theoretical median of the target T . We refer to Figure 1 for two examples when P and Q are known distributions.

Before presenting a rigorous statement of Eq. (5), we provide a result that shows how a gap appears between on one side T_{XX} and T_{YY} and T_{XY} on the other side if the distributions of P and Q are well-separated. This explains the “two bumps” behavior depicted in the left panel of Figure 1. The left mode of the empirical distribution corresponds to realizations of T_{XX} and T_{YY} , that are close to zero by definition, whereas the right mode corresponds to realizations of T_{XY} , which can be arbitrarily far from 0.

Lemma 3.1 (Gap between intra- and inter-distances). *Set μ_X (resp. μ_Y)*

the expectation and Σ_X (resp. Σ_Y) the covariance matrix of X (resp. Y). Assume that there exists $\lambda > 75$ such that

$$\|\mu_X - \mu_Y\|^2 \geq \lambda (\text{tr}(\Sigma_X) + \text{tr}(\Sigma_Y)) .$$

Then, with probability at least $1 - 75/\lambda$,

$$\max(T_{XX}, T_{YY}) + \frac{\|\mu_X - \mu_Y\|^2}{25} < T_{XY} .$$

Remark 3.1. Recall that T is a mixture such that $T \sim T_{XX}$, $T \sim T_{YY}$ and $T \sim T_{XY}$ with respective proportions α^2 , $(1 - \alpha)^2$ and $2\alpha(1 - \alpha)$. Since $\alpha^2 + (1 - \alpha)^2 \geq 2\alpha(1 - \alpha)$ holds, Lemma 3.1 therefore implies that the median is determined by the scales of T_{XX} or T_{YY} , and the influence of T_{XY} is relatively weak if $\alpha < 1/2$.

3.2. Convergence of the empirical cumulative distribution function

It turns out that Eq. (5) is a trivial consequence of a much stronger statement. Indeed, $\widehat{F}_n(t)$ can be seen as a sum of three dependent U -statistics with kernel $h(x, y) = \mathbb{1}_{\|x-y\|^2 \leq t}$, and the following result shows that it follows a central limit theorem. We refer to classical textbooks (Lee, 1990; Korolyuk and Borovskich, 2013) for an introduction to the theory of U -statistics.

Proposition 3.1 (CLT for non-identically distributed triangular array U -statistic). Consider $h : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ such that $h \in L^2(P) \cap L^2(Q) \times L^2(P) \cap L^2(Q)$, and suppose that the $(X_{n,i})_{i=1}^n$ satisfy Assumption 2.1. Define

$$U_n = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} h(X_{n,i}, X_{n,j}) ,$$

and set

$$\theta = \alpha^2 \mathbb{E}[h(X, X')] + 2\alpha(1 - \alpha) \mathbb{E}[h(X, Y)] + (1 - \alpha)^2 \mathbb{E}[h(Y, Y')] . \quad (6)$$

Then

$$\sqrt{n}(U_n - \theta) \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma^2) , \quad (7)$$

where $\sigma = \sigma(h, P, Q)$ is defined as

$$\begin{aligned} \sigma^2 &= \text{Var}_X (\alpha \mathbb{E}[h(X, X')] + (1 - \alpha) \mathbb{E}[h(X, Y)]) \\ &+ \text{Var}_Y (\alpha \mathbb{E}[h(X, Y)] + (1 - \alpha) \mathbb{E}[h(Y, Y')]) , \end{aligned} \quad (8)$$

with all of $X, X' \sim P$ and $Y, Y' \sim Q$ being independent.

Proof (sketch). The idea of the proof is the following: (i) split U_n in three terms depending on the relative position of the indices; (ii) write down the Hoeffding decomposition of each of these terms; (iii) show that the remainders are negligible, and (iv) conclude thanks to the central limit theorem for triangular arrays. The complete proof can be found in Appendix B.1. \square

We make the following remarks.

- Central limit theorems for U -statistics are known since the fundamental article of [Hoeffding \(1948\)](#). Prop. 3.1 is in the line of such results. An asymptotic normality result also exists in the non-identically distributed case; see [Hoeffding \(1948, Th. 8.1\)](#). However, this result was not applicable in our setting. The material in [Jammalamadaka and Janson \(1986\)](#) covers the case of a triangular array scheme, but does not cover the non-identically distributed setting. Results regarding *two-sample* U -statistics are closest in spirit but not directly applicable; see [van der Vaart \(1998, Section 12.2\)](#) for an introduction and [Dehling and Fried \(2012\)](#) for recent developments. With our notation, the two-sample U -statistic is written as $(\alpha n(1 - \alpha)n)^{-1} \sum_{i=1}^{\alpha n} \sum_{j=\alpha n+1}^n h(X_{n,i}, X_{n,j})$. The sole difference is the absence of “intra-segment” interactions: the previous display does not contain terms in $h(X_{n,i}, X_{n,j})$ with i and j in the same segment. It is the existence of these terms in our case which complicates the analysis.
- Suppose that h is degenerate, that is, $\mathbb{E}h(X, y) = \mathbb{E}h(x, Y) = 0$ for $x, y \in \mathcal{X}$. Then the variance term in Eq. (7) is zero, and Proposition 3.1 remains true in the following sense: $\sqrt{n}(U_n - \theta)$ converges towards the constant 0, which is a degenerate Gaussian distribution $\mathcal{N}(0, 0)$. In this case, we believe that the convergence will be faster, but not toward a Gaussian distribution. We refer to [Lee \(1990, Section 3.2.2\)](#) for results in this direction.
- It is possible to prove a version of Prop. 3.1 for the multiple change-point setting. The proof follows the lines of Section B.1, with an additional technical difficulty due to the numerous inter-segment interactions—we only deal with one in the present work.

3.3. Asymptotic normality of the squared sample median

We now turn to the statement of our main result. In the previous section, we only obtained the convergence of the empirical distribution function. It is well-known that such a result implies the convergence of the empirical quantiles towards the theoretical quantiles of the target distribution, if the convergence of the empirical distribution function is “strong enough” ([van der Vaart, 1998, Chapter 21](#)), *i.e.*, if the convergence is uniform or a CLT—as it is the case in our setting.

Proposition 3.2 (Asymptotic normality of H_n). *Suppose that Assumption 2.1 holds, and define T as in Section 3.1. Define $m = \text{med}(T) = F^{-1}(1/2)$ the theoretical median of the target distribution, and suppose that F has a non-zero derivative at m . Define $\sigma = \sigma(h, P, Q)$ as in Eq. (8), where $h(x, y) = \mathbb{1}_{\|x-y\|^2 \leq m}$. Then, for H_n defined in Eq. (4), we have*

$$\sqrt{n}(H_n - m) \xrightarrow{\mathcal{L}} \mathcal{N}\left(0, \frac{\sigma^2}{F'(m)^2}\right). \quad (9)$$

Before providing a sketch of the proof of Prop. 3.2, we make a few remarks.

- Empirical U -quantiles are known to satisfy asymptotic normality in the i.i.d. case (Serfling, 1980). Although a lot of work has been done to relax the independence assumption, however, there is no result regarding the non-identically distributed case. In the two-sample setting, some results exist, both in the independent case (Lehmann, 1951) and with some dependence structure (Dehling and Fried, 2012). Nevertheless, as noted before, in our setting it is necessary to consider the intra-segment interactions, which complicates the analysis.
- As it can be seen for instance in Figure 1, observations lying between $\max(\mathbb{E}T_{XX}, \mathbb{E}T_{YY})$ and $\mathbb{E}T_{XY}$ may be very scarce. In such a case, it is possible for $F'(m)$ to be small, leading to a large variance term in Eq. (9). Note however that $F'(m) \neq 0$ for arbitrary continuous distributions.

Proof sketch. Set $t \in \mathbb{R}$. The general idea of the proof is to rewrite statements about the event $\{\sqrt{n}(H_n - m) \leq t\}$ as statements about a sum of U -statistics. We then control these U -statistics with Prop. 3.1 for conveniently chosen h , and conclude with Slutsky’s Lemma. Throughout this proof, we only suppose that $p \in (0, 1)$ to emphasize that Prop. 3.2 can be extended to *any* quantile, not only the median. The complete proof can be found in Appendix B.2. \square

4. Empirical investigation

We empirically compare the median heuristic and an approach based on the maximization of test power (Gretton et al., 2012b) in the setting of two-sample test. The purpose is to gain insights about when the median heuristic performs reasonably and when it could work poorly.

We consider the following setting. We focus on the use of the Gaussian kernel $k_\nu(x, y) := \exp\left(\frac{-\|x-y\|^2}{2\nu^2}\right)$, and consider the selection of bandwidth $\nu > 0$. We consider the following two scenarios for the distributions P and Q . (i) **Mean:** A *change in the mean* of a univariate Gaussian distribution: $P \sim \mathcal{N}(0, 1)$ and $Q \sim \mathcal{N}(\mu, 1)$ with $\mu > 0$; (ii) **Var:** A *change in the variance* of a univariate Gaussian distribution: $P \sim \mathcal{N}(0, 1)$ and $Q \sim \mathcal{N}(0, \sigma^2)$ with $\sigma > 1$. In both these scenarios, we choose $\alpha = 1/2$ in order to be able to compute the linear-time estimate MMD_ℓ .

Since the kernel k_ν and distributions P and Q are Gaussian, it is possible to compute the cumulative distribution function F_T of T exactly, and thus to have access to the “theoretical” value of the median heuristic given by solving in t the equation $F(t) = 1/2$. This is done by solving this equation $F(t) = 1/2$ in each scenario (see Section D), which gives rise to bandwidth choices $\nu_{\text{med}}^{\text{Mean}}$ and $\nu_{\text{med}}^{\text{Var}}$ —which are functions of μ and σ , respectively. In this way, we suppress the noise coming from choosing empirically the median of the pairwise distance and obtain some fast, precise insights.

Comparisons are made with the bandwidth selected by the maximization of test power (Gretton et al., 2012b), where the criterion is defined as the (squared) MMD divided by a certain variance. We show in Section E of the

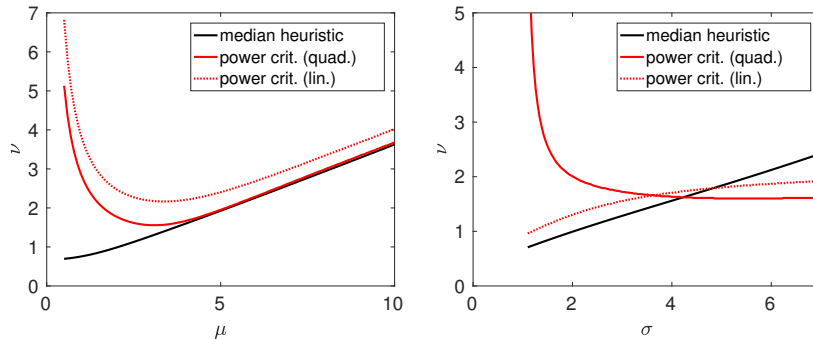


FIG 2. Gaussian kernel bandwidth selected by different means. The left panel corresponds to the **Mean** scenario, where we vary the mean μ of Q . The right panel corresponds to the **Var** scenario, where we vary the variance σ^2 of Q . For each plot, the vertical axis depicts the values of selected bandwidths. The black curves are obtained by computing the theoretical median, the red curves by maximizing R_u (the power criterion with quadratic-time MMD), and the red dotted curves by maximizing R_ℓ (the power criterion with linear-time MMD).

Appendix that, in both scenarios, it is possible to derive the power ratio criterion in closed-form, both for MMD_u and MMD_ℓ . We denote these ratios by R_u and R_ℓ respectively. Let ν_u^{Mean} and ν_u^{Var} be the bandwidths selected the maximization by R_u in scenarios **Mean** and **Var** respectively, and ν_ℓ^{Mean} and ν_ℓ^{Var} be those by the maximization of R_ℓ . We now dispose of two different test statistics, the quadratic-time statistic MMD_u and the linear-time statistic MMD_ℓ . To each of these statistics corresponds two different way to choose the bandwidth: on one side the value given by the theoretical median, on the other side the value that maximizes the corresponding power ratio criterion, giving rise to four different tests.

Comparison of bandwidths selected by different approaches. Figure 2 shows bandwidths selected by different approaches in the two scenarios. In scenario **Mean**, for large values of μ , the values of $\nu_{\text{med}}^{\text{Mean}}$ chosen by the median heuristic are almost identical to the values of ν_u^{Mean} selected by the power ratio maximization with quadratic-time MMD statistics, and are parallel with ν_ℓ^{Mean} of linear-time MMD. We do not have a theoretical explanation for this phenomenon, which looks like a remarkable coincidence in view of the way $\nu_{\text{med}}^{\text{Mean}}$ and ν_u^{Mean} are computed—see Section D and E in the Appendix. On the other hand, in the **Var** scenario, this behavior cannot be observed, supporting the common sense that the median heuristic may not be always the best choice, selecting a bandwidth too large for the scale of the changes.

Comparison of approximate Bahadur slopes. We turn to a measure of performance of a testing procedure, the approximate Bahadur slope (Bahadur, 1960). Intuitively, this slope is the rate of convergence of p -values to 0 as n increases. Therefore, the higher the slope, the faster the p -value vanishes, and thus the lower the sample size required to reject H_0 under θ (the higher the better).

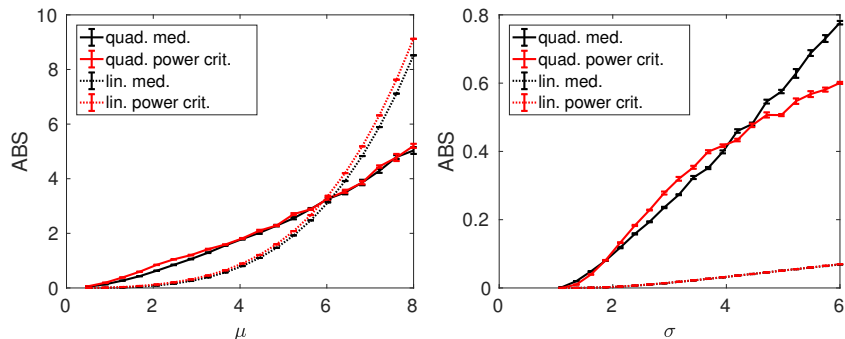


FIG 3. In this figure we plot the ABS of the different tests in the **Mean** (left panel) and **Var** (right panel) scenarios. The black lines correspond to the MMD_u test with the median heuristic, the red lines to the MMD_u test with bandwidth chosen by maximizing the power criterion. Dotted lines correspond to the MMD_ℓ test. We report error bars corresponding to the standard deviation over 5 experiments—the randomness comes from the estimation of λ_1 from 10^3 sample points.

This was first used in the context of kernel two-sample test by [Jitkrittum et al. \(2017b\)](#). We choose the ABS as a measure of quality of the two-sample testing procedures associated to the different bandwidths choices that we exposed.

Let us give a formal definition of ABS for a general testing procedure between $H_0 : \theta \in \Theta_0$ and $H_1 : \theta \in \Theta \setminus \Theta_0$. Let T_n be the associated test statistic, and denote by F the asymptotic distribution of T_n under the null. Assume that F is continuous and common to all $\theta_0 \in \Theta_0$. Suppose that there exists a continuous strictly increasing function $R : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\lim_{n \rightarrow +\infty} R(n) = +\infty$ and $-2 \log(1 - F(T_n))/R(n) \xrightarrow{\mathbb{P}} c(\theta)$ where T_n is drawn under θ , for some function c that is positive on $\Theta \setminus \Theta_0$ and identically zero on Θ_0 . Then the function c is known as the *approximate Bahadur slope* (ABS) of T_n . As it turns out, in our setting, it is possible to compute the ABS associated to MMD_u and MMD_ℓ :

Proposition 4.1 (Approximate Bahadur Slope computations). *The ABS of $nMMD_u^2$ is $\frac{MMD^2}{4\lambda_1}$, where λ_1 is the largest eigenvalue of the centered Gaussian kernel integral operator. The ABS of $\sqrt{n}MMD_\ell^2$ is $\frac{MMD^4}{8\sigma_\ell^2}$, where σ_ℓ^2 is the asymptotic variance of MMD_ℓ^2 (Eq. (16) in the Appendix).*

We refer to Section F of the Appendix for the proof of Prop. 4.1, which is an adaptation of the proof of Th. 5 in [Jitkrittum et al. \(2017b\)](#). In the computation of the ABS for MMD_u^2 , we estimate λ_1 empirically, following [Gretton et al. \(2009, Eq. \(5\)\)](#).

Figure 3 describes the approximate Bahadur slopes for the considered approaches. It shows that, for scenario **Mean** (left panel), the median heuristic provides almost identical slopes to that of the corresponding power-maximization approaches, as suggested by the results in Figure 2. In scenario **Var** (right panel), it is interesting to see that the slope by the median heuristic is almost

identical to that of the power-maximization approach with linear-time MMD statistic. A similar behavior can be observed for the quadratic time statistic, for relatively small values of the variance.

5. Conclusion and future directions

In this article, we partly explained the behavior of the median heuristic for large sample size, providing a convergence analysis in the setting of kernel two-sample test. We believe that our work serves as a step towards a deeper understanding of the median heuristic. In particular, it should open the door to more rigorous statements regarding the optimality of bandwidth choice in kernel two-sample test.

As future work, we first aim to explain the interesting behaviors observed in Figures 2 and 3. Another direction for research is the case where α depends on n . For instance, one could be interested in a situation where we observe asymptotically much less samples of X than samples of Y . Finally, it would be extremely interesting to perform a similar analysis in machine learning settings other than kernel two-sample test. In particular, we suspect that for many tasks where the distributions at hand are not “well-separated,” the median heuristic leads to higher bandwidth than needed and is therefore a poor choice.

References

- Abramowitz, M. and Stegun, I. (1964). *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. Courier Corporation.
- Aizerman, M., Braverman, E., and Rozonoer, L. (1964). Theoretical foundations of the potential function method in pattern recognition. *Automation and Remote Control*, **25**:917–936.
- Arlot, S., Celisse, A., and Harchaoui, Z. (2012). A kernel multiple change-point algorithm via model selection. ArXiv preprint, available at <https://arxiv.org/abs/1202.3878v2>.
- Bach, F. and Jordan, M. I. (2002). Kernel independent component analysis. *Journal of Machine Learning Research*, **3**(7):1–48.
- Bahadur, R. (1960). Stochastic comparison of tests. *The Annals of Mathematical Statistics*, **31**(2):276–295.
- Billingsley, P. (2012). *Probability and Measure*, Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken (NJ), third edition.
- Boser, B. E., Guyon, I. M., and Vapnik, V. N. (1992). A training algorithm for optimal margin classifiers. In *Proceedings of the 5th annual workshop on Computational Learning Theory*, pp. 144–152. ACM Press.
- Dehling, H. and Fried, R. (2012). Asymptotic distribution of two-sample empirical U-quantiles with applications to robust tests for shifts in location. *Journal of Multivariate Analysis*, **105**(1):124–140.
- Flaxman, S., Sejdinovic, D., Cunningham, J. P., and Filippi, S. (2016). Bayesian learning of kernel embeddings. In *Proceedings of the 32th Conference on Uncertainty in Artificial Intelligence*, pp. 182–191.

- Garreau, D. and Arlot, S. (2016). Consistent change-point detection with kernels. ArXiv preprint, available at <http://arxiv.org/abs/1612.04740v3>.
- Genton, M. G. (2001). Classes of kernels for machine learning: a statistics perspective. *Journal of Machine Learning Research*, **2**(12):299–312.
- Gretton, A., Borgwardt, K. M., Rasch, M., Schölkopf, B., and Smola, A. J. (2007). A kernel method for the two-sample-problem. In *Advances in Neural Information Processing Systems 19*, pp. 513–520. MIT Press.
- Gretton, A., Fukumizu, K., Harchaoui, Z., and Sriperumbudur, B. (2009). A fast, consistent kernel two-sample test. In *Advances in Neural Information Processing Systems*, pp. 673–681.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. (2012a). A kernel two-sample test. *Journal of Machine Learning Research*, **13**(1):723–773.
- Gretton, A., Sejdinovic, D., Strathmann, H., Balakrishnan, S., Pontil, M., Fukumizu, K., and Sriperumbudur, B. K. (2012b). Optimal kernel choice for large-scale two-sample tests. In *Advances in Neural Information Processing Systems 25*, pp. 1205–1213. Curran Associates, Inc.
- Harchaoui, Z. and Cappé, O. (2007). Retrospective multiple change-point estimation with kernels. In *Proceedings of the 14th IEEE Workshop on Statistical Signal Processing*, pp. 768–772.
- Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, **19**(3):293–325.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, **12**(1):55–67.
- Jammalamadaka, S. R. and Janson, S. (1986). Limit theorems for a triangular scheme of U-statistics with applications to inter-point distances. *The Annals of Probability*, pp. 1347–1358.
- Jitkrittum, W., Szabó, Z., and Gretton, A. (2017a). An adaptive test of independence with analytic kernel embeddings. In *Proceedings of the 34th International Conference on Machine Learning*.
- Jitkrittum, W., Xu, W., Szabó, Z., Fukumizu, K., and Gretton, A. (2017b). A linear-time kernel goodness-of-fit test. In *Advances in Neural Information Processing Systems*, pp. 261–270.
- Karatzoglou, A., Smola, A., Hornik, K., and Zeileis, A. (2004). kernlab – an S4 package for kernel methods in R. *Journal of Statistical Software*, **11**(9):1–20.
- Korolyuk, V. S. and Borovskich, Y. V. (2013). *Theory of U-statistics*, Mathematics and its Applications, **273**. Springer Science & Business Media.
- Lee, J. (1990). *U-statistics: Theory and Practice*, Statistics: Textbooks and Monographs, **110**. Marcel Dekker, Inc.
- Lehmann, E. L. (1951). Consistency and unbiasedness of certain nonparametric tests. *The Annals of Mathematical Statistics*, **22**(2):165–179.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Poczos, B. (2017). Mmd gan: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems 30*, pp. 2203–2213.
- Long, M., Cao, Y., Wang, J., and Jordan, M. (2015). Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd Inter-*

- national Conference on Machine Learning*, Proceedings of Machine Learning Research, **37**, pp. 97–105. PMLR.
- Mallows, C. L. (1991). Another comment on O’Cinneide. *The American Statistician*, **45**(3):257.
- Marcum, J. (1950). *Table of Q Functions*. Rand Corporation.
- Muandet, K., Sriperumbudur, B., Fukumizu, K., Gretton, A., and Schölkopf, B. (2016). Kernel mean shrinkage estimators. *Journal of Machine Learning Research*, **17**(48):1–41.
- Muandet, K., Fukumizu, K., Sriperumbudur, B. K., and Schölkopf, B. (2017). Kernel mean embedding of distributions : A review and beyond. *Foundations and Trends in Machine Learning*, **10**(1–2):1–141.
- Reddi, S. J., Ramdas, A., Poczos, B., Singh, A., and Wasserman, L. (2015). On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pp. 3571–3577.
- Schölkopf, B. and Smola, A. J. (2002). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, Adaptive Computation and Machine Learning. MIT Press, Cambridge (MA).
- Schölkopf, B., Smola, A., and Müller, K.-R. (1997). Kernel principal component analysis. In *Proceedings of the 7th International Conference on Artificial Neural Networks*, pp. 583–588.
- Schölkopf, B., Smola, A., and Müller, K. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, **10**(5):1299–1319.
- Serfling, R. J. (1980). *Approximation theorems of mathematical statistics*, Wiley Series in Probability and Statistics, **162**. John Wiley & Sons.
- Sriperumbudur, B., Fukumizu, K., Gretton, A., Lanckriet, G. R. G., and Schölkopf, B. (2009). Kernel choice and classifiability for RKHS embeddings of probability distributions. In *Advances in Neural Information Processing Systems 22*, pp. 1750–1758.
- Sutherland, D. J., Tung, H.-Y., Strathmann, H., De, S., Ramdas, A., Smola, A., and Gretton, A. (2017). Generative models and model criticism via optimized maximum mean discrepancy. In *Proceedings of the 5th International Conference on Learning Representations*.
- Takeuchi, I., Le, Q. V., Sears, T. D., and Smola, A. J. (2006). Nonparametric quantile estimation. *Journal of Machine Learning Research*, **7**(7):1231–1264.
- van der Vaart, A. (1998). *Asymptotic Statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, **3**. Cambridge University Press, Cambridge.
- Vert, J.-P., Tsuda, K., and Schölkopf, B. (2004). A primer on kernel methods. In *Kernel Methods in Computational Biology*, pp. 35–70. MIT Press.
- Zhang, Q., Filippi, S., Gretton, A., and Sejdinovic, D. (2017). Large-scale kernel methods for independence testing. *Statistics and Computing*, pp. 1–18.

Appendix

In this Appendix we collect the proofs of all the results stated in the main paper and some further explanations regarding the computation of the bandwidths used in Section 4. It is organized as follows: Section A contains the proof of Lemma 3.1, whereas the proofs of all the other main results can be found in Section B, and the technical lemmas are collected in Section C. The derivation of the cumulative distribution of T in closed form is exposed in Section D. Section E contains further explanation about and computation of the power criteria introduced in Section 4. Finally, Section F is dedicated to the proof of Prop. 4.1.

A. Proof of Lemma 3.1

Set $r := \|\mu_X - \mu_Y\|/5$. According to Chebyshev's inequality,

$$\begin{cases} \mathbb{P}(\|X - \mu_X\| > r) & \leq \frac{25 \operatorname{tr}(\Sigma_X)}{\|\mu_X - \mu_Y\|^2} \\ \mathbb{P}(\|Y - \mu_Y\| > r) & \leq \frac{25 \operatorname{tr}(\Sigma_Y)}{\|\mu_X - \mu_Y\|^2}, \end{cases}$$

and this is also true for any independent copy of X and Y . By the union bound, there exists an event Ω_{XX} with probability greater than $1 - 50 \operatorname{tr}(\Sigma_X) / \|\mu_X - \mu_Y\|^2$ such that $\|X - \mu_X\| \leq r$ and $\|X' - \mu_X\| \leq r$. On this event,

$$\begin{aligned} T_{XX}^{\frac{1}{2}} &= \|X - X'\| && \text{(definition of } T_{XX}) \\ &\leq \|X - \mu_X\| + \|\mu_X - X'\| && \text{(triangle inequality)} \\ T_{XX}^{\frac{1}{2}} &< 2r. \end{aligned}$$

The same reasoning shows that there exists an event Ω_{YY} with probability at least $1 - 50 \operatorname{tr}(\Sigma_Y) / \|\mu_X - \mu_Y\|^2$ such that $T_{YY}^{\frac{1}{2}} < 2r$. There also exists an event Ω_{XY} with probability greater than $1 - 25(\operatorname{tr}(\Sigma_X) + \operatorname{tr}(\Sigma_Y)) / \|\mu_X - \mu_Y\|^2$ on which $\|X - \mu_X\| \leq r$ and $\|Y - \mu_Y\| \leq r$. On this event,

$$\begin{aligned} T_{XY}^{\frac{1}{2}} &= \|X - Y\| && \text{(definition of } T_{XY}) \\ &\geq \|\mu_X - \mu_Y\| - \|X - \mu_X\| - \|Y - \mu_Y\| && \text{(triangle inequality)} \\ &\geq \|\mu_X - \mu_Y\| - 2r && \text{(definition of } \Omega_{XY}) \\ T_{XY}^{\frac{1}{2}} &\geq 3r. && \text{(definition of } r) \end{aligned}$$

Define $\Omega := \Omega_{XX} \cap \Omega_{YY} \cap \Omega_{XY}$. By the union bound, Ω has probability greater than $1 - 75 (\text{tr}(\Sigma_X) + \text{tr}(\Sigma_Y)) / \|\mu_X - \mu_Y\|^2$. Moreover, on this event, $\max\left(T_{XX}^{\frac{1}{2}}, T_{YY}^{\frac{1}{2}}\right) + r < T_{XY}^{\frac{1}{2}}$, from which we deduce the result. \square

B. Proofs of the main results

B.1. Proof of Prop. 3.1

The proof consists of three steps.

Step 1: Decomposition of U_n . We begin by decomposing U_n . To this extent, define

$$A_n = \binom{\alpha n}{2}^{-1} \sum_{1 \leq i < j \leq \alpha n} h(X_{n,i}, X_{n,j}), \quad B_n = \binom{(1-\alpha)n}{2}^{-1} \sum_{\alpha n < i < j \leq n} h(X_{n,i}, X_{n,j}),$$

and

$$C_n = \frac{1}{\alpha n(1-\alpha)n} \sum_{\substack{1 \leq i \leq \alpha n \\ \alpha n < j \leq n}} h(X_{n,i}, X_{n,j}).$$

Note that in C_n there is no term such that $h(X_{n,i}, X_{n,j})$ with $1 \leq j \leq \alpha n$ and $\alpha n < i \leq n$, since the sum in U_n is prescribed to $i < j$. Simple algebra shows that

$$U_n = \alpha^2 A_n + (1-\alpha)^2 B_n + 2\alpha(1-\alpha)C_n + s_n. \quad (10)$$

where $s_n = \frac{\alpha(1-\alpha)}{n-1} A_n - \frac{\alpha(1-\alpha)}{n-1} B_n - \frac{2\alpha(1-\alpha)}{n-1} C_n$. According to Lemma C.1 (see Appendix C), the variance of each of A_n , B_n and C_n is $O(1/n)$, and hence s_n converges to 0 in probability at speed n^{-2} . Therefore we will focus on the terms of Eq. (10) other than s_n .

Step 2: Decompositions of A_n , B_n and C_n . The next step is to obtain the H -decomposition (Lee, 1990, Section 1.6) of A_n and B_n . Let us detail this process for A_n . We set

$$\begin{aligned} \theta_A &= \mathbb{E}[h(X, X')], \\ h_A(x) &= \mathbb{E}[h(x, X')] - \theta_A, \\ g_A(x, y) &= h(x, y) - h_A(x) - h_A(y) - \theta_A \\ L_A &= \frac{2}{\alpha n} \sum_{1 \leq i \leq \alpha n} h_A(X_{n,i}), \\ R_A &= \binom{\alpha n}{2}^{-1} \sum_{1 \leq i < j \leq \alpha n} g_A(X_{n,i}, X_{n,j}). \end{aligned}$$

Then it is possible to write A_n as (Lee, 1990, Th. 1)

$$A_n = \theta_A + L_A + R_A.$$

A totally analogous statement holds for B_n : Define

$$\begin{aligned}\theta_B &= \mathbb{E}[h(Y, Y')], \\ h_B(y) &= \mathbb{E}[h(y, Y')] - \theta_B, \\ g_B(x, y) &= h(x, y) - h_B(x) - h_B(y) - \theta_B, \\ L_B &= \frac{2}{(1-\alpha)n} \sum_{\alpha n < i \leq n} h_B(X_{n,i}), \\ R_B &= \left(\frac{(1-\alpha)n}{2}\right)^{-1} \sum_{\alpha n < i < j \leq n} g_B(X_{n,i}, X_{n,j}).\end{aligned}$$

Then B_n can be written as

$$B_n = \theta_B + L_B + R_B.$$

We decompose C_n in the same fashion, the only difference being the appearance of a second term in the linear part. Namely, set

$$\begin{aligned}\theta_C &= \mathbb{E}[h(X, Y)], \\ h_{C,1}(x) &= \mathbb{E}[h(x, Y)], \\ h_{C,2}(y) &= \mathbb{E}[h(X, y)], \\ g_C(x, y) &= h(x, y) - h_{C,1}(x) - h_{C,2}(y) - \theta_C, \\ L_C &= \frac{1}{\alpha n} \sum_{i=1}^{\alpha n} h_{C,1}(X_{n,i}) + \frac{1}{(1-\alpha)n} \sum_{i=\alpha n+1}^n h_{C,2}(X_{n,i}), \\ R_C &= \frac{1}{\alpha n(1-\alpha)n} \sum_{\substack{1 \leq i \leq \alpha n \\ \alpha n < j \leq n}} g_C(X_{n,i}, X_{n,j}).\end{aligned}$$

Then we have

$$C_n = \theta_C + L_C + R_C.$$

Since θ defined in Eq. (6) can be written as

$$\theta = \alpha^2 \theta_A + (1-\alpha)^2 \theta_B + 2\alpha(1-\alpha) \theta_C,$$

we have by Eq. (10) and the above expressions for A_n , B_n and C_n

$$\begin{aligned}& \sqrt{n}(U_n - \theta) \\ &= \sqrt{n} [\alpha^2(L_A + R_A) + (1-\alpha)^2(L_B + R_B) + 2\alpha(1-\alpha)(L_C + R_C) + s_n] \\ &= \sqrt{n} [\alpha^2 L_A + (1-\alpha)^2 L_B + 2\alpha(1-\alpha) L_C] + \sqrt{n}(r_n + s_n),\end{aligned}\tag{11}$$

where $r_n = \alpha^2 R_A + (1-\alpha)^2 R_B + 2\alpha(1-\alpha) R_C$. According to Lemma C.2 in Appendix C, the variance of each of R_A , R_B and R_C is of order n^{-2} . As mentioned earlier s_n is also of order n^{-2} . Therefore we have $\sqrt{n}(r_n + s_n) \xrightarrow{\mathbb{P}} 0$ as $n \rightarrow \infty$. This shows that we only need to focus on the first term of Eq. (11).

Step 3: Regrouping of the terms in Eq. (11). We now regroup the first term in Eq. (11) that belong to the same segment (*i.e.*, $1 \leq i \leq \alpha n$ or $\alpha n < i \leq n$). That is, a simple calculation yields that

$$\begin{aligned} & \sqrt{n} [\alpha^2 L_A + (1 - \alpha)^2 L_B + 2\alpha(1 - \alpha)L_C] \\ &= \sqrt{n} \left[\frac{2}{n} \sum_{1 \leq i \leq \alpha n} (\alpha h_A(X_{n,i}) + (1 - \alpha)h_{C,1}(X_{n,i})) \right. \\ & \quad \left. + \frac{2}{n} \sum_{\alpha n < i \leq n} (\alpha h_{C,2}(X_{n,i}) + (1 - \alpha)h_B(X_{n,i})) \right]. \end{aligned} \quad (12)$$

For any $1 \leq i \leq \alpha n$, define

$$Z_{n,i}^{(1)} = \alpha h_A(X_{n,i}) + (1 - \alpha)h_{C,1}(X_{n,i}).$$

Since $h \in L^2(P) \cap L^2(Q) \times L^2(P) \cap L^2(Q)$ and $X_{n,i} \sim X$,

$$\sigma_1^2 := \text{Var} \left(Z_{n,i}^{(1)} \right) = \text{Var} (\alpha h_A(X) + (1 - \alpha)h_{C,1}(X))$$

is finite and does not depend on i . Note also the $(Z_{n,i}^{(1)})_{i=1}^{\alpha n}$ are independent. Therefore the Lindeberg's condition is satisfied, and hence we can apply the central limit theorem for triangular arrays of independent random variables [Billingsley \(2012, Th. 27.2\)](#) to obtain

$$\frac{1}{\sqrt{\alpha n}} \sum_{i=1}^{\alpha n} Z_{n,i}^{(1)} \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma_1^2) \quad (n \rightarrow \infty).$$

In a similar fashion, for any $\alpha n < i \leq n$, let

$$Z_{n,i}^{(2)} := \alpha h_{C,2}(X_{n,i}) + (1 - \alpha)h_B(X_{n,i}),$$

and

$$\sigma_2^2 := \text{Var} \left(Z_{n,i}^{(2)} \right) = \text{Var} (\alpha h_{C,2}(Y) + (1 - \alpha)h_B(Y)).$$

Then we have

$$\frac{1}{\sqrt{(1 - \alpha)n}} \sum_{i=\alpha n+1}^n Z_{n,i}^{(2)} \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma_2^2) \quad (n \rightarrow \infty).$$

The two previous sums are independent, and thus by Lévy's theorem and Eq. (12) we have

$$\begin{aligned} & \sqrt{n} [\alpha^2 L_A + (1 - \alpha)^2 L_B + 2\alpha(1 - \alpha)L_C] \\ &= \sqrt{n} \left[\frac{2}{n} \sum_{1 \leq i \leq \alpha n} Z_{n,i}^{(1)} + \frac{2}{n} \sum_{\alpha n < i \leq n} Z_{n,i}^{(2)} \right] \\ &= \frac{2\sqrt{\alpha}}{\sqrt{\alpha n}} \sum_{1 \leq i \leq \alpha n} Z_{n,i}^{(1)} + \frac{2\sqrt{1 - \alpha}}{\sqrt{(1 - \alpha)n}} \sum_{\alpha n < i \leq n} Z_{n,i}^{(2)} \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma^2) \quad (n \rightarrow \infty), \end{aligned}$$

where $\sigma^2 = \sigma_1^2 + \sigma_2^2$, which can be shown to be equal to Eq. (8) by a straightforward computation. Since the remainder term in Eq. (11) converges in probability to 0, we can conclude *via* Slutsky's Lemma (van der Vaart, 1998, Th. 2.7). \square

B.2. Proof of Prop. 3.2

Set $t \in \mathbb{R}$. The general idea of the proof is to rewrite statements about the event $\{\sqrt{n}(H_n - m) \leq t\}$ as statements about a sum of U -statistics. We will then control these U -statistics with Prop. 3.1 for conveniently chosen h , and conclude with Slutsky's Lemma. Throughout this proof, we only suppose that $p \in (0, 1)$ to emphasize that Prop. 3.2 can be extended to *any* quantile, not only the median. Therefore in this proof we define m by

$$m := \hat{F}(p).$$

Note that the median case corresponds to $p = 1/2$.

We use the property of the generalized inverse to obtain

$$\{\sqrt{n}(H_n - m) \leq t\} = \left\{ p \leq \hat{F}_n \left(m + \frac{t}{\sqrt{n}} \right) \right\},$$

and rewrite this event as

$$\left\{ \sqrt{n} \left(\hat{F}_n \left(m + \frac{t}{\sqrt{n}} \right) - F \left(m + \frac{t}{\sqrt{n}} \right) \right) \geq \sqrt{n} \left(p - F \left(m + \frac{t}{\sqrt{n}} \right) \right) \right\}. \quad (13)$$

Since F is differentiable in m , the right-hand side of Eq. (13) converges towards $-tF'(m)$ by Taylor expansion. From Prop. 3.1, it is also true that

$$\sqrt{n} \left(\hat{F}_n(m) - F(m) \right) \xrightarrow{\mathcal{L}} \mathcal{N}(0, \sigma^2).$$

Therefore, if we manage to prove that

$$\sqrt{n} \left[\left(\hat{F}_n \left(m + \frac{t}{\sqrt{n}} \right) - \hat{F}_n(m) \right) - \left(F \left(m + \frac{t}{\sqrt{n}} \right) - F(m) \right) \right] \xrightarrow{\mathbb{P}} 0, \quad (14)$$

Then Eq. (9) will follow by Slutsky's Lemma.

Define $h(x, y) = \mathbf{1}_{m < \|x-y\|^2 \leq m + \frac{t}{\sqrt{n}}}$. Then, with the notations used in the proof of Prop. 3.1, Eq. (14) reads

$$\sqrt{n}(U_n - \theta) \xrightarrow{\mathbb{P}} 0.$$

We dispose of the remainder terms as in the proof of Prop. 3.1, thus we are left to show that

$$\sqrt{n} \left[\alpha^2(A_n - \theta_A) + (1 - \alpha)^2(B_n - \theta_B) + 2\alpha(1 - \alpha)(C_n - \theta_C) \right] \xrightarrow{\mathbb{P}} 0. \quad (15)$$

Let us focus on the first term of the previous display, which can be written

$$\alpha^2 \sqrt{n} (A_n - \theta_A) = \alpha^2 \sqrt{n} (L_A + R_A) .$$

Once again, we use Lemma C.2 to get rid of R_A . The linear term is slightly more tedious to analyze. Recall that $\mathbb{E}[h_A(X)] = 0$. Thanks to Jensen's inequality,

$$\begin{aligned} \text{Var}(h_A(X)) &= \mathbb{E}_X \left[\left(\mathbb{E}_{X'} \left[\mathbb{1}_{m < \|X - X'\|^2 \leq m + \frac{t}{\sqrt{n}}} \right]^2 \right) \right] \\ &\leq \mathbb{P} \left(m < \|X - X'\|^2 \leq m + \frac{t}{\sqrt{n}} \right) . \end{aligned}$$

We recognize

$$F \left(m + \frac{t}{\sqrt{n}} \right) - F(m) ,$$

which goes to 0 when $n \rightarrow \infty$, since we assumed F to have a derivative in m . Furthermore, by independence of the $X_{n,i}$,

$$\text{Var}(\sqrt{n}L_A) = 4 \text{Var}(h_A(X)) \xrightarrow[n \rightarrow \infty]{} 0 .$$

A similar reasoning applies to the other terms in Eq. (15), and the proof is concluded.

C. Technical lemmas

In this section, we state and prove the technical results that are needed in the proofs of the main results. Recall that we denote by $|A|$ the cardinality of any finite set A .

Lemma C.1 (Controlling the variance of A_n , B_n , and C_n). *Let A_n , B_n , and C_n be defined as in the proof of Prop. 3.1. Then $\text{Var}(A_n)$, $\text{Var}(B_n)$, and $\text{Var}(C_n)$ are $O(n^{-1})$.*

The proof is standard in U -statistics (Lee, 1990), we provide a version thereof for completeness' sake.

Proof. We set $m := \alpha n$ in this proof. Recall that

$$A_n = \binom{m}{2}^{-1} \sum_{1 \leq i < j \leq m} h(X_{n,i}, X_{n,j}) ,$$

Define $h_{i,j} = h(X_{n,i}, X_{n,j})$, thus $\mathbb{E}h_{i,j} = \theta_A$ and $\mathbb{E}A_n = \theta_A$. Let us turn to the computation of $\mathbb{E}[A_n^2]$, that is,

$$\mathbb{E}[A_n^2] = \binom{m}{2}^{-2} \sum_{a < b} \sum_{c < d} \mathbb{E}[h_{a,b} h_{c,d}] ,$$

where a, b, c and d range from 1 to m . In this last display, there are three possibilities for the indices in the summation that we detail below.

- The indices a, b, c and d are all distinct. There are $\binom{m}{4} \binom{4}{2} = \frac{m^4}{4} + O(m^3)$ ways to choose such indices, that is, $\binom{m}{4}$ ways to choose the location of the 4 indices among the m possible locations, then $\binom{4}{2}$ choices for, say, $a < b$, and only one possibility left for $c < d$.
- One of the indices is common, that is, $|\{a, b\} \cap \{c, d\}| = 1$. There are $6 \binom{m}{3} = O(m^3)$ ways to do so.
- Both indices are equal, that is, $a = c$ and $b = d$. There are $\binom{m}{2} = O(m^2)$ ways to do so.

Note that when a, b, c and d are all distinct, $\mathbb{E}[h_{a,b}h_{c,d}] = \mathbb{E}[h_{a,b}] \mathbb{E}[h_{c,d}]$ by independence of the $X_{n,i}$ s. Thus

$$\mathbb{E}[A_n^2] = \binom{m}{2}^{-2} \sum_{\substack{a < b \\ c < d \\ \neq}} \mathbb{E}[h_{a,b}] \mathbb{E}[h_{c,d}] + O(m^{-1}),$$

where we sum on *distinct* indices. On the other hand,

$$\mathbb{E}[A_n]^2 = \binom{m}{2}^{-2} \sum_{\substack{a < b \\ c < d}} \mathbb{E}h_{a,b} \mathbb{E}h_{c,d}.$$

By the same combinatorial argument, the terms corresponding to intersecting sets of indices are at most $O(n^3)$ and we have

$$\mathbb{E}[A_n]^2 = \binom{m}{2}^{-2} \sum_{\substack{a < b \\ c < d \\ \neq}} \mathbb{E}h_{a,b} \mathbb{E}h_{c,d} + O(m^{-1}).$$

Since $m = \alpha n$, $O(m^{-1}) = O(n^{-1})$ and we can conclude for A_n :

$$\text{Var}(A_n) = \mathbb{E}[A_n^2] - \mathbb{E}[A_n]^2 = O(n^{-1}).$$

The same proof transfers readily for B_n and C_n . □

Lemma C.2 (Controlling the variance of R_A , R_B , and R_C). *Let R_A , R_B , and R_C be as in the proof of Prop. 3.2. Then $\text{Var}(R_A)$, $\text{Var}(R_B)$, and $\text{Var}(R_C)$ are of order $O(n^{-2})$.*

Proof. Recall that

$$R_A = \binom{m}{2}^{-1} \sum_{1 \leq i < j \leq m} g_{i,j},$$

with $g_{i,j} = g_A(X_{n,i}, X_{n,j})$. By definition of g_A , it holds that $\mathbb{E}g_{i,j} = 0$ for any $i \neq j$, thus $\mathbb{E}R_A = 0$. Hence to control the variance of R_A , we just need to compute $\mathbb{E}[R_A^2]$. As in the proof of Lemma C.1, we have

$$\mathbb{E}[R_A^2] = \binom{m}{2}^{-2} \sum_{a < b} \sum_{c < d} \mathbb{E}[g_{a,b}g_{c,d}].$$

Note that $\mathbb{E}[g_{a,b}g_{c,d}] = 0$ whenever a, b, c and d are all distinct. But a straightforward computation also shows that $\mathbb{E}[g_{a,b}g_{a,c}] = 0$ for any distinct a, b, c . Thus the previous display reduces to

$$\mathbb{E}[R_A^2] = \binom{m}{2}^{-2} \sum_{\substack{a < b \\ c < d \\ \star}} \mathbb{E}[g_{a,b}g_{c,d}],$$

where \star denotes that we sum on indices such that $|\{a, b\} \cap \{c, d\}| \geq 2$. As we have seen in the proof of Lemma C.1, there are only $O(m^2)$ such possibilities, and we can conclude. \square

D. Derivation of the cumulative distribution function of T

In this section, we derive the population cumulative distribution of the pairwise distances in closed-form in both scenarios considered in Section 4. We first introduce some additional notations. We let γ be the *lower* incomplete gamma function:

$$\gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt,$$

and Q_M the Marcum Q -function (Marcum, 1950), defined by

$$Q_M(a, b) = \int_b^{+\infty} x \left(\frac{x}{a}\right)^{M-1} \exp\left(-\frac{x^2 + a^2}{2}\right) I_{M-1}(ax) dx,$$

where I_{M-1} is the modified Bessel function of order $M-1$ (Abramowitz and Stegun, 1964).

Change in the mean. In the **Mean** scenario, T_{XX} and T_{YY} have the law of $2\chi_1^2$, whereas T_{XY} has the distribution of a non-central chi-squared random variable $\chi_1^2(\mu/2)$. It is well-known that the cumulative distribution function of a chi-squared random variable with one degree of freedom is given by $\frac{\gamma(\frac{1}{2}, \frac{t}{2})}{\Gamma(\frac{1}{2})}$, whereas that of a non-central chi-squared with non-centrality parameter λ is given by $1 - Q_{1/2}(\sqrt{\lambda}, \sqrt{t})$. Thus, according to the definition of T as a mixture, we have

$$F_T(t) = [\alpha^2 + (1 - \alpha)^2] \frac{\gamma\left(\frac{1}{2}, \frac{t}{4}\right)}{\Gamma\left(\frac{1}{2}\right)} + 2\alpha(1 - \alpha) \left(1 - Q_{\frac{1}{2}}\left(\frac{\mu}{\sqrt{2}}, \sqrt{\frac{t}{2}}\right)\right).$$

We plot F_T for a few values of μ in Figure 4. Finding the value of the population median of T is equivalent to solving $F_T(t) = 1/2$ for t . Since F_T is an increasing function, this can be achieved by dichotomy, or, even faster, *via* Newton method. A straightforward computation yields the derivative of F_T ,

$$f_T(t) = \frac{[\alpha^2 + (1 - \alpha)^2] e^{-t/4} + 2\alpha(1 - \alpha) e^{-(\mu^2+t)/4} \cosh\left(\frac{\mu\sqrt{t}}{2}\right)}{2\sqrt{\pi t}}.$$




FIG 4. In this figure, we plot the cumulative distribution function of T . Left panel: change in the mean scenario. Right panel: change in the variance. Note that F_T is becoming very “flat” when μ increases in the left panel, leading to numerical problems when solving $F_T(t) = 1/2$.

Change in the variance. In the **Var** scenario, T_{XX} has the law of $2\chi_1^2$, T_{YY} has the law of $2\sigma^2\chi_1^2$, and T_{XY} has the law of $(1+\sigma^2)\chi_1^2$. By the same reasoning, we obtain the cumulative distribution function of T ,

$$F_T(t) = \frac{\alpha^2 \gamma\left(\frac{1}{2}, \frac{t}{4}\right) + (1-\alpha)^2 \gamma\left(\frac{1}{2}, \frac{t}{4\sigma^2}\right) + 2\alpha(1-\alpha) \gamma\left(\frac{1}{2}, \frac{t}{2(\sigma^2+1)}\right)}{\Gamma\left(\frac{1}{2}\right)}.$$

We plot F_T for a few values of σ in Figure 4. A straightforward computation yields

$$f_T(t) = \frac{\alpha^2 \exp\left(\frac{-t}{4}\right)}{2\sqrt{\pi t}} + \frac{(1-\alpha)^2 \exp\left(\frac{-t}{4\sigma^2}\right)}{2\sigma\sqrt{\pi t}} + \frac{2\alpha(1-\alpha) \exp\left(\frac{-t}{2(\sigma^2+1)}\right)}{\sqrt{2(\sigma^2+1)\pi t}}.$$

E. Power criterion for kernel two-sample test

In this section, we set $z = (x, y)$ with $x \sim X$ and $y \sim Y$, where it should be clear that X and Y depend on the scenario that we are investigating. As before, x' and y' denote independent copies of x and y , thus z' is an independent copy of z . Note that the proportion of X and Y is fixed to $\alpha = 1/2$. Recall that we focus on $k_\nu(x, y) = \exp(-(x-y)^2/(2\nu^2))$, the Gaussian kernel with bandwidth ν , that we denote by k for concision. We also set

$$h(z, z') := k(x, x') + k(y, y') - k(x, y') - k(x', y).$$

According to [Gretton et al. \(2012b\)](#), the asymptotic probability of a type II error for the linear-time test statistic MMD_ℓ at level a is given by

$$\Phi\left(\Phi^{-1}(1-a) - \frac{\text{MMD}^2 \sqrt{n}}{\sigma_\ell \sqrt{2}}\right),$$

where MMD^2 is the maximum mean discrepancy between P and Q and σ_ℓ^2 is an asymptotic variance term given by (Gretton et al., 2012a, Section 6)

$$\sigma_\ell^2 := 2 \left(\mathbb{E}_{z, z'} [h(z, z')^2] - (\mathbb{E}_{z, z'} [h(z, z')])^2 \right). \quad (16)$$

Thus a natural way to select the kernel—in our case the bandwidth—is to maximize the *power ratio criterion*

$$R_\ell = R_\ell(P, Q, k) := \frac{\text{MMD}^2}{\sigma_\ell},$$

keeping in mind that both the maximum mean discrepancy and σ_ℓ are depending on P , Q , and k .

For the quadratic-time test statistic MMD_u , the picture is not so clear since the null distribution is more complicated and the type II error depends on a quantile of this distribution, namely

$$\frac{c_\alpha}{\sqrt{n}} - \frac{\text{MMD}^2 \sqrt{n}}{\sigma_u \sqrt{2}},$$

where σ_u^2 is an asymptotic variance term given by (Gretton et al., 2012a, Section 6)

$$\sigma_u^2 := 4 \left(\mathbb{E}_z \left[(\mathbb{E}_{z'} [h(z, z')])^2 \right] - (\mathbb{E}_{z, z'} [h(z, z')])^2 \right). \quad (17)$$

Nevertheless, we choose the bandwidth that maximizes the following power ratio criterion

$$R_u = R_u(P, Q, k) := \frac{\text{MMD}^2}{\sigma_u}.$$

We now proceed to compute the ratio R_ℓ and R_u in both the change of mean and the change of variance scenarios. Note that, though the MMD expressions are well-known, the closed-form expressions we obtain for the variance terms are new up to the best of our knowledge.

Change in the mean. In this case, the MMD is given by

$$\mathbb{E}_{z, z'} [h(z, z')] = \frac{2\nu}{\sqrt{\nu^2 + 2}} \left(1 - \exp \left(\frac{-\mu^2}{2(\nu^2 + 2)} \right) \right).$$

See, for instance, Reddi et al. (2015). We now turn to the variance terms. The second term in Eq. (16) and Eq. (17) is the MMD, given by the previous display. The remaining computation is, in both the linear-time and the quadratic-time case, a lengthy derivation that uses intensively the following lemma:

Lemma E.1 (Gaussian integral computation). *Let $a, b, c, d \in \mathbb{R}$ with $b, d > 0$. Then*

$$\frac{1}{\sqrt{2\pi}} \int \exp \left(\frac{-(x-a)^2}{b} + \frac{-(x-c)^2}{d} \right) dx = \sqrt{\frac{bd}{2(b+d)}} \exp \left(\frac{-(a-c)^2}{b+d} \right).$$

Proof. We write

$$\frac{-(x-a)^2}{b} + \frac{-(x-c)^2}{d} = \frac{1}{bd} \{ (b+d)x^2 + (2ad+2bc)x - (a^2d+bc^2) \},$$

and then use the well-known identity

$$\int_{-\infty}^{+\infty} e^{-\alpha x^2 + \beta x + \gamma} dx = \sqrt{\frac{\pi}{\alpha}} e^{\frac{\beta^2}{4\alpha} + \gamma}.$$

□

For the linear-time statistic, we obtain

$$\begin{aligned} \mathbb{E}_{z,z'} [h(z, z')^2] &= \frac{2\nu}{\sqrt{\nu^2+4}} \left(1 + \exp\left(\frac{-\mu^2}{\nu^2+4}\right) \right) \\ &+ \frac{2\nu^2}{\nu^2+2} \left(1 + \exp\left(\frac{-\mu^2}{\nu^2+2}\right) \right) \\ &- \frac{8\nu^2}{\sqrt{(\nu^2+1)(\nu^2+3)}} \exp\left(\frac{-(\nu^2+2)\mu^2}{2(\nu^2+1)(\nu^2+3)}\right). \end{aligned}$$

For the quadratic-time statistic, we obtain

$$\begin{aligned} \mathbb{E}_z \left[\left(\mathbb{E}_{z'} [h(z, z')] \right)^2 \right] &= \frac{2\nu^2}{\sqrt{(\nu^2+1)(\nu^2+3)}} \left(1 + \exp\left(\frac{-\mu^2}{\nu^2+3}\right) \right) \\ &+ \frac{2\nu^2}{\nu^2+2} \left(1 + \exp\left(\frac{-\mu^2}{\nu^2+2}\right) \right) \\ &- \frac{4\nu^2}{\sqrt{(\nu^2+1)(\nu^2+3)}} \exp\left(\frac{-(\nu^2+2)\mu^2}{2(\nu^2+1)(\nu^2+3)}\right) \\ &- \frac{4\nu^2}{\nu^2+2} \exp\left(\frac{-\mu^2}{2(\nu^2+2)}\right). \end{aligned}$$

In Figure 5, we plot both R_ℓ and R_u for various values of the change parameter μ .

Change in the variance. In this case, the MMD is given by

$$\mathbb{E}_{z,z'} [h(z, z')] = \nu \left(\frac{1}{\sqrt{\nu^2+2}} + \frac{1}{\sqrt{\nu^2+2\sigma^2}} - \frac{2}{\sqrt{\nu^2+\sigma^2+1}} \right).$$

As in the previous paragraph, the second term in Eq. (16) and Eq. (17) is the MMD, which is given by the previous display. As for the first term, we obtain for the linear-time statistic:

$$\begin{aligned} \mathbb{E}_{z,z'} [h(z, z')^2] &= \frac{\nu}{\sqrt{\nu^2+4}} + \frac{\nu}{\sqrt{\nu^2+4\sigma^2}} + \frac{2\nu}{\sqrt{\nu^2+2\sigma^2+2}} \\ &+ \frac{2\nu^2}{\sqrt{(\nu^2+2)(\nu^2+2\sigma^2)}} - \frac{4\nu^2}{\sqrt{(\nu^2+1)(\nu^2+\sigma^2)+2\nu^2+\sigma^2+1}} \\ &- \frac{4\nu^2}{\sqrt{(\nu^2+1)(\nu^2+\sigma^2)+\sigma^2(2\nu^2+\sigma^2+1)}} + \frac{2\nu^2}{\nu^2+\sigma^2+1}. \end{aligned}$$




FIG 5. In this figure, we plot R_ℓ (left panel) and R_u (right panel) as a function of the bandwidth in the change of mean scenario. Note that both R_ℓ and R_u are becoming very “flat” when μ goes to 0, leading to numerical problems when maximizing with respect to ν .




FIG 6. In this figure, we plot R_ℓ (left panel) and R_u (right panel) as a function of the bandwidth in the change of variance scenario. Note that both R_ℓ and R_u are becoming very “flat” when σ goes to 1, leading to numerical problems when maximizing with respect to ν . This is similar to the phenomenon observed in the change in the mean scenario when $\mu \rightarrow 0$.

And for the the quadratic-time statistic:

$$\begin{aligned} \mathbb{E}_z \left[\left(\mathbb{E}_{z'} [h(z, z')] \right)^2 \right] = & \nu^2 \times \left\{ \frac{1}{\sqrt{(\nu^2 + 1)(\nu^2 + 3)}} + \frac{1}{\sqrt{(\nu^2 + \sigma^2)(\nu^2 + 3\sigma^2)}} \right. \\ & + \frac{1}{\sqrt{(\nu^2 + \sigma^2)(\nu^2 + \sigma^2 + 2)}} + \frac{1}{\sqrt{(\nu^2 + 1)(\nu^2 + 2\sigma^2 + 1)}} \\ & + \frac{2}{\sqrt{(\nu^2 + 2\sigma^2)(\nu^2 + 2)}} - \frac{2}{\sqrt{(\nu^2 + 1)(\nu^2 + \sigma^2) + 2\nu^2 + \sigma^2 + 1}} \\ & - \frac{2}{\sqrt{(\nu^2 + 2)(\nu^2 + \sigma^2 + 1)}} - \frac{2}{\sqrt{(\nu^2 + \sigma^2 + 1)(\nu^2 + 2\sigma^2)}} \\ & \left. - \frac{2}{\sqrt{(\nu^2 + 1)(\nu^2 + \sigma^2) + \sigma^2(2\nu^2 + \sigma^2 + 1)}} + \frac{2}{\nu^2 + \sigma^2 + 1} \right\}. \end{aligned}$$

F. Proof of Prop. 4.1

Set $T_n^u := n\text{MMD}_u^2$. According to Th. 12 in [Gretton et al. \(2012a\)](#), with our notations,

$$n\text{MMD}_u^2 \xrightarrow{\mathcal{L}} \sum_{\ell=1}^{+\infty} \lambda_{\ell} \left[\left(\sqrt{2}a_{\ell} - \sqrt{2}b_{\ell} \right)^2 - 4 \right],$$

where a_{ℓ}, b_{ℓ} are i.i.d. standard Gaussian random variables and the λ_{ℓ} are the eigenvalues of the centered Gaussian kernel operator. We deduce that $n\text{MMD}_u^2 \xrightarrow{\mathcal{L}} 4 \sum_{\ell=1}^{+\infty} \lambda_{\ell} (c_{\ell}^2 - 1)$, where the c_{ℓ} are i.i.d. standard Gaussian random variables. Denote by F the cumulative distribution function of this random variables. According to [Bahadur \(1960\)](#), the cumulative distribution function of c_{ℓ}^2 belongs to $\mathcal{D}(1, 1)$. It follows from Th. 11 in [Jitkrittum et al. \(2017b\)](#) that F belongs to $\mathcal{D}\left(\frac{1}{4\lambda_1}, 1\right)$. We now take $R(n) = n$. According to Th. 10 in [Gretton et al. \(2012a\)](#), $T_n/R(n) = \text{MMD}_u^2 \xrightarrow{\mathbb{P}} \text{MMD}^2$. We conclude via Th. 9 in [Jitkrittum et al. \(2017b\)](#).

Set $T_n^{\ell} := \sqrt{n}\text{MMD}_{\ell}^2$. According to [Gretton et al. \(2012b, Eq. 5\)](#) the limiting distribution of $\sqrt{n}\text{MMD}_{\ell}$ under the null is $\mathcal{N}(0, 8\sigma_{\ell}^2)$. The cumulative distribution function of a standard Gaussian random variable belongs to $\mathcal{D}(1, 2)$ according to [Bahadur \(1960\)](#) and we deduce that the cumulative distribution function of $\mathcal{N}(0, 8\sigma_{\ell}^2)$ belongs to $\mathcal{D}(1/(8\sigma_{\ell}^2), 2)$. Set $R(n) := \sqrt{n}$. Since $T_n^{\ell}/R(n) = \text{MMD}_{\ell}^2$ converges in probability towards the squared MMD, we conclude again with Th. 9 in [Jitkrittum et al. \(2017b\)](#). \square