

Generalized Kernel Thinning

Raaz Dwivedi^{1,2}, Lester Mackey³

¹ Department of Computer Science, Harvard University

² Department of EECS, Massachusetts Institute of Technology

³ Microsoft Research New England

raaz@mit.edu, lmackey@microsoft.com

Abstract

The kernel thinning (KT) algorithm of Dwivedi and Mackey (2021) compresses a probability distribution more effectively than independent sampling by targeting a reproducing kernel Hilbert space (RKHS) and leveraging a less smooth square-root kernel. Here we provide four improvements. First, we show that KT applied directly to the target RKHS yields tighter, dimension-free guarantees for any kernel, any distribution, and any fixed function in the RKHS. Second, we show that, for analytic kernels like Gaussian, inverse multiquadric, and sinc, target KT admits maximum mean discrepancy (MMD) guarantees comparable to or better than those of square-root KT without making explicit use of a square-root kernel. Third, we prove that KT with a fractional power kernel yields better-than-Monte-Carlo MMD guarantees for non-smooth kernels, like Laplace and Matérn, that do not have square-roots. Fourth, we establish that KT applied to a sum of the target and power kernels (a procedure we call KT+) simultaneously inherits the improved MMD guarantees of power KT and the tighter individual function guarantees of target KT. In our experiments with target KT and KT+, we witness significant improvements in integration error even in 100 dimensions and when compressing challenging differential equation posteriors.

1. Introduction

A core task in probabilistic inference is learning a compact representation of a probability distribution \mathbb{P} . This problem is usually solved by sampling points x_1, \dots, x_n independently from \mathbb{P} or, if direct sampling is intractable, generating n points from a Markov chain converging to \mathbb{P} . The benefit of these approaches is that they provide asymptotically exact sample estimates $\mathbb{P}_{\text{in}} f \triangleq \frac{1}{n} \sum_{i=1}^n f(x_i)$ for intractable expectations $\mathbb{P} f \triangleq \mathbb{E}_{X \sim \mathbb{P}}[f(X)]$. However, they also suffer from a serious drawback: the learned representations are unnecessarily large, requiring n points to achieve $|\mathbb{P} f - \mathbb{P}_{\text{in}} f| = \Theta(n^{-\frac{1}{2}})$ integration error. These inefficient representations quickly become prohibitive for expensive downstream tasks and function evaluations: for example, in computational cardiology, each function evaluation $f(x_i)$ initiates a heart or tissue simulation that consumes 1000s of CPU hours (Niederer et al., 2011; Augustin et al., 2016; Strocchi et al., 2020).

To reduce the downstream computational burden, a standard practice is to *thin* the initial sample by discarding every t -th sample point (Owen, 2017). Unfortunately, standard thinning often results in a substantial loss of accuracy: for example, thinning an i.i.d. or fast-mixing Markov chain sample from n points to $n^{\frac{1}{2}}$ points increases integration error from $\Theta(n^{-\frac{1}{2}})$ to $\Theta(n^{-\frac{1}{4}})$.

The recent *kernel thinning* (KT) algorithm of Dwivedi and Mackey (2021) addresses this issue by producing thinned coresets with better-than-i.i.d. integration error in a reproducing kernel Hilbert space (RKHS, Berlinet and Thomas-Agnan, 2011). Given a target kernel¹ \mathbf{k} and a suitable sequence

1. A kernel \mathbf{k} is any function that yields positive semi-definite matrices $(\mathbf{k}(z_i, z_j))_{i,j=1}^l$ for all inputs $(z_i)_{i=1}^l$. See Tab. 1 in App. B for expressions of popular kernels.

of input points $\mathcal{S}_{\text{in}} = (x_i)_{i=1}^n$ approximating \mathbb{P} , KT returns a subsequence \mathcal{S}_{out} of \sqrt{n} points with better-than-i.i.d. *maximum mean discrepancy* (MMD, Gretton et al., 2012),²

$$\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{P}_{\text{out}}) \triangleq \sup_{\|f\|_{\mathbf{k}} \leq 1} |\mathbb{P}f - \mathbb{P}_{\text{out}}f| \quad \text{for} \quad \mathbb{P}_{\text{out}} \triangleq \frac{1}{\sqrt{n}} \sum_{x \in \mathcal{S}_{\text{out}}} \delta_x, \quad (1)$$

where $\|\cdot\|_{\mathbf{k}}$ denotes the norm for the RKHS \mathcal{H} associated with \mathbf{k} . That is, the KT output admits $\mathcal{O}(n^{-\frac{1}{4}})$ worst-case integration error across the unit ball of \mathcal{H} . KT achieves its improvement with high probability using non-uniform randomness and a less smooth *square-root kernel* \mathbf{k}_{rt} satisfying

$$\mathbf{k}(x, y) = \int_{\mathbb{R}^d} \mathbf{k}_{\text{rt}}(x, z) \mathbf{k}_{\text{rt}}(z, y) dz. \quad (2)$$

When the input points are sampled i.i.d. or from a fast-mixing Markov chain on \mathbb{R}^d , Dwivedi and Mackey prove that the KT output has, with high probability, $\mathcal{O}_d(n^{-\frac{1}{2}} \sqrt{\log n})$ -MMD $_{\mathbf{k}}$ error for \mathbb{P} and \mathbf{k}_{rt} with bounded support, $\mathcal{O}_d(n^{-\frac{1}{2}} (\log^{d+1} n \log \log n)^{\frac{1}{2}})$ -MMD $_{\mathbf{k}}$ error for \mathbb{P} and \mathbf{k}_{rt} with light tails, and $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2} + \frac{d}{2\rho}} \sqrt{\log n \log \log n}))$ -MMD $_{\mathbf{k}}$ error for \mathbb{P} and \mathbf{k}_{rt}^2 with $\rho > 2d$ moments. Meanwhile, an i.i.d. coreset of the same size suffers $\Omega(n^{-\frac{1}{4}})$ MMD $_{\mathbf{k}}$. We refer to the original KT algorithm as ROOT KT hereafter.

Our contributions In this work, we offer four improvements over the original KT algorithm. First, we show in Thm. 1 that a generalization of KT that uses only the target kernel \mathbf{k} , henceforth called TARGET KT, provides a tighter $\mathcal{O}(n^{-\frac{1}{2}} \sqrt{\log n})$ integration error guarantee for each function f in the RKHS. This TARGET KT guarantee (a) applies to **any kernel \mathbf{k} on any domain** (even kernels that do not admit a square-root and kernels defined on non-Euclidean spaces), (b) applies to **any target distribution \mathbb{P}** (even heavy-tailed \mathbb{P} not covered by ROOT KT guarantees), and (c) is **dimension-free**, eliminating the exponential dimension dependence and $(\log n)^{d/2}$ factors of prior ROOT KT guarantees.

Second, we prove in Thm. 2 that, for analytic kernels, like Gaussian, inverse multiquadric (IMQ), and sinc, TARGET KT admits MMD guarantees comparable to or better than those of Dwivedi and Mackey (2021) without making explicit use of a square-root kernel. Third, we establish in Thm. 3 that generalized KT with a fractional α -power kernel \mathbf{k}_{α} yields improved MMD guarantees for kernels that do not admit a square-root, like Laplace and non-smooth Matérn. Fourth, we show in Thm. 4 that, remarkably, applying generalized KT to a sum of \mathbf{k} and \mathbf{k}_{α} —a procedure we call *kernel thinning+* (KT+)—simultaneously inherits the improved MMD of POWER KT and the dimension-free individual function guarantees of TARGET KT.

In Sec. 4, we use our new tools to generate substantially compressed representations of both i.i.d. samples in dimensions $d = 2$ through 100 and Markov chain Monte Carlo samples targeting challenging differential equation posteriors. In line with our theory, we find that TARGET KT and KT+ significantly improve both single function integration error and MMD, even for kernels without fast-decaying square-roots.

Related work For bounded \mathbf{k} , both i.i.d. samples (Tolstikhin et al., 2017, Prop. A.1) and thinned geometrically ergodic Markov chains (Dwivedi and Mackey, 2021, Prop. 1) deliver $n^{\frac{1}{2}}$ points with $\mathcal{O}(n^{-\frac{1}{4}})$ MMD with high probability. The *online Haar strategy* of Dwivedi et al. (2019) and low discrepancy *quasi-Monte Carlo* methods (see, e.g., Hickernell, 1998; Novak and Wozniakowski, 2010;

2. MMD is a metric for *characteristic* \mathbf{k} , like those in Tab. 1, and controls integration error for all bounded continuous f when \mathbf{k} *determines convergence*, like each \mathbf{k} in Tab. 1 except SINC (Simon-Gabriel et al., 2020).

Dick et al., 2013) provide improved $\mathcal{O}_d(n^{-\frac{1}{2}} \log^d n)$ MMD guarantees but are tailored specifically to the uniform distribution on $[0, 1]^d$. Alternative coreset constructions for more general \mathbb{P} include *kernel herding* (Chen et al., 2010), *discrepancy herding* (Harvey and Samadi, 2014), *super-sampling with a reservoir* (Paige et al., 2016), *support points convex-concave procedures* (Mak and Joseph, 2018), *greedy sign selection* (Karnin and Liberty, 2019, Sec. 3.1), *Stein point MCMC* (Chen et al., 2019), and *Stein thinning* (Riabiz et al., 2020a). While some admit better-than-i.i.d. MMD guarantees for finite-dimensional kernels on \mathbb{R}^d (Chen et al., 2010; Harvey and Samadi, 2014), none apart from KT are known to provide better-than-i.i.d. MMD or integration error for the infinite-dimensional kernels covered in this work. The minimax lower bounds of Phillips and Tai (2020, Thm. 3.1) and Tolstikhin et al. (2017, Thm. 1) respectively establish that any procedure outputting $n^{\frac{1}{2}}$ -sized coresets and any procedure estimating \mathbb{P} based only on n i.i.d. sample points must incur $\Omega(n^{-\frac{1}{2}})$ MMD in the worst case. Our guarantees in Sec. 2 match these lower bounds up to logarithmic factors.

Notation We define the norm $\|\mathbf{k}\|_\infty = \sup_{x,y} |\mathbf{k}(x,y)|$ and the shorthand $[n] \triangleq \{1, \dots, n\}$, $\mathbb{R}_+ \triangleq \{x \in \mathbb{R} : x \geq 0\}$, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, $\mathcal{B}_\mathbf{k} \triangleq \{f \in \mathcal{H} : \|f\|_\mathbf{k} \leq 1\}$, and $\mathcal{B}_2(r) = \{y \in \mathbb{R}^d : \|y\|_2 \leq r\}$. We write $a \lesssim b$ and $a \gtrsim b$ to mean $a = \mathcal{O}(b)$ and $a = \Omega(b)$, use \lesssim_d when masking constants dependent on d , and write $a = \mathcal{O}_p(b)$ to mean a/b is bounded in probability. For any distribution \mathbb{Q} and point sequences $\mathcal{S}, \mathcal{S}'$ with empirical distributions $\mathbb{Q}_n, \mathbb{Q}'_n$, we define $\text{MMD}_\mathbf{k}(\mathbb{Q}, \mathcal{S}) \triangleq \text{MMD}_\mathbf{k}(\mathbb{Q}, \mathbb{Q}_n)$ and $\text{MMD}_\mathbf{k}(\mathcal{S}, \mathcal{S}') \triangleq \text{MMD}_\mathbf{k}(\mathbb{Q}_n, \mathbb{Q}'_n)$.

2. Generalized Kernel Thinning

Our generalized kernel thinning algorithm (Alg. 1) for compressing an input point sequence $\mathcal{S}_{\text{in}} = (x_i)_{i=1}^n$ proceeds in two steps: **KT-SPLIT** and **KT-SWAP** detailed in App. A. First, given a thinning parameter m and an auxiliary kernel $\mathbf{k}_{\text{split}}$, **KT-SPLIT** divides the input sequence into 2^m candidate coresets of size $n/2^m$ using non-uniform randomness. Next, given a target kernel \mathbf{k} , **KT-SWAP** selects a candidate coreset with smallest $\text{MMD}_\mathbf{k}$ to \mathcal{S}_{in} and iteratively improves that coreset by exchanging coreset points for input points whenever the swap leads to reduced $\text{MMD}_\mathbf{k}$. When $\mathbf{k}_{\text{split}}$ is a square-root kernel \mathbf{k}_{rt} (2) of the target kernel \mathbf{k} , generalized KT recovers the original **ROOT KT** algorithm of Dwivedi and Mackey (2021). In this section, we establish performance guarantees for more general $\mathbf{k}_{\text{split}}$ with special emphasis on the practical choice of $\mathbf{k}_{\text{split}} = \mathbf{k}$. Like **ROOT KT**, for any m , generalized KT has time complexity dominated by $\mathcal{O}(n^2)$ evaluations of $\mathbf{k}_{\text{split}}$ and \mathbf{k} and $\mathcal{O}(n \min(d, n))$ space complexity from storing either \mathcal{S}_{in} or the kernel matrices.

Algorithm 1: Generalized Kernel Thinning – Return coreset of size $\lfloor n/2^m \rfloor$ with small $\text{MMD}_\mathbf{k}$

Input: split kernel $\mathbf{k}_{\text{split}}$, target kernel \mathbf{k} , point sequence $\mathcal{S}_{\text{in}} = (x_i)_{i=1}^n$, thinning parameter $m \in \mathbb{N}$, probabilities $(\delta_i)_{i=1}^{bn/2^c}$

$(\mathcal{S}^{(m,\ell)})_{\ell=1}^{2^m} \leftarrow \text{KT-SPLIT}(\mathbf{k}_{\text{split}}, \mathcal{S}_{\text{in}}, m, (\delta_i)_{i=1}^{bn/2^c})$ // Split \mathcal{S}_{in} into 2^m candidate coresets of size $\lfloor \frac{n}{2^m} \rfloor$

$\mathcal{S}_{\text{KT}} \leftarrow \text{KT-SWAP}(\mathbf{k}, \mathcal{S}_{\text{in}}, (\mathcal{S}^{(m,\ell)})_{\ell=1}^{2^m})$ // Select best coreset and iteratively refine

return coreset \mathcal{S}_{KT} of size $\lfloor n/2^m \rfloor$

Our first main result, proved in App. C, bounds the **KT-SPLIT** integration error for any fixed function in the RKHS $\mathcal{H}_{\text{split}}$ generated by $\mathbf{k}_{\text{split}}$. Throughout, the term *oblivious* indicates that a sequence is generated independently of any randomness in **KT**.

Theorem 1 (Single function guarantees for KT-SPLIT) Consider KT-SPLIT (Alg. 1a) with oblivious \mathcal{S}_{in} and $(\delta_i)_{i=1}^{n/2}$ and $\delta^* \triangleq \min_i \delta_i$. If $\frac{n}{2^m} \in \mathbb{N}$, then, for any fixed $f \in \mathcal{H}_{\text{split}}$, index $\ell \in [2^m]$, and scalar $\delta' \in (0, 1)$, the output coresets $\mathcal{S}^{(m, \ell)}$ with $\mathbb{P}_{\text{split}}^{(\ell)} \triangleq \frac{1}{n/2^m} \sum_{x \in \mathcal{S}^{(m, \ell)}} x$ satisfies

$$|\mathbb{P}_{\text{in}} f - \mathbb{P}_{\text{split}}^{(\ell)} f| \leq \|f\|_{\mathbf{k}_{\text{split}}} \cdot \sigma_m \sqrt{2 \log\left(\frac{2}{\delta'}\right)} \quad \text{for} \quad \sigma_m \triangleq \frac{2}{\sqrt{3}} \frac{2^m}{n} \sqrt{\|\mathbf{k}_{\text{split}}\|_{\infty, \text{in}} \cdot \log\left(\frac{4}{\delta^*}\right)}$$

with probability at least $p_{\text{sg}} \triangleq 1 - \sum_{j=1}^m \sum_{i=1}^{n/2^j} \delta_i - \delta'$. Here, $\|\mathbf{k}_{\text{split}}\|_{\infty, \text{in}} \triangleq \max_{x \in \mathcal{S}_{\text{in}}} \mathbf{k}_{\text{split}}(x, x)$.

To ensure that the success probability p_{sg} is at least $1 - \delta$, we can set $\delta' = \frac{\delta}{2}$ and $\delta_i = \frac{\delta}{2n(1-2^{-m})}$ for a stopping time n known *a priori* or $\delta_i = \delta / (4m(i+1) \log^2(i+1))$ for an arbitrary oblivious stopping time n . When compressing heavily from n to \sqrt{n} points, Thm. 1 guarantees $\mathcal{O}(n^{-\frac{1}{2}} \sqrt{\log n})$ integration error with high probability for any fixed function $f \in \mathcal{H}_{\text{split}}$. This represents a near-quadratic improvement over the $\Omega(n^{-\frac{1}{4}})$ integration error of \sqrt{n} i.i.d. points.

Our second main result bounds the $\text{MMD}_{\mathbf{k}}$ (1)—the worst-case integration error across the unit ball of \mathcal{H} —for generalized KT applied to the target kernel, i.e., $\mathbf{k}_{\text{split}} = \mathbf{k}$. The proof of this result in App. D is based on Thm. 1 and an appropriate covering number for the unit ball $\mathcal{B}_{\mathbf{k}}$ of the \mathbf{k} RKHS.

Definition 1 (k covering number) For a set \mathcal{A} and scalar $\varepsilon > 0$, we define the \mathbf{k} covering number $\mathcal{N}_{\mathbf{k}}(\mathcal{A}, \varepsilon)$ with $\mathcal{M}_{\mathbf{k}}(\mathcal{A}, \varepsilon) \triangleq \log \mathcal{N}_{\mathbf{k}}(\mathcal{A}, \varepsilon)$ as the minimum cardinality of a set $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}}$ satisfying

$$\mathcal{B}_{\mathbf{k}} \subseteq \bigcup_{h \in \mathcal{C}} \{g \in \mathcal{B}_{\mathbf{k}} : \sup_{x \in \mathcal{A}} |h(x) - g(x)| \leq \varepsilon\}. \quad (3)$$

Theorem 2 (MMD guarantee for TARGET KT) Consider generalized KT (Alg. 1) with $\mathbf{k}_{\text{split}} = \mathbf{k}$, oblivious \mathcal{S}_{in} and $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$, and $\delta^* \triangleq \min_i \delta_i$. If $\frac{n}{2^m} \in \mathbb{N}$, then for any $\delta' \in (0, 1)$, the output coresets \mathcal{S}_{KT} is of size $\frac{n}{2^m}$ and satisfies

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \leq \inf_{\varepsilon \in (0, 1), \mathcal{S}_{\text{in}} \subset \mathcal{A}} 2\varepsilon + \frac{2^m}{n} \cdot \sqrt{\frac{8}{3} \|\mathbf{k}\|_{\infty, \text{in}} \log\left(\frac{4}{\delta^*}\right) \cdot [\log\left(\frac{4}{\delta'}\right) + \mathcal{M}_{\mathbf{k}}(\mathcal{A}, \varepsilon)]} \quad (4)$$

with probability at least p_{sg} , where $\|\mathbf{k}\|_{\infty, \text{in}}$ and p_{sg} were defined in Thm. 1.

When compressing heavily from n to \sqrt{n} points, Thm. 2 with $\varepsilon = \sqrt{\frac{\|\mathbf{k}\|_{\mathcal{J}, \text{in}}}{n}}$ and $\mathcal{A} = \mathcal{B}_2(\mathbb{R}_{\text{in}})$ for $\mathbb{R}_{\text{in}} \triangleq \max_{x \in \mathcal{S}_{\text{in}}} \|x\|_2$ guarantee

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \lesssim_{\delta} \sqrt{\frac{\|\mathbf{k}\|_{\mathcal{J}, \text{in}} \log n}{n} \cdot \mathcal{M}_{\mathbf{k}}(\mathcal{B}_2(\mathbb{R}_{\text{in}}), \sqrt{\frac{\|\mathbf{k}\|_{\mathcal{J}, \text{in}}}{n}})} \quad (5)$$

with high probability, thereby giving an MMD guarantee for any kernel \mathbf{k} with a covering number bound and a comparable guarantee for \mathbb{P} since $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{KT}}) \leq \text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{in}}) + \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}})$. We unpack the results further in Tabs. 2 and 3 in App. B for a range of \mathbb{P} and \mathbf{k} .

3. Kernel Thinning+

We next introduce and analyze two new generalized KT variants: (i) POWER KT which leverages a power kernel \mathbf{k}_{α} that interpolates between \mathbf{k} and \mathbf{k}_{r} to improve upon the MMD guarantees of target KT even when \mathbf{k}_{r} is not available and (ii) KT+ which uses a sum of \mathbf{k} and \mathbf{k}_{α} to retain both the improved MMD guarantee of \mathbf{k}_{α} and the superior single function guarantees of \mathbf{k} .

First, we define a generalization of the square-root kernel (2) for shift-invariant \mathbf{k} . Let \widehat{f} denote the generalized Fourier transform (GFT, Wendland, 2004, Def. 8.9) of order 0 of $f : \mathbb{R}^d \rightarrow \mathbb{R}$.

Definition 2 (α -power kernel) Define $\mathbf{k}_1 \triangleq \mathbf{k}$. We say a kernel $\mathbf{k}_{\frac{1}{2}}$ is a $\frac{1}{2}$ -power kernel for \mathbf{k} if $\mathbf{k}(x, y) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \mathbf{k}_{\frac{1}{2}}(x, z) \mathbf{k}_{\frac{1}{2}}(z, y) dz$. For $\alpha \in (\frac{1}{2}, 1)$, a kernel $\mathbf{k}_\alpha(x, y) = \kappa_\alpha(x - y)$ on \mathbb{R}^d is an α -power kernel for $\mathbf{k}(x, y) = \kappa(x - y)$ if $\widehat{\kappa}_\alpha = \widehat{\kappa}^\alpha$.

By design, $\mathbf{k}_{\frac{1}{2}}$ matches \mathbf{k}_τ (2) up to an immaterial constant rescaling. Given a power kernel \mathbf{k}_α we define POWER KT as generalized KT with $\mathbf{k}_{\text{split}} = \mathbf{k}_\alpha$. Our next result (with proof in App. E) provides an MMD guarantee for POWER KT.

Theorem 3 (MMD guarantee for POWER KT) Consider generalized KT (Alg. 1) with $\mathbf{k}_{\text{split}} = \mathbf{k}_\alpha$ for some $\alpha \in [\frac{1}{2}, 1]$, oblivious sequences \mathcal{S}_{in} and $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$, and $\delta^* \triangleq \min_i \delta_i$. If $\frac{n}{2^m} \in \mathbb{N}$, then for any $\delta' \in (0, 1)$, the output coreset \mathcal{S}_{KT} is of size $\frac{n}{2^m}$ and satisfies

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \leq \left(\frac{2^m}{n} \|\mathbf{k}_\alpha\|_\infty \right)^{\frac{1}{2\alpha}} (2 \cdot \widetilde{\mathcal{M}}_\alpha)^{1 - \frac{1}{2\alpha}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2} + 1)}} \cdot \mathcal{R}_{\text{max}}^{\frac{d}{2}} \cdot \widetilde{\mathcal{M}}_\alpha \right)^{\frac{1}{\alpha} - 1}, \quad (6)$$

with probability at least p_{sg} (defined in Thm. 1), where the parameters $\widetilde{\mathcal{M}}_\alpha$ and \mathcal{R}_{max} are defined in App. E. When each $\delta_i = \frac{\delta^0}{2^n}$ these parameters satisfy $\widetilde{\mathcal{M}}_\alpha = \mathcal{O}_d(\sqrt{\log n})$ and $\mathcal{R}_{\text{max}} = \mathcal{O}_d(1)$ for \mathbb{P} and \mathbf{k}_α with compact support and $\widetilde{\mathcal{M}}_\alpha = \mathcal{O}_d(\sqrt{\log n \log \log n})$ and $\mathcal{R}_{\text{max}} = \mathcal{O}_d(\log n)$ for \mathbb{P} and \mathbf{k}_α with subexponential tails.

Thm. 3 reproduces the ROOT KT guarantee of Dwivedi and Mackey (2021, Thm. 1) when $\alpha = \frac{1}{2}$ and more generally accommodates any power kernel via an MMD interpolation result (Prop. 1) that may be of independent interest. This generalization is especially valuable for less-smooth kernels like LAPLACE and MATÉRN(ν, γ) with $\nu \in (\frac{d}{2}, d]$ that have no square-root kernel. Our TARGET KT MMD guarantees are no better than i.i.d. for these kernels, but, as shown in App. J, these kernels have MATÉRN kernels as α -power kernels, which yield $o(n^{-\frac{1}{4}})$ MMD in conjunction with Thm. 3.

Our final KT variant, *kernel thinning+*, runs KT-SPLIT with a scaled sum of the target and power kernels, $\mathbf{k}^\dagger \triangleq \mathbf{k} / \|\mathbf{k}\|_\infty + \mathbf{k}_\alpha / \|\mathbf{k}_\alpha\|_\infty$. Remarkably, this choice simultaneously provides the improved MMD guarantees of Thm. 3 and the dimension-free single function guarantees of Thm. 1.

Theorem 4 (Single function & MMD guarantees for KT+) Consider generalized KT (Alg. 1) with $\mathbf{k}_{\text{split}} = \mathbf{k}^\dagger$, oblivious \mathcal{S}_{in} and $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$, $\delta^* \triangleq \min_i \delta_i$, and $\frac{n}{2^m} \in \mathbb{N}$. For any fixed function $f \in \mathcal{H}$, index $\ell \in [2^m]$, and scalar $\delta' \in (0, 1)$, the KT-SPLIT coreset $\mathcal{S}^{(m, \ell)}$ satisfies

$$|\mathbb{P}_{\text{in}} f - \mathbb{P}_{\text{split}}^{(\ell)} f| \leq \frac{2^m}{n} \cdot \sqrt{\frac{16}{3} \log\left(\frac{4}{\delta^*}\right) \log\left(\frac{2}{\delta^0}\right)} \|f\|_{\mathbf{k}} \sqrt{\|\mathbf{k}\|_\infty}, \quad (7)$$

with probability at least p_{sg} (for p_{sg} and $\mathbb{P}_{\text{split}}^{(\ell)}$ defined in Thm. 1). Moreover,

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \leq \min \left[\sqrt{2} \cdot \overline{\mathcal{M}}_{\text{targetKT}}(\mathbf{k}), \quad 2^{\frac{1}{2\alpha}} \cdot \overline{\mathcal{M}}_{\text{powerKT}}(\mathbf{k}_\alpha) \right] \quad (8)$$

with probability at least p_{sg} where $\overline{\mathcal{M}}_{\text{targetKT}}(\mathbf{k})$ denotes the right hand side of (4) with $\|\mathbf{k}\|_{\infty, \text{in}}$ replaced by $\|\mathbf{k}\|_\infty$, and $\overline{\mathcal{M}}_{\text{powerKT}}(\mathbf{k}_\alpha)$ denotes the right hand side of (6).

As shown in Tab. 3 in App. B, KT+ provides better-than-i.i.d. MMD guarantees for several popular kernels—even the Laplace, non-smooth Matérn, and odd B-spline kernels neglected by prior analyses—while matching or improving upon the guarantees of TARGET KT and ROOT KT in each case.

4. Experiments

Dwivedi and Mackey (2021) illustrated the MMD benefits of ROOT KT over i.i.d. sampling and standard MCMC thinning with a series of vignettes focused on the Gaussian kernel. We revisit those vignettes with the broader range of kernels covered by generalized KT and demonstrate significant improvements in both MMD and single-function integration error. We let \mathbb{P}_{in} and \mathbb{P}_{out} denote the empirical distribution of each input and output coresets respectively and set output size as \sqrt{n} given n input points. Besides MMD error, we evaluate the integration error for four test functions both inside and outside of \mathcal{H} , (a) a random element of the target kernel RKHS ($f(x) = \mathbf{k}(X', x)$), (b) moments ($f(x) = x_1$ and $f(x) = x_1^2$), and (c) a standard numerical integration benchmark test function from the *continuous integrand family* (CIF, Genz, 1984), $f_{CIF}(x) = \exp(-\frac{1}{d} \sum_{j=1}^d |x_j - u_j|)$. App. H provides additional details besides several other experiments with other MCMC and i.i.d. targets.

In Fig. 1, we evaluate results for three MCMC chains targeting the posteriors of three distinct coupled ordinary differential equation model: the Goodwin (1965) model of oscillatory enzymatic control ($d = 4$), the Lotka (1925) model of oscillatory predator-prey evolution ($d = 4$), the Hinch et al. (2004) model of calcium signalling in cardiac cells ($d = 38$), and a tempered Hinch posterior. We employ KT+ ($\alpha = 0.81$) with LAPLACE \mathbf{k} for Goodwin and Lotka-Volterra and KT+ ($\alpha = 0.5$) with IMQ \mathbf{k} ($\nu = 0.5$) for Hinch. Notably, neither kernel has a square-root with fast-decaying tails. We observe that KT+ uniformly improves upon the MMD error of standard thinning (ST), even when ST exhibits better-than-i.i.d. accuracy. In nearly every setting, KT+ provides significantly smaller integration error for functions inside of the RKHS (like $\mathbf{k}(X', \cdot)$) and outside of the RKHS (like the first and second moments and the benchmark CIF function).




Figure 1: Kernel thinning+ (KT+) vs. standard MCMC thinning (ST). For kernels *without* fast-decaying square-roots, KT+ improves *both* MMD and integration error decay rates in each inference task.

References

- Christoph M Augustin, Aurel Neic, Manfred Liebmann, Anton J Prassl, Steven A Niederer, Gundolf Haase, and Gernot Plank. Anatomically accurate high resolution modeling of human whole heart electromechanics: A strongly scalable algebraic multigrid solver method for nonlinear deformation. *Journal of computational physics*, 305:622–646, 2016.
- Necdet Batir. Bounds for the gamma function. *Results in Mathematics*, 72(1):865–874, 2017. doi: 10.1007/s00025-017-0698-0. URL <https://doi.org/10.1007/s00025-017-0698-0>.
- Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- Wilson Ye Chen, Alessandro Barp, François-Xavier Briol, Jackson Gorham, Mark Girolami, Lester Mackey, and Chris Oates. Stein point Markov chain Monte Carlo. In *International Conference on Machine Learning*, pages 1011–1021. PMLR, 2019.
- Yutian Chen, Max Welling, and Alex Smola. Super-samples from kernel herding. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, UAI’10*, page 109–116, Arlington, Virginia, USA, 2010. AUAI Press. ISBN 9780974903965.
- Josef Dick, Frances Y Kuo, and Ian H Sloan. High-dimensional integration: The quasi-Monte Carlo way. *Acta Numerica*, 22:133–288, 2013.
- Raaz Dwivedi and Lester Mackey. Kernel thinning. *arXiv preprint arXiv:2105.05842*, 2021.
- Raaz Dwivedi, Ohad N Feldheim, Ori Gurel-Gurevich, and Aaditya Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 174(1):103–131, 2019.
- Alan Genz. Testing multidimensional integration routines. In *Proc. of international conference on Tools, methods and languages for scientific and engineering computation*, pages 81–94, 1984.
- Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.
- Brian C Goodwin. Oscillatory behavior in enzymatic control process. *Advances in Enzyme Regulation*, 3: 318–356, 1965.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(25):723–773, 2012.
- Heikki Haario, Eero Saksman, and Johanna Tamminen. Adaptive proposal distribution for random walk Metropolis algorithm. *Computational Statistics*, 14(3):375–395, 1999.
- Nick Harvey and Samira Samadi. Near-optimal herding. In *Conference on Learning Theory*, pages 1165–1182, 2014.
- Fred Hickernell. A generalized discrepancy and quadrature error bound. *Mathematics of computation*, 67(221):299–322, 1998.
- Robert Hinch, JL Greenstein, AJ Tanskanen, L Xu, and RL Winslow. A simplified local control model of calcium-induced calcium release in cardiac ventricular myocytes. *Biophysical journal*, 87(6):3723–3736, 2004.
- Zohar Karnin and Edo Liberty. Discrepancy, coresets, and sketches in machine learning. In *Conference on Learning Theory*, pages 1975–1993. PMLR, 2019.

- Alfred James Lotka. *Elements of physical biology*. Williams & Wilkins, 1925.
- Simon Mak and V Roshan Joseph. Support points. *The Annals of Statistics*, 46(6A):2562–2592, 2018.
- Whitney K. Newey and Daniel McFadden. Chapter 36: Large sample estimation and hypothesis testing. In *Handbook of Econometrics*, volume 4, pages 2111–2245. Elsevier, 1994. doi: [https://doi.org/10.1016/S1573-4412\(05\)80005-4](https://doi.org/10.1016/S1573-4412(05)80005-4). URL <https://www.sciencedirect.com/science/article/pii/S1573441205800054>.
- Steven A Niederer, Lawrence Mitchell, Nicolas Smith, and Gernot Plank. Simulating human cardiac electrophysiology on clinical time-scales. *Frontiers in Physiology*, 2:14, 2011.
- E Novak and H Wozniakowski. Tractability of multivariate problems, volume ii: Standard information for functionals, european math. *Soc. Publ. House, Zürich*, 3, 2010.
- Art B Owen. Statistically efficient thinning of a Markov chain sampler. *Journal of Computational and Graphical Statistics*, 26(3):738–744, 2017.
- Brooks Paige, Dino Sejdinovic, and Frank Wood. Super-sampling with a reservoir. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pages 567–576, 2016.
- Jeff M Phillips and Wai Ming Tai. Near-optimal coresets of kernel density estimates. *Discrete & Computational Geometry*, 63(4):867–887, 2020.
- Marina Riabiz, Wilson Chen, Jon Cockayne, Pawel Swietach, Steven A Niederer, Lester Mackey, and Chris Oates. Optimal thinning of MCMC output. *arXiv preprint arXiv:2005.03952*, 2020a.
- Marina Riabiz, Wilson Ye Chen, Jon Cockayne, Pawel Swietach, Steven A. Niederer, Lester Mackey, and Chris J. Oates. Replication Data for: Optimal Thinning of MCMC Output, 2020b. URL <https://doi.org/10.7910/DVN/MDKNWM>. Accessed on Mar 23, 2021.
- Gareth O Roberts and Richard L Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- Alessandro Rudi, Ulysse Marteau-Ferey, and Francis Bach. Finding global minima via kernel approximations. *arXiv preprint arXiv:2012.11978*, 2020.
- Carl-Johann Simon-Gabriel, Alessandro Barp, and Lester Mackey. Metrizing weak convergence with maximum mean discrepancies. *arXiv preprint arXiv:2006.09268*, 2020.
- Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- Ingo Steinwart and Simon Fischer. A closer look at covering number bounds for Gaussian kernels. *Journal of Complexity*, 62:101513, 2021.
- Marina Stocchi, Matthias AF Gsell, Christoph M Augustin, Orod Razeghi, Caroline H Roney, Anton J Prassl, Edward J Vigmond, Jonathan M Behar, Justin S Gould, Christopher A Rinaldi, Martin J Bishop, Gernot Plank, and Steven A Niederer. Simulating ventricular systolic motion in a four-chamber heart model with spatially varying robin boundary conditions to model the effect of the pericardium. *Journal of Biomechanics*, 101:109645, 2020.
- Hong-Wei Sun and Ding-Xuan Zhou. Reproducing kernel Hilbert spaces associated with analytic translation-invariant Mercer kernels. *Journal of Fourier Analysis and Applications*, 14(1):89–101, 2008.

Ilya Tolstikhin, Bharath K Sriperumbudur, and Krikamol Muandet. Minimax estimation of kernel mean embeddings. *The Journal of Machine Learning Research*, 18(1):3002–3048, 2017.

Vito Volterra. Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. 1926.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.

Haizhang Zhang and Liang Zhao. On the inclusion relation of reproducing kernel hilbert spaces. *Analysis and Applications*, 11(02):1350014, 2013.

Ding-Xuan Zhou. Capacity of reproducing kernel spaces in learning theory. *IEEE Transactions on Information Theory*, 49(7):1743–1752, 2003.

Appendix

A	Details of KT-SPLIT and KT-SWAP	11
B	Consequences of Thm. 2 and Thm. 4	11
C	Proof of Thm. 1: Single function guarantees for KT-SPLIT	13
D	Proof of Thm. 2: MMD guarantee for TARGET KT	14
E	Proof of Thm. 3: MMD guarantee for POWER KT	15
F	Proof of Thm. 4: Single function & MMD guarantees for KT+	16
G	Proof of Prop. 1: An interpolation result for MMD	18
H	Additional experimental results	19
I	Upper bounds on RKHS covering numbers	23
J	Proof of Tab. 3 results	29

Appendix A. Details of KT-SPLIT and KT-SWAP

Algorithm 1a: KT-SPLIT – Divide points into candidate coresets of size $\lfloor n/2^m \rfloor$

Input: kernel $\mathbf{k}_{\text{split}}$, point sequence $\mathcal{S}_{\text{in}} = (x_i)_{i=1}^n$, thinning parameter $m \in \mathbb{N}$, probabilities $(\delta_i)_{i=1}^n$
 $\mathcal{S}^{(j,\ell)} \leftarrow \{\}$ for $0 \leq j \leq m$ and $1 \leq \ell \leq 2^j$ // Empty coresets: $\mathcal{S}^{(j,\ell)}$ has size $\lfloor \frac{i}{2^j} \rfloor$ after round i
 $\sigma_{j,\ell} \leftarrow 0$ for $1 \leq j \leq m$ and $1 \leq \ell \leq 2^{j-1}$ // Swapping parameters
for $i = 1, \dots, \lfloor n/2 \rfloor$ **do**
 $\mathcal{S}^{(0,1)}$.append(x_i); $\mathcal{S}^{(0,1)}$.append(x_{2i})
 // Every 2^j rounds, add one point from parent coreset $\mathcal{S}^{(j-1,\ell)}$ to each child coreset $\mathcal{S}^{(j,2\ell-1)}, \mathcal{S}^{(j,2\ell)}$
 for ($j = 1; j \leq m$ and $i/2^{j-1} \in \mathbb{N}; j = j + 1$) **do**
 for $\ell = 1, \dots, 2^{j-1}$ **do**
 $(\mathcal{S}, \mathcal{S}^\theta) \leftarrow (\mathcal{S}^{(j-1,\ell)}, \mathcal{S}^{(j,2\ell-1)})$; $(x, \tilde{x}) \leftarrow \text{get_last_two_points}(\mathcal{S})$
 // Compute swapping threshold a
 $a, \sigma_{j,\ell} \leftarrow \text{get_swap_params}(\sigma_{j,\ell}, \mathbf{b}, \delta_{j\mathcal{S}^{(j-1,\ell)}} \frac{2^j}{m})$ with
 $\mathbf{b}^2 = \mathbf{k}_{\text{split}}(x, x) + \mathbf{k}_{\text{split}}(\tilde{x}, \tilde{x}) - 2\mathbf{k}_{\text{split}}(x, \tilde{x})$
 // Assign one point to each child after probabilistic swapping
 $\alpha \leftarrow$
 $\mathbf{k}_{\text{split}}(\tilde{x}, \tilde{x}) - \mathbf{k}_{\text{split}}(x, x) + \sum_{y \in \mathcal{S}} (\mathbf{k}_{\text{split}}(y, x) - \mathbf{k}_{\text{split}}(y, \tilde{x})) - 2\sum_{z \in \mathcal{S}^\theta} (\mathbf{k}_{\text{split}}(z, x) - \mathbf{k}_{\text{split}}(z, \tilde{x}))$
 $(x, \tilde{x}) \leftarrow (\tilde{x}, x)$ with probability $\min(1, \frac{1}{2}(1 - \frac{\alpha}{a})_+)$
 $\mathcal{S}^{(j,2\ell-1)}$.append(x); $\mathcal{S}^{(j,2\ell)}$.append(\tilde{x})
 end
 end
end
return $(\mathcal{S}^{(m,\ell)})_{\ell=1}^{2^m}$, candidate coresets of size $\lfloor n/2^m \rfloor$

function $\text{get_swap_params}(\sigma, \mathbf{b}, \delta)$:
 $a \leftarrow \max(\mathbf{b}\sigma\sqrt{2\log(2/\delta)}, \mathbf{b}^2)$
 $\sigma \leftarrow \sigma + \mathbf{b}^2(1 + (\mathbf{b}^2 - 2a)\sigma^2/a^2)_+$
return (a, σ)

Algorithm 1b: KT-SWAP – Identify and refine the best candidate coreset

Input: kernel \mathbf{k} , point sequence $\mathcal{S}_{\text{in}} = (x_i)_{i=1}^n$, candidate coresets $(\mathcal{S}^{(m,\ell)})_{\ell=1}^{2^m}$
 $\mathcal{S}^{(m,0)}$ baseline.thinning($\mathcal{S}_{\text{in}}, \text{size} = bn/2^m c$) // Compare to baseline
 $\mathcal{S}_{\text{KT}} = \mathcal{S}^{(m,\ell^*)}$ for $\ell^* = \text{argmin}_{\ell \in \{0, 1, \dots, 2^m\}} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}^{(m,\ell)})$ // Select best candidate coreset
 // Swap out each point in \mathcal{S}_{KT} for best alternative in \mathcal{S}_{in}
for $i = 1, \dots, bn/2^m c$ **do**
 $\mathcal{S}_{\text{KT}}[i] = \text{argmin}_{z \in \mathcal{S}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}} \text{ with } \mathcal{S}_{\text{KT}}[i] = z)$
end
return \mathcal{S}_{KT} , refined coreset of size $bn/2^m c$

Appendix B. Consequences of Thm. 2 and Thm. 4

Tab. 2 summarizes the MMD guarantees of Thm. 2 under common growth conditions on the log covering number $\mathcal{M}_{\mathbf{k}}$ and the input point radius $\mathcal{R}_{\mathcal{S}_{\text{in}}} \triangleq \max_{x \in \mathcal{S}_{\text{in}}} \|x\|_2$. In Props. 2 and 3 of

GAUSS(σ) $\sigma > 0$	LAPLACE(σ) $\sigma > 0$	MATÉRN(ν, γ) $\nu > \frac{d}{2}, \gamma > 0$	IMQ(ν, γ) $\nu > 0, \gamma > 0$	SINC(θ) $\theta \neq 0$	B-SPLINE($2\beta+1, \gamma$) $\beta \in \mathbb{N}$
$\exp\left(-\frac{\ z\ _2^2}{2\sigma^2}\right)$	$\exp\left(-\frac{\ z\ _2}{\sigma}\right)$	$c_{\nu-\frac{d}{2}}(\gamma\ z\ _2)^{\nu-\frac{d}{2}} \cdot K_{\nu-\frac{d}{2}}(\gamma\ z\ _2)$	$\frac{1}{(1+\ z\ _2^2/\gamma^2)^\nu}$	$\prod_{j=1}^d \frac{\sin(\theta z_j)}{\theta z_j}$	$\mathbb{B}_{2\beta+2}^{-d} \prod_{j=1}^d h_\beta(\gamma z_j)$

Table 1: Common kernels $\mathbf{k}(x, y)$ on \mathbb{R}^d with $z = x - y$. In each case, $\|\mathbf{k}\|_1 = 1$. Here, $c_a \triangleq \frac{2^1 a}{\Gamma(a)}$, K_a is the modified Bessel function of the third kind of order a (Wendland, 2004, Def. 5.10), h_β is the recursive convolution of $2\beta + 2$ copies of $\mathbf{1}_{[\frac{1}{2}, \frac{1}{2}]}$, and $\mathbb{B}_\beta \triangleq \frac{1}{(\beta-1)!} \sum_{j=0}^{\beta-2} (-1)^j \binom{\beta}{j} \left(\frac{\beta-j}{2}\right)^\beta$.

App. I, we show that analytic kernels, like Gaussian, inverse multiquadric (IMQ), and sinc, have **LOGGROWTH** $\mathcal{M}_\mathbf{k}$ (i.e., $\mathcal{M}_\mathbf{k}(\mathcal{B}_2(r), \varepsilon) \lesssim_d r^d \log^\omega(\frac{1}{\varepsilon})$) while finitely differentiable kernels (like Matérn and B-spline) have **POLYGROWTH** $\mathcal{M}_\mathbf{k}$ (i.e., $\mathcal{M}_\mathbf{k}(\mathcal{B}_2(r), \varepsilon) \lesssim_d r^d \varepsilon^{-\omega}$).

Our conditions on $\mathcal{R}_{\mathcal{S}_{\text{in}}}$ arise from four forms of target distribution tail decay: (1) **COMPACT** ($\mathcal{R}_{\mathcal{S}_{\text{in}}} \lesssim_d 1$), (2) **SUBGAUSS** ($\mathcal{R}_{\mathcal{S}_{\text{in}}} \lesssim_d \sqrt{\log n}$), (3) **SUBEXP** ($\mathcal{R}_{\mathcal{S}_{\text{in}}} \lesssim_d \log n$), and (4) **HEAVYTAIL** ($\mathcal{R}_{\mathcal{S}_{\text{in}}} \lesssim_d n^{1/\rho}$). The first setting arises with a compactly supported \mathbb{P} (e.g., on the unit cube $[0, 1]^d$), and the other three settings arise in expectation and with high probability when \mathcal{S}_{in} is generated i.i.d. from \mathbb{P} with sub-Gaussian tails, sub-exponential tails, or ρ moments respectively.

Substituting these conditions into (5) yields the eight entries of Tab. 2. We find that, for LOGGROWTH $\mathcal{M}_\mathbf{k}$, TARGET KT MMD is within log factors of the $\Omega(n^{-1/2})$ lower bounds of Sec. 1 for light-tailed \mathbb{P} and is $o(n^{-1/4})$ (i.e., better than i.i.d.) for any distribution with $\rho > 4d$ moments. Meanwhile, for POLYGROWTH $\mathcal{M}_\mathbf{k}$, TARGET KT MMD is $o(n^{-1/4})$ whenever $\omega < \frac{1}{2}$ for light-tailed \mathbb{P} or whenever \mathbb{P} has $\rho > 2d/(\frac{1}{2} - \omega)$ moments.

Next, for each of the popular convergence-determining kernels of Tab. 1, we compare the ROOT KT MMD guarantees of Dwivedi and Mackey (2021) with the TARGET KT guarantees of Thm. 2 combined with covering number bounds derived in Apps. I and J. We see in Tab. 3 that Thm. 2 provides better-than-i.i.d. and better-than-ROOT KT guarantees for kernels with slowly decaying or non-existent square-roots (e.g., IMQ with $\nu < \frac{d}{2}$, sinc, and B-spline) and nearly matches known ROOT KT guarantees for analytic kernels like Gauss and IMQ with $\nu \geq \frac{d}{2}$, even though TARGET KT makes no explicit use of a square-root kernel.

	COMPACT \mathbb{P} $\mathcal{R}_{\text{in}} \lesssim_d 1$	SUBGAUSS \mathbb{P} $\mathcal{R}_{\text{in}} \lesssim_d \sqrt{\log n}$	SUBEXP \mathbb{P} $\mathcal{R}_{\text{in}} \lesssim_d \log n$	HEAVYTAIL \mathbb{P} $\mathcal{R}_{\text{in}} \lesssim_d n^{1/\rho}$
LOGGROWTH $\mathcal{M}_\mathbf{k}$ $\mathcal{M}_\mathbf{k}(\mathcal{B}_2(r), \varepsilon) \lesssim_d r^d \log^\omega(\frac{1}{\varepsilon})$	$\sqrt{\frac{(\log n)^{\omega+1}}{n}}$	$\sqrt{\frac{(\log n)^{d+\omega+1}}{n}}$	$\sqrt{\frac{(\log n)^{2d+\omega+1}}{n}}$	$\sqrt{\frac{(\log n)^{\omega+1}}{n^{1-2d/\rho}}}$
POLYGROWTH $\mathcal{M}_\mathbf{k}$ $\mathcal{M}_\mathbf{k}(\mathcal{B}_2(r), \varepsilon) \lesssim_d r^d \varepsilon^{-\omega}$	$\sqrt{\frac{\log n}{n^{1-\omega}}}$	$\sqrt{\frac{(\log n)^{d+1}}{n^{1-\omega}}}$	$\sqrt{\frac{(\log n)^{2d+1}}{n^{1-\omega}}}$	$\sqrt{\frac{\log n}{n^{1-\omega-2d/\rho}}}$

Table 2: MMD guarantees for TARGET KT under $\mathcal{M}_\mathbf{k}$ (3) growth and \mathbb{P} tail decay. We report the $\text{MMD}_\mathbf{k}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}})$ bound (5) for target KT with n input points and \sqrt{n} output points, up to constants depending on d and $\|\mathbf{k}\|_1$. Here $\mathcal{R}_{\text{in}} \triangleq \max_{x \in \mathcal{S}_{\text{in}}} \|x\|_2$.

Kernel \mathbf{k}	TARGET KT	ROOT KT	KT+
GAUSS(σ)	$\frac{(\log n)^{\frac{3d}{4}+1}}{\sqrt{n} c_n^d}$	$\frac{(\log n)^{\frac{d}{4}+\frac{1}{2}} \sqrt{c_n}}{\sqrt{n}}$	$\frac{(\log n)^{\frac{d}{4}+\frac{1}{2}} \sqrt{c_n}}{\sqrt{n}}$
LAPLACE(σ)	$n^{-\frac{1}{4}}$	N/A	$\left(\frac{c_n(\log n)^{1+2d(1-\alpha)}}{n}\right)^{\frac{1}{4\alpha}}$
MATÉRN(ν, γ) $\nu \in (\frac{d}{2}, d]$	$n^{-\frac{1}{4}}$	N/A	$\left(\frac{c_n(\log n)^{1+2d(1-\alpha)}}{n}\right)^{\frac{1}{4\alpha}}$
MATÉRN(ν, γ) $\nu > d$	$\min(n^{-\frac{1}{4}}, \frac{(\log n)^{\frac{d+1}{2}}}{n^{(\nu-d)/(2\nu-d)}}$	$\frac{(\log n)^{\frac{d+1}{2}} \sqrt{c_n}}{\sqrt{n}}$	$\frac{(\log n)^{\frac{d+1}{2}} \sqrt{c_n}}{\sqrt{n}}$
IMQ(ν, γ) $\nu < \frac{d}{2}$	$\frac{(\log n)^{d+1}}{\sqrt{n}}$	$n^{-\frac{1}{4}}$	$\frac{(\log n)^{d+1}}{\sqrt{n}}$
IMQ(ν, γ) $\nu \geq \frac{d}{2}$	$\frac{(\log n)^{d+1}}{n}$	$\frac{(\log n)^{\frac{d+1}{2}} \sqrt{c_n}}{\sqrt{n}}$	$\frac{(\log n)^{\frac{d+1}{2}} \sqrt{c_n}}{\sqrt{n}}$
SINC(θ)	$\frac{(\log n)^2}{\sqrt{n}}$	$n^{-\frac{1}{4}}$	$\frac{(\log n)^2}{\sqrt{n}}$
B-SPLINE($2\beta + 1, \gamma$) $\beta \in 2\mathbb{N}$	$\sqrt{\frac{\log n}{n^{2\beta/(2\beta+1)}}}$	N/A	$\sqrt{\frac{\log n}{n}}$
B-SPLINE($2\beta + 1, \gamma$) $\beta \in 2\mathbb{N}_0 + 1$	$\sqrt{\frac{\log n}{n^{2\beta/(2\beta+1)}}}$	$\sqrt{\frac{\log n}{n}}$	$\sqrt{\frac{\log n}{n}}$

Table 3: $\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}; \mathcal{S}_{\text{KT}})$ guarantees for commonly used kernels. For n input points and \sqrt{n} thinned points, we report the MMD bounds from Thm. 2 for TARGET KT, from Dwivedi and Mackey (2021, Thm. 1) for ROOT KT, and from Thm. 4 for KT+ (with $\alpha = \frac{1}{2}$ wherever feasible). We assume a SUBGAUSS \mathbb{P} for the GAUSS kernel, a COMPACT \mathbb{P} for the B-SPLINE kernel, and a SUBEXP \mathbb{P} for all other kernels (see Tab. 2 for a definition of each \mathbb{P} class). Here, $c_n \triangleq \log \log n$, $\delta_i = \frac{\delta}{2n}$, $\delta^\ell = \frac{\delta}{2}$, and error is reported up to constants depending on $(\mathbf{k}, d, \delta, \alpha)$. The KT+ guarantee for LAPLACE applies with $\alpha > \frac{d}{d+1}$ and for MATÉRN with $\alpha > \frac{d}{2\nu}$. The TARGET KT guarantee for MATÉRN with $\nu > 3d/2$ assumes $\nu - d/2 \in \mathbb{N}$ to simplify the presentation (see (45) for a more explicit expression). The best rate is highlighted in blue.

Appendix C. Proof of Thm. 1: Single function guarantees for KT-SPLIT

The proof is identical for each index ℓ , so, without loss of generality, we prove the result for the case $\ell = 1$. Define

$$\widetilde{\mathcal{W}}_m \triangleq \mathcal{W}_{1,m} = \mathbb{P}_{\text{in}} \mathbf{k} - \mathbb{P}_{\text{out}}^{(1)} \mathbf{k} = \frac{1}{n} \sum_{x \in \mathcal{S}_{\text{in}}} \mathbf{k}(x, \cdot) - \frac{1}{n/2^m} \sum_{x \in \mathcal{S}^{(m,1)}} \mathbf{k}(x, \cdot).$$

Next, we use the results about an intermediate algorithm, kernel halving (Dwivedi and Mackey, 2021, Alg. 3) that was introduced for the analysis of kernel thinning. Using the arguments from Dwivedi and Mackey (2021, Sec. 5.2), we conclude that KT-SPLIT with $\mathbf{k}_{\text{split}}$ set as \mathbf{k} and thinning parameter m , is equivalent to repeated kernel halving with kernel \mathbf{k} for m rounds (with no Failure in any rounds of kernel halving). On this event of equivalence, denoted by $\mathcal{E}_{\text{equi}}$, Dwivedi and Mackey (2021, Eqns. (50, 51)) imply that the function $\widetilde{\mathcal{W}}_m \in \mathcal{H}$ is equal in distribution to another random function \mathcal{W}_m , where \mathcal{W}_m is unconditionally sub-Gaussian with parameter

$$\sigma_m = \frac{2}{\sqrt{3}} \frac{2^m}{n} \sqrt{\|\mathbf{k}\|_\infty \log\left(\frac{6m}{2^m \delta^*}\right)},$$

that is,

$$\mathbb{E}[\exp(\langle \mathcal{W}_m, f \rangle_{\mathbf{k}})] \leq \exp\left(\frac{1}{2}\sigma_m^2 \|f\|_{\mathbf{k}}^2\right) \quad \text{for all } f \in \mathcal{H}, \quad (9)$$

where we note that the analysis of [Dwivedi and Mackey \(2021\)](#) remains unaffected when we replace $\|\mathbf{k}\|_{\infty}$ by $\|\mathbf{k}\|_{\infty, \text{in}}$ in all the arguments. Applying the sub-Gaussian Hoeffding inequality ([Wainwright, 2019](#), Prop. 2.5) along with (9), we obtain that

$$\mathbb{P}[|\langle \mathcal{W}_m, f \rangle_{\mathbf{k}}| > t] \leq 2 \exp\left(-\frac{1}{2}t^2/(\sigma_m^2 \|f\|_{\mathbf{k}}^2)\right) \leq \delta' \text{ for } t \triangleq \sigma_m \|f\|_{\mathbf{k}} \sqrt{2 \log\left(\frac{2}{\delta'}\right)}.$$

Call this event \mathcal{E}_{sg} . As noted above, conditional to the event $\mathcal{E}_{\text{equi}}$, we also have

$$\mathcal{W}_m \stackrel{d}{=} \widetilde{\mathcal{W}}_m \implies \langle \mathcal{W}_m, f \rangle_{\mathbf{k}} \stackrel{d}{=} \mathbb{P}_{\text{in}} f - \mathbb{P}_{\text{out}}^{(1)} f,$$

where $\stackrel{d}{=}$ denotes equality in distribution. Furthermore, [Dwivedi and Mackey \(2021, Eqn. 48\)](#) implies that

$$\mathbb{P}(\mathcal{E}_{\text{equi}}) \geq 1 - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{n/2^j} \delta_i.$$

Putting the pieces together, we have

$$\mathbb{P}[|\mathbb{P}_{\text{in}} f - \mathbb{P}_{\text{out}}^{(1)} f| \leq t] \geq \mathbb{P}(\mathcal{E}_{\text{equi}} \cap \mathcal{E}_{\text{sg}}^c) \geq \mathbb{P}(\mathcal{E}_{\text{equi}}) - \mathbb{P}(\mathcal{E}_{\text{sg}}) \geq 1 - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{n/2^j} \delta_i - \delta' = p_{\text{sg}},$$

as claimed. The proof is now complete.

Appendix D. Proof of Thm. 2: MMD guarantee for TARGET KT

First, we note that by design, KT-SWAP ensures

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \leq \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}^{(m,1)}),$$

where $\mathcal{S}^{(m,1)}$ denotes the first coreset returned by KT-SPLIT. Thus it suffices to show that $\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}^{(m,1)})$ is bounded by the term stated on the right hand side of (4). Let $\mathbb{P}_{\text{out}}^{(1)} \triangleq \frac{1}{n/2^m} \sum_{x \in \mathcal{S}^{(m,1)}} x$. By design of KT-SPLIT, $\text{supp}(\mathbb{P}_{\text{out}}^{(1)}) \subseteq \text{supp}(\mathbb{P}_{\text{in}})$. Recall the set \mathcal{A} is such that $\text{supp}(\mathbb{P}_{\text{in}}) \subseteq \mathcal{A}$.

Proof of (4) Let $\mathcal{C} \triangleq \mathcal{C}_{\mathbf{k}, \varepsilon}(\mathcal{A})$ denote the cover of minimum cardinality satisfying (3). Fix any $f \in \mathcal{B}_{\mathbf{k}}$. By the triangle inequality and the covering property (3) of \mathcal{C} , we have

$$\begin{aligned} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})f \right| &\leq \inf_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(f - g) \right| + \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \\ &\leq \inf_{g \in \mathcal{C}} \left| \mathbb{P}_{\text{in}}(f - g) \right| + \left| \mathbb{P}_{\text{out}}^{(1)}(f - g) \right| + \sup_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \\ &\leq \inf_{g \in \mathcal{C}} 2 \sup_{x \in \mathcal{A}} |f(x) - g(x)| + \sup_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \\ &\leq 2\varepsilon + \sup_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right|. \end{aligned} \quad (10)$$

Applying Thm. 1, we have

$$\left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \leq \frac{2^m}{n} \|g\|_{\mathbf{k}} \sqrt{\frac{8}{3} \|\mathbf{k}\|_{\infty, \text{in}} \cdot \log\left(\frac{4}{\delta^*}\right) \log\left(\frac{4}{\delta'}\right)} \quad (11)$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^j - 1}{m} \sum_{i=1}^{n/2^j} \delta_i = p_{\text{sg}} - \delta'$. A standard union bound then yields that

$$\sup_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \leq \frac{2^m}{n} \sup_{g \in \mathcal{C}} \|g\|_{\mathbf{k}} \sqrt{\frac{8}{3} \|\mathbf{k}\|_{\infty, \text{in}} \cdot \log\left(\frac{4}{\delta^*}\right) [\log |\mathcal{C}| + \log\left(\frac{4}{\delta^0}\right)]}$$

probability at least $p_{\text{sg}} - \delta'$. Since $f \in \mathcal{B}_{\mathbf{k}}$ was arbitrary, and $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}}$ and thus $\sup_{g \in \mathcal{C}} \|g\|_{\mathbf{k}} \leq 1$, we therefore have

$$\begin{aligned} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}^{(m,1)}) &= \sup_{\|f\|_{\mathbf{k}} \leq 1} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})f \right| \stackrel{(10)}{\leq} 2\varepsilon + \sup_{g \in \mathcal{C}} \left| (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}^{(1)})(g) \right| \\ &\leq 2\varepsilon + \sqrt{\frac{8\|\mathbf{k}\|_{\infty, \text{in}}}{3}} \cdot \frac{2^m}{n} \sqrt{\log\left(\frac{4}{\delta^*}\right) [\log |\mathcal{C}| + \log\left(\frac{4}{\delta^0}\right)]}, \end{aligned}$$

with probability at least $p_{\text{sg}} - \delta'$ as claimed.

Appendix E. Proof of Thm. 3: MMD guarantee for POWER KT

Definition of \widetilde{M}_{α} and R_{max} Define the \mathbf{k}_{α} tail radii,

$$\begin{aligned} R_{\mathbf{k}_{\alpha}, n}^{\dagger} &\triangleq \min\{r : \tau_{\mathbf{k}_{\alpha}}(r) \leq \frac{\|\mathbf{k}_{\alpha}\|_{\infty}}{n}\}, \quad \text{where } \tau_{\mathbf{k}_{\alpha}}(R) \triangleq \left(\sup_x \int_{\|y\|_2 \geq R} \mathbf{k}_{\alpha}^2(x, x-y) dy\right)^{\frac{1}{2}}, \\ R_{\mathbf{k}_{\alpha}, n} &\triangleq \min\{r : \sup_{\|x-y\|_2 \geq r} |\mathbf{k}_{\alpha}(x, y)| \leq \frac{\|\mathbf{k}_{\alpha}\|_{\infty}}{n}\}, \end{aligned} \quad (12)$$

and the \mathcal{S}_{in} tail radii

$$R_{\mathcal{S}_{\text{in}}} \triangleq \max_{x \in \mathcal{S}_{\text{in}}} \|x\|_2, \quad \text{and} \quad R_{\mathcal{S}_{\text{in}}, \mathbf{k}_{\alpha}, n} \triangleq \min\left(R_{\mathcal{S}_{\text{in}}}, n^{1+\frac{1}{d}} R_{\mathbf{k}_{\alpha}, n} + n^{\frac{1}{d}} \|\mathbf{k}_{\alpha}\|_{\infty} / L_{\mathbf{k}_{\alpha}}\right). \quad (13)$$

Furthermore, define the inflation factor

$$M_{\mathbf{k}_{\alpha}}(n, m, d, \delta, \delta', R) \triangleq 37 \sqrt{\log\left(\frac{6m}{2^m \delta}\right)} \left[\sqrt{\log\left(\frac{4}{\delta^0}\right)} + 5 \sqrt{d \log\left(2 + 2 \frac{L_{\mathbf{k}_{\alpha}}}{\|\mathbf{k}_{\alpha}\|_{\infty}} (R_{\mathbf{k}_{\alpha}, n} + R)\right)} \right],$$

where $L_{\mathbf{k}_{\alpha}}$ denotes a Lipschitz constant satisfying $|\mathbf{k}_{\alpha}(x, y) - \mathbf{k}_{\alpha}(x, z)| \leq L_{\mathbf{k}_{\alpha}} \|y - z\|_2$ for all $x, y, z \in \mathbb{R}^d$. With the notations in place, we can define the quantities appearing in Thm. 3:

$$\widetilde{M}_{\alpha} \triangleq M_{\mathbf{k}_{\alpha}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_{\text{in}}, \mathbf{k}_{\alpha}, n}) \quad \text{and} \quad R_{\text{max}} \triangleq \max(R_{\mathcal{S}_{\text{in}}}, R_{\mathbf{k}_{\alpha}, n/2^m}^{\dagger}). \quad (14)$$

The scaling of these two parameters depends on the tail behavior of \mathbf{k}_{α} and the growth of the radii $R_{\mathcal{S}_{\text{in}}}$ (which in turn would typically depend on the tail behavior of \mathbb{P}). The scaling of \widetilde{M}_{α} and R_{max} stated in Thm. 3 under the compactly supported or subexponential tail conditions follows directly from Dwivedi and Mackey (2021, Tab. 2, App. I).

Proof of Thm. 3 The KT-SWAP step ensures that

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\alpha \text{KT}}) \leq \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\alpha}^{(m,1)}),$$

where $\mathcal{S}_{\alpha}^{(m,1)}$ denotes the first coreset output by KT-SPLIT with $\mathbf{k}_{\text{split}} = \mathbf{k}_{\alpha}$. Next, we state a key interpolation result for $\text{MMD}_{\mathbf{k}}$ that relates it to the MMD of its power kernels (Def. 2) (see App. G for the proof).

Proposition 1 (An interpolation result for MMD) *Consider a shift-invariant kernel \mathbf{k} that admits valid α and 2α -power kernels \mathbf{k}_α and $\mathbf{k}_{2\alpha}$ respectively for some $\alpha \in [\frac{1}{2}, 1]$. Then for any two discrete measures \mathbb{P} and \mathbb{Q} supported on finitely many points, we have*

$$\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{Q}) \leq (\text{MMD}_{\mathbf{k}_\alpha}(\mathbb{P}, \mathbb{Q}))^{2-\frac{1}{\alpha}} \cdot (\text{MMD}_{\mathbf{k}_{2\alpha}}(\mathbb{P}, \mathbb{Q}))^{\frac{1}{\alpha}-1}. \quad (15)$$

Given Prop. 1, it remains to establish suitable upper bounds on MMDs of \mathbf{k}_α and $\mathbf{k}_{2\alpha}$. To this end, first we note that for any reproducing kernel \mathbf{k} and any two distributions \mathbb{P} and \mathbb{Q} , Hölder's inequality implies that

$$\begin{aligned} \text{MMD}_{\mathbf{k}}^2(\mathbb{P}, \mathbb{Q}) &= \|(\mathbb{P} - \mathbb{Q})\mathbf{k}\|_{\mathbf{k}}^2 = (\mathbb{P} - \mathbb{Q})(\mathbb{P} - \mathbb{Q})\mathbf{k} \leq \|\mathbb{P} - \mathbb{Q}\|_1 \|(\mathbb{P} - \mathbb{Q})\mathbf{k}\|_\infty \\ &\leq 2\|(\mathbb{P} - \mathbb{Q})\mathbf{k}\|_\infty. \end{aligned}$$

Now, let \mathbb{P}_{in} and $\mathbb{P}_\alpha^{(m,1)}$ denote the empirical distributions of \mathcal{S}_{in} and $\mathcal{S}_\alpha^{(m,1)}$. Now applying [Dwivedi and Mackey \(2021, Thm. 4\(b\)\)](#), we find that

$$\text{MMD}_{\mathbf{k}_\alpha}(\mathcal{S}_{\text{in}}, \mathcal{S}_\alpha^{(m,1)}) \leq \sqrt{2\|(\mathbb{P}_{\text{in}} - \mathbb{P}_\alpha^{(m,1)})\mathbf{k}_\alpha\|_{\infty, \text{in}}} \leq \sqrt{2 \cdot \frac{2^m}{n} \|\mathbf{k}_\alpha\|_{\infty, \text{in}} \widetilde{\mathcal{M}}_{\mathbf{k}_\alpha}} \quad (16)$$

with probability $p_{\text{sg}} - \delta'$, where $\widetilde{\mathcal{M}}_{\mathbf{k}_\alpha}$ was defined in (14). We note that while [Dwivedi and Mackey \(2021, Thm. 4\(b\)\)](#) uses $\|\mathbf{k}_\alpha\|_\infty$ in their bounds, we can replace it by $\|\mathbf{k}_\alpha\|_{\infty, \text{in}}$, and verifying that all the steps of the proof continue to be valid (noting that $\|\mathbf{k}_\alpha\|_{\infty, \text{in}}$ is deterministic given \mathcal{S}_{in}). Furthermore, [Dwivedi and Mackey \(2021, Thm. 4\(b\)\)](#) yields that

$$\text{MMD}_{\mathbf{k}_{2\alpha}}(\mathcal{S}_{\text{in}}, \mathcal{S}_\alpha^{(m,1)}) \leq \frac{2^m}{n} \|\mathbf{k}_\alpha\|_{\infty, \text{in}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} \cdot \mathbf{R}_{\text{max}}^{\frac{d}{2}} \cdot \widetilde{\mathcal{M}}_\alpha \right), \quad (17)$$

with probability $p_{\text{sg}} - \delta'$, where we have once again replaced the term $\|\mathbf{k}_\alpha\|_\infty$ with $\|\mathbf{k}_\alpha\|_{\infty, \text{in}}$ for the same reasons as stated above. We note that the two bounds (16) and (17) apply under the same high probability event as noted in [Dwivedi and Mackey \(2021, proof of Thm. 1, eqn. \(18\)\)](#). Putting together the pieces, we find that

$$\begin{aligned} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_\alpha^{(m,1)}) &\stackrel{(15)}{\leq} (\text{MMD}_{\mathbf{k}_\alpha}(\mathcal{S}_{\text{in}}, \mathcal{S}_\alpha^{(m,1)})^{2-\frac{1}{\alpha}} \cdot (\text{MMD}_{\mathbf{k}_{2\alpha}}(\mathcal{S}_{\text{in}}, \mathcal{S}_\alpha^{(m,1)}))^{\frac{1}{\alpha}-1} \\ &\stackrel{(16,17)}{\leq} \left[2 \cdot \frac{2^m}{n} \|\mathbf{k}_\alpha\|_{\infty, \text{in}} \widetilde{\mathcal{M}}_\alpha \right]^{1-\frac{1}{2\alpha}} \left[\frac{2^m}{n} \|\mathbf{k}_\alpha\|_{\infty, \text{in}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} \cdot \mathbf{R}_{\text{max}}^{\frac{d}{2}} \cdot \widetilde{\mathcal{M}}_\alpha \right) \right]^{\frac{1}{\alpha}-1} \\ &= \left(\frac{2^m}{n} \|\mathbf{k}_\alpha\|_{\infty, \text{in}} \right)^{\frac{1}{2\alpha}} (2 \cdot \widetilde{\mathcal{M}}_\alpha)^{1-\frac{1}{2\alpha}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} \cdot \mathbf{R}_{\text{max}}^{\frac{d}{2}} \cdot \widetilde{\mathcal{M}}_\alpha \right)^{\frac{1}{\alpha}-1}, \end{aligned}$$

as claimed. The proof is now complete.

Appendix F. Proof of Thm. 4: Single function & MMD guarantees for KT+

Proof of (7) First, we note that the RKHS \mathcal{H} of \mathbf{k} is contained in the RKHS \mathcal{H}^\dagger of \mathbf{k}^\dagger [Berlinet and Thomas-Agnan \(2011, Thm. 5\)](#). Now, applying Thm. 1 with $\mathbf{k}_{\text{split}} = \mathbf{k}^\dagger$ for any fixed function

$f \in \mathcal{H} \subset \mathcal{H}^\dagger$ and $\delta' \in (0, 1)$, we obtain that

$$\begin{aligned} \left| \mathbb{P}_{\text{in}} f - \mathbb{P}_{\text{split}}^{(\ell)} f \right| &\leq \|f\|_{\mathbf{k}^\vee} \cdot \frac{2}{\sqrt{3}} \frac{2^m}{n} \sqrt{\|\mathbf{k}^\dagger\|_{\infty, \text{in}} \cdot \log\left(\frac{6m}{2^m \delta^*}\right)} \sqrt{2 \log\left(\frac{2}{\delta'}\right)} \\ &\stackrel{(i)}{\leq} \|f\|_{\mathbf{k}^\vee} \cdot \frac{2^m}{n} \sqrt{\frac{16}{3} \log\left(\frac{6m}{2^m \delta^*}\right) \log\left(\frac{2}{\delta'}\right)}, \\ &\stackrel{(ii)}{\leq} \|f\|_{\mathbf{k}} \cdot \frac{2^m}{n} \sqrt{\frac{16}{3} \|\mathbf{k}\|_{\infty} \log\left(\frac{6m}{2^m \delta^*}\right) \log\left(\frac{2}{\delta'}\right)}, \end{aligned}$$

with probability at least p_{sg} . Here step (i) follows from the inequality $\|\mathbf{k}^\dagger\|_{\infty} \leq 2$, and step (ii) follows from the inequality $\|f\|_{\mathbf{k}^\vee} \leq \sqrt{\|\mathbf{k}\|_{\infty}} \|f\|_{\mathbf{k}}$, which in turn follows from the standard facts that

$$\|f\|_{\lambda \mathbf{k}} \stackrel{(iii)}{=} \frac{\|f\|_{\mathbf{k}}}{\sqrt{\lambda}}, \quad \text{and} \quad \|f\|_{\lambda \mathbf{k} + \mathbf{k}_\alpha} \stackrel{(iv)}{\leq} \|f\|_{\lambda \mathbf{k}} \quad \text{for } \lambda > 0, f \in \mathcal{H},$$

see, e.g., [Zhang and Zhao \(2013, Proof of Prop. 2.5\)](#) for a proof of step (iii), [Berlinet and Thomas-Agnan \(2011, Thm. 5\)](#) for step (iv). The proof for the bound (7) is now complete. \square

Proof of (8) Repeating the proof of Thm. 2 with the bound (11) replaced by (7) yields that

$$\begin{aligned} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}+}) &\leq \inf_{\varepsilon, \mathcal{S}_{\text{in}} \subset \mathcal{A}} 2\varepsilon + \frac{2^m}{n} \sqrt{\frac{16}{3} \|\mathbf{k}\|_{\infty} \log\left(\frac{6m}{2^m \delta^*}\right) \cdot \left[\log\left(\frac{4}{\delta'}\right) + \mathcal{M}_{\mathbf{k}}(\mathcal{A}, \varepsilon)\right]} \\ &\leq \sqrt{2} \cdot \overline{\text{M}}_{\text{targetKT}}(\mathbf{k}) \end{aligned} \tag{18}$$

with probability at least p_{sg} . Let us denote this event by \mathcal{E}_1 .

To establish the other bound, first we note that KT-SWAP step ensures that

$$\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}+}) \leq \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}+}^{(m,1)}), \tag{19}$$

where $\mathcal{S}_{\text{KT}+}^{(m,1)}$ denotes the first coresnet output by KT-SPLIT with $\mathbf{k}_{\text{split}} = \mathbf{k}^\dagger$. We can now repeat the proof of Thm. 3, using the sub-Gaussian tail bound (7), and with a minor substitution, namely, $\|\mathbf{k}_\alpha\|_{\infty, \text{in}}$ replaced by $2\|\mathbf{k}_\alpha\|_{\infty}$. Putting it together with (19), we conclude that

$$\begin{aligned} \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}+}) &\leq \left(\frac{2^m}{n} 2\|\mathbf{k}_\alpha\|_{\infty, \text{in}}\right)^{\frac{1}{2\alpha}} (2\widetilde{\text{M}}_\alpha)^{1-\frac{1}{2\alpha}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} \cdot \mathbf{R}_{\text{max}}^{\frac{d}{2}} \cdot \widetilde{\text{M}}_\alpha\right)^{\frac{1}{\alpha}-1} \\ &= 2^{\frac{1}{2\alpha}} \cdot \overline{\text{M}}_{\text{powerKT}}(\mathbf{k}_\alpha), \end{aligned} \tag{20}$$

with probability at least p_{sg} . Let us denote this event by \mathcal{E}_2 .

Note that the quantities on the right hand side of the bounds (18) and (20) are deterministic given \mathcal{S}_{in} , and thus can be computed apriori. Consequently, we apply the high probability bound only for one of the two events \mathcal{E}_1 or \mathcal{E}_2 depending on which of the two quantities (deterministically) attains the minimum. Thus, the bound (8) holds with probability at least p_{sg} as claimed. \square

Appendix G. Proof of Prop. 1: An interpolation result for MMD

For two arbitrary distributions \mathbb{P} and \mathbb{Q} , and any reproducing kernel \mathbf{k} , [Gretton et al. \(2012, Lem. 4\)](#) yields that

$$\text{MMD}_{\mathbf{k}}^2(\mathbb{P}, \mathbb{Q}) = \|(\mathbb{P} - \mathbb{Q})\mathbf{k}\|_{\mathbf{k}}^2. \quad (21)$$

Let \mathcal{F} denote the generalized Fourier transform (GFT) operator ([Wendland \(2004, Def. 8.9\)](#)). Since $\mathbf{k}(x, y) = \kappa(x - y)$, [Wendland \(2004, Thm. 10.21\)](#) yields that

$$\|f\|_{\mathbf{k}}^2 = \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \frac{(\mathcal{F}(f)(\omega))^2}{\mathcal{F}(\kappa)(\omega)} d\omega, \quad \text{for } f \in \mathcal{H}. \quad (22)$$

Let $\widehat{\kappa} \triangleq \mathcal{F}(\kappa)$, and consider a discrete measure $\mathbb{D} = \sum_{i=1}^n w_i \delta_{x_i}$ supported on finitely many points, and let $\mathbb{D}\mathbf{k}(x) \triangleq \sum w_i \mathbf{k}(x, x_i) = \sum w_i \kappa(x - x_i)$. Now using the linearity of the GFT operator \mathcal{F} , we find that for any $\omega \in \mathbb{R}^d$,

$$\begin{aligned} \mathcal{F}(\mathbb{D}\mathbf{k})(\omega) &= \mathcal{F}\left(\sum_{i=1}^n w_i \kappa(\cdot - x_i)\right) = \sum_{i=1}^n w_i \mathcal{F}(\kappa(\cdot - x_i)) = \left(\sum_{i=1}^n w_i e^{-\langle \omega, x_i \rangle}\right) \cdot \widehat{\kappa}(\omega) \\ &= \widehat{D}(\omega) \widehat{\kappa}(\omega) \end{aligned} \quad (23)$$

where we used the time-shifting property of GFT that $\mathcal{F}(\kappa(\cdot - x_i))(\omega) = e^{-\langle \omega, x_i \rangle} \widehat{\kappa}(\omega)$ (proven for completeness in [Lem. 1](#)), and used the shorthand $\widehat{D}(\omega) \triangleq \left(\sum_{i=1}^n w_i e^{-\langle \omega, x_i \rangle}\right)$ in the last step. Putting together (21) to (23) with $\mathbb{D} = \mathbb{P} - \mathbb{Q}$, we find that

$$\begin{aligned} \text{MMD}_{\mathbf{k}}^2(\mathbb{P}, \mathbb{Q}) &= \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \widehat{D}^2(\omega) \widehat{\kappa}(\omega) d\omega \quad (24) \\ &= \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \widehat{D}^2(\omega) \widehat{\kappa}^\alpha(\omega) (\widehat{\kappa}^\alpha(\omega))^{-\frac{1}{\alpha}} d\omega \\ &= \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \widehat{D}^2(\omega') \widehat{\kappa}^\alpha(\omega') d\omega' \int_{\mathbb{R}^d} \frac{\widehat{D}^2(\omega) \widehat{\kappa}^\alpha(\omega)}{\int_{\mathbb{R}^d} \widehat{D}^2(\omega^\theta) \widehat{\kappa}^\alpha(\omega^\theta) d\omega^\theta} (\widehat{\kappa}^\alpha(\omega))^{-\frac{1}{\alpha}} d\omega \\ &\stackrel{(i)}{\leq} \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \widehat{D}^2(\omega') \widehat{\kappa}^\alpha(\omega') d\omega' \left(\int_{\mathbb{R}^d} \frac{\widehat{D}^2(\omega) \widehat{\kappa}^\alpha(\omega)}{\int_{\mathbb{R}^d} \widehat{D}^2(\omega^\theta) \widehat{\kappa}^\alpha(\omega^\theta) d\omega^\theta} \widehat{\kappa}^\alpha(\omega) d\omega \right)^{\frac{1}{\alpha}} \\ &= \frac{1}{(2\pi)^{d/2}} \left(\int_{\mathbb{R}^d} \widehat{D}^2(\omega') \widehat{\kappa}^\alpha(\omega') d\omega' \right)^{2-\frac{1}{\alpha}} \left(\int_{\mathbb{R}^d} \frac{\widehat{D}^2(\omega) \widehat{\kappa}^{2\alpha}(\omega)}{d} \omega \right)^{\frac{1}{\alpha}} \\ &= \left(\frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \widehat{D}^2(\omega') \widehat{\kappa}^\alpha(\omega') d\omega' \right)^{2-\frac{1}{\alpha}} \left(\frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} \frac{\widehat{D}^2(\omega) \widehat{\kappa}^{2\alpha}(\omega)}{d} \omega \right)^{\frac{1}{\alpha}} \\ &\stackrel{(ii)}{=} (\text{MMD}_{\mathbf{k}_\alpha}^2(\mathbb{P}, \mathbb{Q}))^{2-\frac{1}{\alpha}} \cdot (\text{MMD}_{\mathbf{k}_{2\alpha}}^2(\mathbb{P}, \mathbb{Q}))^{\frac{1}{\alpha}-1}, \end{aligned}$$

where step (i) makes use of Jensen's inequality and the fact that the function $t \mapsto t^{-\frac{1}{\alpha}}$ for $t \geq 0$ is concave for $\alpha \in [\frac{1}{2}, 1]$, and step (ii) follows by applying (24) for kernels \mathbf{k}_α and $\mathbf{k}_{2\alpha}$ and noting that by definition $\mathcal{F}(\mathbf{k}_\alpha) = \widehat{\kappa}^\alpha$, and $\mathcal{F}(\mathbf{k}_{2\alpha}) = \widehat{\kappa}^{2\alpha}$. Noting MMD is a non-negative quantity, and taking square-root establishes the claim (15).

Lemma 1 (Shifting property of the generalized Fourier transform) *If $\widehat{\kappa}$ denotes the generalized Fourier transform (GFT) ([Wendland, 2004, Def. 8.9](#)) of the function $\kappa : \mathbb{R}^d \rightarrow \mathbb{R}$, then $e^{-\langle \cdot, x_i \rangle} \widehat{\kappa}$ denotes the GFT of the shifted function $\kappa(\cdot - x_i)$, for any $x_i \in \mathbb{R}^d$.*

Proof Note that by definition of the GFT $\widehat{\kappa}$ (Wendland, 2004, Def. 8.9), we have

$$\int \kappa(x)\widehat{\gamma}(x)dx = \int \widehat{\kappa}(\omega)\gamma(\omega)d\omega, \quad (25)$$

for all suitable Schwartz functions γ (Wendland, 2004, Def. 5.17), where $\widehat{\gamma}$ denotes the Fourier transform (Wendland, 2004, Def. 5.15) of γ since GFT and FT coincide for these functions (as noted in the discussion after Wendland (2004, Def. 8.9)). Thus to prove the lemma, we need to verify that

$$\int \kappa(x - x_i)\widehat{\gamma}(x)dx = \int e^{-\langle \omega, x_i \rangle} \widehat{\kappa}(\omega)\gamma(\omega)d\omega, \quad (26)$$

for all suitable Schwartz functions γ . Starting with the right hand side of the display (26), we have

$$\int e^{-\langle \omega, x_i \rangle} \widehat{\kappa}(\omega)\gamma(\omega)d\omega = \int \widehat{\kappa}(\omega)(e^{-\langle \omega, x_i \rangle} \gamma(\omega))d\omega \stackrel{(i)}{=} \int \kappa(x)\widehat{\gamma}(x + x_i)dx \stackrel{(ii)}{=} \int \kappa(z - x_i)\widehat{\gamma}(z)dz,$$

where step (i) follows from the shifting property of the FT (Wendland, 2004, Thm. 5.16(4)), and the fact that the GFT condition (25) holds for the shifted function $\gamma(\cdot + x_i)$ function as well since it is still a Schwartz function (recall that $\widehat{\gamma}$ is the FT), and step (ii) follows from a change of variable. We have thus established (26), and the proof is complete. \blacksquare

Appendix H. Additional experimental results

This section provides additional experimental details and results deferred from Sec. 4.

Details for target distributions and kernels We consider three classes of target distributions on \mathbb{R}^d : (i) mixture of Gaussians (MoG) $\mathbb{P} = \frac{1}{M} \sum_{j=1}^M \mathcal{N}(\mu_j, \mathbf{I}_2)$ with M component means $\mu_j \in \mathbb{R}^2$ defined in (27), (ii) Gaussian $\mathbb{P} = \mathcal{N}(0, \mathbf{I}_d)$ for $d \in \{2, 10, 20, 50, 100\}$, and (iii) the posteriors of four distinct coupled ordinary differential equation models: the *Goodwin (1965) model* of oscillatory enzymatic control ($d = 4$), the *Lotka (1925) model* of oscillatory predator-prey evolution ($d = 4$), the *Hinch et al. (2004) model* of calcium signalling in cardiac cells ($d = 38$), and a tempered Hinch posterior. For settings (i) and (ii), we use an i.i.d. input sequence \mathcal{S}_{in} from \mathbb{P} and kernel bandwidths $\sigma = 1/\gamma = \sqrt{2d}$. For setting (iii), we use MCMC input sequences \mathcal{S}_{in} from 12 posterior inference experiments of Riabiz et al. (2020a) and set the bandwidths $\sigma = 1/\gamma$ as specified by Dwivedi and Mackey (2021, Sec. K.2).

Details of test functions For f_{CIF} , the shift parameters u_j are drawn i.i.d. and uniformly from $[0, 1]$ (once and then fixed across all experiments). Furthermore, we note the following: (a) For Gaussian targets, the error with CIF function and i.i.d. input is measured across the sample mean over the n input points and \sqrt{n} output points obtained by standard thinning the input sequence, since $\mathbb{P}f_{\text{CIF}}$ does not admit a closed form. (b) To define the function $f : x \mapsto \mathbf{k}(X', x)$, first we draw a sample $X \sim \mathbb{P}$, independent of the input, and then set $X' = 2X$. For the MCMC targets, we draw a point uniformly from a held out data point not used as input for KT. For each target, the sample is drawn exactly once and then fixed throughout all sample sizes and repetitions.

Common settings and error computation To obtain an output coreset of size $n^{\frac{1}{2}}$ with n input points, we (a) take every $n^{\frac{1}{2}}$ -th point for standard thinning (ST) and (b) run KT with $m = \frac{1}{2} \log_2 n$ using an ST coreset as the base coreset in KT-SWAP. For Gaussian and MoG target we use i.i.d. points as input, and for MCMC targets we use an ST coreset after burn-in as the input. We compute errors with respect to \mathbb{P} whenever available in closed form and otherwise use \mathbb{P}_{in} . For each input sample size $n \in \{2^4, 2^6, \dots, 2^{14}\}$ with $\delta_i = \frac{1}{2n}$, we report the mean MMD or function integration error ± 1 standard error across 10 independent replications of the experiment (the standard errors are too small to be visible in all experiments). We also plot the ordinary least squares fit (for log mean error vs log coreset size), with the slope of the fit denoted as the empirical decay rate, e.g., for an OLS fit with slope -0.25 , we display the decay rate of $n^{-0.25}$.

H.1. TARGET KT vs. i.i.d. sampling

For Gaussian \mathbb{P} and Gaussian \mathbf{k} , Fig. 2 quantifies the improvements in distributional approximation obtained when using TARGET KT in place of a more typical i.i.d. summary. Remarkably, TARGET KT significantly improves the rate of decay and order of magnitude of mean $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{P}_{\text{out}})$, even in $d = 100$ dimensions with as few as 4 output points. Moreover, in line with our theory, TARGET KT MMD closely tracks that of ROOT KT without using \mathbf{k}_{r} . Finally, TARGET KT delivers improved single-function integration error, both of functions in the RKHS (like $\mathbf{k}(X', \cdot)$) and those outside (like the first moment and CIF benchmark function), even with large d and relatively small sample sizes.

H.2. Mixture of Gaussians Experiments

Our mixture of Gaussians target is given by $\mathbb{P} = \frac{1}{M} \sum_{j=1}^M \mathcal{N}(\mu_j, \mathbf{I}_d)$ for $M \in \{4, 6, 8\}$ where

$$\begin{aligned} \mu_1 &= [-3, 3]^\top, & \mu_2 &= [-3, 3]^\top, & \mu_3 &= [-3, -3]^\top, & \mu_4 &= [3, -3]^\top, \\ \mu_5 &= [0, 6]^\top, & \mu_6 &= [-6, 0]^\top, & \mu_7 &= [6, 0]^\top, & \mu_8 &= [0, -6]^\top. \end{aligned} \quad (27)$$

For an 8-component mixture of Gaussians target \mathbb{P} , the top row of Fig. 3 highlights the visual differences between i.i.d. coresets and coresets generated using generalized KT. We consider ROOT KT with GAUSS \mathbf{k} , TARGET KT with GAUSS \mathbf{k} , and KT+ ($\alpha = 0.7$) with LAPLACE \mathbf{k} , KT+ ($\alpha = \frac{1}{2}$) with IMQ \mathbf{k} ($\nu = 0.5$), and KT+ ($\alpha = \frac{2}{3}$) with B-SPLINE(5) \mathbf{k} , and note that the B-SPLINE(5) and LAPLACE \mathbf{k} do not admit square-root kernels. In each case, even for small n , generalized KT provides a more even distribution of points across components with fewer within-component gaps and clumps. Moreover, as suggested by our theory, TARGET KT and ROOT KT coresets for GAUSS \mathbf{k} have similar quality despite TARGET KT making no explicit use of a square-root kernel. The MMD error plots in the bottom row of Fig. 3 provide a similar conclusion quantitatively, where we observe that for both variants of KT, the MMD error decays as $n^{-\frac{1}{2}}$, a significant improvement over the $n^{-\frac{1}{4}}$ rate of i.i.d. sampling. We also observe that the empirical MMD decay rates are in close agreement with the rates guaranteed by our theory in Tab. 3 ($n^{-\frac{1}{2}}$ for GAUSS, B-SPLINE, and IMQ and $n^{-\frac{1}{4\alpha}} = n^{-0.36}$ for LAPLACE).

Finally, we display mean MMD (± 1 standard error across ten independent experiment replicates) as a function of coreset size in Fig. 4 for $M = 4, 6$ component MoG targets. The conclusions from Fig. 4 are identical to those from the bottom row of Fig. 3: TARGET KT and ROOT KT provide similar MMD errors with GAUSS \mathbf{k} , and all variants of KT provide a significant improvement over

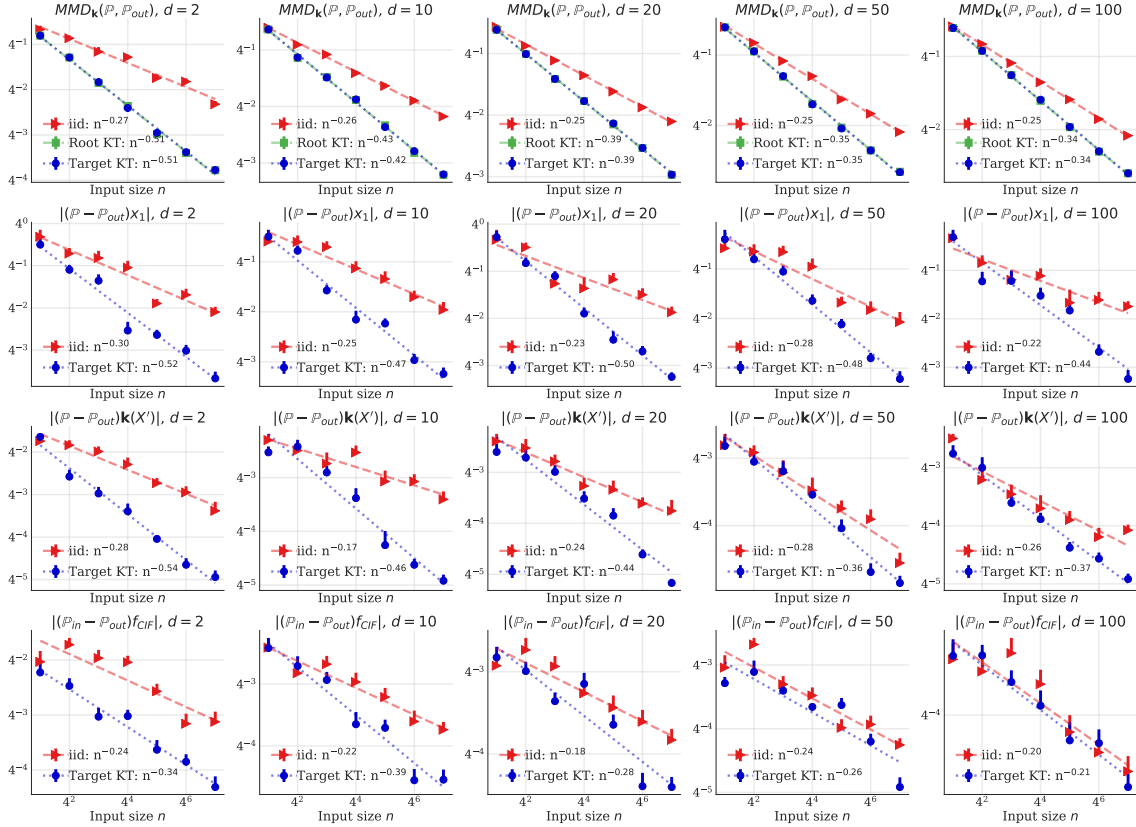


Figure 2: MMD and single-function integration error for Gaussian k and standard Gaussian \mathbb{P} in \mathbb{R}^d . Without using a square-root kernel, TARGET KT matches the MMD performance of ROOT KT and improves upon i.i.d. MMD and single-function integration error, even in $d = 100$ dimensions.

i.i.d. sampling both in terms of magnitude and decay rate with input size. Moreover the observed decay rates for KT+ closely match the rates guaranteed by our theory in Tab. 3.

H.3. MCMC experiments

Our set-up for MCMC experiments follows closely that of Dwivedi and Mackey (2021). For complete details on the targets and sampling algorithms we refer the reader to Riabiz et al. (2020a, Sec. 4).

Goodwin and Lotka-Volterra experiments From Riabiz et al. (2020b), we use the output of four distinct MCMC procedures targeting each of two $d = 4$ -dimensional posterior distributions \mathbb{P} : (1) a posterior over the parameters of the *Goodwin model* of oscillatory enzymatic control (Goodwin, 1965) and (2) a posterior over the parameters of the *Lotka-Volterra model* of oscillatory predator-prey evolution (Lotka, 1925; Volterra, 1926). For each of these targets, Riabiz et al. (2020b) provide 2×10^6 sample points from the following four MCMC algorithms: Gaussian random walk (RW), adaptive Gaussian random walk (adaRW, Haario et al., 1999), Metropolis-adjusted Langevin algorithm (MALA, Roberts and Tweedie, 1996), and pre-conditioned MALA (pMALA, Girolami and Calderhead, 2011).

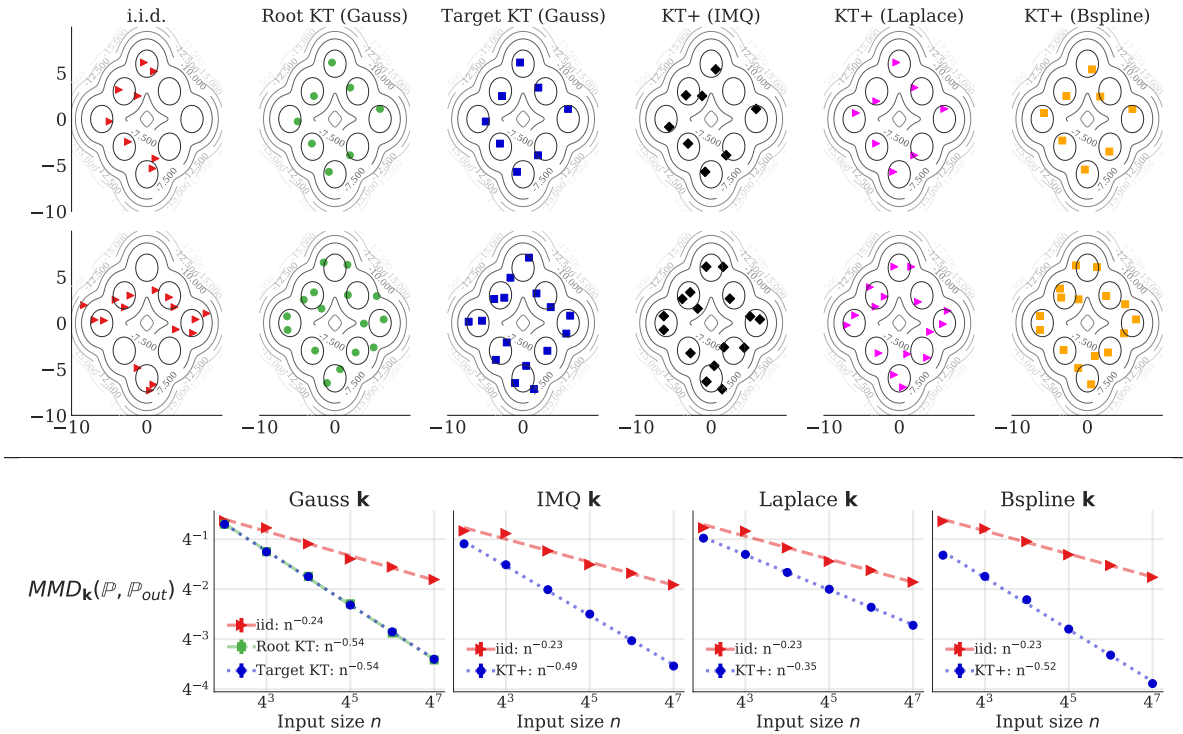


Figure 3: Generalized kernel thinning (KT) vs i.i.d. sampling for various kernels k and an 8-component mixture of Gaussian target \mathbb{P} : Scatter plots of various coresets with equidensity contours overlaid in the top row and the MMD errors for the respective coresets versus the input size on the bottom row.

Hinch experiments [Riabiz et al. \(2020b\)](#) also provide the output of two independent Gaussian random walk MCMC chains targeting each of two $d = 38$ -dimensional posterior distributions \mathbb{P} : (1) a posterior over the parameters of the Hinch model of calcium signalling in cardiac cells ([Hinch et al., 2004](#)) and (2) a tempered version of the same posterior, as defined by [Riabiz et al. \(2020a, App. S5.4\)](#).

In total we run 12 MCMC experiments. While [Fig. 1](#) provided results for 3 of these experiments, the results for the other 9 MCMC settings are available in [Fig. 5](#).

Burn-in and standard thinning We discard the initial burn-in points of each chain using the maximum burn-in period reported in [Riabiz et al. \(2020a, Tabs. S4 & S6, App. S5.4\)](#). Furthermore, we also normalize each Hinch chain by subtracting the post-burn-in sample mean and dividing each coordinate by its post-burn-in sample standard deviation. To obtain an input sequence \mathcal{S}_{in} of length n to be fed into a thinning algorithm, we downsample the remaining even indices of points using standard thinning (odd indices are held out). When applying standard thinning to any Markov chain output, we adopt the convention of keeping the final sample point.

The selected burn-in periods for the Goodwin task were 820,000 for RW; 824,000 for adaRW; 1,615,000 for MALA; and 1,475,000 for pMALA. The respective numbers for the Lotka-Volterra task were 1,512,000 for RW; 1,797,000 for adaRW; 1,573,000 for MALA; and 1,251,000 for pMALA.

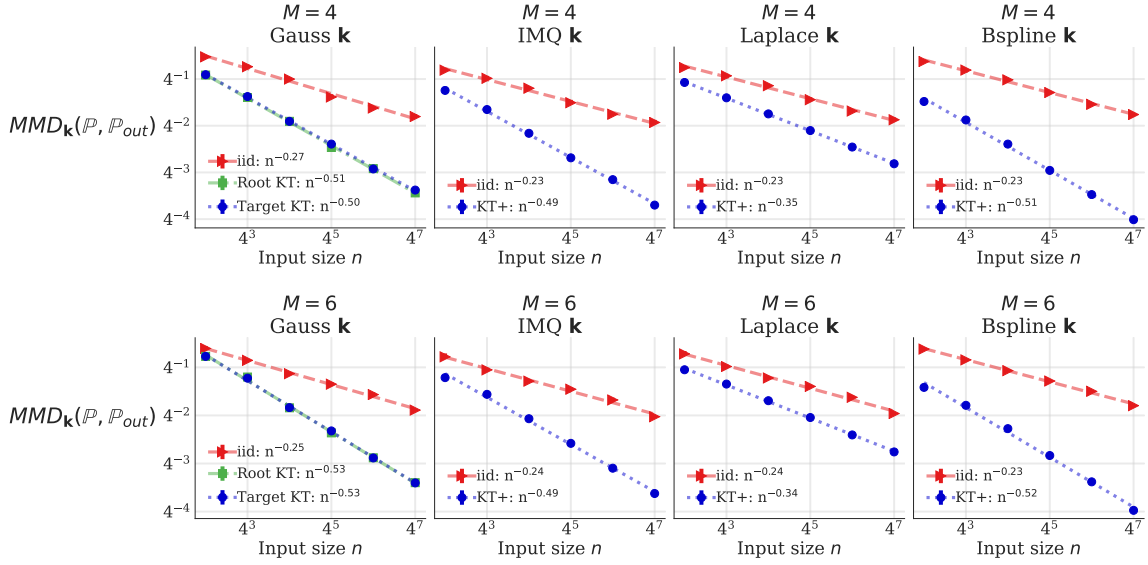


Figure 4: Kernel thinning versus i.i.d. sampling. For mixture of Gaussians \mathbb{P} with $M \in \{4, 6\}$ components and the kernel choices of Sec. 4, the TARGET KT with GAUSS \mathbf{k} provides comparable $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{P}_{\text{out}})$ error to the ROOT KT, and both provide an $n^{-\frac{1}{2}}$ decay rate improving significantly over the $n^{-\frac{1}{4}}$ decay rate from i.i.d. sampling. For the other kernels, KT+ provides a decay rate close to $n^{-\frac{1}{2}}$ for IMQ and B-SPLINE \mathbf{k} , and $n^{-0.35}$ for LAPLACE \mathbf{k} , providing an excellent agreement with the MMD guarantees provided by our theory. See Sec. 4 for further discussion.

Appendix I. Upper bounds on RKHS covering numbers

In this section, we state several results on covering bounds for RKHSes for both generic and specific kernels. We then use these bounds with Thm. 2 (or Tab. 2) to establish MMD guarantees for the output of generalized kernel thinning as summarized in Tab. 3.

We first state covering number bounds for RKHS associated with generic kernels, that are either (a) analytic, or (b) finitely many times differentiable. These results follow essentially from Sun and Zhou (2008); Steinwart and Christmann (2008), but we provide a proof in App. I.2 for completeness.

Proposition 2 (Covering numbers for analytic and differentiable kernels) *The following results hold true.*

- (a) **Analytic kernels:** Suppose that $\mathbf{k}(x, y) = \kappa(\|x - y\|_2^2)$ for $\kappa : \mathbb{R}_+ \rightarrow \mathbb{R}$ real-analytic with convergence radius R_κ , that is,

$$\left| \frac{1}{j!} \kappa_+^{(j)}(0) \right| \leq C_\kappa (2/R_\kappa)^j \quad \text{for all } j \in \mathbb{N}_0 \quad (28)$$

for some constant C_κ , where $\kappa_+^{(j)}$ denotes the right-sided j -th derivative of κ . Then for any set $\mathcal{A} \subset \mathbb{R}^d$ and any $\varepsilon \in (0, \frac{1}{2})$, we have

$$\mathcal{M}_{\mathbf{k}}(\mathcal{A}, \varepsilon) \leq \mathcal{N}_2(\mathcal{A}, r^\dagger/2) \cdot (4 \log(1/\varepsilon) + 2 + 4 \log(16\sqrt{C_\kappa} + 1))^{d+1}, \quad (29)$$

$$\text{where } r^\dagger \triangleq \min\left(\frac{\sqrt{R_\kappa}}{2^d}, \sqrt{R_\kappa + D_{\mathcal{A}}^2} - D_{\mathcal{A}}\right), \text{ and } D_{\mathcal{A}} \triangleq \max_{x, y \in \mathcal{A}} \|x - y\|_2. \quad (30)$$

(b) **Differentiable kernels:** Suppose that for $\mathcal{X} \subset \mathbb{R}^d$, the kernel $\mathbf{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is s -times continuously differentiable, i.e., all partial derivatives $\partial^{\alpha, \alpha} \mathbf{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ exist and are continuous for all multi-indices $\alpha \in \mathbb{N}_0^d$ with $|\alpha| \leq s$. Then, for any closed Euclidean ball $\bar{\mathcal{B}}_2(r)$ contained in \mathcal{X} and any $\varepsilon > 0$, we have

$$\mathcal{M}_{\mathbf{k}}(\bar{\mathcal{B}}_2(r), \varepsilon) \leq c_{s,d,\mathbf{k}} \cdot r^d \cdot (1/\varepsilon)^{d/s}, \quad (31)$$

for some constant $c_{s,d,\mathbf{k}}$ that depends only on s, d and \mathbf{k} .

Next, we state several explicit bounds on covering numbers for several popular kernels. See App. I.3 for the proof.

Proposition 3 (Covering numbers for specific kernels) *The following statements hold true.*

(a) When $\mathbf{k} = \text{GAUSS}(\sigma)$, we have

$$\mathcal{M}_{\mathbf{k}}(\mathcal{B}_2(r), \varepsilon) \leq C_{\text{GAUSS},d} \cdot \left(\frac{\log(4/\varepsilon)}{\log \log(4/\varepsilon)} \right)^d \log(1/\varepsilon) \cdot \begin{cases} 1 & \text{when } r \leq \frac{1}{\sqrt{2}\sigma}, \\ (3\sqrt{2}r\sigma)^d & \text{otherwise,} \end{cases} \quad (32)$$

$$\text{where } C_{\text{GAUSS},d} \triangleq \binom{4e+d}{d} e^{-d} \leq \begin{cases} 4.3679 & \text{for } d = 1 \\ 0.05 \cdot d^{4e} e^{-d} & \text{for } d \geq 2 \end{cases} \leq 30 \text{ for all } d \geq 1. \quad (33)$$

(b) When $\mathbf{k} = \text{MATÉRN}(\nu, \gamma)$, $\nu \geq \frac{d}{2} + 1$, then for some constant $C_{\text{MATÉRN},\nu,\gamma,d}$, we have

$$\mathcal{M}_{\mathbf{k}}(\mathcal{B}_2(r), \varepsilon) \leq C_{\text{MATÉRN},\nu,\gamma,d} \cdot r^d \cdot (1/\varepsilon)^{d/\lfloor \nu - \frac{d}{2} \rfloor}. \quad (34)$$

(c) When $\mathbf{k} = \text{IMQ}(\nu, \gamma)$, we have

$$\mathcal{M}_{\mathbf{k}}(\mathcal{B}_2(r), \varepsilon) \leq \left(1 + \frac{4r}{\tilde{r}}\right)^d \cdot (4 \log(1/\varepsilon) + 2 + C_{\text{IMQ},\nu,\gamma})^{d+1}, \quad (35)$$

$$\text{where } C_{\text{IMQ},\nu,\gamma} \triangleq 4 \log \left(16 \frac{(2\nu+1)^{\nu+1}}{\gamma^{2\nu}} + 1 \right), \text{ and } \tilde{r} \triangleq \min \left(\frac{\gamma}{2d}, \sqrt{\gamma^2 + 4r^2} - 2r \right). \quad (36)$$

(d) When $\mathbf{k} = \text{SINC}(\theta)$, then for $\varepsilon \in (0, \frac{1}{2})$, we have

$$\mathcal{M}_{\mathbf{k}}([-r, r]^d, \varepsilon) \leq d \cdot \left(1 + \frac{4r}{\tilde{r}_\theta}\right) \cdot (4 \log(d/\varepsilon) + 2 + 4 \log 17)^2, \quad (37)$$

$$\text{where } \tilde{r}_\theta \triangleq \min \left(\frac{\sqrt{3}}{|\theta|}, \sqrt{\frac{12}{\theta^2} + 4r^2} - 2r \right). \quad (38)$$

(e) When $\mathbf{k} = \text{B-SPLINE}(2\beta + 1, \gamma)$, then for some universal constant $C_{\text{B-SPLINE}}$, we have

$$\mathcal{M}_{\mathbf{k}}\left(\left[-\frac{1}{2}, \frac{1}{2}\right]^d, \varepsilon\right) \leq d \cdot \max(\gamma, 1) \cdot C_{\text{B-SPLINE}} \cdot (d/\varepsilon)^{\frac{1}{\beta + \frac{1}{2}}}. \quad (39)$$

I.1. Auxiliary results about RKHS and Euclidean covering numbers

In this section, we collect several results regarding the covering numbers of Euclidean and RKHS spaces that come in handy for our proofs. These results can also be of independent interest.

We start by defining the notion of restricted kernel and its unit ball (Rudi et al. (2020, Prop. 8)). For $\mathcal{X} \subset \mathbb{R}^d$, let $|_{\mathcal{X}}$ denotes the restriction operator. That is, for any function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we have $f|_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{R}$ such that $f|_{\mathcal{A}}(x) = f(x)$ for $x \in \mathcal{X}$.

Definition 3 (Restricted kernel and its RKHS) Consider a kernel \mathbf{k} defined on $\mathbb{R}^d \times \mathbb{R}^d$ with the corresponding RKHS \mathcal{H} , any set $\mathcal{X} \subset \mathbb{R}^d$. The restricted kernel $\mathbf{k}|_{\mathcal{X}}$ is defined as

$$\mathbf{k}|_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R} \quad \text{such that} \quad \mathbf{k}|_{\mathcal{X}}(x, y) \triangleq \mathbf{k}|_{\mathcal{X} \times \mathcal{X}}(x, y) = \mathbf{k}(x, y) \quad \text{for all } x, y \in \mathcal{X},$$

and $\mathcal{H}|_{\mathcal{X}}$ denotes its RKHS. For $f \in \mathcal{H}|_{\mathcal{X}}$, the restricted RKHS norm is defined as follows:

$$\|f\|_{\mathbf{k}|_{\mathcal{X}}} = \inf_{h \in \mathcal{H}} \|h\|_{\mathbf{k}} \quad \text{such that} \quad h|_{\mathcal{X}} = f.$$

Furthermore, we use $\mathcal{B}_{\mathbf{k}|_{\mathcal{X}}} \triangleq \{f \in \mathcal{H}|_{\mathcal{X}} : \|f\|_{\mathbf{k}|_{\mathcal{X}}} \leq 1\}$ to denote the unit ball of the RKHS corresponding to this restricted kernel.

In this notation, the unit ball of unrestricted kernel satisfies $\mathcal{B}_{\mathbf{k}} \triangleq \mathcal{B}_{\mathbf{k}|_{\mathbb{R}^d}}$. Now, recall the RKHS covering number definition from Def. 1. In the sequel, we also use the covering number of the restricted kernel defined as follows:

$$\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{X}, \varepsilon) = \mathcal{N}_{\mathbf{k}|_{\mathcal{X}}}(\mathcal{X}, \varepsilon), \quad (40)$$

that is $\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{X}, \varepsilon)$ denotes the minimum cardinality over all possible covers $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}|_{\mathcal{X}}}$ that satisfy

$$\mathcal{B}_{\mathbf{k}|_{\mathcal{X}}} \subset \bigcup_{h \in \mathcal{C}} \left\{ g \in \mathcal{B}_{\mathbf{k}|_{\mathcal{X}}} : \sup_{x \in \mathcal{X}} |h(x) - g(x)| \leq \varepsilon \right\}.$$

With this notation in place, we now state a result that relates the covering numbers \mathcal{N}^{\dagger} (40) and \mathcal{N} Def. 1.

Lemma 2 (Relation between restricted and unrestricted RKHS covering numbers) We have

$$\mathcal{N}_{\mathbf{k}, \varepsilon}(\mathcal{X}) \leq \mathcal{N}_{\mathbf{k}, \varepsilon}^{\dagger}(\mathcal{X})$$

Proof Rudi et al. (2020, Prop. 8(d,f)) imply that there exists a bounded linear extension operator $E : \mathcal{H}|_{\mathcal{X}} \rightarrow \mathcal{H}$ with operator norm bounded by 1, which when combined with Steinwart and Christmann (2008, eqns. (A.38), (A.39)) yields the claim. \blacksquare

Next, we state results that relate RKHS covering numbers for a change of domain for a shift-invariant kernel. We use $\mathcal{B}_{\|\cdot\|}(x; r) \triangleq \{y \in \mathbb{R}^d : \|x - y\| \leq r\}$ to denote the r radius ball in \mathbb{R}^d defined by the metric induced by a norm $\|\cdot\|$.

Definition 4 (Euclidean covering numbers) Given a set $\mathcal{X} \subset \mathbb{R}^d$, a norm $\|\cdot\|$, and a scalar $\varepsilon > 0$, we use $\mathcal{N}_{\|\cdot\|}(\mathcal{X}, \varepsilon)$ to denote the ε -covering number of \mathcal{X} with respect to $\|\cdot\|$ -norm. That is, $\mathcal{N}_{\|\cdot\|}(\mathcal{X}, \varepsilon)$ denotes the minimum cardinality over all possible covers $\mathcal{C} \subset \mathcal{X}$ that satisfy

$$\mathcal{X} \subset \bigcup_{z \in \mathcal{C}} \mathcal{B}_{\|\cdot\|}(z; \varepsilon).$$

When $\|\cdot\| = \|\cdot\|_q$ for some $q \in [1, \infty]$, we use the shorthand $\mathcal{N}_q \triangleq \mathcal{N}_{\|\cdot\|_q}$.

Lemma 3 (Relation between RKHS covering numbers on different domains) *Given a shift-invariant kernel \mathbf{k} , a norm $\|\cdot\|$ on \mathbb{R}^d , and any set $\mathcal{X} \subset \mathbb{R}^d$, we have*

$$\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{X}, \varepsilon) \leq \left[\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{B}_{\|\cdot\|}, \varepsilon) \right]^{\mathcal{N}_{\kappa}(\mathcal{X}, 1)}.$$

Proof Let $\mathcal{C} \subset \mathcal{X}$ denote the cover of minimum cardinality such that

$$\mathcal{X} \subseteq \bigcup_{z \in \mathcal{C}} \mathcal{B}_{\|\cdot\|}(z, 1).$$

We then have

$$\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{X}, \varepsilon) \stackrel{(i)}{\leq} \prod_{z \in \mathcal{C}} \mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{B}_{\|\cdot\|}(z, 1), \varepsilon) \stackrel{(ii)}{\leq} \prod_{z \in \mathcal{C}} \mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{B}_{\|\cdot\|}, \varepsilon) \leq \left[\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{B}_{\|\cdot\|}, \varepsilon) \right]^{|\mathcal{C}|},$$

where step (i) follows by applying [Steinwart and Fischer \(2021, Lem. 3.11\)](#),³ and step (ii) follows by applying [Steinwart and Fischer \(2021, Lem. 3.10\)](#). The claim follows by noting that \mathcal{C} denotes a cover of minimum cardinality, and hence by definition $|\mathcal{C}| = \mathcal{N}_{\|\cdot\|}(\mathcal{X}, 1)$. \blacksquare

Lemma 4 (Covering number for product kernel) *Given $\mathcal{X} \subset \mathbb{R}$ and a reproducing kernel $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, consider the product kernel $\mathbf{k} \triangleq \kappa^{\otimes d} : \mathcal{X}^{\otimes 2d} \rightarrow \mathbb{R}$ defined as*

$$\mathbf{k}(x, y) = \prod_{i=1}^d \kappa(x_i, y_i) \quad \text{for } x, y \in \mathcal{X}^{\otimes d} \triangleq \underbrace{\mathcal{X} \times \mathcal{X} \dots \times \mathcal{X}}_{d \text{ times}} \subset \mathbb{R}^d.$$

Then the covering numbers of the two kernels are related as follows:

$$\mathcal{N}_{\mathbf{k}}^{\dagger}(\mathcal{X}^{\otimes d}, \varepsilon) \leq \left[\mathcal{N}_{\kappa}^{\dagger}(\mathcal{X}, \varepsilon / (d \|\kappa\|_{\infty}^{\frac{d-1}{2}})) \right]^d. \quad (41)$$

Proof Let \mathcal{H} denote the RKHS corresponding to κ . Then the RKHS corresponding to the kernel \mathbf{k} is given by the tensor product $\mathcal{H}_{\mathbf{k}} \triangleq \mathcal{H} \times \mathcal{H} \times \dots \times \mathcal{H}$ [Berlinet and Thomas-Agnan \(2011, Sec. 4.6\)](#), i.e., for any $f \in \mathcal{H}_{\mathbf{k}}$, there exists $f_1, f_2, \dots, f_d \in \mathcal{H}$ such that

$$f(x) = \prod_{i=1}^d f_i(x_i) \quad \text{for all } x \in \mathcal{X}^{\otimes d}. \quad (42)$$

Let $\mathcal{C}_{\kappa}(\mathcal{X}, \varepsilon) \subset \mathcal{B}_{\kappa}$ denote an ε -cover of \mathcal{B}_{κ} in L^{∞} -norm ([Def. 1](#)). Then for each $f_i \in \mathcal{H}$, we have $\tilde{f}_i \in \mathcal{C}_{\kappa}(\mathcal{X}, \varepsilon)$ such that

$$\sup_{z \in \mathcal{X}} |f_i(z) - \tilde{f}_i(z)| \leq \varepsilon. \quad (43)$$

Now, we claim that for every $f \in \mathcal{B}_{\mathbf{k}}$, there exists $g \in \mathcal{C}_{\mathbf{k}} \triangleq (\mathcal{C}_{\kappa}(\mathcal{X}, \varepsilon))^{\otimes d}$ such that

$$\sup_{x \in \mathcal{X}^{\otimes d}} |f(x) - g(x)| \leq d\varepsilon \|\kappa\|_{\infty}^{\frac{d-1}{2}}, \quad (44)$$

3. [Steinwart and Fischer \(2021, Lem. 3.11\)](#) is stated for disjoint partition of X in two sets, but the argument can be repeated for any finite cover of X .

which immediately implies the claimed bound (41) on the covering number. We now prove the claim (44). For any fixed $f \in \mathcal{H}_k$, let f_i, \tilde{f}_i denote the functions satisfying (42) and (43) respectively. Then, we prove our claim (44) with $g = \prod_{i=1}^d \tilde{f}_i \in \mathcal{C}_k$. Using the convention $\prod_{k=1}^0 \tilde{f}_k(x_k) = 1$, we find that

$$\begin{aligned} |f(x) - g(x)| &= \left| \prod_{i=1}^d f_i(x_i) - \prod_{i=1}^d \tilde{f}_i(x_i) \right| \\ &\leq \sum_{i=1}^d \left| f_i(x_i) - \tilde{f}_i(x_i) \right| \left| \prod_{j=i+1}^d f_j(x_j) \prod_{k=1}^{i-1} \tilde{f}_k(x_k) \right| \\ &\stackrel{(43)}{\leq} d\varepsilon \cdot \sup_{h \in \mathcal{B}_\kappa} \|h\|_\infty^{d-1} \leq d\varepsilon \|\kappa\|_\infty^{\frac{d-1}{2}}, \end{aligned}$$

where in the last step we have used the following argument:

$$\sup_{z \in \mathcal{X}} h(x) = \sup_{z \in \mathcal{X}} h, \kappa(z, \cdot)_\kappa \leq \|h\|_\kappa \sqrt{\kappa(z, z)} \leq \sqrt{\|\kappa\|_\infty} \quad \text{for any } h \in \mathcal{B}_\kappa.$$

The proof is now complete. ■

Lemma 5 (Relation between Euclidean covering numbers) *We have*

$$\mathcal{N}_\infty(\mathcal{B}_2(r), 1) \leq \frac{1}{\sqrt{\pi d}} \cdot \left[\left(1 + \frac{2r}{\sqrt{d}}\right) \sqrt{2\pi e} \right]^d \quad \text{for all } d \geq 1.$$

Proof We apply [Wainwright \(2019, Lem. 5.7\)](#) with $\mathcal{B} = \mathcal{B}_2(r)$ and $\mathcal{B}' = \mathcal{B}_\infty(1)$ to conclude that

$$\mathcal{N}_\infty(\mathcal{B}_2(r), 1) \leq \frac{\text{Vol}(2\mathcal{B}_2(r) + \mathcal{B}_\infty(1))}{\text{Vol}(\mathcal{B}_\infty(1))} \leq \text{Vol}(\mathcal{B}_2(2r + \sqrt{d})) \leq \frac{\pi^{d/2}}{\Gamma(\frac{d}{2}+1)} \cdot (2r + \sqrt{d})^d,$$

where $\text{Vol}(\mathcal{X})$ denotes the d -dimensional Euclidean volume of $\mathcal{X} \subset \mathbb{R}^d$, and $\Gamma(a)$ denotes the Gamma function. Next, we apply the following bounds on the Gamma function from [Batir \(2017, Thm. 2.2\)](#):

$$\Gamma(b+1) \geq (b/e)^b \sqrt{2\pi b} \quad \text{for any } b \geq 1, \quad \text{and} \quad \Gamma(b+1) \leq (b/e)^b \sqrt{e^2 b} \quad \text{for any } b \geq 1.1.$$

Thus, we have

$$\mathcal{N}_\infty(\mathcal{B}_2(r), 1) \leq \frac{\pi^{d/2}}{\sqrt{2\pi d} (\frac{d}{2e})^{d/2}} \cdot (2r + \sqrt{d})^d \leq \frac{1}{\sqrt{\pi d}} \cdot \left[\left(1 + \frac{2r}{\sqrt{d}}\right) \sqrt{2e\pi} \right]^d,$$

as claimed, and we are done. ■

I.2. Proof of Prop. 2: Covering numbers for analytic and differentiable kernels

First we apply [Lem. 2](#) so that it remains to establish the stated bounds simply on $\log \mathcal{N}_k^\dagger(\mathcal{X}, \varepsilon)$.

Proof of bound (29) in part (a) The bound (29) for the real-analytic kernel is a restatement of [Sun and Zhou \(2008, Thm. 2\)](#) in our notation (in particular, after making the following substitutions in their notation: $R \leftarrow 1, C_0 \leftarrow C_\kappa, r \leftarrow R_\kappa, \mathcal{X} \leftarrow \mathcal{A}, \tilde{r} \leftarrow r^\dagger, \eta \leftarrow \varepsilon, D \leftarrow D_{\mathcal{A}}^2, n \leftarrow d$). □

Proof of bound (31) for part (b): Under these assumptions, [Steinwart and Christmann \(2008, Thm. 6.26\)](#) states that the i -th dyadic entropy number [Steinwart and Christmann \(2008, Def. 6.20\)](#) of the identity inclusion mapping from $\mathcal{H}|_{\bar{B}_2(r)}$ to $L^\infty_{\bar{B}_2(r)}$ is bounded by $c'_{s,d,\mathbf{k}} \cdot r^s i^{-s/d}$ for some constant $c'_{s,d,\mathbf{k}}$ independent of ε and r . Given this bound on the entropy number, and applying [Steinwart and Christmann \(2008, Lem. 6.21\)](#), we conclude that the log-covering number $\log \mathcal{N}_{\mathbf{k}}^\dagger(\bar{B}_2(r), \varepsilon)$ is bounded by $\ln 4 \cdot (c'_{s,d,\mathbf{k}} r^s / \varepsilon)^{d/s} = c_{s,d,\mathbf{k}} r^d \cdot (1/\varepsilon)^{d/s}$ as claimed. \square

I.3. Proof of Prop. 3: Covering numbers for specific kernels

First we apply [Lem. 2](#) so that it remains to establish the stated bounds in each part on the corresponding $\log \mathcal{N}_{\mathbf{k}}$.

Proof for GAUSS kernel: Part (a) The bound (32) for the Gaussian kernel follows directly from [Steinwart and Fischer \(2021, Eqn. 11\)](#) along with the discussion stated just before it. Furthermore, the bound (33) for $C_{\text{Gauss},d}$ are established in [Steinwart and Fischer \(2021, Eqn. 6\)](#), and in the discussion around it. \square

Proof for MATÉRN kernel: Part (b) We claim that $\text{MATÉRN}(\nu, \gamma)$ is $\lfloor \nu - \frac{d}{2} \rfloor$ -times continuously differentiable which immediately implies the bound (34) using [Prop. 2\(b\)](#).

To prove the differentiability, we use Fourier transform of Matérn kernels. For $\mathbf{k} = \text{MATÉRN}(\nu, \gamma)$, let $\kappa : \mathbb{R}^d \rightarrow \mathbb{R}$ denote the function such that noting that $\mathbf{k}(x, y) = \kappa(x - y)$. Then using the Fourier transform of κ from [Wendland \(2004, Thm 8.15\)](#), and noting that κ is real-valued, we can write

$$\mathbf{k}(x, y) = c_{\mathbf{k},d} \int \cos(\omega^\top(x - y))(\gamma^2 + \|\omega\|_2^2)^{-\nu} d\omega$$

for some constant $c_{\mathbf{k},d}$ depending only on the kernel parameter, and d (due to the normalization of the kernel, and the Fourier transform convention). Next, for any multi-index $a \in \mathbb{N}_0^d$, we have

$$\left| \partial^{a,a} \cos(\omega^\top(x - y))(\gamma^2 + \|\omega\|_2^2)^{-\nu} \right| \leq \prod_{j=1}^d \omega_j^{2a_j} (\gamma^2 + \|\omega\|_2^2)^{-\nu} \leq \frac{\|\omega\|_2^{2 \sum_{j=1}^d a_j}}{(\gamma^2 + \|\omega\|_2^2)^\nu},$$

where $\partial^{a,a}$ denotes the partial derivative of order a . Moreover, we have

$$\int \frac{\|\omega\|_2^{2 \sum_{j=1}^d a_j}}{(\gamma^2 + \|\omega\|_2^2)^\nu} d\omega = c_d \int_{r>0} r^{d-1} \frac{r^{2 \sum_{j=1}^d a_j}}{(\gamma^2 + r^2)^\nu} dr \leq c_d \int_{r>0} r^{d-1+2 \sum_{j=1}^d a_j - 2\nu} \stackrel{(i)}{<} \infty,$$

where step (i) holds whenever

$$d - 1 + 2 \sum_{j=1}^d a_j - 2\nu < -1 \iff \sum_{j=1}^d a_j < \nu - \frac{d}{2}.$$

Then applying [Newey and McFadden \(1994, Lemma 3.6\)](#), we conclude that for all multi-indices a such that $\sum_{j=1}^d a_j \leq \lfloor \nu - \frac{d}{2} \rfloor$, the partial derivative $\partial^{a,a} \mathbf{k}$ exists and is given by

$$c_{\mathbf{k},d} \int \partial^{a,a} \cos(\omega^\top(x - y))(\gamma^2 + \|\omega\|_2^2)^{-\nu} d\omega,$$

and we are done. \square

Proof for IMQ kernel: Part (c) The bounds (35) and (36) follow from Sun and Zhou (2008, Ex. 3), and noting that $\mathcal{N}_2(\mathcal{B}_2(r), \tilde{r}/2)$ is bounded by $(1 + \frac{4r}{\tilde{r}})^d$ (Wainwright, 2019, Lem. 5.7). \square

Proof for SINC kernel: Part (d) For $\mathbf{k} = \text{SINC}(\theta)$, we can write $\mathbf{k}(x, y) = \prod_{i=1}^d \kappa_\theta(x_i - y_i)$ for $\kappa_\theta : \mathbb{R} \rightarrow \mathbb{R}$ defined as $\kappa_\theta(t) = \frac{\sin(\theta t)}{\theta t} \stackrel{(i)}{=} \frac{\sin(|\theta t|)}{|\theta t|}$, where step (i) follows from the fact that $t \mapsto \sin t/t$ is an even function. Thus, we can apply Lem. 4. Given the bound (41), and noting that $\|\kappa_\theta\|_\infty = 1$, it suffices to establish the univariate version of the bound (37), namely,

$$\mathcal{M}_{\mathbf{k}}([-r, r], \varepsilon) \leq (1 + \frac{4r}{r_\theta}) \cdot (4 \log(1/\varepsilon) + 2 + 4 \log 17)^2.$$

To do so, we claim that univariate SINC kernel is an analytic kernel that satisfies the condition (28) of Prop. 2(a) with $\kappa(t) = \text{SINC}(\theta\sqrt{t})$, $R_\kappa = \frac{12}{\theta^2}$, and $C_\kappa = 1$; and thus applying the bounds (29) and (30) from Prop. 2(a) with $\mathcal{A} = \mathcal{B}_2^d(r)$ yields the claimed bound (37) and (38). To verify the condition (28) with the stated parameters, we note that

$$\begin{aligned} \kappa(t) = \text{SINC}(\theta\sqrt{t}) &= \frac{1}{|\theta|\sqrt{t}} \sum_{j=0}^{\infty} \frac{1}{(2j+1)!} \cdot (\theta\sqrt{t})^{2j+1} = \sum_{j=0}^{\infty} \frac{1}{(2j+1)!} \cdot (\theta\sqrt{t})^{2j} \\ &= \sum_{j=0}^{\infty} \frac{1}{(2j+1)!} \cdot \theta^{2j} \cdot t^j \end{aligned}$$

which implies

$$\left| \kappa_+^{(j)}(0) \right| = \frac{1}{(2j+1)!} \cdot \theta^{2j} j! \leq (2/R_\kappa)^j j! \quad \text{for} \quad R_\kappa \triangleq \frac{2}{\theta^2} \cdot \inf_{j \geq 1} ((2j+1)!)^{1/j} = \frac{12}{\theta^2},$$

and we are done. \square

Proof for B-SPLINE kernel: Part (e) For $\mathbf{k} = \text{B-SPLINE}(2\beta + 1, \gamma)$, we can write $\mathbf{k}(x, y) = \prod_{i=1}^d \kappa_{\beta, \gamma}(x_i - y_i)$ for $\kappa_{\beta, \gamma} : \mathbb{R} \rightarrow \mathbb{R}$ defined as $\kappa_{\beta, \gamma}(t) = \mathbf{B}_{2\beta+2}^{-1} \otimes^{2\beta+2} \mathbf{1}_{[-\frac{1}{2}, \frac{1}{2}]}(\gamma \cdot t)$, and thus we can apply Lem. 4. Given the bound (41), and noting that $\|\kappa_{\beta, \gamma}\|_\infty \leq 1$ (Dwivedi and Mackey (2021, Eqn. 107)), it suffices to establish the univariate version of the bound (39). Abusing notation and using $\kappa_{\beta, \gamma}$ to denote the univariate B-SPLINE($2\beta + 1, \gamma$) kernel, we find that

$$\begin{aligned} \log \mathcal{N}_{\kappa_{\beta, \gamma}}^\dagger([-\frac{1}{2}, \frac{1}{2}], \varepsilon) &\stackrel{(i)}{\leq} \mathcal{N}_1([0, \gamma], 1) \cdot \log \mathcal{N}_{\kappa_{\beta, 1}}^\dagger([-\frac{1}{2}, \frac{1}{2}], \varepsilon) \\ &\stackrel{(ii)}{\leq} \max(\gamma, 1) \cdot C_{\text{B-SPLINE}} \cdot (1/\varepsilon)^{\frac{1}{\beta+\frac{1}{2}}}, \end{aligned}$$

where step (i) follows from Steinwart and Fischer (2021, Thm. 2.4, Sec. 3.3), and for step (ii) we use the fact that the unit-covering number of $[0, \gamma]$ is bounded by $\max(\gamma, 1)$, and apply the covering number bound for the univariate B-SPLINE kernel from Zhou (2003, Ex. 4) (by substituting $m = 2\beta + 2$ in their notation) along with the fact that $\log \mathcal{N}_{\kappa_{\beta, 1}}^\dagger([-\frac{1}{2}, \frac{1}{2}], \varepsilon) = \log \mathcal{N}_{\kappa_{\beta, 1}}^\dagger([0, 1], \varepsilon)$ since κ_β is shift-invariant. \square

Appendix J. Proof of Tab. 3 results

In Tab. 3, the stated results for all the entries in the TARGET KT column follow directly by substituting the covering number bounds from Prop. 3 in the corresponding entry along with the stated

radii growth conditions for the target \mathbb{P} . (We substitute $m = \frac{1}{2} \log_2 n$ since we thin to \sqrt{n} output size.) For the KT+ column, the stated result follows by either taking the minimum of the first two columns (whenever the ROOT KT guarantee applies) or using the POWER KT guarantee. First we remark how to always ensure a rate of at least $\mathcal{O}(n^{-\frac{1}{4}})$ even when the guarantee from our theorems are larger, using a suitable baseline procedure and then proceed with our proofs.

Remark 1 (Improvement over baseline thinning) *First we note that the KT-SWAP step ensures that, deterministically, $\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}}) \leq \text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{base}})$ and $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{KT}}) \leq 2 \text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{in}}) + \text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{base}})$ for $\mathcal{S}_{\text{base}}$ a baseline thinned coreset of size $\frac{n}{2^m}$ and any target \mathbb{P} . For example if the input and baseline coresets are drawn i.i.d. and \mathbf{k} is bounded, then $\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}})$ and $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}_{\text{KT}})$ are $\mathcal{O}(\sqrt{2^m/n})$ with high probability (Tolstikhin et al., 2017, Thm. A.1), even if the guarantee of Thm. 2 is larger. As a consequence, in all well-defined KT variants, we can guarantee a rate of $n^{-\frac{1}{4}}$ for $\text{MMD}_{\mathbf{k}}(\mathcal{S}_{\text{in}}, \mathcal{S}_{\text{KT}})$ when the output size is \sqrt{n} simply by using baseline as i.i.d. thinning in the KT-SWAP step.*

GAUSS kernel The TARGET KT guarantee follows by substituting the covering number bound for the Gaussian kernel from Prop. 3(a) in (5), and the ROOT KT guarantee follows directly from Dwivedi and Mackey (2021, Tab. 2). Putting the guarantees for the ROOT KT and TARGET KT together (and taking the minimum of the two) yields the guarantee for KT+.

IMQ kernel The TARGET KT guarantee follows by putting together the covering bound Prop. 3(c) and the MMD bounds (5).

For the ROOT KT guarantee, we use a square-root dominating kernel $\tilde{\mathbf{k}}_{\text{rt}}$ IMQ(ν', γ') Dwivedi and Mackey (2021, Def.2) as suggested by Dwivedi and Mackey (2021). Dwivedi and Mackey (2021, Eqn.(117)) shows that $\tilde{\mathbf{k}}_{\text{rt}}$ is always defined for appropriate choices of ν', γ' . The best ROOT KT guarantees are obtained by choosing largest possible ν' (to allow the most rapid decay of tails), and Dwivedi and Mackey (2021, Eqn.(117)) implies with $\nu < \frac{d}{2}$, the best possible parameter satisfies $\nu' \leq \frac{d}{4} + \frac{\nu}{2}$. For this parameter, some algebra shows that $\max(\mathbb{R}_{\tilde{\mathbf{k}}_{\text{rt}}, n}^{\dagger}, \mathbb{R}_{\tilde{\mathbf{k}}_{\text{rt}}, n}^{\sim}) \lesssim_{d, \nu, \gamma} n^{1/2\nu}$, leading to a guarantee worse than $n^{-\frac{1}{4}}$, so that the guarantee degenerates to $n^{-\frac{1}{4}}$ using Rem. 1 for ROOT KT. When $\nu \geq \frac{d}{2}$, we can use a MATÉRN kernel as a square-root dominating kernel from Dwivedi and Mackey (2021, Prop. 3), and then applying the bounds for the kernel radii (12), and the inflation factor (14) for a generic Matérn kernel from Dwivedi and Mackey (2021, Tab. 3) leads to the entry for the ROOT KT stated in Tab. 2. The guarantee for KT+ follows by taking the minimum of the two.

MATÉRN kernel For TARGET KT, substituting the covering number bound from Prop. 3(b) in (5) with $R = \log n$ yields the MMD bound of order

$$\sqrt{\frac{\log n \cdot (\log n)^d \cdot n^{2b\nu} \cdot \frac{d}{2} c}{n}}, \quad (45)$$

which is better than $n^{-\frac{1}{4}}$ only when $\nu > 3d/2$, and simplified to the entry in the Tab. 3 when we assume $\nu - \frac{d}{2}$ is an integer. When $\nu \leq 3d/2$, we can simply use baseline as i.i.d. thinning to obtain an order $n^{-\frac{1}{4}}$ MMD error as in Rem. 1.

The ROOT KT (and thereby KT+) guarantees for $\nu > d$ follow from Dwivedi and Mackey (2021, Tab. 2).

When $\nu \in (\frac{d}{2}, d]$, we use POWER KT with a suitable α to establish the KT+ guarantee. For MATÉRN(ν, γ) kernel, the α -power kernel is given by MATÉRN($\alpha\nu, \gamma$) if $\alpha\nu > \frac{d}{2}$ (a proof of this follows from Def. 2 and Dwivedi and Mackey (2021, Eqns (71-72))). Since LAPLACE(σ) = MATÉRN($\frac{d+1}{2}, \sigma^{-1}$), we conclude that its α -power kernel is defined for $\alpha > \frac{d}{d+1}$. And using the various tail radii (12), and the inflation factor (14) for a generic Matérn kernel from Dwivedi and Mackey (2021, Tab. 3), we conclude that $\widetilde{M}_\alpha \lesssim_{d, \mathbf{k}_\alpha, \delta} \sqrt{\log n \log \log n}$, and $\max(\mathbf{R}_{\mathbf{k}_\alpha, n}^\dagger, \mathbf{R}_{\mathbf{k}_\alpha, n}) \lesssim_{d, \mathbf{k}_\alpha} \log n$, so that $\mathbf{R}_{\max} = \mathcal{O}_{d, \mathbf{k}_\alpha}(\log n)$ (13) for SUBEXP \mathbb{P} setting. Thus for this case, the MMD guarantee for \sqrt{n} thinning with POWER KT (tracking only scaling with n) is

$$\begin{aligned} & \left(\frac{2^m}{n} \|\mathbf{k}_\alpha\|_\infty \right)^{\frac{1}{2\alpha}} (2 \cdot \widetilde{M}_\alpha)^{1 - \frac{1}{2\alpha}} \left(2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} \cdot \mathbf{R}_{\max}^{\frac{d}{2}} \cdot \widetilde{M}_\alpha \right)^{\frac{1}{\alpha} - 1} \\ & \lesssim_{d, \mathbf{k}_\alpha, \delta} \left(\frac{1}{\sqrt{n}} \right)^{\frac{1}{2\alpha}} (\sqrt{c_n \log n})^{1 - \frac{1}{2\alpha}} \cdot ((\log n)^{\frac{d}{2} + \frac{1}{2}} \sqrt{c_n})^{\frac{1}{\alpha} - 1} = \left(\frac{c_n (\log n)^{1+2d(1-\alpha)}}{n} \right)^{\frac{1}{4\alpha}} \end{aligned}$$

where $c_n = \log \log n$; and we thus obtain the corresponding entry (for KT+) stated in Tab. 3.

SINC kernel The guarantee for TARGET KT follows directly from substituting the covering number bounds from Prop. 3(d) in (5).

For the ROOT KT guarantee, we note that the square-root kernel construction of Dwivedi and Mackey (2021, Prop.2) implies that SINC(θ) itself is a square-root of SINC(θ) since the Fourier transform of SINC is a rectangle function on a bounded domain. However, the tail of the SINC kernel does not decay fast enough for the guarantee of Dwivedi and Mackey (2021, Thm. 1) to improve beyond the $n^{-\frac{1}{4}}$ bound of Dwivedi and Mackey (2021, Rem. 2) obtained when running ROOT KT with i.i.d. baseline thinning.

In this case, TARGET KT and KT+ are identical since $\mathbf{k}_{\text{rt}} = \mathbf{k}$.

B-SPLINE kernel The guarantee for TARGET KT follows directly from substituting the covering number bounds from Prop. 3(d) in (5).

For B-SPLINE($2\beta + 1, \gamma$) kernel, using arguments similar to that in Dwivedi and Mackey (2021, Tab.4), we conclude that (up to a constant scaling) the α -power kernel is defined to be B-SPLINE($A + 1, \gamma$) whenever $A \triangleq 2\alpha\beta + 2\alpha - 2 \in 2\mathbb{N}_0$. For odd β we can always take $\alpha = \frac{1}{2}$ and B-SPLINE($\beta + 1, \gamma$) is a valid (up to a constant scaling) square-root kernel (Dwivedi and Mackey, 2021). For even β , we have to choose $\alpha \triangleq \frac{p+1}{\beta+1} \in (\frac{1}{2}, 1)$ by taking $p \in \mathbb{N}$ suitably, and the smallest suitable choice is $p = \lceil \frac{\beta-1}{2} \rceil = \frac{\beta}{2} \in \mathbb{N}$, which is feasible as long as $\beta \geq 2$. And, thus B-SPLINE($\beta + 1, \gamma$) is a suitable \mathbf{k}_α for B-SPLINE($2\beta + 1$) for even $\beta \geq 2$ with $\alpha = \frac{\beta+2}{2\beta+2} \in (\frac{1}{2}, 1)$. Whenever the α -power kernel is defined, we can then apply the various tail radii (12), and the inflation factor (14) for the power B-SPLINE kernel from Dwivedi and Mackey (2021, Tab. 3) to obtain the MMD rates for POWER KT from Dwivedi and Mackey (2021, Tab. 2) (which remains the same as ROOT KT upto factors depending on α and β).

The guarantee for KT+ follows by taking the minimum MMD error for TARGET KT and ROOT KT for even β , and α -POWER KT for odd β .

GENERALIZED KERNEL THINNING

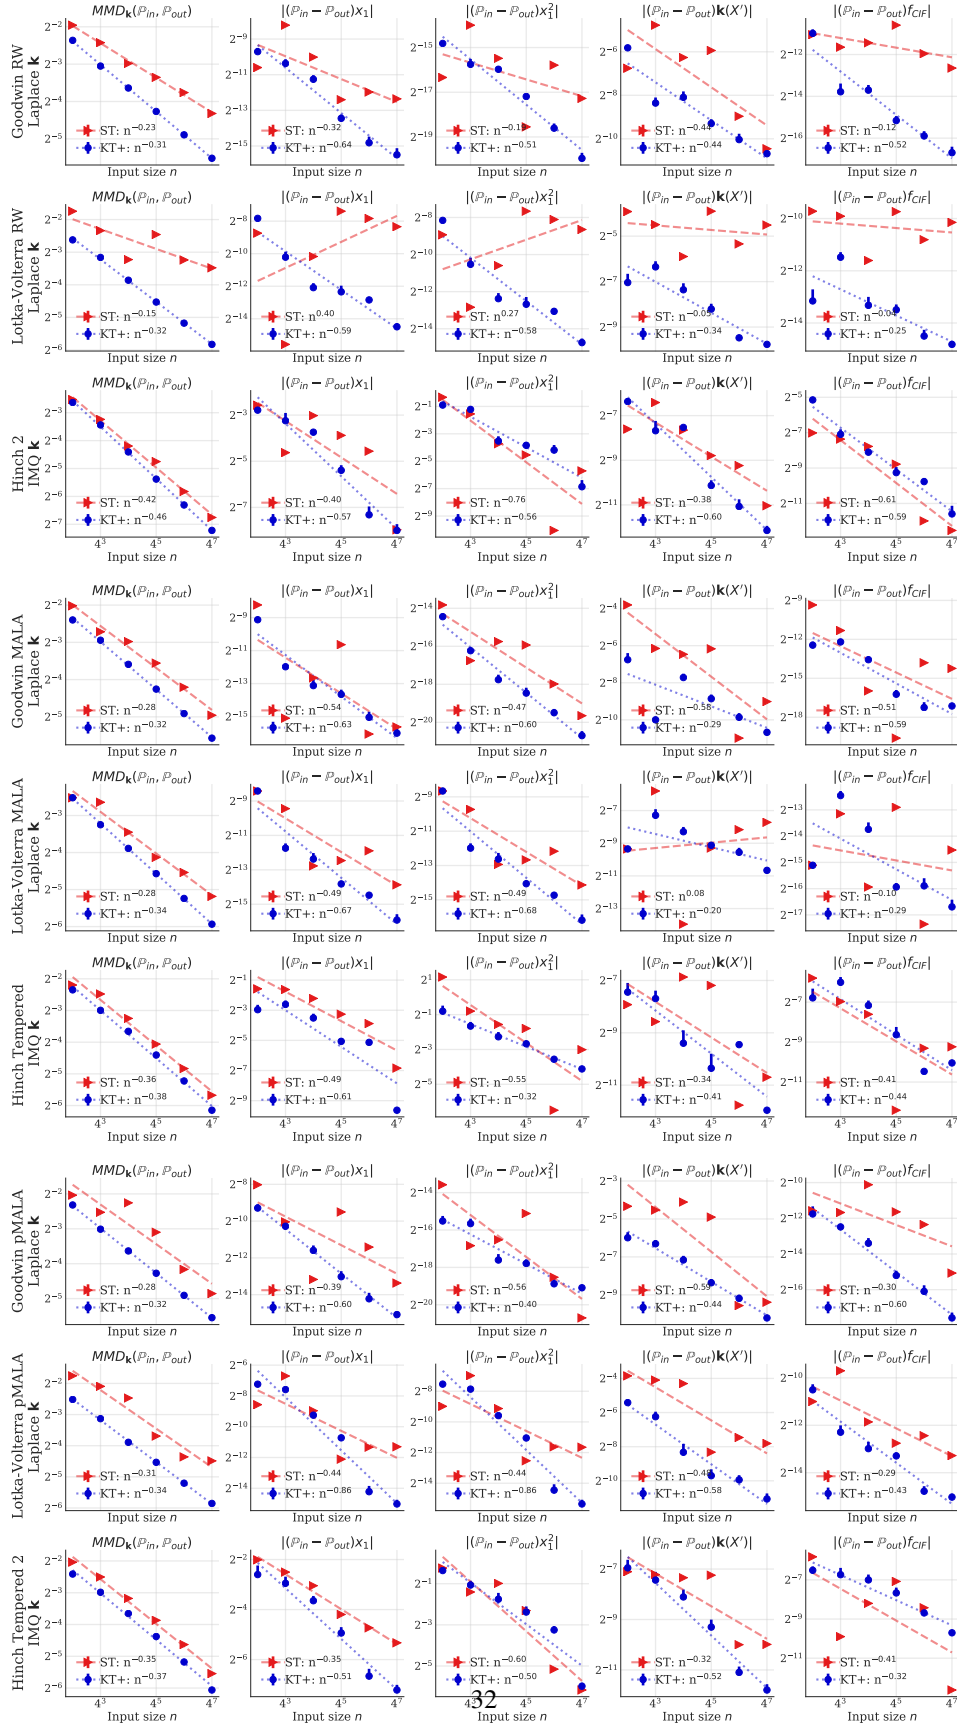


Figure 5: Kernel thinning+ (KT+) vs. standard MCMC thinning (ST). For kernels without fast-decaying square-roots, KT+ improves MMD and integration error decay rates in each inference task.