

Kernel Thinning

Raaz Dwivedi

Cornell Tech

DWIVEDI@CORNELL.EDU

Lester Mackey

Microsoft Research New England

LMACKEY@MICROSOFT.COM

Editor: Ingo Steinwart

Abstract

We introduce kernel thinning, a new procedure for compressing a distribution \mathbb{P} more effectively than i.i.d. sampling or standard thinning. Given a suitable reproducing kernel \mathbf{k}_* and $\mathcal{O}(n^2)$ time, kernel thinning compresses an n -point approximation to \mathbb{P} into a \sqrt{n} -point approximation with comparable worst-case integration error across the associated reproducing kernel Hilbert space. The maximum discrepancy in integration error is $\mathcal{O}_d(n^{-1/2}\sqrt{\log n})$ in probability for compactly supported \mathbb{P} and $\mathcal{O}_d(n^{-\frac{1}{2}}(\log n)^{(d+1)/2}\sqrt{\log \log n})$ for sub-exponential \mathbb{P} on \mathbb{R}^d . In contrast, an equal-sized i.i.d. sample from \mathbb{P} suffers $\Omega(n^{-1/4})$ integration error. Our sub-exponential guarantees resemble the classical quasi-Monte Carlo error rates for uniform \mathbb{P} on $[0, 1]^d$ but apply to general distributions on \mathbb{R}^d and a wide range of common kernels. Moreover, the same construction delivers near-optimal L^∞ coresets in $\mathcal{O}(n^2)$ time. We use our results to derive explicit non-asymptotic maximum mean discrepancy bounds for Gaussian, Matérn, and B-spline kernels and present two vignettes illustrating the practical benefits of kernel thinning over i.i.d. sampling and standard Markov chain Monte Carlo thinning, in dimensions $d = 2$ through 100.

Keywords: coresets, distribution compression, Markov chain Monte Carlo, maximum mean discrepancy, reproducing kernel Hilbert space, thinning

1. Introduction

Monte Carlo and Markov chain Monte Carlo (MCMC) methods (Brooks et al., 2011) are commonly used to approximate intractable target expectations $\mathbb{P}f \triangleq \mathbb{E}_{X \sim \mathbb{P}}[f(X)]$ of \mathbb{P} -integrable functions f with asymptotically exact averages $\mathbb{P}_n f \triangleq \frac{1}{n} \sum_{i=1}^n f(x_i)$ based on points $(x_i)_{i=1}^n$ generated from a Markov chain. A standard practice, to minimize the expense of downstream function evaluation, is to *thin* the Markov chain output down to a smaller size n_{out} by keeping only every (n/n_{out}) -th sample point (Owen, 2017). We call this approach *standard thinning*, and such sample compression is critical in fields like computational cardiology in which each function evaluation triggers an organ or tissue simulation consuming thousands of CPU hours (Niederer et al., 2011; Augustin et al., 2016; Stocchi et al., 2020). Unfortunately, standard thinning also leads to a significant reduction in accuracy. For example, thinning one’s chain down to $n_{\text{out}} = \sqrt{n}$ sample points increases integration error from $\mathcal{O}(n^{-\frac{1}{2}})$ in probability to $\Omega(n^{-\frac{1}{4}})$ by the Markov chain central limit theorem (Roberts and Rosenthal, 2004, Prop. 29). Our primary contribution is a more effective thinning strategy, which provides $\mathcal{O}_p(n^{-\frac{1}{4}})$ -integration error when $n^{\frac{1}{2}}$ points are returned.

1.1 Thinned MMD coresets

We focus on integration error in a reproducing kernel Hilbert space (RKHS, Steinwart and Christmann, 2008, Def. 4.18) of bounded, measurable functions with a target kernel $\mathbf{k}_\star : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ for $d \in \mathbb{N}$ and RKHS norm $\|\cdot\|_{\mathbf{k}_\star}$.

Assumption 1 (RKHS of bounded, measurable functions). *The RKHS $\mathcal{H}_{\mathbf{k}}$ of a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ contains only bounded measurable functions. Equivalently, \mathbf{k} is bounded with $\mathbf{k}(x, \cdot)$ measurable for all $x \in \mathbb{R}^d$ (Steinwart and Christmann, 2008, Lems. 4.23, 4.24).¹*

The worst-case integration error over the RKHS unit ball is given by the kernel *maximum mean discrepancy* (MMD, Gretton et al., 2012).

Definition 1 (Maximum mean discrepancy (Gretton et al., 2012)). *For a kernel \mathbf{k} satisfying Assump. 1, we define the kernel maximum mean discrepancy,*

$$\text{MMD}_{\mathbf{k}}(\mu, \nu) \triangleq \sup_{f \in \mathcal{H}_{\mathbf{k}}: \|f\|_{\mathbf{k}} \leq 1} |\mu f - \nu f| \quad \text{for all probability measures } \mu, \nu \text{ on } \mathbb{R}^d. \quad (1)$$

For sequences of points \mathcal{S} and \mathcal{S}' in \mathbb{R}^d with empirical distributions \mathbb{Q} and \mathbb{Q}' , we overload this notation to write $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathcal{S}) \triangleq \text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{Q})$ and $\text{MMD}_{\mathbf{k}}(\mathcal{S}, \mathcal{S}') \triangleq \text{MMD}_{\mathbf{k}}(\mathbb{Q}, \mathbb{Q}')$.

Given \mathbf{k}_\star satisfying Assump. 1, a target distribution \mathbb{P} on \mathbb{R}^d , and a sequence of \mathbb{R}^d -valued points $(x_i)_{i=1}^n$ generated to approximate \mathbb{P} , our aim is to identify a *thinned MMD coreset*, a shorter subsequence that continues to approximate \mathbb{P} well in $\text{MMD}_{\mathbf{k}_\star}$.

Definition 2 (MMD coreset). *We call a sequence of n_{out} points in \mathbb{R}^d with empirical measure \mathbb{Q} an $(n_{\text{out}}, \varepsilon)$ -MMD coreset for $(\mathbf{k}_\star, \mathbb{P})$ if $\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathbb{Q}) \leq \varepsilon$.*

Notably, when the initial sequence is drawn i.i.d. or from a fast-mixing Markov chain targeting \mathbb{P} , standard thinning down to size $n_{\text{out}} = n^{\frac{1}{2}}$ yields an order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ -MMD coreset in probability (see Prop. 1). A benchmark for improvement is provided by the online Haar strategy of Dwivedi et al. (2019), which generates an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}} \log^{2d} n))$ -MMD coreset in probability from $2n^{\frac{1}{2}}$ i.i.d. sample points when \mathbb{P} is specifically the uniform distribution on the unit cube $[0, 1]^d$.² Our goal is to develop thinned coresets of improved quality for any target \mathbb{P} with sufficiently fast tail decay.

1.2 Our contributions

To this end, we introduce *kernel thinning* (Alg. 1), a new, practical solution to the thinned MMD coreset problem that takes as input an $(n, \mathcal{O}_p(n^{-\frac{1}{2}}))$ -MMD coreset and outputs an $(n^{\frac{1}{2}}, \mathcal{O}_p(n^{-\frac{1}{4}}))$ -MMD coreset for a wide-range of $(\mathbf{k}_\star, \mathbb{P})$. Kernel thinning uses non-uniform randomness and evaluations of a less smooth *square-root kernel* \mathbf{k}_{rt} (see Def. 5) to partition the input into subsets of comparable quality and then greedily refines the best of these subsets using \mathbf{k}_\star . Our primary contributions include:

1. Throughout, we use \mathbf{k} for statements involving a generic kernel that is potentially distinct from the target kernel \mathbf{k}_\star .
2. Dwivedi et al. (2019) specifically control the *star discrepancy*, a quantity which in turn upper bounds a Sobolev space MMD called the *L^2 discrepancy* (Hickernell, 1998; Novak and Wozniakowski, 2010).

1. **Better-than-i.i.d. MMD coresets:** Given n input points sampled i.i.d. or from a fast-mixing Markov chain, kernel thinning yields, in probability, an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}}\sqrt{\log n}))$ -MMD coreset for \mathbb{P} and \mathbf{k}_{rt} with bounded support, an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}}\sqrt{\log^{d+1} n \log \log n}))$ -MMD coreset for \mathbb{P} and \mathbf{k}_{rt} with light tails, and an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2} + \frac{d}{2\rho}}\sqrt{\log n \log \log n}))$ -MMD coreset for \mathbb{P} and \mathbf{k}_{rt}^2 with $\rho > 2d$ moments (Thm. 1 and Cor. 1). For compactly supported or light-tailed \mathbb{P} and \mathbf{k}_{rt} , these results compare favorably with known $\Omega_d(n^{-\frac{1}{2}})$ lower bounds (see Sec. 8.1). Our guarantees extend to more general input point sequences, including deterministic sequences based on quadrature or kernel herding (Chen et al., 2010), and give rise to explicit, non-asymptotic error bounds for a wide variety of popular kernels including Gaussian, Matérn, and B-spline kernels. While $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}} \log^{\frac{d-1}{2}} n))$ -MMD coresets have been developed for specific $(\mathbf{k}_\star, \mathbb{P})$ pairings like the uniform distribution on $[0, 1]^d$ and an L^2 discrepancy kernel \mathbf{k}_\star (see Sec. 8.1), to the best of our knowledge, no prior $(n^{\frac{1}{2}}, o_p(n^{-\frac{1}{4}}))$ -MMD coreset constructions were known for the range of \mathbb{P} and \mathbf{k}_\star studied in this work.
2. **MMD error from square-root L^∞ error:** To derive our MMD guarantees for kernel thinning, we first establish an important link between MMD coresets for \mathbf{k}_\star and L^∞ coresets for \mathbf{k}_{rt} .

Definition 3 (L^∞ coreset). *For any kernel \mathbf{k} satisfying Assump. 1, probability measure μ on \mathbb{R}^d , and $z \in \mathbb{R}^d$, let $\mu\mathbf{k}(z) \triangleq \mathbb{E}_{X \sim \mu}[\mathbf{k}(X, z)]$. We call a sequence of n_{out} points in \mathbb{R}^d with empirical measure \mathbb{Q} an $(n_{\text{out}}, \varepsilon)$ - L^∞ coreset for (\mathbf{k}, \mathbb{P}) if $\|\mathbb{P}\mathbf{k} - \mathbb{Q}\mathbf{k}\|_\infty \leq \varepsilon$.*

Thm. 2 and Cor. 2 show that *any* L^∞ coreset for $(\mathbf{k}_{\text{rt}}, \mathbb{P})$ is also an MMD coreset for $(\mathbf{k}_\star, \mathbb{P})$ with quality depending on the tail decay of \mathbf{k}_{rt} and \mathbb{P} .
3. **Online vector balancing in Hilbert spaces:** As a building block for constructing high-quality coresets, we introduce and analyze a Hilbert space generalization of the self-balancing walk of Alweiss et al. (2021) to partition a sequence of functions (like $(\mathbf{k}_{\text{rt}}(x_i, \cdot))_{i=1}^n$) into nearly equal halves. Our analysis of this *self-balancing Hilbert walk* (SBHW, Alg. 3) in Thm. 3 may be of independent interest for solving the online vector balancing problem of Spencer (1977) in Hilbert spaces (Cor. 4).
4. **Efficient, near-optimal L^∞ coresets:** We then design a symmetrized version of SBHW for RKHSes—kernel halving—that delivers 2-thinned coresets with small L^∞ error (Alg. 2 and Thm. 4). The first stage of kernel thinning, KT-SPLIT, recursively applies kernel halving to \mathbf{k}_{rt} to obtain near-minimax-optimal L^∞ coresets in $\mathcal{O}(n^2)$ time with $\mathcal{O}(n \min(d, n))$ space (Cors. 5 and 6).

After describing our kernel and input point requirements in Sec. 2, we detail the kernel thinning and kernel halving algorithms in Sec. 3. Sec. 4 houses our main MMD guarantees, both for kernel thinning and for generic L^∞ square-root kernel coresets. We introduce and analyze the self-balancing Hilbert walk in Sec. 5 and present our main L^∞ guarantees for kernel halving and KT-SPLIT in Sec. 6. Sec. 7 complements our theoretical contributions with two vignettes illustrating the practical benefits of kernel thinning over (a) i.i.d. sampling in dimensions $d = 2$ through 100 and (b) standard MCMC thinning across twelve

experiments targeting challenging differential equation posterior distributions. We conclude with a discussion of our results, related work, and future directions in Sec. 8 and defer all proofs to the appendices.

Notation We define the shorthand $[n] \triangleq \{1, \dots, n\}$ for $n \in \mathbb{N}$, $a \wedge b \triangleq \min(a, b)$ for $a, b \in \mathbb{R}$, $\mathbb{R}_+ \triangleq \{x \in \mathbb{R} : x \geq 0\}$, and $\mathcal{B}(x; r) \triangleq \{y \in \mathbb{R}^d \mid \|x - y\|_2 < r\}$ for $r \in \mathbb{R}$. We use $\text{Vol}(\mathcal{B})$ to denote the volume of a compact set $\mathcal{B} \subset \mathbb{R}^d$. We let \mathcal{A}^c denote the complement of a set $\mathcal{A} \subset \mathbb{R}^d$ and $\mathbb{I}_{\mathcal{A}}(x) = 1$ if $x \in \mathcal{A}$ and 0 otherwise. We use $\Pr(\mathcal{E})$ to denote the probability of an event \mathcal{E} . For real-valued kernels \mathbf{k} and functions f on \mathbb{R}^d , we make frequent use of the norms $\|\mathbf{k}\|_\infty = \sup_{x, y \in \mathbb{R}^d} |\mathbf{k}(x, y)|$ and $\|f\|_\infty = \sup_{x \in \mathbb{R}^d} |f(x)|$. For $x > 0$, we use $\Gamma(x) = \int_0^\infty t^{x-1} e^{-t} dt$ to denote the Gamma function (with $\Gamma(n) = (n-1)!$ for $n \in \mathbb{N}$).

For two sequences of real numbers $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}}$, we say that a_n is of order b_n and write $a_n = \mathcal{O}(b_n)$ or $a_n \lesssim b_n$ to denote that $a_n \leq cb_n$ for all $n \in \mathbb{N}$ and some constant $c > 0$. We write $a_n = \Omega(b_n)$ if $b_n = \mathcal{O}(a_n)$ and $a_n = \Theta(b_n)$ when $a = \Omega(b_n)$ and $a = \mathcal{O}(b_n)$. Moreover, we use $a_n = \mathcal{O}_d(b_n)$, $a_n \lesssim_d b_n$, $a_n = \Omega_d(b_n)$, $a_n \gtrsim_d b_n$ to indicate dependency of underlying universal constant on d . We say $a_n = o(b_n)$ if $\lim_{n \rightarrow \infty} a_n/b_n = 0$. For a sequence of real-valued random variables $(X_n)_{n \in \mathbb{N}}$, we write $X_n = \mathcal{O}_P(a_n)$ or $X_n = \mathcal{O}(a_n)$ in probability, when $\frac{X_n}{a_n}$ is stochastically bounded, i.e., for all $\delta > 0$, there exists finite c_δ and n_δ such that $\Pr(|\frac{X_n}{a_n}| > c_\delta) < \delta$, for all $n > n_\delta$. We write $X_n = \Omega_P(a_n)$ if $1/X_n = \mathcal{O}_P(1/b_n)$ and $X_n = o_p(a_n)$ when $\frac{X_n}{a_n} \rightarrow 0$ in probability, i.e., for all $\varepsilon > 0$, $\Pr(|\frac{X_n}{a_n}| \geq \varepsilon) \rightarrow 0$.

We write *order (n, ε) -MMD (or L^∞) coresets* to mean an $(n, \mathcal{O}(\varepsilon))$ -MMD (or L^∞) coresets and append *in probability* to mean an $(n, \mathcal{O}_p(\varepsilon))$ -MMD (or L^∞) coresets.

2. Input Point and Kernel Requirements

Given a target distribution \mathbb{P} on \mathbb{R}^d , a kernel \mathbf{k}_\star satisfying Assump. 1, and a sequence of \mathbb{R}^d -valued input points $\mathcal{S}_n = (x_i)_{i=1}^n$ generated either randomly or deterministically, our goal is to identify a better-than-i.i.d. thinned MMD coresets, that is, a subsequence \mathcal{S}_{out} of size $n^{\frac{1}{2}}$ satisfying $\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathcal{S}_{\text{out}}) = o_p(n^{-\frac{1}{4}})$. When drawing asymptotic conclusions, we will view d as fixed and \mathcal{S}_n as a prefix of an infinite sequence of points $\mathcal{S}_\infty \triangleq (x_i)_{i=1}^\infty$.

2.1 Input point requirements

Our algorithms are designed to return high quality MMD coresets for the *input* \mathcal{S}_n . To translate these into high quality coresets for the target \mathbb{P} , it suffices, by the triangle inequality, for the input points to have quality $\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathcal{S}_n) = \mathcal{O}_p(n^{-\frac{1}{2}})$. As we discuss in Sec. 8.1, input sequences generated by i.i.d. sampling, kernel herding (Chen et al., 2010), Stein Point MCMC (Chen et al., 2019), and greedy sign selection (Karnin and Liberty, 2019) all satisfy this property. Moreover, we prove in App. B that an analogous guarantee holds for the iterates of a fast-mixing Markov chain.

Proposition 1 (MMD guarantee for MCMC). *Consider a homogeneous ϕ -irreducible geometrically ergodic Markov chain (Gallegos-Herrada et al., 2023, Thm. 1xi) with initial state x_0 , subsequent iterates \mathcal{S}_∞ , and stationary distribution \mathbb{P} . If \mathbf{k}_\star satisfies Assump. 1, then there exists a \mathbb{P} -almost everywhere finite function $c : \mathbb{R}^d \rightarrow (0, \infty]$ such that, for any given $n \in \mathbb{N}$ and $\delta \in (0, 1)$, $\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathcal{S}_n) \leq \sqrt{\frac{c(x_0)\|\mathbf{k}_\star\|_\infty \log(\varepsilon/\delta)}{n}}$ with probability $1 - \delta$ given x_0 .*

The *input radius*,

$$R_{\mathcal{S}_n} \triangleq \max_{x \in \mathcal{S}_n} \|x\|_2, \quad (2)$$

will also play an important role in our results. In particular, the growth rate of this radius as a function of n impacts the growth rate of our MMD bounds. Our next definition assigns familiar names to the most commonly encountered growth rates.

Definition 4 (Input radius growth rates). *We say the point sequence \mathcal{S}_∞ with prefixes \mathcal{S}_n for $n \in \mathbb{N}$ is COMPACT if $R_{\mathcal{S}_n} = \mathcal{O}_d(1)$, SUBGAUSS if $R_{\mathcal{S}_n} = \mathcal{O}_d(\sqrt{\log n})$, SUBEXP if $R_{\mathcal{S}_n} = \mathcal{O}_d(\log n)$, and HEAVYTAIL(ρ) with $\rho > 0$ if $R_{\mathcal{S}_n} = \mathcal{O}_d(n^{1/\rho})$.*

These growth rates are exactly those which arise with probability 1 when an input sequence is generated identically from \mathbb{P} with corresponding tail behavior or from a fast-mixing Markov chain targeting \mathbb{P} . Our proof of this result is given in App. C.

Proposition 2 (Almost sure radius growth). *Consider either (i) points \mathcal{S}_∞ sampled identically (but not necessarily independently) from \mathbb{P} with x_0 independent or (ii) a homogeneous ϕ -irreducible geometrically ergodic Markov chain with initial state x_0 , subsequent iterates \mathcal{S}_∞ , and stationary distribution \mathbb{P} . Then the following statements hold true for any non-negative c and ρ and \mathbb{P} -almost every x_0 .*

- (a) *If \mathbb{P} is compactly supported, then, with probability 1 conditional on x_0 , \mathcal{S}_∞ is COMPACT.*
- (b) *If $\mathbb{E}_{X \sim \mathbb{P}}[e^{c\|X\|_2^2}] < \infty$, then, with probability 1 conditional on x_0 , \mathcal{S}_∞ is SUBGAUSS.*
- (c) *If $\mathbb{E}_{X \sim \mathbb{P}}[e^{c\|X\|_2}] < \infty$, then, with probability 1 conditional on x_0 , \mathcal{S}_∞ is SUBEXP.*
- (d) *If $\mathbb{E}_{X \sim \mathbb{P}}[\|X\|_2^\rho] < \infty$, then, with probability 1 conditional on x_0 , \mathcal{S}_∞ is HEAVYTAIL(ρ).*

Finally, we will also require \mathcal{S}_∞ to be *oblivious*, that is, generated independently of any randomness in the thinning algorithm. To capture this assumption, we treat \mathcal{S}_∞ as fixed and deterministic hereafter. This treatment is without loss of generality since our results hold conditional on the observed values of $(x_i)_{i=1}^\infty$ when the points are random and oblivious.

2.2 Kernel requirements

We use the terms *reproducing kernel* and *kernel* interchangeably to indicate that \mathbf{k}_\star is symmetric and positive definite, i.e., that the kernel matrix $(\mathbf{k}_\star(z_i, z_j))_{i,j=1}^l$ is symmetric and positive semidefinite for any evaluation points $(z_i)_{i=1}^l$ in \mathbb{R}^d . In addition to \mathbf{k}_\star , our algorithm takes as input a *square-root kernel* for \mathbf{k}_\star .

Definition 5 (Square-root kernel). *We say a kernel $\mathbf{k}_{\text{rt}} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a square-root kernel for $\mathbf{k}_\star : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ if $\mathbf{k}_{\text{rt}}(x, \cdot)$ is square integrable for all $x \in \mathbb{R}^d$ with*

$$\mathbf{k}_\star(x, y) = \int_{\mathbb{R}^d} \mathbf{k}_{\text{rt}}(x, z) \mathbf{k}_{\text{rt}}(y, z) dz \quad \text{for all } x, y \in \mathbb{R}^d. \quad (3)$$

We highlight that a square-root kernel need not be unique and that its existence is an indication of a certain degree of smoothness in the target kernel \mathbf{k}_\star . One convenient tool for deriving square-root kernels is the notion of a *spectral density*.

Name of kernel $\mathbf{k}_*(x, y) = \kappa(x - y)$	Expression for $\kappa(z)$	Fourier transform $\widehat{\kappa}(\omega)$	Square-root kernel \mathbf{k}_{rt}
Gaussian (σ) : $\sigma > 0$	$\exp\left(-\frac{\ z\ _2^2}{2\sigma^2}\right)$	$\sigma^d \exp\left(-\frac{\sigma^2\ \omega\ _2^2}{2}\right)$	$\left(\frac{2}{\pi\sigma^2}\right)^{\frac{d}{4}}$ Gaussian $\left(\frac{\sigma}{\sqrt{2}}\right)$
Matérn (ν, γ) : $\nu > d, \gamma > 0$	$c_{\nu-\frac{d}{2}}(\gamma\ z\ _2)^{\nu-\frac{d}{2}} K_{\nu-\frac{d}{2}}(\gamma\ z\ _2)$	$\phi_{d,\nu,\gamma}(\gamma^2 + \ \omega\ _2^2)^{-\nu}$	$A_{\nu,\gamma,d}$ Matérn $\left(\frac{\nu}{2}, \gamma\right)$
B-spline ($2\beta + 1$) : $\beta \in 2\mathbb{N} + 1$	$S_{2\beta+2,d} \prod_{j=1}^d \otimes^{2\beta+2} \mathbb{I}_{[-\frac{1}{2}, \frac{1}{2}]}(z_j)$	$S'_{2\beta+2,d} \prod_{j=1}^d \frac{\sin^{2\beta+2}(\frac{\omega_j}{2})}{\omega_j^{2\beta+2}}$	$\widetilde{S}_{\beta,d}$ B-spline (β)

Table 1: Square-root kernels \mathbf{k}_{rt} for common target kernels \mathbf{k}_* . Each \mathbf{k}_* satisfies $\|\mathbf{k}_*\|_\infty = 1$, and the parameter range ensures the existence of \mathbf{k}_{rt} . Above, \otimes^ℓ denotes recursive convolution with ℓ function copies, K_a denotes the modified Bessel function of the third kind (Wendland, 2004, Def. 5.10), $c_b \triangleq \frac{2^{1-b}}{\Gamma(b)}$, $\phi_{d,\nu,\gamma} = \frac{c_{\nu-d/2}}{c_\nu} \gamma^{2\nu-d}$, $A_{\nu,\gamma,d} \triangleq \left(\frac{1}{4\pi} \gamma^2\right)^{d/4} \sqrt{\frac{\Gamma(\nu)}{\Gamma(\nu-d/2)}} \cdot \frac{\Gamma((\nu-d)/2)}{\Gamma(\nu/2)}$, $S'_{2\beta+2,d} \triangleq S_{2\beta+2,d} \cdot \left(\frac{4^{\beta+1}}{\sqrt{2\pi}}\right)^d$, and $\widetilde{S}_{\beta,d} \triangleq \frac{\sqrt{S_{2\beta+2,d}}}{S_{\beta+1,d}}$ where $S_{\beta,d}$ is defined in (89). See App. N for our derivation.

Definition 6 (Shift invariance and spectral density). *We call a kernel of the form $\mathbf{k}(x, y) = \kappa(x - y)$ for $\kappa : \mathbb{R}^d \rightarrow \mathbb{R}$ shift-invariant and say \mathbf{k} has spectral density $\widehat{\kappa}$ if κ is the Fourier transform of a finite measure with Lebesgue density $\widehat{\kappa}$, i.e., $\kappa(z) = \frac{1}{(2\pi)^{d/2}} \int e^{-i(\omega,z)} \widehat{\kappa}(\omega) d\omega$.*

As we show in Apps. N and Q, many familiar kernels admit spectral densities, including Gaussian, Matérn, B-spline, inverse multiquadric, sech, and Wendland’s compactly supported kernels. Moreover, by Bochner’s theorem (Bochner, 1933; Wendland, 2004, Thm. 6.6) and the Fourier inversion theorem (Wendland, 2004, Cor. 5.24), any continuous $\mathbf{k}_*(x, y) = \kappa(x - y)$ with absolutely integrable κ has a spectral density equal to the Fourier transform of κ . Our next result (proved in App. D) derives a square-root kernel for any shift-invariant \mathbf{k}_* with a square-root integrable spectral density.

Proposition 3 (Shift-invariant square-root kernels). *If a kernel $\mathbf{k}_*(x, y) = \kappa(x - y)$ admits a spectral density (Def. 6) $\widehat{\kappa}$ with $\int \sqrt{\widehat{\kappa}(\omega)} d\omega < \infty$, then $\mathbf{k}_{\text{rt}}(x, y) = \frac{\kappa_{\text{rt}}(x - y)}{(2\pi)^{d/4}}$ is a square-root kernel of \mathbf{k}_* for κ_{rt} the Fourier transform of $\sqrt{\widehat{\kappa}}$.*

Tab. 1 gives several examples of common kernels satisfying the conditions of Prop. 3 along with their associated square-root kernels. For example, if \mathbf{k}_* is Gaussian with bandwidth σ , then a rescaled Gaussian kernel with bandwidth $\frac{\sigma}{\sqrt{2}}$ is a valid choice for \mathbf{k}_{rt} . For simplicity, our results in the sequel assume the use of an exact square-root kernel \mathbf{k}_{rt} , but, as we detail in App. Q, it suffices to use the square-root of any kernel that dominates \mathbf{k}_* in the positive-definite order (see Def. 8). For example, we show in Prop. 4 of App. Q that a standard Matérn kernel is a suitable square-root dominating kernel for any sufficiently-smooth shift-invariant \mathbf{k}_* with absolutely integrable κ . In Tab. 5 of App. Q, we also derive convenient tailored square-root dominating kernels for inverse multiquadric, sech, and Wendland’s compactly supported kernels.

Finally, we define several kernel growth and decay properties that will be explicitly assumed in some of our results.

Assumption 2 (Lipschitz kernel). *The kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ admits a Lipschitz constant*

$$L_{\mathbf{k}} \triangleq \sup_{x,y,z} \frac{|\mathbf{k}(x,y) - \mathbf{k}(x,z)|}{\|y-z\|_2} < \infty. \quad (4)$$

Assumption 3 (Kernel tail decay). *The kernel \mathbf{k} satisfies Assump. 1 and, for each $\varepsilon > 0$,*

$$\begin{aligned} \max\{ \inf\{r : \sup_{\|x-y\|_2 \geq r} \sup_{x,y} |\mathbf{k}(x,y)| \leq \varepsilon\}, \quad \inf\{r : \tau_{\mathbf{k}}(r) \leq \varepsilon\} \} < \infty, \\ \text{where } \tau_{\mathbf{k}}(r) \triangleq (\sup_x \int_{\|y\|_2 \geq r} \mathbf{k}^2(x, x-y) dy)^{\frac{1}{2}} \quad \text{for } r \geq 0. \end{aligned} \quad (5)$$

The following definition gives familiar names to commonly encountered tail decay rates.

Definition 7 (Kernel tail decay rate). *For a kernel \mathbf{k} satisfying Assump. 3, define*

$$\begin{aligned} R_{\mathbf{k},n} \triangleq \inf\{r : \sup_{\|x-y\|_2 \geq r} \sup_{x,y} |\mathbf{k}(x,y)| \leq \frac{\|\mathbf{k}\|_{\infty}}{n}\}, \quad R'_{\mathbf{k},n} \triangleq \inf\{r : \tau_{\mathbf{k}}(r) \leq \frac{\|\mathbf{k}\|_{\infty}}{\sqrt{n}}\}, \\ \text{and } R_{\mathbf{k},n}^{\dagger} \triangleq \max\{R_{\mathbf{k},n}, R'_{\mathbf{k},n}\} \end{aligned} \quad (6)$$

for $\tau_{\mathbf{k}}$ defined in (5). We say \mathbf{k} is (a) COMPACT if $R_{\mathbf{k},n}^{\dagger} = \mathcal{O}_d(1)$, (b) SUBGAUSS if $R_{\mathbf{k},n}^{\dagger} = \mathcal{O}_d(\sqrt{\log n})$, (c) SUBEXP if $R_{\mathbf{k},n}^{\dagger} = \mathcal{O}_d(\log n)$, (d) HEAVYTAIL(ρ) for $\rho > 0$ if $R_{\mathbf{k},n}^{\dagger} = \mathcal{O}_d(n^{1/\rho})$, and (e) SUBPOLY if $\log R_{\mathbf{k},n}^{\dagger} = \mathcal{O}_d(\log n)$. Notably, any COMPACT, SUBGAUSS, SUBEXP, or HEAVYTAIL(ρ) \mathbf{k} is also SUBPOLY.

Remark 1 (\mathbf{k}_{rt} tail decay implies \mathbf{k}_{\star} boundedness). *If \mathbf{k}_{rt} satisfying Assump. 3 is a square-root kernel of \mathbf{k}_{\star} , then there exists a finite r for which $\mathbf{k}_{\star}(x,x) = \int \mathbf{k}_{\text{rt}}^2(x, x-y) dy \leq \|\mathbf{k}_{\text{rt}}\|_{\infty}^2 \text{Vol}(\mathcal{B}(0; r)) + \tau_{\mathbf{k}_{\text{rt}}}^2(r) < \infty$.*

Popular examples of COMPACT, SUBGAUSS, and SUBEXP \mathbf{k} are B-spline, Gaussian, and Matérn kernels respectively (see Tab. 2). Moreover, one can directly verify that an inverse multiquadric $\mathbf{k}(x,y) = (\gamma^2 + \|x-y\|_2^2)^{-\nu}$ with $\gamma > 0$ and $\nu > \frac{d}{4}$ is HEAVYTAIL($2\nu \wedge (4\nu-d)$).

3. Kernel Thinning

Our solution to the thinned coreset problem is *kernel thinning*, described in Alg. 1. Given a thinning parameter $m \in \mathbb{N}$, kernel thinning proceeds in two stages: KT-SPLIT and KT-SWAP.

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

Input: kernels $(\mathbf{k}_{\star}, \mathbf{k}_{\text{rt}})$, input points $\mathcal{S}_n = (x_i)_{i=1}^n$, thinning parameter $m \in \mathbb{N}$, probabilities $(\delta_i)_{i=1}^{\lfloor \frac{n}{2} \rfloor}$

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

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

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




Figure 1: Overview of KT-SPLIT. (Left) KT-SPLIT recursively partitions its input \mathcal{S}_n into 2^m balanced coresets $\mathcal{S}^{(m,\ell)}$ of size $\lfloor \frac{n}{2^m} \rfloor$. (Right) In Sec. 6, we interpret each coreset $\mathcal{S}^{(m,\ell)}$ as the output of repeated *kernel halving*: on each halving round, remaining points are paired, and one point from each pair is selected using non-uniform randomness.

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

Input: kernel \mathbf{k}_{rt} , point sequence $\mathcal{S}_n = (x_i)_{i=1}^n$, thinning parameter $m \in \mathbb{N}$, probabilities $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$

$\mathcal{S}^{(j,\ell)} \leftarrow \{\}$ for $0 \leq j \leq m$ and $1 \leq \ell \leq 2^j$ // Empty coresets: $\mathcal{S}^{(j,\ell)}$ has size $\lfloor \frac{i}{2^{j-1}} \rfloor$ after i rounds

$\sigma_{j,\ell} \leftarrow 0$ 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_{2i-1}); $\mathcal{S}^{(0,1)}$.append(x_{2i})

 // Every 2^{j-1} rounds add point from parent coreset $\mathcal{S}^{(j-1,\ell)}$ to each child $\mathcal{S}^{(j,2\ell-1)}$, $\mathcal{S}^{(j,2\ell)}$

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**

$(\mathcal{S}, \mathcal{S}') \leftarrow (\mathcal{S}^{(j-1,\ell)}, \mathcal{S}^{(j,2\ell-1)})$; $(x, x') \leftarrow \text{get_last_two_points}(\mathcal{S})$

 // Compute swapping threshold \mathbf{a}

$\mathbf{a}, \sigma_{j,\ell} \leftarrow \text{get_swap_params}(\sigma_{j,\ell}, \mathbf{b}, \delta_{|S|/2} \cdot \frac{2^{j-1}}{m})$ for
 $\mathbf{b}^2 = \mathbf{k}_{\text{rt}}(x, x) + \mathbf{k}_{\text{rt}}(x', x') - 2\mathbf{k}_{\text{rt}}(x, x')$

 // Assign one point to each child after probabilistic swapping

$\alpha \leftarrow \mathbf{k}_{\text{rt}}(x', x') - \mathbf{k}_{\text{rt}}(x, x) + \sum_{y \in \mathcal{S}} (\mathbf{k}_{\text{rt}}(y, x) - \mathbf{k}_{\text{rt}}(y, x')) - 2\sum_{z \in \mathcal{S}'} (\mathbf{k}_{\text{rt}}(z, x) - \mathbf{k}_{\text{rt}}(z, x'))$

$(x, x') \leftarrow (x', x)$ with probability $\min(1, \frac{1}{2}(1 - \frac{\alpha}{\mathbf{a}})_+)$

$\mathcal{S}^{(j,2\ell-1)}$.append(x); $\mathcal{S}^{(j,2\ell)}$.append(x')

end

end

end

return $(\mathcal{S}^{(m,\ell)})_{\ell=1}^{2^m}$, candidate coresets of size $\lfloor n/2^m \rfloor$

function `get_swap_params`($\sigma, \mathbf{b}, \delta$):

$\mathbf{a} \leftarrow \max(\mathbf{b}\sigma\sqrt{2\log(2/\delta)}, \mathbf{b}^2)$

$\sigma^2 \leftarrow \sigma^2 + \mathbf{b}^2(1 + (\mathbf{b}^2 - 2\mathbf{a})\sigma^2/\mathbf{a}^2)_+$

return (\mathbf{a}, σ)

KT-SPLIT The first stage, KT-SPLIT, is an initialization stage that partitions the input sequence $\mathcal{S}_n = (x_i)_{i=1}^n$ into 2^m balanced candidate coresets, each of size $\lfloor \frac{n}{2^m} \rfloor$.³ As depicted in Fig. 1, this partitioning is carried out recursively in m rounds, first dividing the input sequence in half, then halving those halves into quarters, and so on until coresets of size $\lfloor \frac{n}{2^m} \rfloor$ are produced. The details of KT-SPLIT can appear a bit complicated as, in practice, all m halving rounds are carried out concurrently in an online manner. However, under the hood, each candidate coreset is generated by recursively applying a new simple subroutine called *kernel halving*.

Kernel halving Kernel halving (KH, Alg. 2) is a simple randomized procedure for dividing an input sequence \mathcal{S}_n into two balanced, equal-sized coresets $\mathcal{S}^{(1)}$ and $\mathcal{S}^{(2)}$ using a kernel \mathbf{k} . KH begins with two empty coresets $\mathcal{S}^{(1)}$ and $\mathcal{S}^{(2)}$ and adds points (x, x') from the input sequence two at a time, assigning one point from each pair to each coreset. To encourage balance between the coresets during generation, KH effectively computes which assignment of (x, x') leads to a smaller $\text{MMD}_{\mathbf{k}}(\mathcal{S}^{(1)}, \mathcal{S}^{(2)})$ and then favors that assignment using non-uniform randomness. More precisely, on the i -th step with $(x, x') = (x_{2i-1}, x_{2i})$, KH computes the imbalance contrast

$$\alpha_i = i^2 (\text{MMD}_{\mathbf{k}}^2(\mathcal{S}^{(1)} \cup \{x\}, \mathcal{S}^{(2)} \cup \{x'\}) - \text{MMD}_{\mathbf{k}}^2(\mathcal{S}^{(1)} \cup \{x'\}, \mathcal{S}^{(2)} \cup \{x\}))$$

and then adds x to $\mathcal{S}^{(1)}$ and x' to $\mathcal{S}^{(2)}$ or x' to $\mathcal{S}^{(1)}$ and x to $\mathcal{S}^{(2)}$ with probability biased toward the more balanced outcome. We refer to this step as *probabilistic swapping* in the algorithm statements. The exact value of this probability depends on a *swapping threshold* α_i that is produced automatically each round based on the user-supplied inputs $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$. In Sec. 4.1, we will learn how to set these inputs to achieve better balance than standard thinning or uniform subsampling, and in Sec. 6.1 we will discuss the role of α_i and its generation procedure `get_swap_params` in achieving this balance.

Notably, in the context of KT-SPLIT, KH is run specifically with the square-root kernel \mathbf{k}_{rt} rather than the target kernel \mathbf{k}_* . This choice enables us to take advantage of the strong L^∞ balance properties established for KH in Sec. 6 and the close connection between square-root L^∞ error and target MMD error revealed in Sec. 4.2.

KT-SWAP The second stage, KT-SWAP, refines the candidate coresets produced by KT-SPLIT in three steps. First, KT-SWAP adds a baseline coreset of size $\lfloor \frac{n}{2^m} \rfloor$ to the candidate list (for example, one produced by standard thinning or uniform subsampling) to ensure that the KT-SWAP output is never worse than that of the baseline. Next, it selects the candidate coreset closest to \mathcal{S}_n in terms of $\text{MMD}_{\mathbf{k}_*}$. Finally, it refines the selected coreset by replacing each coreset point in turn with the best alternative in \mathcal{S}_n , as measured by $\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \cdot)$. This stage serves to greedily improve upon the MMD of the initial KT-SPLIT candidates, and, when computable, $\text{MMD}_{\mathbf{k}_*}$ to the target distribution \mathbb{P} can be substituted for the surrogate $\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \cdot)$ throughout.

Complexity For any m , the time complexity of kernel thinning is dominated by $\mathcal{O}(n^2)$ kernel evaluations, while the space complexity is $\mathcal{O}(n \min(d, n))$, achieved by storing the smaller of the input sequence $(x_i)_{i=1}^n$ and the kernel matrix $(\mathbf{k}_{\text{rt}}(x_i, x_j))_{i,j=1}^n$. In addition,

3. When 2^m does not evenly divide n , the final $n - 2^m \lfloor \frac{n}{2^m} \rfloor$ points are discarded.

Algorithm 2: Kernel Halving

Input: kernel \mathbf{k} , point sequence $\mathcal{S}_n = (x_i)_{i=1}^n$, probability sequence $(\delta_i)_{i=1}^{\lfloor n/2 \rfloor}$
 $\mathcal{S}^{(1)}, \mathcal{S}^{(2)} \leftarrow \{\}; \quad \psi_0 \leftarrow \mathbf{0} \in \mathcal{H} \quad //$ Initialize empty coresets: $\mathcal{S}^{(1)}, \mathcal{S}^{(2)}$ have size i after round i
 $\sigma_0 \leftarrow 0 \quad //$ Swapping parameter
for $i = 1, 2, \dots, \lfloor n/2 \rfloor$ **do**
 $//$ Construct kernel difference function using next two points
 $(x, x') \leftarrow (x_{2i-1}, x_{2i}); \quad f_i \leftarrow \mathbf{k}(x_{2i-1}, \cdot) - \mathbf{k}(x_{2i}, \cdot); \quad \eta_i \leftarrow -1$
 $//$ Compute swapping threshold \mathbf{a}_i
 $\mathbf{a}_i, \sigma_i \leftarrow \text{get_swap_params}(\sigma_{i-1}, \mathbf{b}, \delta_i)$ with $\mathbf{b}^2 = \|f_i\|_{\mathbf{k}}^2 = \mathbf{k}(x, x) + \mathbf{k}(x', x') - 2\mathbf{k}(x, x')$
 $//$ Compute RKHS inner product $\langle \psi_{i-1}, f_i \rangle_{\mathbf{k}}$, which has a simple form
 $\alpha_i \leftarrow \sum_{j=1}^{2i-2} (\mathbf{k}(x_j, x) - \mathbf{k}(x_j, x')) - 2 \sum_{z \in \mathcal{S}^{(1)}} (\mathbf{k}(z, x) - \mathbf{k}(z, x'))$
 $//$ Assign one point to each coreset after probabilistic swapping
 $(x, x') \leftarrow (x', x)$ and $\eta_i \leftarrow 1$ with probability $\min(1, \frac{1}{2}(1 - \frac{\alpha_i}{\mathbf{a}_i})_+)$
 $\mathcal{S}^{(1)}.append(x); \quad \mathcal{S}^{(2)}.append(x'); \quad \psi_i \leftarrow \psi_{i-1} + \eta_i f_i \quad //$
 $\psi_i = \sum_{x' \in \mathcal{S}^{(2)}} \mathbf{k}(x', \cdot) - \sum_{x \in \mathcal{S}^{(1)}} \mathbf{k}(x, \cdot)$
end
return $\mathcal{S}^{(1)}$, coreset of size $\lfloor n/2 \rfloor$

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

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

scaling either \mathbf{k}_* or \mathbf{k}_{t} by a positive multiplier has no impact on Alg. 1, so the kernels need only be specified up to arbitrary rescalings.

4. MMD Guarantees

We are now prepared to present our main MMD guarantees.

4.1 MMD guarantees for kernel thinning

Our first main result, proved in App. E, bounds the MMD of a kernel thinning coreset in terms of the input (2) and kernel (6) radii, the combined radii

$$R_{\mathcal{S}_n, \mathbf{k}, n} \triangleq \min(R_{\mathcal{S}_n}, n^{1+\frac{1}{d}} R_{\mathbf{k}, n} + n^{\frac{1}{d}} \frac{\|\mathbf{k}\|_{\infty}}{L_{\mathbf{k}}}) \quad \text{and} \quad R'_{\mathcal{S}_n, \mathbf{k}, n} \triangleq \max(R_{\mathcal{S}_n}, R'_{\mathbf{k}, n}), \quad (7)$$

and the kernel thinning inflation factor

$$\mathfrak{M}_{\mathbf{k}}(n, m, d, \delta, \delta', R) \triangleq 2 \mathbb{I}(\frac{n}{2^m} \notin \mathbb{N}) + 37 \sqrt{\log(\frac{6m}{2^m \delta})} \left[\sqrt{\log(\frac{4}{\delta'})} + 5 \sqrt{d \log(2 + \frac{2L_{\mathbf{k}}}{\|\mathbf{k}\|_{\infty}} (R_{\mathbf{k}, n} + R))} \right], \quad (8)$$

defined for any kernel \mathbf{k} satisfying Assumps. 2 and 3, $n, m, d \in \mathbb{N}$, $\delta \in (0, \frac{6m}{2^m}]$, $\delta' \in (0, 1]$, and $R \geq 0$.

Theorem 1 (MMD guarantee for kernel thinning). *Consider kernel thinning (Alg. 1) with \mathbf{k}_{\star} satisfying Assump. 1, \mathbf{k}_{rt} a square-root kernel of \mathbf{k}_{\star} , $\delta^{\star} \triangleq \min_i \delta_i$, and $n_{\text{out}} \triangleq \lfloor \frac{n}{2^m} \rfloor$ for $m \leq \lfloor \log_2 n \rfloor$. If \mathbf{k}_{rt} satisfies Assumps. 2 and 3, then, for any fixed $\delta' \in (0, 1)$, we have*

$$\text{MMD}_{\mathbf{k}_{\star}}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) \leq \frac{\|\mathbf{k}_{\text{rt}}\|_{\infty}}{n_{\text{out}}} \left[2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (R'_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n_{\text{out}}})^{\frac{d}{2}} \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, m, d, \delta^{\star}, \delta', R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n}) \right], \quad (9)$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} n_{\text{out}}} \delta_i$.

Remark 2 (Guarantee for target \mathbb{P}). *A guarantee for any target distribution \mathbb{P} follows directly from the triangle inequality, $\text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathcal{S}_{\text{KT}}) \leq \text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathcal{S}_n) + \text{MMD}_{\mathbf{k}_{\star}}(\mathcal{S}_n, \mathcal{S}_{\text{KT}})$.*

Remark 3 (Comparison with baseline thinning). *The KT-SWAP step ensures that, deterministically, $\text{MMD}_{\mathbf{k}_{\star}}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) \leq \text{MMD}_{\mathbf{k}_{\star}}(\mathcal{S}_n, \mathcal{S}_{\text{base}})$ for $\mathcal{S}_{\text{base}}$ a baseline thinned core-set of size n_{out} . Therefore, we additionally have $\text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathcal{S}_{\text{KT}}) \leq 2 \text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathcal{S}_n) + \text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathcal{S}_{\text{base}})$.*

Remark 4 (Finite-time and anytime guarantees). *To obtain a success probability of at least $1 - \delta$ with $\delta' = \frac{\delta}{2}$, it suffices to choose $\delta_i = \frac{\delta}{n}$ when the input size n is known in advance and $\delta_i = \frac{m\delta}{2^{m+2}(i+1)\log^2(i+1)}$ when the input size n is not known in advance (but is chosen independently of the randomness used in kernel thinning). In either case, $\delta^{\star} \leq \frac{6m}{2^m}$ is a valid argument to $\mathfrak{M}_{\mathbf{k}_{\text{rt}}}$ (8). See App. F for our proof.*

Our next corollary, proved in App. G, translates Thm. 1 into specific rates of MMD decay depending on the radius growth of \mathcal{S}_{∞} and the tail decay of \mathbf{k}_{rt} .

Corollary 1 (MMD rates for kernel thinning). *Under the notation and assumptions of Thm. 1, consider a sequence of kernel thinning runs (Alg. 1), indexed by $n \in \mathbb{N}$, with $m = \lfloor \frac{1}{2} \log_2 n \rfloor$, $\log(1/\delta^{\star}) = \mathcal{O}(\log n)$, and $\sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\lfloor n/2^j \rfloor} \delta_i = o(1)$ as $n \rightarrow \infty$. If \mathcal{S}_{∞} and \mathbf{k}_{rt} respectively satisfy one of the radius growth (Def. 4) and tail decay (Def. 7) conditions in the table below, then $\text{MMD}_{\mathbf{k}_{\star}}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) = \mathcal{O}_d(\varepsilon_{\text{MMD}, n})$ in probability where $\varepsilon_{\text{MMD}, n}$ is the corresponding table entry.*

$\varepsilon_{\text{MMD}, n}$	COMPACT \mathcal{S}_{∞} $R_{\mathcal{S}_n} \lesssim_d 1$	SUBGAUSS \mathcal{S}_{∞} $R_{\mathcal{S}_n} \lesssim_d \sqrt{\log n}$	SUBEXP \mathcal{S}_{∞} $R_{\mathcal{S}_n} \lesssim_d \log n$	HEAVYTAIL(ρ) \mathcal{S}_{∞} $R_{\mathcal{S}_n} \lesssim_d n^{1/\rho}$
COMPACT \mathbf{k}_{rt} $R_{\mathbf{k}_{\text{rt}}, n}^{\dagger} \lesssim_d 1$	$\sqrt{\frac{\log n}{n}}$	$(\log n)^{\frac{d+2}{4}} \sqrt{\frac{\log \log n}{n}}$	$(\log n)^{\frac{d+1}{2}} \sqrt{\frac{\log \log n}{n}}$	$\frac{\log n}{\sqrt{n^{1-d/\rho}}}$
SUBGAUSS \mathbf{k}_{rt} $R_{\mathbf{k}_{\text{rt}}, n}^{\dagger} \lesssim_d \sqrt{\log n}$	$(\log n)^{\frac{d+2}{4}} \sqrt{\frac{\log \log n}{n}}$	$(\log n)^{\frac{d+2}{4}} \sqrt{\frac{\log \log n}{n}}$	$(\log n)^{\frac{d+1}{2}} \sqrt{\frac{\log \log n}{n}}$	$\frac{\log n}{\sqrt{n^{1-d/\rho}}}$
SUBEXP \mathbf{k}_{rt} $R_{\mathbf{k}_{\text{rt}}, n}^{\dagger} \lesssim_d \log n$	$(\log n)^{\frac{d+1}{2}} \sqrt{\frac{\log \log n}{n}}$	$(\log n)^{\frac{d+1}{2}} \sqrt{\frac{\log \log n}{n}}$	$(\log n)^{\frac{d+1}{2}} \sqrt{\frac{\log \log n}{n}}$	$\frac{\log n}{\sqrt{n^{1-d/\rho}}}$
HEAVYTAIL(ρ') \mathbf{k}_{rt} $R_{\mathbf{k}_{\text{rt}}, n}^{\dagger} \lesssim_d n^{1/\rho'}$	$\frac{\log n}{\sqrt{n^{1-d/\rho'}}$	$\frac{\log n}{\sqrt{n^{1-d/\rho'}}$	$\frac{\log n}{\sqrt{n^{1-d/\rho'}}$	$\frac{\log n}{\sqrt{n^{1-d/(\rho \wedge \rho')}}}$

Square-root kernel \mathbf{k}_{rt}	$\ \mathbf{k}_{\text{rt}}\ _\infty$	$(R_{\mathbf{k}_{\text{rt}},n}, R'_{\mathbf{k}_{\text{rt}},n})$ $\underset{\sim}{\approx}$	$\mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, \frac{1}{2} \log_2 n, d, \frac{\delta}{n}, \delta', R)$ $\underset{\sim}{\approx}$
$(\frac{2}{\pi\sigma^2})^{\frac{d}{4}}$ Gaussian $(\frac{\sigma}{\sqrt{2}})$	$(\frac{2}{\pi\sigma^2})^{\frac{d}{4}}$	$(\sigma\sqrt{\log n}, \sigma\sqrt{d + \log n})$	$\sqrt{\log(\frac{n}{\delta})[\log(\frac{1}{\delta'}) + d \log(\sqrt{\log n + \frac{R}{\sigma}})]}$
$A_{\nu,\gamma,d}$ Matérn $(\frac{\nu}{2}, \gamma)$	$\nu(\frac{\gamma^2}{2\pi(a-1)})^{\frac{d}{4}}$	$(\gamma^{-1}(\log n + a \log(1+a)), \gamma^{-1}(a + \log n + E))$	$\sqrt{\log(\frac{n}{\delta})[\log(\frac{1}{\delta'}) + d \log(\log n + B + \gamma R)]}$
$\tilde{S}_{\beta,d}$ B-spline (β)	c_β^d	$(\frac{1}{2}\sqrt{d}(\beta+1), \frac{1}{2}\sqrt{d}(\beta+1))$	$\sqrt{\log(\frac{n}{\delta})[\log(\frac{1}{\delta'}) + d \log(d\beta + \sqrt{d}R)]}$

Table 2: Explicit bounds on Thm. 1 quantities for common kernels. Here, $A_{\nu,\gamma,d}$, and $\tilde{S}_{\beta,d}$ are as in Tab. 1, $a \triangleq \frac{1}{2}(\nu-d)(> 1)$, $B \triangleq a \log(1+a)$, $E \triangleq d \log(\frac{\sqrt{2e\pi}}{\gamma}) + \log(\frac{(\nu-2)^{\nu-\frac{3}{2}}}{(2(a-1))^{2a-1}d^{\frac{d}{2}+1}})$, $c_1 = \frac{2}{\sqrt{3}}$, and $c_\beta < 1$ for $\beta > 1$ (see (120)). See App. O for our derivation.

Remark 5 (Probability parameters). *The condition $(\star) \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\lfloor n/2^j \rfloor} \delta_i = o(1)$ is satisfied when $\delta_1 = \dots = \delta_{\lfloor \frac{n}{2} \rfloor} = o(\frac{1}{n})$. Hence, both $\log(1/\delta^*) = \mathcal{O}(\log n)$ and (\star) are satisfied when, for example, $\delta_1 = \dots = \delta_{\lfloor \frac{n}{2} \rfloor} = \frac{1}{n \log \log n}$.*

Cor. 1 shows that kernel thinning returns an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}}\sqrt{\log n}))$ -MMD coresets in probability when \mathcal{S}_∞ and \mathbf{k}_{rt} are compactly supported. For fixed d , this guarantee significantly improves upon the baseline $\Omega_p(n^{-\frac{1}{4}})$ rates of i.i.d. sampling and standard MCMC thinning and matches the minimax lower bounds of Sec. 8.1 up to a $\sqrt{\log n}$ term and constants depending on d . For example, when \mathcal{S}_∞ is drawn i.i.d. from \mathbb{P} , kernel thinning is nearly minimax optimal amongst *all* distributional approximations (even weighted coresets and non-coreset approximations) that depend on \mathbb{P} only through n i.i.d. input points (Tolstikhin et al., 2017, Thms. 1 and 6).

More generally, when \mathcal{S}_∞ and \mathbf{k}_{rt} are SUBGAUSS, SUBEXP, or HEAVYTAIL(ρ) with $\rho > 2d$, Cor. 1 shows that the kernel thinning provides an MMD error of $\mathcal{O}_d(n^{-\frac{1}{2}}\sqrt{(\log n)^{d/2+1} \log \log n})$, $\mathcal{O}_d(n^{-\frac{1}{2}}\sqrt{(\log n)^{d+1} \log \log n})$, and $\mathcal{O}_d(n^{-\frac{1}{2}}n^{\frac{d}{2\rho}} \log n)$ in probability with output coresets of size \sqrt{n} . In each case, we find that kernel thinning significantly improves upon an $\Omega_p(n^{-\frac{1}{4}})$ baseline when n is sufficiently large relative to d and, by Rem. 3, is never significantly worse than the baseline when n is small. Our SUBEXP guarantees also resemble the classical quasi-Monte Carlo guarantees for the uniform distribution on $[0, 1]^d$ (see Sec. 8.1) but allow for non-uniform and unbounded target distributions \mathbb{P} .

Thm. 1 also allows us to derive more precise, explicit error bounds for specific kernels. For example, for the popular Gaussian, Matérn, and B-spline kernels, Tab. 2 provides explicit bounds on each kernel-dependent quantity in Thm. 1: $\|\mathbf{k}_{\text{rt}}\|_\infty$, the kernel radii $(R_{\mathbf{k}_{\text{rt}},n}, R'_{\mathbf{k}_{\text{rt}},\sqrt{n}})$, and the inflation factor $\mathfrak{M}_{\mathbf{k}_{\text{rt}}}$.

4.2 MMD coresets from square-root L^∞ coresets

Thm. 1 builds on a second key result, proved in App. H, that bounds MMD error for \mathbf{k}_\star in terms of L^∞ error for the *square-root* kernel \mathbf{k}_{rt} .

Theorem 2 (MMD guarantee for square-root L^∞ approximations). *Suppose \mathbf{k}_\star satisfying Assump. 1 has a square-root kernel \mathbf{k}_{rt} satisfying Assump. 3. Then for any distributions μ*

and ν on \mathbb{R}^d and scalars $r, a, b \geq 0$ with $a + b = 1$,

$$\text{MMD}_{\mathbf{k}_\star}(\mu, \nu) \leq e_d r^{\frac{d}{2}} \cdot \|\mu \mathbf{k}_{\text{rt}} - \nu \mathbf{k}_{\text{rt}}\|_\infty + 2\tau_{\mathbf{k}_{\text{rt}}}(ar) + 2\|\mathbf{k}_\star\|_\infty^{\frac{1}{2}} \cdot \max\{\tau_\mu(br), \tau_\nu(br)\}, \quad (10)$$

where $e_d \triangleq (\pi^{d/2}/\Gamma(d/2 + 1))^{\frac{1}{2}}$ decreases super-exponentially in d and $\tau_\mu(r) \triangleq \mu(\mathcal{B}^c(0, r))$.

Importantly, Thm. 2 implies that *any* L^∞ coreset for the square-root kernel \mathbf{k}_{rt} , even one not produced by kernel thinning, is also an MMD coreset for the target kernel \mathbf{k}_\star with MMD error depending on the tail decay of $(\mathbf{k}_{\text{rt}}, \mu, \nu)$. Cor. 2 summarizes the implications of Thm. 2 for common classes of tail decay. See App. I for the proof with explicit constants.

Corollary 2 (MMD error from square-root L^∞ error). *Under the setting and assumptions of Thm. 2, define the L^∞ error $\varepsilon \triangleq \|\mu \mathbf{k}_{\text{rt}} - \nu \mathbf{k}_{\text{rt}}\|_\infty$ and the tail decay function $\tilde{\tau}(r) \triangleq \tau_{\mathbf{k}_{\text{rt}}}(r) + \|\mathbf{k}_\star\|_\infty^{1/2} \max\{\tau_\mu(r), \tau_\nu(r)\}$ for $r \geq 0$. Then the following implications hold for any nonnegative c and ρ .*

TAIL DECAY	COMPACT	SUBGAUSS	SUBEXP	HEAVYTAIL(ρ)
$\tilde{\tau}(r) \lesssim$	$\mathbb{I}(r \leq c)$	e^{-cr^2}	e^{-cr}	$r^{-\rho}$
$\Rightarrow \text{MMD}_{\mathbf{k}_\star}(\mu, \nu) \lesssim_d$	ε	$\varepsilon(\log \frac{1}{\varepsilon})^{\frac{d}{4}}$	$\varepsilon(\log \frac{1}{\varepsilon})^{\frac{d}{2}}$	$\varepsilon^{\frac{2\rho}{d+2\rho}}$

Remark 6 (Tail decay from finite moments). *By Markov's inequality (Durrett, 2019, Thm. 1.6.4), $\tilde{\tau}$ has (i) COMPACT decay when \mathbf{k}_{rt} is COMPACT and μ and ν have compact support; (ii) SUBGAUSS decay when \mathbf{k}_{rt} is SUBGAUSS and $\mathbb{E}_{X \sim \mu}[e^{c\|X\|_2^2}], \mathbb{E}_{X \sim \nu}[e^{c\|X\|_2^2}] < \infty$; (iii) SUBEXP decay when \mathbf{k}_{rt} is SUBEXP and $\mathbb{E}_{X \sim \mu}[e^{c\|X\|_2}], \mathbb{E}_{X \sim \nu}[e^{c\|X\|_2}] < \infty$; and (iv) HEAVYTAIL(ρ) decay when \mathbf{k}_{rt} is HEAVYTAIL(ρ) and $\mathbb{E}_{X \sim \mu}[\|X\|_2^\rho], \mathbb{E}_{X \sim \nu}[\|X\|_2^\rho] < \infty$.*

Cor. 2 highlights that MMD quality for (\mathbf{k}_\star, μ) is of the same order as L^∞ quality for $(\mathbf{k}_{\text{rt}}, \mu)$ when $\mathbf{k}_{\text{rt}}, \mu$, and the approximation ν have compact support. MMD quality then degrades naturally as the tail behavior worsens. In Secs. 5 and 6, we show that, with high probability, KT-SPLIT provides a high-quality L^∞ coreset for $(\mathbf{k}_{\text{rt}}, \mathcal{S}_n)$ and hence, by Cor. 2, also provides a high-quality MMD coreset for $(\mathbf{k}_\star, \mathcal{S}_n)$.

5. Self-balancing Hilbert Walk

To exploit the L^∞ -MMD connection revealed in Thm. 2, we now turn our attention to constructing high-quality thinned L^∞ coresets. Our strategy relies on a new Hilbert space generalization of the self-balancing walk of Alweiss et al. (2021). We dedicate this section to defining and analyzing this *self-balancing Hilbert walk*, and we detail its connection to kernel thinning in Sec. 6.

Alweiss et al. (2021, Thm. 1.2) introduced a randomized algorithm called the *self-balancing walk* that takes as input a streaming sequence of Euclidean vectors $x_i \in \mathbb{R}^d$ with $\|x_i\|_2 \leq 1$ and outputs a online sequence of random assignments $\eta_i \in \{-1, 1\}$ satisfying

$$\|\sum_{i=1}^n \eta_i x_i\|_\infty \lesssim \sqrt{\log(d/\delta) \log(n/\delta)} \quad \text{with probability at least } 1 - \delta. \quad (11)$$

Algorithm 3: Self-balancing Hilbert Walk

Input: sequence of functions $(f_i)_{i=1}^n$ in a Hilbert space \mathcal{H} , threshold sequence $(\mathbf{a}_i)_{i=1}^n$
 $\psi_0 \leftarrow \mathbf{0} \in \mathcal{H}$
for $i = 1, 2, \dots, n$ **do**
 $\alpha_i \leftarrow \langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}$ // Compute Hilbert space inner product
 if $|\alpha_i| > \mathbf{a}_i$:
 $\psi_i \leftarrow \psi_{i-1} - f_i \cdot \alpha_i / \mathbf{a}_i$
 else:
 $\eta_i \leftarrow 1$ with probability $\frac{1}{2}(1 - \alpha_i / \mathbf{a}_i)$ and $\eta_i \leftarrow -1$ otherwise
 $\psi_i \leftarrow \psi_{i-1} + \eta_i f_i$
end
return ψ_n , combination of signed input functions

Since our ultimate aim is to combine kernel functions, we define a suitable Hilbert space generalization in Alg. 3.

Given a streaming sequence of functions f_i in an arbitrary Hilbert space \mathcal{H} with a norm $\|\cdot\|_{\mathcal{H}}$, this *self-balancing Hilbert walk* (SBHW) outputs a streaming sequence of signed function combinations ψ_i satisfying the following desirable properties established in App. J.

Theorem 3 (Self-balancing Hilbert walk properties). *Consider the self-balancing Hilbert walk (Alg. 3) with each $f_i \in \mathcal{H}$ and $\mathbf{a}_i > 0$ and define the sub-Gaussian constants*

$$\sigma_0^2 \triangleq 0 \quad \text{and} \quad \sigma_i^2 \triangleq \sigma_{i-1}^2 + \|f_i\|_{\mathcal{H}}^2 \left(1 + \frac{\sigma_{i-1}^2}{\mathbf{a}_i^2} (\|f_i\|_{\mathcal{H}}^2 - 2\mathbf{a}_i)\right)_+ \quad \forall i \geq 1. \quad (12)$$

The following properties hold true.

(i) **Functional sub-Gaussianity:** For each $i \in [n]$, ψ_i is σ_i sub-Gaussian:

$$\mathbb{E}[\exp(\langle \psi_i, u \rangle_{\mathcal{H}})] \leq \exp\left(\frac{\sigma_i^2 \|u\|_{\mathcal{H}}^2}{2}\right) \quad \text{for all } u \in \mathcal{H}. \quad (13)$$

(ii) **Signed sum representation:** If $\mathbf{a}_i \geq \sigma_{i-1} \|f_i\|_{\mathcal{H}} \sqrt{2 \log(2/\delta_i)}$ for $\delta_i \in (0, 1]$, then, with probability at least $1 - \sum_{i=1}^n \delta_i$,

$$|\alpha_i| \leq \mathbf{a}_i, \forall i \in [n], \quad \text{and} \quad \psi_n = \sum_{i=1}^n \eta_i f_i. \quad (14)$$

(iii) **Exact halving via symmetrization:** If $\mathbf{a}_i \geq \sigma_{i-1} \|f_i\|_{\mathcal{H}} \sqrt{2 \log(2/\delta_i)}$ for $\delta_i \in (0, 1]$ and each $f_i = g_{2i-1} - g_{2i}$ for $g_1, \dots, g_{2n} \in \mathcal{H}$, then, with prob. at least $1 - \sum_{i=1}^n \delta_i$,

$$|\alpha_i| \leq \mathbf{a}_i, \forall i \in [n], \quad \text{and} \quad \frac{1}{2n} \psi_n = \frac{1}{2n} \sum_{i=1}^{2n} g_i - \frac{1}{n} \sum_{i \in \mathcal{I}} g_i \quad \text{for } \mathcal{I} = \{2i + \frac{n_i-1}{2} : i \in [n]\}.$$

(iv) **Pointwise sub-Gaussianity in RKHS:** If \mathcal{H} is the RKHS of a kernel $\mathbf{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, then, for each $i \in [n]$ and $x \in \mathcal{X}$, $\psi_i(x)$ is $\sigma_i \sqrt{\mathbf{k}(x, x)}$ sub-Gaussian:

$$\mathbb{E}[\exp(\psi_i(x))] \leq \exp\left(\frac{\sigma_i^2 \mathbf{k}(x, x)}{2}\right).$$

(v) **Sub-Gaussian constant bound:** Fix any $q \in [0, 1)$, and suppose $\frac{1}{2}\|f_i\|_{\mathcal{H}}^2 \leq \mathbf{a}_i$ for all $i \in [n]$. If $\frac{\|f_i\|_{\mathcal{H}}^2}{1+q} \leq \mathbf{a}_i \leq \frac{\|f_i\|_{\mathcal{H}}^2}{1-q}$ whenever both $\sigma_{i-1}^2 < \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \|f_i\|_{\mathcal{H}}^2}$ and $\|f_i\|_{\mathcal{H}} > 0$, then

$$\sigma_i^2 \leq \frac{\max_{j \in [i]} \|f_j\|_{\mathcal{H}}^2}{1-q^2} \quad \text{for all } i \in [n]. \quad (15)$$

(vi) **Adaptive thresholding:** If $\mathbf{a}_i = \max(c_i \sigma_{i-1} \|f_i\|_{\mathcal{H}}, \|f_i\|_{\mathcal{H}}^2)$ for $c_i \geq 0$, then

$$\sigma_n^2 \leq \frac{\max_{i \in [n]} \|f_i\|_{\mathcal{H}}^2}{4} (c^* + 1/c^*)^2 \quad \text{for } c^* \triangleq \max_{i \in [n]} c_i.$$

Remark 7. The kernel \mathbf{k} in Property (iv) can be arbitrary and need not be bounded.

Property (i) ensures that the functions ψ_i produced by Alg. 3 are mean zero and unlikely to be large in any particular direction u . Property (ii) builds on this functional sub-Gaussianity to ensure that ψ_n is precisely a sum of the signed input functions $\pm f_i$ with high probability. The two properties together imply that, with high probability and an appropriate setting of \mathbf{a}_i , Alg. 3 partitions the input functions f_i into two groups such that the function sums are nearly balanced across the two groups. Property (iii) uses the signed sum representation to construct a two-thinned coreset for any input function sequence $(g_i)_{i=1}^{2n}$. This is achieved by offering the consecutive function differences $g_{2i-1} - g_{2i}$ as the inputs f_i to Alg. 3. Property (iv) highlights that functional sub-Gaussianity also implies sub-Gaussianity of the function values $\psi_i(x)$ whenever the Hilbert space \mathcal{H} is an RKHS. Finally, Properties (v) and (vi) provide explicit bounds on the sub-Gaussian constants σ_i when adaptive settings of the thresholds \mathbf{a}_i are employed. In Sec. 6, we will connect the SBHW to kernel halving and use Properties (iii), (iv), and (vi) together to show that kernel halving and hence also KT-SPLIT coresets have provably small L^∞ kernel error. Specifically, we will boost the pointwise sub-Gaussianity (Property (iv)) of the output function ψ_n into a high probability bound for $\|\psi_n\|_\infty = \sup_x |\psi_n(x)| = \sup_x |\langle \psi_n, \mathbf{k}(x, \cdot) \rangle|$ by constructing a finite cover for $(\psi_n(x))_{x \in \mathbb{R}^d}$ based on the decay and smoothness of the kernel \mathbf{k} .

Comparison with i.i.d. signs A simple alternative to Alg. 3 is to assign signs uniformly at random to each vector, that is, to output $\psi'_n = \sum_{i=1}^n \eta'_i f_i$ with independent Rademacher $\eta'_i \in \{\pm 1\}$. Since the minimal squared sub-Gaussian constant of a sum of independent weighted Rademachers is equal to its variance (Buldygin and Kozachenko, 1980, Lem. 5 & Ex. 1), the minimal squared sub-Gaussian constant of ψ'_n satisfies

$$\sigma_{\text{IID}}^2 = \sup_{u \in \mathcal{H}} \text{Var}(\langle \psi'_n, u \rangle_{\mathcal{H}}) / \|u\|_{\mathcal{H}}^2 = \sup_{u \in \mathcal{H}} \sum_{i=1}^n \langle f_i, u \rangle_{\mathcal{H}}^2 / \|u\|_{\mathcal{H}}^2.$$

In the best case, all f_i are orthogonal and bounded and $\sigma_{\text{IID}}^2 = \max_{i \in [n]} \|f_i\|_{\mathcal{H}}^2$ does not grow with n ; the reader can check that Alg. 3 also reduces to i.i.d. signing in this case. However, in the worst case, all f_i are equal, and $\sigma_{\text{IID}}^2 = n \max_{i \in [n]} \|f_i\|_{\mathcal{H}}^2$. In contrast, if we choose \mathbf{a}_i as in Property (vi) with $c_i = \sqrt{2 \log(2n/\delta)}$, then the SBHW output $\psi_n = \sum_{i=1}^n \eta_i f_i$ with probability $1 - \delta$ by Property (ii) and has squared sub-Gaussian constant $\sigma_{\text{SBHW}}^2 = \mathcal{O}(\log(2n/\delta) \max_{i \in [n]} \|f_i\|_{\mathcal{H}}^2)$ in every case by Property (vi). This drop from $\Omega(n)$ to $\mathcal{O}(\log n)$ represents an exponential improvement in worst-case balance over employing i.i.d. signs.

We can attribute this gain to the carefully chosen updates of Alg. 3. Notice that, on round i , the function ψ_{i-1} is updated only in the f_i direction, so it suffices to examine the evolution of $\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}$. We show in App. J.1 that this evolution takes the form

$$\langle \psi_i, f_i \rangle_{\mathcal{H}} = (1 - \|f_i\|_{\mathcal{H}}^2/\mathbf{a}_i)\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}} + \varepsilon_i \|f_i\|_{\mathcal{H}}^2$$

where $\varepsilon_i \triangleq \mathbb{I}[|\alpha_i| \leq \mathbf{a}_i](\eta_i + \alpha_i/\mathbf{a}_i)$ is mean-zero and 1-sub-Gaussian given ψ_{i-1} . In other words, whenever $\mathbf{a}_i \geq \|f_i\|_{\mathcal{H}}^2$ (as recommended in Property (vi)), Alg. 3 first shrinks the magnitude of ψ_{i-1} in the f_i direction before adding a sub-Gaussian variable in this direction. This targeted shrinkage is absent in the i.i.d. signing update,

$$\langle \psi'_i, f_i \rangle_{\mathcal{H}} = \langle \psi'_{i-1}, f_i \rangle_{\mathcal{H}} + \eta'_i \|f_i\|_{\mathcal{H}}^2,$$

which simply adds a sub-Gaussian variable in the f_i direction, and allows the SBHW to maintain a substantially smaller sub-Gaussian constant.

General recipe for exact halving The symmetrization construction introduced in Property (iii) can be used to convert any vector balancing algorithm (i.e., any algorithm which assigns ± 1 signs to a sequence of vectors) into an exact halving algorithm (i.e., one which assigns -1 to exactly half of the points) simply by running the algorithm on paired vector differences. We will use this property in the sequel to painlessly construct coresets of an exact target size.

Comparison with the self-balancing walk of Alweiss et al. (2021) In the Euclidean setting with $\mathcal{H} = \mathbb{R}^d$, constant thresholds $\mathbf{a}_i = 30 \log(n/\delta)$, and $\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}$ the usual Euclidean dot product, Alg. 3 recovers a slight variant of the Euclidean self-balancing walk of Alweiss et al. (2021, Proof of Thm. 1.2). The original algorithm differs only superficially by terminating with failure whenever $|\alpha_i| > \mathbf{a}_i$. We allow the walk to continue with the update $\psi_i \leftarrow \psi_{i-1} - f_i \cdot \alpha_i/\mathbf{a}_i$, as it streamlines our sub-Gaussianity analysis and avoids the reliance on distributional symmetry present in Sec. 2.1 of Alweiss et al. (2021). We show in App. R that Thm. 3 recovers the guarantee (11) of Alweiss et al. (2021, Thm. 1.2) with improved constants and a less conservative setting of \mathbf{a}_i .

6. L^∞ Guarantees

In this section, we derive near-optimal L^∞ coreset guarantees for kernel halving (Alg. 2) and KT-SPLIT by relating the two algorithms to the self-balancing Hilbert walk (Alg. 3).

6.1 L^∞ guarantees for kernel halving

To make the connection between KH and SBHW more apparent, we have translated each line of Alg. 2 into the notation of Alg. 3. In this notation, we see that Alg. 2 forms signed combinations ψ_i of paired kernel differences $f_i = \mathbf{k}(x_{2i-1}, \cdot) - \mathbf{k}(x_{2i}, \cdot)$; that the inner product $\alpha_i = \langle \psi_{i-1}, f_i \rangle_{\mathbf{k}}$ has a simple explicit form in terms of kernel evaluations; and that, under the event $\mathcal{E} = \{|\alpha_i| \leq \mathbf{a}_i \text{ for all } i = 1, \dots, \lfloor n/2 \rfloor\}$, the function $\psi_{\lfloor n/2 \rfloor}$ of Alg. 2 exactly matches the output of SBHW. Indeed, the function `get_swap_params` serves to compute the sub-Gaussian constants σ_i exactly as defined in Thm. 3 and to adaptively select the thresholds \mathbf{a}_i exactly as recommended in Thm. 3(iii) and (vi): $\mathbf{a}_i = \max(\|f_i\|_{\mathcal{H}} \sigma_{i-1} \sqrt{2 \log(2/\delta_i)}, \|f_i\|_{\mathcal{H}}^2)$.

This choice of \mathbf{a}_i simultaneously ensures that the KH-SBHW equivalence event \mathcal{E} occurs with high probability by Thm. 3(iii) and that the SBHW sub-Gaussian constant remains small by Thm. 3(vi). Hence, we can invoke the pointwise SBHW sub-Gaussianity revealed in Thm. 3(iv) to control the KH L^∞ coresets error with high probability.

Theorem 4 (L^∞ guarantees for kernel halving). *Let $\mathcal{S}_{\text{KH}}(\mathbf{k}, \mathcal{S}, \Delta)$ denote the output of kernel halving (Alg. 2) with kernel \mathbf{k} satisfying Assumps. 2 and 3, input point sequence \mathcal{S} , and probability sequence Δ . Let $\mathbb{P}_n \triangleq \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$, and recall the definitions of $\mathfrak{M}_{\mathbf{k}}$ (8) and $R_{\mathcal{S}_n, \mathbf{k}, n}$ (7). The following statements hold for any $m \in \mathbb{N}$ with $m \leq \log_2 n$ and $\delta' \in (0, 1)$.*

(a) **Kernel halving yields a 2-thinned L^∞ coresets:** *The output $\mathcal{S}^{(1)} \triangleq \mathcal{S}_{\text{KH}}(\mathbf{k}, \mathcal{S}_n, (\delta_i)_{i=1}^{\lfloor \frac{n}{2} \rfloor})$ has size $n_{\text{out}} = \lfloor \frac{n}{2} \rfloor$ with $\mathbb{Q}_{\text{KH}}^{(1)} \triangleq \frac{1}{n_{\text{out}}} \sum_{x \in \mathcal{S}^{(1)}} \delta_x$ satisfying*

$$\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_\infty \leq \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}) \quad (16)$$

with probability at least $1 - \delta' - \sum_{i=1}^{n_{\text{out}}} \delta_i$ for $\delta^ \triangleq \min_{i \leq n_{\text{out}}} \delta_i$.*

(b) **Repeated kernel halving yields a 2^m -thinned L^∞ coresets:** *For each $j > 1$, let $\mathcal{S}^{(j)} \triangleq \mathcal{S}_{\text{KH}}(\mathbf{k}, \mathcal{S}^{(j-1)}, (\frac{2^{j-1}}{m} \delta_i)_{i=1}^{\lfloor \frac{n}{2^j} \rfloor})$ be the output of kernel halving recursively applied for j rounds. Then $\mathcal{S}^{(m)}$ has size $n_{\text{out}} = \lfloor \frac{n}{2^m} \rfloor$ with $\mathbb{Q}_{\text{KH}}^{(m)} \triangleq \frac{1}{n_{\text{out}}} \sum_{x \in \mathcal{S}^{(m)}} \delta_x$ satisfying*

$$\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty \leq \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}) \quad (17)$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} n_{\text{out}}} \delta_i$ for $\delta^ \triangleq \min_{i \leq 2^m n_{\text{out}}} \delta_i$.*

Thm. 4, proved in App. K, shows that L^∞ error for KH scales simply as the kernel thinning inflation factor $\mathfrak{M}_{\mathbf{k}}$ (8) divided by the size of the output. Our next corollary, an immediate consequence of Thm. 4(b) and the definition of $\mathfrak{M}_{\mathbf{k}}$ (8), translates these bounds into rates of decay depending on the radius growth of \mathcal{S}_∞ and the tail decay of \mathbf{k} .

Corollary 3 (L^∞ rates for kernel halving). *Under the notation and assumptions of Thm. 4, consider a sequence of repeated kernel halving runs (Alg. 2), indexed by $n \in \mathbb{N}$, with $m = \lfloor \frac{1}{2} \log_2 n \rfloor$ rounds, $\log(1/\delta^*) = \mathcal{O}(\log n)$, and $\sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\lfloor \frac{n}{2^j} \rfloor} \delta_i = o(1)$ as $n \rightarrow \infty$. If \mathcal{S}_∞ and \mathbf{k} respectively satisfy one of the radius growth (Def. 4) and tail decay (Def. 7) conditions in the table below, then $\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty = \mathcal{O}(\varepsilon_{\infty, n})$ in probability where $\varepsilon_{\infty, n}$ is the corresponding table entry and all hidden constants are independent of the dimension d .*

TYPE OF \mathcal{S}_∞	COMPACT with $R_{\mathcal{S}_n} = \mathcal{O}(\sqrt{d})$	ARBITRARY
TYPE OF \mathbf{k}	COMPACT with $\frac{L_{\mathbf{k}}(R_{\mathbf{k}, n} \wedge R_{\mathcal{S}_n})}{\ \mathbf{k}\ _\infty} = \mathcal{O}(1)$	SUBPOLY with $\log(\frac{L_{\mathbf{k}} R_{\mathbf{k}, n}}{\ \mathbf{k}\ _\infty}) = \mathcal{O}(\log n)$
$\varepsilon_{\infty, n}$	$\sqrt{\frac{d \log n}{n}}$	$\sqrt{\frac{d}{n}} \log n$

Remark 8 (Example settings for KH L^∞ rates). *Rem. 5 specifies settings of $(\delta_i)_{i=1}^{\lfloor \frac{n}{2} \rfloor}$ that meet the conditions of Cor. 3. For any radial kernel $\mathbf{k}(x, y) = \kappa(\|x - y\|_2 / \sigma)$ with L -Lipschitz $\kappa : \mathbb{R} \rightarrow \mathbb{R}$, bandwidth $\sigma > 0$, and $\|\mathbf{k}\|_\infty = 1$, we have $\frac{L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty} \leq \frac{L}{\sigma}$, $R_{\mathbf{k}, n} \leq \sigma \kappa^\dagger(1/n)$,*

and therefore $\frac{L_{\mathbf{k}} R_{\mathbf{k},n}}{\|\mathbf{k}\|_{\infty}} \leq L \kappa^{\dagger}(1/n)$ where L and $\kappa^{\dagger}(u) \triangleq \sup\{r : \kappa(r) \geq u\}$ are independent of d and σ . Hence Gaussian, Matérn (with $\nu > \frac{d}{2} + 1$), and IMQ kernels satisfy the SUBPOLY requirements of Cor. 3 with any choice of bandwidth and satisfy the COMPACT requirements when restricted to a compact domain with $\sigma = \sqrt{d}$. See App. L for our proof.

Near-optimal L^{∞} coresets For any bounded, radial \mathbf{k} satisfying mild decay and smoothness conditions, Phillips and Tai (2020, Thm. 3.1) proved that any procedure outputting a coreset of size $n^{\frac{1}{2}}$ must suffer $\Omega(\min(\sqrt{dn}^{-\frac{1}{2}}, n^{-\frac{1}{4}}))$ L^{∞} error with constant probability for some \mathbb{P}_n . Hence, the repeated KH quality guarantees from Cor. 3 are within a $\log n$ factor of minimax optimality for this kernel family which includes Gaussian, Matérn, Wendland, and IMQ kernels.

Online vector balancing in an RKHS In the online vector balancing problem of Spencer (1977), one must assign signs η_i to Euclidean vectors f_i in an online fashion while keeping the norm of the signed sum $\|\sum_{i=1}^n \eta_i f_i\|_{\infty}$ as small as possible. As an immediate consequence of Thm. 4(a), we find that kernel halving solves an RKHS generalization of the online vector balancing problem.

Corollary 4 (Online vector balancing in an RKHS). *Consider a sequence of kernel halving runs (Alg. 2), indexed by $n \in \mathbb{N}$, with $\log(1/\delta^*) = \mathcal{O}(\log n)$ and $\sum_{i=1}^{\lfloor n/2^j \rfloor} \delta_i = o(1)$ as $n \rightarrow \infty$. If the kernel \mathbf{k} is SUBPOLY (Def. 7) with $\log(\frac{L_{\mathbf{k}} R_{\mathbf{k},n}}{\|\mathbf{k}\|_{\infty}}) = \mathcal{O}(\log n)$, then $\|\sum_{i=1}^{2^n} \epsilon_i \mathbf{k}(x_i, \cdot)\|_{\infty} = \mathcal{O}(\sqrt{d} \log n)$ in probability for the generated signs $\epsilon_i \triangleq \eta_{\lceil i/2 \rceil} (-1)^i$.*

Proof By Thm. 4(a) and the definition of $\mathfrak{M}_{\mathbf{k}}$ (8), $\|\sum_{i=1}^n \eta_i f_i\|_{\infty} = \mathcal{O}(\sqrt{d} \log n)$ in probability. Noting that $-\eta_i f_i = \eta_i (\mathbf{k}(x_{2i}, \cdot) - \mathbf{k}(x_{2i-1}, \cdot)) = \epsilon_{2i-1} \mathbf{k}(x_{2i-1}, \cdot) + \epsilon_{2i} \mathbf{k}(x_{2i}, \cdot)$, the claim follows. \blacksquare

Cor. 4 applies to *any* fixed input point sequence \mathcal{S}_{∞} and to a broad range of kernels including, by Rem. 8, Gaussian, Matérn, and IMQ kernels, as well as B-spline kernel since (122) and (123) imply that $\log(\frac{L_{\mathbf{k}} R_{\mathbf{k},n}}{\|\mathbf{k}\|_{\infty}}) = \mathcal{O}(\log d) = \mathcal{O}(\log n)$ for B-spline \mathbf{k} .

6.2 L^{∞} and MMD guarantees for KT-SPLIT

We finally extend our L^{∞} and MMD guarantees to the KT-SPLIT step of kernel thinning by observing that each candidate coreset generated by KT-SPLIT is the product of repeated kernel halving (Alg. 2) with \mathbf{k}_{rt} as the chosen kernel \mathbf{k} . Hence, Thm. 4(a) applies equally to the coreset $\mathcal{S}^{(1,1)}$ returned by KT-SPLIT with $m = 1$, and, when $m > 1$, Thm. 4(b) applies to the coreset $\mathcal{S}^{(m,1)}$ produced by KT-SPLIT. Combining these L^{∞} bounds with our L^{∞} to MMD conversion theorem (Thm. 2) yields the following corollary proved in App. M.

Corollary 5 (L^{∞} and MMD guarantees for KT-SPLIT). *Consider KT-SPLIT with \mathbf{k}_{rt} satisfying Assumps. 2 and 3, $\delta^* \triangleq \min_i \delta_i$, and $\mathbb{P}_n \triangleq \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$. The guarantees of Thm. 4 with $\mathbf{k} = \mathbf{k}_{\text{rt}}$ hold if $\mathbb{Q}_{\text{KT}}^{(m,1)} \triangleq \frac{1}{n_{\text{out}}} \sum_{x \in \mathcal{S}^{(m,1)}} \delta_x$ is substituted for $\mathbb{Q}_{\text{KH}}^{(m)}$, and the guarantees of Thm. 1 hold if the output coreset $\mathcal{S}^{(m,1)}$ is substituted for \mathcal{S}_{KT} ,*

Cor. 5 ensures that, with high probability, at least one KT-SPLIT candidate is a high-quality L^{∞} coreset for $(\mathbf{k}_{\text{rt}}, \mathbb{P}_n)$ and hence also a high-quality MMD coreset for $(\mathbf{k}_{\star}, \mathbb{P}_n)$.

The proof of Thm. 1 in App. E establishes the same MMD guarantee for kernel thinning by noting that the subsequent KT-SWAP step directly minimizes the $\text{MMD}_{\mathbf{k}_*}$ to \mathbb{P}_n and hence can only improve or maintain this MMD quality. Finally, exactly as in the proofs of Cors. 1 and 3 we can deduce both L^∞ and MMD rate bounds for KT-SPLIT coresets.

Corollary 6 (L^∞ and MMD rates for KT-SPLIT). *Under the notation and assumptions of Cor. 5, consider a sequence of KT-SPLIT runs, indexed by $n \in \mathbb{N}$, with $m = \lfloor \frac{1}{2} \log_2 n \rfloor$ rounds, $\log(1/\delta^*) = \mathcal{O}(\log n)$, and $\sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\lfloor n/2^j \rfloor} \delta_i = o(1)$ as $n \rightarrow \infty$. The guarantees of Cor. 3 with $\mathbf{k} = \mathbf{k}_{\text{rt}}$ hold if $\mathbb{Q}_{\text{KT}}^{(m,1)}$ is substituted for $\mathbb{Q}_{\text{KH}}^{(m)}$, and the guarantees of Cor. 1 hold if $\mathcal{S}^{(m,1)}$ is substituted for \mathcal{S}_{KT} .*

7. Vignettes

We complement our primary methodological and theoretical development with two vignettes illustrating the promise of kernel thinning for improving upon (a) i.i.d. sampling in dimensions $d = 2$ through 100 and (b) standard MCMC thinning when targeting challenging differential equation posteriors. See App. P for supplementary details and

<https://github.com/microsoft/goodpoints>

for a Python implementation of kernel thinning and code replicating each vignette.

7.1 Common settings

Throughout, we adopt a **Gaussian**(σ) target kernel $\mathbf{k}_*(x, y) = \exp(-\frac{1}{2\sigma^2} \|x - y\|_2^2)$ and the corresponding square-root kernel \mathbf{k}_{rt} from Tab. 1. To output a coreset of size $n^{\frac{1}{2}}$ with n input points, we (a) take every $n^{\frac{1}{2}}$ -th point for standard thinning and (b) run kernel thinning (KT) with $m = \frac{1}{2} \log_2 n$ using a standard thinning coreset as the base coreset in KT-SWAP. For each input sample size $n \in \{2^4, 2^6, \dots, 2^{14}, 2^{16}\}$ with $\delta_i = \frac{1}{2n}$, we report the mean coreset error $\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathcal{S}) \pm 1$ standard error across 10 independent replications of the experiment (the standard errors are too small to be visible in all experiments). We additionally regress the log mean MMD onto the log input size using ordinary least squares and display both the best linear fit and an empirical decay rate based on the slope of that fit, e.g., for a slope of -0.25 , we report an empirical decay rate of $n^{-0.25}$ for the mean MMD.

7.2 Kernel thinning versus i.i.d. sampling

We first illustrate the benefits of kernel thinning over i.i.d. sampling from a target \mathbb{P} . We generate each input sequence \mathcal{S}_n i.i.d. from \mathbb{P} , use squared kernel bandwidth $\sigma^2 = 2d$, and consider both Gaussian targets $\mathbb{P} = \mathcal{N}(0, \mathbf{I}_d)$ for $d \in \{2, 4, 10, 100\}$ and mixture of Gaussians (MoG) targets $\mathbb{P} = \frac{1}{M} \sum_{j=1}^M \mathcal{N}(\mu_j, \mathbf{I}_2)$ with $M \in \{4, 6, 8\}$ component locations $\mu_j \in \mathbb{R}^2$ defined in App. P.1.

Fig. 2 highlights the visible differences between the KT and i.i.d. coresets for the MoG targets. Even for small sample sizes, the KT coresets achieves better stratification across components with less clumping and fewer gaps within components, suggestive of a better approximation to \mathbb{P} . Indeed, when we examine MMD error as a function of coreset size in Fig. 3, we observe that kernel thinning provides a significant improvement across all




Figure 2: Kernel thinning (KT) and i.i.d. coresets for 4- and 8-component mixture of Gaussian targets with equidensity contours of the target underlaid. See Sec. 7.2 for more details.

settings. For the Gaussian $d = 2$ target and each MoG target, the KT MMD error scales as $n^{-\frac{1}{2}}$, a quadratic improvement over the $\Omega_p(n^{-\frac{1}{4}})$ MMD error of i.i.d. sampling. Notably, we would not expect to see an exact empirical rate of $n^{-\frac{1}{2}}$ for larger d and small n due to the logarithmic factors in our MMD bounds. However, even for small sample sizes and high dimensions, we observe in Fig. 3(b) that KT significantly improves both the magnitude and the decay rate of MMD relative to i.i.d. sampling.

7.3 Kernel thinning versus standard MCMC thinning

Next, we illustrate the benefits of kernel thinning over standard Markov chain Monte Carlo (MCMC) thinning on twelve posterior inference experiments conducted by Riabiz et al. (2021). We briefly describe each experiment here and refer the reader to Riabiz et al. (2021, Sec. 4) for more details.

Goodwin and Lotka-Volterra experiments From Riabiz et al. (2020), we obtain 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 target, Riabiz et al. (2020) provide 2×10^6 sample points from each of 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).

Hinch experiments From Riabiz et al. (2020), we also obtain the output of two independent Gaussian random walk MCMC chains for 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. (2021, App. S5.4). In computational cardiology, the calcium signalling model represents one component of a heart simulator, and one aims to propagate uncertainty in the signalling model through the whole heart simulation, an operation which




 (a) Mixture of Gaussians \mathbb{P}

 (b) Gaussian \mathbb{P}

Figure 3: Kernel thinning versus i.i.d. sampling. For (a) mixture of Gaussian targets with $M \in \{4, 6, 8\}$ components and (b) standard multivariate Gaussian targets in dimension $d \in \{2, 4, 10, 100\}$, kernel thinning (KT) reduces $\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathcal{S})$ significantly more quickly than the standard $n^{-\frac{1}{4}}$ rate for $n^{\frac{1}{2}}$ i.i.d. points, even in high dimensions. For reference, decay rates of $\sqrt{\log n/n}$ and $\log n/\sqrt{n}$ are plotted in each of the figures in panel (b).

requires thousands of CPU hours per sample point (Riabiz et al., 2021). In this setting, the costs of running kernel thinning are dwarfed by the time required to generate the input sample (two weeks) and more than offset by the cost savings in the downstream uncertainty propagation task.

Preprocessing and kernel settings We discard the initial points of each chain as burn-in using the maximum burn-in period reported in Riabiz et al. (2021, Tabs. S4 & S6, App. S5.4). and 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 form an input sequence \mathcal{S}_n of length n for coreset construction, we downsample the remaining points using standard thinning. Since exact computation of $\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathcal{S})$ is intractable for these posterior targets, we report $\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \mathcal{S})$ —the error that is controlled directly in our theoretical results—for these experiments. We select the kernel bandwidth σ using the popular median heuristic (see, e.g., Garreau et al., 2017). Additional details can be found in App. P.3.

Results Fig. 4 compares the mean $\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \mathcal{S})$ error of the generated kernel thinning and standard thinning coresets. In each of the twelve experiments, KT significantly improves both the rate of decay and the order of magnitude of mean MMD, in line with the guarantees




Figure 4: Kernel thinning versus standard MCMC thinning. Kernel thinning (KT) significantly improves both the rate of decay and the order of magnitude of mean $\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \mathcal{S})$ in each posterior inference task, including eight tasks with 4-dimensional targets (Goodwin and Lotka-Volterra) and four tasks with 38-dimensional targets (Hinch). See Sec. 7.3 for more details.

of Thm. 1. Notably, in the $d = 38$ -dimensional Hinch experiments, standard thinning already improves upon the $n^{-\frac{1}{4}}$ rate of i.i.d. subsampling but is outpaced by KT which consistently provides further improvements.

8. Discussion

We introduced kernel thinning (Alg. 1), a new, practical solution to the thinned MMD coreset problem that, given $\mathcal{O}(n^2)$ time and $\mathcal{O}(n \min(d, n))$ storage, improves upon the integration error of i.i.d. sampling and standard MCMC thinning. To achieve this we first showed that any L^∞ coreset for a square-root kernel \mathbf{k}_{rt} also provides an MMD coreset for its associated target kernel \mathbf{k}_* (Thm. 2). We next introduced and analyzed a self-balancing Hilbert walk for solving the online vector balancing problem in Hilbert spaces (Alg. 3 and Thm. 3). We then designed a symmetrized version of SBHW for RKHSes—kernel halving—that delivers 2-thinned coresets with small L^∞ error (Alg. 2 and Thm. 4). Our online algorithm, KT-SPLIT, recursively applies kernel halving to a square-root kernel to obtain near-optimal L^∞ coresets in $\mathcal{O}(n^2)$ time with $\mathcal{O}(n \min(d, n))$ space (Cors. 3 and 5). Kernel thinning then combines KT-SPLIT with a greedy refinement step (KT-SWAP) to yield

coresets with better-than-i.i.d. MMD for a broad range of kernels and target distributions (Thm. 1 and Cor. 1). While our analysis is restricted to kernels that admit square-root dominating kernels, Dwivedi and Mackey (2022) recently generalized the KT algorithm and analysis to support arbitrary kernels. Separately, Shetty et al. (2022) have developed a distribution compression meta-algorithm, Compress++, which reduces the runtime of KT to near-linear $\mathcal{O}(n \log^3 n)$ time with MMD error that is worse by at most a factor of 4. Hence, KT-Compress++ can be practically deployed even for very large input sizes.

8.1 Related work on MMD coresets

While $(n^{\frac{1}{2}}, o_p(n^{-\frac{1}{4}}))$ -MMD coresets have been developed for specific $(\mathbb{P}, \mathbf{k}_\star)$ pairings like the uniform distribution on the unit cube paired with a Sobolev kernel \mathbf{k}_\star , to the best of our knowledge, no prior $(n^{\frac{1}{2}}, o_p(n^{-\frac{1}{4}}))$ -MMD coreset constructions were known for the range of \mathbb{P} and \mathbf{k}_\star studied in this work. For comparison, we review here both lower bounds and prior strategies for generating coresets with small MMD.

Lower bounds For any bounded and radial (i.e., $\mathbf{k}_\star(x, y) = \kappa(\|x - y\|_2^2)$) kernel satisfying mild decay and smoothness conditions, Phillips and Tai (2020, Thm. 3.1) showed that any procedure outputting coresets of size $n^{\frac{1}{2}}$ must suffer $\Omega(\min(\sqrt{d}n^{-\frac{1}{2}}, n^{-\frac{1}{4}}))$ $\text{MMD}_{\mathbf{k}_\star}$ for some (discrete) target distribution \mathbb{P} . This lower bound applies, for example, to Matérn kernels and to infinitely smooth Gaussian kernels. For any continuous and shift-invariant (i.e., $\mathbf{k}_\star(x, y) = \kappa(x - y)$) kernel taking on at least two values, Tolstikhin et al. (2017, Thm. 1) showed that *any* estimator (including non-coreset estimators) based only on n i.i.d. draws from \mathbb{P} must suffer $C_{\mathbf{k}_\star} n^{-\frac{1}{2}}$ $\text{MMD}_{\mathbf{k}_\star}$ with probability at least 1/4 for some discrete target \mathbb{P} and a constant $C_{\mathbf{k}_\star}$ depending only \mathbf{k}_\star . If, in addition, \mathbf{k}_\star is characteristic (i.e., $\text{MMD}_{\mathbf{k}_\star}(\mu, \nu) \neq 0$ when $\mu \neq \nu$), then Tolstikhin et al. (2017, Thm. 6) establish the same lower bound for some continuous target \mathbb{P} with infinitely differentiable density. These last two lower bounds hold, for example, for Gaussian, Matérn, and B-spline kernels and apply in particular to any thinning algorithm that compresses n i.i.d. sample points without additional knowledge of \mathbb{P} . For light-tailed \mathbb{P} and \mathbf{k}_{rt} , the kernel thinning guarantees of Thm. 1 match each of these lower bounds up to factors of $\sqrt{\log n}$ and constants depending on d .

Order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ -MMD coresets for general target \mathbb{P} By Prop. A.1 of Tolstikhin et al. (2017), an i.i.d. sample from \mathbb{P} yields an order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ -MMD coreset in probability. Chen et al. (2010) showed that *kernel herding* with a finite-dimensional kernel (like the linear $\mathbf{k}_\star(x, y) = \langle x, y \rangle$) finds an $(n^{\frac{1}{2}}, (C_{\mathbb{P}, \mathbf{k}_\star, d} n)^{-\frac{1}{2}})$ -MMD coreset for an inexplicit parameter $C_{\mathbb{P}, \mathbf{k}_\star, d}$. However, Bach et al. (2012) showed that their analysis does not apply to any infinite-dimensional kernel (like the Gaussian, Matérn, and B-spline kernels studied in this work), as $C_{\mathbb{P}, \mathbf{k}_\star, d}$ would necessarily equal 0. The best known rate for kernel herding with bounded infinite-dimensional kernels (Lacoste-Julien et al., 2015, Thm. G.1) guarantees an order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ -MMD coreset, matching the i.i.d. guarantee. For bounded kernels, the same guarantee is available for *Stein Point MCMC* (Chen et al., 2019, Thm. 1) which greedily minimizes MMD^4 over random draws from \mathbb{P} and for a variant of the *greedy sign selection*

4. To bound $\text{MMD}_{\mathbf{k}_\star}$ using Chen et al. (2019, Thm. 1), choose $\mathbf{k}_{\star 0}(x, y) = \mathbf{k}_\star(x, y) - \mathbb{P}\mathbf{k}_\star(x) - \mathbb{P}\mathbf{k}_\star(y) + \mathbb{P}\mathbb{P}\mathbf{k}_\star$.

algorithm described in Karnin and Liberty (2019, Sec. 3.1).⁵ Slightly inferior guarantees were established for Stein points (Chen et al., 2018, Thm. 1) and Stein thinning (Riabiz et al., 2021, Thm. 1), both of which accommodate unbounded kernels as well.

Finite-dimensional kernels Harvey and Samadi (2014) construct $(n^{\frac{1}{2}}, \sqrt{d}n^{-\frac{1}{2}} \log^{2.5} n)$ -MMD coresets for finite-dimensional linear kernels on \mathbb{R}^d but do not address infinite-dimensional kernels.

Uniform distribution on $[0, 1]^d$ The explicit low discrepancy *quasi-Monte Carlo (QMC)* construction of Chen and Skrganov (2002) provides a $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}} \log^{\frac{d-1}{2}} n))$ -MMD coreset for an L^2 discrepancy kernel when \mathbb{P} is the uniform distribution on the unit cube $[0, 1]^d$. For the same target, the *online Haar strategy* of Dwivedi et al. (2019) yields an $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}} \log^{2d} n))$ -MMD coreset in probability. Dwivedi et al. (2019) also conjecture that a greedy variant of their Haar strategy would provide an improved $(n^{\frac{1}{2}}, \mathcal{O}_d(n^{-\frac{1}{2}} \log^d n))$ -MMD coreset. These constructions satisfy our quality criteria but are tailored specifically to the uniform distribution on the unit cube.

Unknown coreset quality On compact manifolds, optimal coresets of size $n^{\frac{1}{2}}$ minimize the weighted Riesz energy (a form of relative MMD with a weighted Riesz kernel) at known rates (Borodachov et al., 2014); however, practical *minimum Riesz energy* (Borodachov et al., 2014) and *minimum energy design* (Joseph et al., 2015, 2019) constructions have not been analyzed. When \mathbf{k}_\star is nonnegative and the kernel matrix $(\mathbf{k}_\star(x_i, x_j))_{i,j=1}^n$ satisfies a strong diagonal dominance condition, Kim et al. (2016, Cor. 3, Thm. 6) show that greedy optimization of $\text{MMD}_{\mathbf{k}_\star}$ yields an *MMD-critic* coreset $\hat{\mathcal{S}}$ of size $n^{\frac{1}{2}}$ satisfying

$$\text{MMD}_{\mathbf{k}_\star}^2(\mathbb{P}_n, \hat{\mathcal{S}}) \leq (1 - \frac{1}{e}) \text{MMD}_\star^2 + \frac{1}{e} \mathbb{P}_n \mathbb{P}_n \mathbf{k}_\star \quad \text{for} \quad \text{MMD}_\star = \min_{|\mathcal{S}|=\sqrt{n}} \text{MMD}_{\mathbf{k}_\star}(\mathbb{P}_n, \mathcal{S}).$$

In the usual case when $\mathbb{P}_n \mathbb{P}_n \mathbf{k}_\star = \Omega(1)$, this error bound does not decay to 0 with n . Paige et al. (2016) analyze the impact of approximating a kernel in *super-sampling with a reservoir* but do not analyze the quality of the constructed MMD coreset. For the conditionally positive definite energy distance kernel, Mak and Joseph (2018) establish that an optimal coreset of size $n^{\frac{1}{2}}$ has $o(n^{-\frac{1}{4}})$ MMD but do not provide a construction; in addition, Mak and Joseph (2018) propose two *support points convex-concave procedures* for constructing MMD coresets but do not establish their optimality and do not analyze their quality.

8.2 Related work on weighted MMD coresets

While coresets satisfy a number of valuable constraints that are critical for some downstream applications—exact approximation of constants, automatic preservation of convex integrand constraints, compatibility with unweighted downstream tasks, easy visualization, straightforward sampling, and increased numerical stability against errors in integral evaluations (Karvonen et al., 2019)—some applications also support *weighted coreset* approximations of \mathbb{P} of the form $\sum_{i=1}^{\sqrt{n}} w_i \delta_{x_i}$ for weights $w_i \in \mathbb{R}$ that need not be equal, need not be non-negative, or need not sum to 1. Notably, weighted coresets that depend on \mathbb{P} only through an i.i.d. sample of size n are subject to the same $\Omega(n^{-\frac{1}{2}})$ MMD lower bounds of Tolstikhin

5. The statement of Karnin and Liberty (2019, Thm. 24) bounds $\|\cdot\|_\infty$, but the proof bounds $\text{MMD}_{\mathbf{k}_\star}$.

et al. (2017) described in Sec. 8.1. Any constructions that violate these bounds do so only by exploiting additional information about \mathbb{P} (for example, exact knowledge of $\mathbb{P}\mathbf{k}_*$) that is not generally available and not required for our kernel thinning guarantees. Moreover, while weighted coresets need not provide satisfactory solutions to the unweighted coreset problem studied in this work, kernel thinning coreset points can be converted into an optimally weighted coreset of no worse quality by explicitly minimizing $\text{MMD}_{\mathbf{k}_*}(\mathbb{P}_n, \sum_{i=1}^{\sqrt{n}} w_i \delta_{x_i})$ or, if computable, $\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \sum_{i=1}^{\sqrt{n}} w_i \delta_{x_i})$ over the weights w_i in $\mathcal{O}(n^{3/2})$ time.

With this context, we now review known weighted MMD coreset guarantees. We highlight that only one of the weighted $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ -MMD guarantees covers the unbounded distributions addressed in this work and that the single unbounded guarantee relies on a restrictive uniformly bounded eigenfunction assumption that is typically not satisfied. In other words, our analysis establishes MMD improvements for practical $(\mathbf{k}_*, \mathbb{P})$ pairings not covered by prior weighted analyses.

\mathbb{P} with bounded support If the target \mathbb{P} has bounded density and bounded, regular support and \mathbf{k}_* is a Gaussian or Matérn kernel, then *Bayesian quadrature* (O’Hagan, 1991) and *Bayes-Sard cubature* (Karvonen et al., 2018) with quasi-uniform unisolvent point sets yield weighted $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ -MMD coresets by Wendland (2004, Thm. 11.22 and Cor. 11.33). If \mathbb{P} has bounded support, and \mathbf{k}_* has more than d continuous derivatives, then the *P-greedy* algorithm (De Marchi et al., 2005) also yields weighted $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ -MMD coresets by Santin and Haasdonk (2017, Thm. 4.1). For $(\mathbf{k}_*, \mathbb{P})$ pairs with compact support and sufficiently rapid eigenvalue decay, approximate *continuous volume sampling kernel quadrature* (Belhadji et al., 2020) using the Gibbs sampler of Rezaei and Gharan (2019) yields weighted coresets with $o(n^{-\frac{1}{4}})$ root mean squared MMD.

Finite-dimensional kernels with compactly supported \mathbb{P} For compactly supported \mathbb{P} , Briol et al. (2015, Thm. 1) and Bach et al. (2012, Prop. 1) show that *Frank-Wolfe Bayesian quadrature* and weighted variants of kernel herding respectively yield weighted $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ -MMD coresets for continuous finite-dimensional kernels, but, by Bach et al. (2012, Prop. 2), these analyses do not extend to infinite-dimensional kernels, like the Gaussian, Matérn, and B-spline kernels studied in this work.

Eigenfunction restrictions For $(\mathbf{k}_*, \mathbb{P})$ pairs with known Mercer eigenfunctions, Belhadji et al. (2019) bound the expected squared MMD of *determinantal point process (DPP) kernel quadrature* in terms of kernel eigenvalue decay and provide explicit rates for univariate Gaussian \mathbb{P} and uniform \mathbb{P} on $[0, 1]$. Their construction makes explicit use of the kernel eigenfunctions which are not available for most $(\mathbf{k}_*, \mathbb{P})$ pairings. For $(\mathbf{k}_*, \mathbb{P})$ pairs with $\mathbb{P}\mathbf{k}_* = \mathbf{0}$, uniformly bounded eigenfunctions, and rapidly decaying eigenvalues, Liu and Lee (2017, App. B.2) prove that *black-box importance sampling* generates probability-weighted coresets with $o(n^{-\frac{1}{4}})$ root mean squared MMD but do not provide any examples verifying their assumptions. The uniformly bounded eigenfunction condition is considered particularly difficult to check (Steinwart and Scovel, 2012), does not hold for Gaussian kernels with Gaussian \mathbb{P} (Minh, 2010, Thm. 1), and need not hold even for infinitely univariate smooth kernels on $[0, 1]$ (Zhou, 2002, Ex. 1).

Unknown coreset quality Huszár and Duvenaud (2012, Prop. 2) bound the MMD error of weighted *sequential Bayesian quadrature* coresets using weak submodularity, but this bound does not decay to zero with n . Khanna and Mahoney (2019, Thm. 2) prove that *weighted kernel herding* yields a weighted $(n^{\frac{1}{2}}, \exp(-n^{\frac{1}{2}}/\kappa_n))$ -MMD coreset. However, the κ_n term in Khanna and Mahoney (2019, Thm.3, Assum. 2) is at least as large as the condition number of an $\sqrt{n} \times \sqrt{n}$ kernel matrix, which for typical kernels (including the Gaussian and Matérn kernels) is $\Omega(\sqrt{n})$ (Koltchinskii and Giné, 2000; El Karoui, 2010); the resulting MMD error bound therefore does not decay with n . The *ProtoGreedy* and *ProtoDash* algorithms of Gurumoorthy et al. (2019, Thm. IV.3, IV.5) yield nonnegative weighted coresets $\hat{\mathcal{S}}$ of size $n^{\frac{1}{2}}$ satisfying $\text{MMD}_{\mathbf{k}_*}^2(\mathbb{P}_n, \hat{\mathcal{S}}) \leq \text{MMD}_{\mathbf{k}_*}^2 + (\mathbb{P}_n \mathbb{P}_n \mathbf{k}_* - \text{MMD}_{\mathbf{k}_*}^2) e^{-\lambda \sqrt{n}}$ where $\text{MMD}_{\mathbf{k}_*}$ is the optimal MMD error to \mathbb{P}_n for a nonnegatively weighted coreset of size $n^{\frac{1}{2}}$. However, careful inspection reveals that $\lambda \sqrt{n} \leq 1$ for any kernel and any n . Hence, in the usual case in which $\mathbb{P}_n \mathbb{P}_n \mathbf{k}_* = \Omega(1)$, this error bound does not decay to 0 with n . Campbell and Broderick (2019, Thm. 4.4) prove that *Hilbert coresets via Frank-Wolfe* with n input points yield weighted order $(n^{\frac{1}{2}}, \nu_n \sqrt{n})$ -MMD coresets for some $\nu_n < 1$ but do not analyze the dependence of ν_n on n .

Non-MMD guarantees For \mathbb{P} with continuously differentiable Lebesgue density and \mathbf{k}_* a bounded Langevin Stein kernel with $\mathbb{P} \mathbf{k}_* = \mathbf{0}$, Thm. 2 of Oates et al. (2017) does not bound MMD but does prove that a randomized *control functionals* weighted coreset satisfies $\sqrt{\mathbb{E}[(\mathbb{E}f - \sum_{i=1}^{\sqrt{n}} w_i f(x_i))^2]} \leq C_{\mathbb{P}, \mathbf{k}_*, d, f} / n^{\frac{7}{24}}$ for each f in the RKHS of \mathbf{k}_* and an unspecified $C_{\mathbb{P}, \mathbf{k}_*, d, f}$. This bound is asymptotically better than the $\Omega(n^{-\frac{1}{4}})$ guarantee for unweighted i.i.d. coresets but worse than the unweighted kernel thinning guarantees of Thm. 1. On compact domains, Thm. 1 of Oates et al. (2019) establishes improved rates for the same weighted coreset when both \mathbb{P} and \mathbf{k}_* are sufficiently smooth. Bardenet and Hardy (2020) establish an $n^{-\frac{1}{4} - \frac{1}{4d}}$ asymptotic decay of $\mathbb{E}f - \sum_{i=1}^{\sqrt{n}} w_i f(x_i)$ for DPP kernel quadrature with \mathbb{P} on $[-1, 1]^d$ and each f in the RKHS of a particular kernel.

8.3 Related work on L^∞ coresets

A number of alternative strategies are available for constructing coresets with L^∞ guarantees. For example, for any bounded \mathbf{k}_{rt} , Cauchy-Schwarz and the reproducing property imply that

$$\|(\mathbb{P} - \mathbb{P}_n) \mathbf{k}_*\|_\infty = \sup_{z \in \mathbb{R}^d} |\langle \mathbf{k}_*(z, \cdot), \mathbb{P} \mathbf{k}_* - \mathbb{P}_n \mathbf{k}_* \rangle_{\mathbf{k}_*}| \leq \text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{P}_n) \cdot \|\mathbf{k}_*\|_\infty^{\frac{1}{2}},$$

so that all of the order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ -MMD coreset constructions discussed in Sec. 8.1 also yield order $(n^{\frac{1}{2}}, n^{-\frac{1}{4}})$ - L^∞ coresets. However, none of those constructions is known to provide a $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ - L^∞ coreset.

A series of breakthroughs due to Joshi et al. (2011); Phillips (2013); Phillips and Tai (2018, 2020); Tai (2020) has led to a sequence of increasingly compressed $(n^{\frac{1}{2}}, o(n^{-\frac{1}{4}}))$ - L^∞ coreset constructions, with the best known guarantees currently due to Phillips and Tai (2020) and Tai (2020). Given n input points, Phillips and Tai (2020) developed an offline, polynomial-time construction to find an $(n^{\frac{1}{2}}, \mathcal{O}_p(\sqrt{dn}^{-\frac{1}{2}} \sqrt{\log n}))$ - L^∞ coreset for Lipschitz kernels exhibiting suitable decay, while Tai (2020) developed an offline construction for

Gaussian kernels that runs in $\Omega(d^{5d})$ time and yields an $(n^{\frac{1}{2}}, \mathcal{O}_p(2^d n^{-\frac{1}{2}} \sqrt{\log(d \log n)}))$ - L^∞ coresset. More details on these constructions based on the Gram-Schmidt walk of Bansal et al. (2018) can be found in App. S. Notably, the Phillips and Tai (hereafter, PT) guarantee is tighter than that of Thm. 4 by a factor of $\sqrt{\log \log n}$ for sub-Gaussian kernels and input points and $\sqrt{\log n}$ for heavy-tailed kernels and input points. Similarly, the Tai guarantee provides an improvement when n is doubly-exponential in the dimension, that is, when $\sqrt{d \log n} = \Omega(2^d)$.

Moreover, by Thm. 2, we may apply the PT and Tai constructions to a square-root kernel \mathbf{k}_t to obtain comparable MMD guarantees for the target kernel \mathbf{k}_* with high probability. However, kernel thinning has a number of practical advantages that lead us to recommend it. First with n input points, using standard matrix multiplication, the PT and Tai constructions have $\Omega(n^4)$ computational complexity and $\Omega(n^2)$ storage costs, a substantial increase over the $\mathcal{O}(n^2)$ running time and $\mathcal{O}(n \min(d, n))$ storage of kernel thinning. Second, KT-SPLIT is an online algorithm while the PT and Tai constructions require the entire set of input points to be available a priori. Finally, each halving round of KT-SPLIT splits the sample size exactly in half, allowing the user to run all m halving rounds simultaneously; the PT and Tai constructions require a rebalancing step after each round forcing the halving rounds to be conducted sequentially.

8.4 Future directions

Several other opportunities for future development recommend themselves. First, since our results cover any target \mathbb{P} with at least $2d$ moments—even discrete and other non-smooth targets—a natural question is whether tighter error bounds with better sample complexities are available when \mathbb{P} is also known to have a smooth Lebesgue density. Second, the MMD to L^∞ reduction in Thm. 2 applies also to weighted L^∞ coresets, and, in applications in which weighted point sets are supported, we would expect either quality or compression improvements from employing non-uniform weights (see, e.g. Turner et al., 2021).

Appendix

Table of contents

1	Introduction	1
2	Input Point and Kernel Requirements	4
3	Kernel Thinning	7
4	MMD Guarantees	10
5	Self-balancing Hilbert Walk	13
6	L^∞ Guarantees	16
7	Vignettes	19
8	Discussion	22

A	Appendix Notation	28
B	Proof of Prop. 1: MMD guarantee for MCMC	29
C	Proof of Prop. 2: Almost sure radius growth	30
D	Proof of Prop. 3: Shift-invariant square-root kernels	33
E	Proof of Thm. 1: MMD guarantee for kernel thinning	34
F	Proof of Rem. 4: Finite-time and anytime guarantees	34
G	Proof of Cor. 1: MMD rates for kernel thinning	34
H	Proof of Thm. 2: MMD guarantee for square-root L^∞ approximations	35
I	Proof of Cor. 2: MMD error from square-root L^∞ error	37
J	Proof of Thm. 3: Self-balancing Hilbert walk properties	37
K	Proof of Thm. 4: L^∞ guarantees for kernel halving	41
L	Proof of Rem. 8: Example settings for KH L^∞ rates	53
M	Proof of Cor. 5: L^∞ and MMD guarantees for KT-SPLIT	53
N	Derivation of Tab. 1: Square-root kernels \mathbf{k}_{rt} for common target kernels \mathbf{k}_*	54
O	Derivation of Tab. 2: Explicit bounds on Thm. 1 quantities for common kernels	56
P	Supplementary Details for Vignettes of Sec. 7	64
Q	Kernel Thinning with Square-root Dominating Kernels	66
R	Online Vector Balancing in Euclidean Space	69
S	L^∞ Coresets of Phillips and Tai (2020) and Tai (2020)	70

Appendix A. Appendix Notation

For each $p \geq 1$, we define L^p as the set of measurable $g : \mathbb{R}^d \rightarrow \mathbb{R}$ with $\|g\|_{L^p} \triangleq (\int |g(x)|^p dx)^{1/p} < \infty$ and C^p as the set of $g : \mathbb{R}^d \rightarrow \mathbb{R}$ for which all partial derivatives of order p exist and are continuous. For a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, we also write $\mathbf{k} \in L^{2,\infty}$ to indicate that \mathbf{k} is measurable with finite

$$\|\mathbf{k}\|_{L^{2,\infty}} \triangleq \sup_{x \in \mathbb{R}^d} (\int \mathbf{k}^2(x, y) dy)^{\frac{1}{2}} = \sup_{x \in \mathbb{R}^d} \|\mathbf{k}(x, \cdot)\|_{L^2}. \quad (18)$$

Throughout, we follow the unitary angular frequency convention of Wendland (2004, Def. 5.15) and define the Fourier transform $\mathcal{F}(f)$ of an integrable complex function $f : \mathbb{R}^d \rightarrow \mathbb{C}$ via

$$\mathcal{F}(f)(\omega) \triangleq \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} f(x) e^{-i\langle x, \omega \rangle} dx \quad \text{for all } \omega \in \mathbb{R}^d. \quad (19)$$

Appendix B. Proof of Prop. 1: MMD guarantee for MCMC

By Douc et al. (2018, Lem. 9.3.9, Cor. 9.2.16), a homogeneous ϕ -irreducible geometrically ergodic Markov chain with stationary distribution \mathbb{P} is also aperiodic with a unique stationary distribution.⁶ Since $(x_i)_{i=0}^\infty$ are the iterates of such a chain, there exist, by Gallegos-Herrada et al. (2023, Thm. 1xi), constants $\rho \in (0, 1)$ and $\tau < \infty$ and a measurable \mathbb{P} -almost everywhere finite function $V : \mathbb{R}^d \rightarrow [1, \infty]$ satisfying $\mathbb{P}V < \infty$ and

$$\sup_{\text{measurable } h: \frac{|h(x)|}{V(x)} \leq 1, \forall x \in \mathbb{R}^d} |\mathbb{E}[h(x_i) \mid x_0 = x] - \mathbb{P}h| \leq \tau V(x) \rho^i, \text{ for all } x \in \mathbb{R}^d \text{ and } i \in \mathbb{N}. \quad (20)$$

Since V is finite \mathbb{P} -almost everywhere, we will choose $c(x) = \infty \Leftrightarrow V(x) = \infty$ to ensure that our claim is (vacuously) true whenever $V(x_0) = \infty$.

Hereafter, suppose $V(x_0) < \infty$. Since the Markov chain is irreducible and aperiodic with a unique stationary distribution \mathbb{P} , Assump. H1 of Havet et al. (2020) is satisfied. Hence, by an application of Havet et al. (2020, Prop. 2.1) with $V' = V\rho^g$ and $\zeta = 2/\rho^g$ for sufficiently large $g \in \mathbb{N}$, there exists a set $C(x_0) \subseteq \mathbb{R}^d$ that contains x_0 and satisfies Assumps. H2 and H3 of Havet et al. (2020).

Now fix any $y_1, \dots, y_n, z_1, \dots, z_n \in \mathbb{R}^d$. We invoke the definition of MMD (1), the triangle inequality, the reproducing property of an RKHS (Steinwart and Christmann, 2008, Def. 4.18), and Cauchy-Schwarz in turn to deduce a bounded differences property for MMD:

$$\begin{aligned} & \text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \frac{1}{n} \sum_{i=1}^n \delta_{y_i}) - \text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \frac{1}{n} \sum_{i=1}^n \delta_{z_i}) \\ &= \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} |\mathbb{P}f - \frac{1}{n} \sum_{i=1}^n f(y_i)| - \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} |\mathbb{P}f - \frac{1}{n} \sum_{i=1}^n f(z_i)| \\ &\leq \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} |\frac{1}{n} \sum_{i=1}^n f(y_i) - f(z_i)| = \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} \frac{1}{n} \sum_{i=1}^n |\langle \mathbf{k}_\star(y_i, \cdot) - \mathbf{k}_\star(z_i, \cdot), f \rangle_{\mathbf{k}_\star}| \\ &\leq \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} \frac{1}{n} \sum_{i=1}^n \|\mathbf{k}_\star(y_i, \cdot) - \mathbf{k}_\star(z_i, \cdot)\|_{\mathbf{k}_\star} \|f\|_{\mathbf{k}_\star} \\ &= \sup_{\|f\|_{\mathbf{k}_\star} \leq 1} \frac{1}{n} \sum_{i=1}^n \sqrt{\mathbf{k}_\star(y_i, y_i) + \mathbf{k}_\star(z_i, z_i) - 2\mathbf{k}_\star(y_i, z_i)} \|f\|_{\mathbf{k}_\star} \leq \frac{2}{n} \|\mathbf{k}_\star\|_\infty^{\frac{1}{2}} \sum_{i=1}^n \mathbb{I}[y_i \neq z_i]. \end{aligned}$$

Since x_0 belongs to a set $C(x_0)$ satisfying Assumps. H2 and H3 of Havet et al. (2020), McDiarmid's inequality for geometrically ergodic Markov chains (Havet et al., 2020, Thm. 3.1) implies that, with probability at least $1 - \delta$ conditional on x_0 ,

$$\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathbb{P}_n) \leq \mathbb{E}[\text{MMD}_{\mathbf{k}_\star}(\mathbb{P}, \mathbb{P}_n) \mid x_0] + \sqrt{c_1(x_0) \|\mathbf{k}_\star\|_\infty \log(1/\delta)/n}$$

where $c_1(x_0)$ is a finite value depending only on the transition probabilities of the chain and the set $C(x_0)$.

Now, define the \mathbb{P} centered kernel $\mathbf{k}_{\star\mathbb{P}}(x, y) = \mathbf{k}_\star(x, y) - \mathbb{P}\mathbf{k}_\star(x) - \mathbb{P}\mathbf{k}_\star(y) + \mathbb{P}\mathbb{P}\mathbf{k}_\star$. To bound the expectation, we will use a slight modification of Lem. 3 of Riabiz et al. (2021). The original lemma used the assumption of V -uniform ergodicity (Meyn and Tweedie, 2012, Defn. (16.0.1)) and the assumption $V(x) \geq \sqrt{\mathbf{k}_{\star\mathbb{P}}(x, x)}$ solely to argue that, for some $R > 0$,

$$|\mathbb{E}[f(x_i) \mid x_0 = x] - \mathbb{P}f| \leq RV(x) \rho^i \text{ for all } x \in \mathbb{R}^d \text{ and } f \in \mathcal{H}_{\mathbf{k}_{\star\mathbb{P}}} \text{ with } \|f\|_{\mathbf{k}_{\star\mathbb{P}}} = 1.$$

6. In Havet et al. (2020); Douc et al. (2018, Def. 9.2.1) the term *irreducible* is synonymous with ϕ -irreducible as defined by Gallegos-Herrada et al. (2023, Sec. 2).

In our case, since $\mathbf{k}_{\star\mathbb{P}}$ is bounded and any $f \in \mathcal{H}_{\mathbf{k}_{\star\mathbb{P}}}$ with $\|f\|_{\mathbf{k}_{\star\mathbb{P}}} = 1$ satisfies

$$|f(x)| = |(\mathbf{k}_{\star\mathbb{P}}(x, \cdot), f)_{\mathbf{k}_{\star\mathbb{P}}}| \leq \|\mathbf{k}_{\star\mathbb{P}}(x, \cdot)\|_{\mathbf{k}_{\star\mathbb{P}}} \|f\|_{\mathbf{k}_{\star\mathbb{P}}} = \sqrt{\mathbf{k}_{\star\mathbb{P}}(x, x)} \leq \sqrt{\|\mathbf{k}_{\star\mathbb{P}}\|_{\infty}}$$

by the reproducing property and Cauchy-Schwarz, the geometric ergodicity property (27) implies the analogous bound

$$|\mathbb{E}[f(x_i) \mid x_0 = x] - \mathbb{P}f| \leq \tau \sqrt{\|\mathbf{k}_{\star\mathbb{P}}\|_{\infty}} V(x) \rho^i \text{ for all } x \in \mathbb{R}^d \text{ and } f \in \mathcal{H}_{\mathbf{k}_{\star\mathbb{P}}} \text{ with } \|f\|_{\mathbf{k}_{\star\mathbb{P}}} = 1.$$

Hence, the conclusions of Riabiz et al. (2021, Lem. 3) with $R = \tau \sqrt{\|\mathbf{k}_{\star\mathbb{P}}\|_{\infty}}$ hold under our assumptions. Jensen's inequality and the conclusion of Lem. 3 of Riabiz et al. (2021) now yield the sure bound

$$\begin{aligned} \mathbb{E}[\text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathbb{P}_n) \mid x_0]^2 &\leq \mathbb{E}[\text{MMD}_{\mathbf{k}_{\star}}(\mathbb{P}, \mathbb{P}_n)^2 \mid x_0] \\ &= \mathbb{E}\left[\frac{1}{n^2} \sum_{i=1}^n \mathbf{k}_{\star\mathbb{P}}(x_i, x_i) + \frac{1}{n^2} \sum_{i=1}^n \sum_{j \neq i} \mathbf{k}_{\star\mathbb{P}}(x_i, x_j) \mid x_0\right] \\ &\leq \frac{1}{n} \|\mathbf{k}_{\star\mathbb{P}}\|_{\infty} \left(1 + \frac{2\tau\rho}{1-\rho} \frac{1}{n} \sum_{i=1}^{n-1} \mathbb{E}[V(x_i) \mid x_0]\right) \leq \frac{4}{n} \|\mathbf{k}_{\star}\|_{\infty} \left(1 + \frac{2\tau\rho}{1-\rho} \frac{1}{n} \sum_{i=1}^{n-1} \mathbb{E}[V(x_i) \mid x_0]\right). \end{aligned}$$

Now, define $c_2(x_0) \triangleq \sup_{n \in \mathbb{N}} \frac{1}{n} \sum_{i=1}^{n-1} \mathbb{E}[V(x_i) \mid x_0]$. Since $V(x_0) < \infty$, the geometric ergodicity property (27) and the fact that $\mathbb{P}V < \infty$ imply

$$c_2(x_0) \leq \mathbb{P}V + V(x_0) \sup_{n \in \mathbb{N}} \frac{1}{n} \sum_{i=1}^{n-1} \rho^i \leq \mathbb{P}V + V(x_0) \rho < \infty.$$

Taking $c(x_0) = \sqrt{2} \max(c_1(x_0), 4c_2(x_0)(1 + \frac{2\tau\rho}{1-\rho}))$ completes the proof.

Appendix C. Proof of Prop. 2: Almost sure radius growth

We prove this result for the identically distributed case in App. C.1 and for the Markov chain case in App. C.2.

C.1 Radius growth for identically distributed sequences

Suppose x_0 and \mathcal{S}_{∞} are drawn identically from \mathbb{P} . Claim (a) is true by definition. To establish the remaining claims, we use the following more general result proved in App. C.1.1.

Lemma 5 (Growth rate for identically distributed sequence). *Consider a sequence of identically distributed random variables $(Y_i)_{i=1}^{\infty}$ on \mathbb{R} and a measurable function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ with an increasing inverse function ψ^{-1} . If $\psi(Y_1) \geq 0$ almost surely, then*

$$\Pr(\max_{i \leq n} Y_i > \psi^{-1}(n) \text{ for some } n \in \mathbb{N}) \leq \mathbb{E}(\psi(Y_1)).$$

Consequently, if $\mathbb{E}(\psi(Y_1)) < \infty$, then, for any $\delta \in (0, 1]$,

$$\Pr(\max_{i \leq n} Y_i \leq \psi^{-1}\left(\frac{n\mathbb{E}[\psi(Y_1)]}{\delta}\right) \text{ for all } n \in \mathbb{N}) \geq 1 - \delta.$$

Fix any $\delta \in (0, 1]$. Thm. 5 with $Y_i = \|x_i\|_2$ implies that, with probability $1 - \delta$,

$$R_{\mathcal{S}_n} \leq \psi^{-1}\left(\frac{n\mathbb{E}[\psi(\|x_1\|_2)]}{\delta}\right) \text{ for all } n \in \mathbb{N}$$

for $\psi^{-1}(r) = \frac{\sqrt{\log r}}{\sqrt{c}}$ in case (b), $\psi^{-1}(r) = \frac{\log r}{c}$ in case (c), and $\psi^{-1}(r) = r^{1/\rho}$ in case (d).

As a result we have, with probability $1 - \delta$, $R_{\mathcal{S}_n} = \mathcal{O}_d(\sqrt{\log n})$ in case (b), $R_{\mathcal{S}_n} = \mathcal{O}_d(\log n)$ in case (c), and $R_{\mathcal{S}_n} = \mathcal{O}_d(n^{1/\rho})$ in case (d). Since δ is arbitrary, these orders hold with probability 1 as claimed.

C.1.1 PROOF OF THM. 5: GROWTH RATE FOR IDENTICALLY DISTRIBUTED SEQUENCE

We make use of three lemmas. The first rewrites maximum exceedance events in terms of individual variable exceedance events when thresholds are nondecreasing.

Lemma 6 (Exceedance equivalence). *For any real-valued $(a_i)_{i=1}^\infty$ and nondecreasing $(b_i)_{i=1}^\infty$,*

$$\max_{i \leq n} a_i > b_n \text{ for some } n \in \mathbb{N} \iff a_i > b_i \text{ for some } i \in \mathbb{N}. \quad (21)$$

Proof The \Leftarrow part follows immediately. To prove the \Rightarrow part, suppose $\max_{i \leq n^*} a_i > b_{n^*}$ for some n^* . Then there exists an $i \leq n^*$ with $a_i > b_{n^*} \geq b_i$ since $(b_i)_{i=1}^\infty$ is nondecreasing. ■

The second bounds the probability of growth rate violation for *any* sequence of random variables in terms of a sum of exceedance probabilities.

Lemma 7 (Growth rate for arbitrary sequence). *For any sequence of random variables $(Y_i)_{i=1}^\infty$ on \mathbb{R} and a nondecreasing real-valued sequence $(b_i)_{i=1}^\infty$, we have*

$$\Pr(\max_{i \leq n} Y_i > b_n \text{ for some } n \in \mathbb{N}) = \Pr(Y_i > b_i \text{ for some } i \in \mathbb{N}) \leq \sum_{i=1}^\infty \Pr(Y_i > b_i).$$

Proof The result follows from immediately from Thm. 6 and the union bound. ■

The third lemma bounds a sum of exceedance probabilities whenever the random variables are identically distributed and nonnegative.

Lemma 8 (Bounding exceedances with expectations). *If the random variables $(Z_i)_{i=1}^\infty$ are identically distributed and almost surely nonnegative, then*

$$\sum_{i=1}^\infty \Pr(Z_i > i) \leq \mathbb{E}(Z_1). \quad (22)$$

Proof Since Z_1 is almost surely nonnegative, we have $Z_1 = \int_0^{Z_1} dt = \int_0^\infty \mathbb{I}(Z_1 > t) dt$ almost surely. Tonelli's theorem (Mukherjea, 1972, Thm. 1) therefore implies that

$$\begin{aligned} \mathbb{E}(Z_1) &= \mathbb{E}(\int_0^\infty \mathbb{I}(Z_1 > t) dt) = \int_0^\infty \Pr(Z_1 > t) dt \\ &\geq \int_0^\infty \Pr(Z_1 > \lceil t \rceil) dt = \sum_{i=1}^\infty \Pr(Z_1 > i) = \sum_{i=1}^\infty \Pr(Z_i > i) \end{aligned}$$

where the final inequality uses the identically distributed assumption. ■

Since $\psi(Y_1) \geq 0$ almost surely and ψ^{-1} is increasing, we invoke Thm. 8 with $Z_i = \psi(Y_i)$, the invertibility of ψ , and Thm. 7 with $b_i = \psi^{-1}(i)$ in turn to conclude

$$\begin{aligned} \mathbb{E}(\psi(Y_1)) &\stackrel{(26)}{\geq} \sum_{i=1}^\infty \Pr(\psi(Y_i) > i) = \sum_{i=1}^\infty \Pr(Y_i > \psi^{-1}(i)) \\ &\stackrel{(21)}{\geq} \Pr(\max_{i \leq n} Y_i > \psi^{-1}(n) \text{ for some } n \in \mathbb{N}). \end{aligned}$$

C.2 Radius growth for MCMC

Now suppose x_0 and \mathcal{S}_∞ are the iterates of a homogeneous ϕ -irreducible geometrically ergodic Markov chain with initial state x_0 , subsequent iterates \mathcal{S}_∞ , and stationary distribution \mathbb{P} . Our claims will follow from the following more detailed result proved in App. C.2.1.

Lemma 9 (Growth rate for MCMC). *Consider a homogeneous ϕ -irreducible geometrically ergodic Markov chain with initial state x_0 , subsequent iterates $(x_i)_{i=1}^\infty$, and stationary distribution \mathbb{P} . There exist constants $\rho \in (0, 1)$ and $\tau < \infty$ and a measurable \mathbb{P} -almost everywhere finite function $V : \mathbb{R}^d \rightarrow [1, \infty]$ such that, for any index $j \in \mathbb{N}$, measurable function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, measurable nonnegative function ψ on \mathbb{R} with increasing inverse function ψ^{-1} , and $X \sim \mathbb{P}$,*

$$\begin{aligned} \Pr(\max_{i \leq n} g(x_i) > \psi^{-1}(n) \text{ for some } n \in \mathbb{N} \mid x_0) &\leq \\ \mathbb{E}(\psi(g(X))) + \frac{\tau \rho^{j+1}}{1-\rho} V(x_0) + \sum_{i=1}^j \Pr(g(x_i) > \psi^{-1}(i) \mid x_0). \end{aligned} \quad (23)$$

Now suppose $V(x_0) < \infty$, and fix any measurable nonnegative ψ on \mathbb{R} with increasing ψ^{-1} and any measurable $g : \mathbb{R}^d \rightarrow \mathbb{R}$. If $\mathbb{E}(\psi(g(X))) < \infty$, then, for any $\delta \in (0, 1]$, there exists a constant $c_{\delta, \psi \circ g}(x_0) \in (0, \infty)$ such that

$$\Pr(\max_{i \leq n} g(x_i) \leq \psi^{-1}(c_{\delta, \psi \circ g}(x_0)n) \text{ for all } n \in \mathbb{N} \mid x_0) \geq 1 - \delta. \quad (24)$$

Moreover, if \mathbb{P} is compactly supported, then for any $\delta \in (0, 1]$, there exists a constant $c_\delta(x_0) \in (0, \infty)$ such that

$$\Pr(\sup_{i \in \mathbb{N}} \|x_i\|_2 \leq c_\delta(x_0) \mid x_0) \geq 1 - \delta. \quad (25)$$

Instantiate the function V from Thm. 9, and suppose that $V(x_0) < \infty$, an event that holds for \mathbb{P} -almost every x_0 . Claims (b) to (d) then follow by invoking the time-uniform tail bound (24) with $g = \|\cdot\|_2$ and proceeding as in App. C.1. Finally, claim (a) follows from the bound (25), which establishes $\|x_i\|_2 = \mathcal{O}_d(1)$ with probability 1 conditional on x_0 .

C.2.1 PROOF OF THM. 9: GROWTH RATE FOR MCMC

The proof closely parallels that of Thm. 5 except that we substitute the following estimate for Thm. 8.

Lemma 10 (Bounding MCMC exceedances with expectations). *Consider a homogeneous ϕ -irreducible geometrically ergodic Markov chain with initial state x_0 , subsequent iterates $(x_i)_{i=1}^\infty$, and stationary distribution \mathbb{P} . There exist constants $\rho \in (0, 1)$ and $\tau < \infty$ and a measurable \mathbb{P} -almost everywhere finite function $V : \mathbb{R}^d \rightarrow [1, \infty]$ such that, for any index $j \in \mathbb{N}$, measurable nonnegative function f on \mathbb{R}^d , and $X \sim \mathbb{P}$,*

$$\sum_{i=j}^\infty \mathbb{P}(f(x_i) > i \mid x_0) \leq \mathbb{E}(f(X)) + \frac{\tau \rho^j}{1-\rho} V(x_0). \quad (26)$$

Proof By Douc et al. (2018, Lem. 9.3.9), a homogeneous ϕ -irreducible geometrically ergodic Markov chain with stationary distribution \mathbb{P} is also aperiodic. By Gallegos-Herrada et al.

(2023, Thm. 1xi), there exist constants $\rho \in (0, 1)$ and $\tau < \infty$ and a measurable \mathbb{P} -almost everywhere finite function $V : \mathbb{R}^d \rightarrow [1, \infty]$ satisfying

$$\sup_{h: |h(x)| \leq V(x), \forall x \in \mathbb{R}^d} |\mathbb{E}[h(x_i) \mid x_0 = x] - \mathbb{P}h| \leq \tau V(x) \rho^i \quad \text{for all } x \in \mathbb{R}^d. \quad (27)$$

Applying this result to the functions $h_i(x) \triangleq \mathbb{I}[f(x) > i]$, we find that

$$\sum_{i=j}^{\infty} \mathbb{P}(f(x_i) > i \mid x_0) \leq \sum_{i=j}^{\infty} \mathbb{P}(f(X) > i) + \sum_{i=j}^{\infty} \tau V(x_0) \rho^i \leq \mathbb{E}[f(X)] + \frac{\tau \rho^j}{1-\rho} V(x_0)$$

where the final inequality uses Thm. 8. \blacksquare

Fix any V , ρ , and τ satisfying the conclusions of Thm. 10, any measurable $g : \mathbb{R}^d \rightarrow \mathbb{R}$, and any measurable nonnegative ψ on \mathbb{R} with increasing ψ^{-1} . The first claim (23) follows by applying Thm. 7 with $b_i = \psi^{-1}(i)$, the assumed invertibility and strict monotonicity of ψ^{-1} , and Thm. 10 with $f = \psi \circ g$ in turn to find that

$$\begin{aligned} & \mathbb{P}(\max_{i \leq n} g(x_i) > \psi^{-1}(n) \text{ for some } n \mid x_0) \leq \sum_{i=1}^{\infty} \mathbb{P}(g(x_i) > \psi^{-1}(i) \mid x_0) \\ & = \sum_{i=1}^{\infty} \mathbb{P}(\psi(g(x_i)) > i \mid x_0) \leq \sum_{i=1}^j \mathbb{P}(\psi(g(x_i)) > i \mid x_0) + \mathbb{E}(\psi(g(X))) + \frac{\tau \rho^{j+1}}{1-\rho} V(x_0). \end{aligned}$$

Now suppose $V(x_0) < \infty$, fix any $\delta \in (0, 1]$, and let j be the smallest positive index with $\frac{\tau \rho^{j+1}}{1-\rho} V(x_0) < \frac{\delta}{3}$. Since each $\psi(g(x_i))$ is a tight random variable given x_0 , we can additionally choose a constant $c_{\delta, \psi \circ g}(x_0) \in (0, \infty)$ satisfying $\sum_{i=1}^j \mathbb{P}(\psi(g(x_i)) / c_{\delta, \psi \circ g}(x_0) > i \mid x_0) < \frac{\delta}{3}$ and $\mathbb{E}(\psi(g(X))) / c_{\delta, \psi \circ g}(x_0) < \frac{\delta}{3}$. The claim (24) now follows by applying the initial result (23) to the function $\psi / c_{\delta, \psi \circ g}(x_0)$.

To establish the final claim (25), suppose that \mathbb{P} is compactly supported. Since \mathbb{P} has compact support and each $\|x_i\|_2$ is a tight random variable, there exists a constant $c_{\delta} \in (0, \infty)$ satisfying $\Pr_{X \sim \mathbb{P}}(\|X\|_2 > c_{\delta}) = 0$ and $\sum_{i=1}^j \mathbb{P}(\psi(\|x_i\|_2) > c_{\delta} \mid x_0) < \frac{\delta}{2}$. The union bound and the geometric ergodicity property (27) applied to the function $h(x) = \mathbb{I}(\|x\|_2 > c_{\delta})$ with $\mathbb{P}h = 0$ now imply

$$\begin{aligned} \Pr(\sup_{i \in \mathbb{N}} \|x_i\|_2 > c_{\delta} \mid x_0) & \leq \sum_{i=1}^{\infty} \Pr(\|x_i\|_2 > c_{\delta} \mid x_0) \\ & \leq \sum_{i=1}^j \Pr(\|x_i\|_2 > c_{\delta} \mid x_0) + \tau V(x_0) \sum_{i=j+1}^{\infty} \rho^i \\ & = \sum_{i=1}^j \Pr(\|x_i\|_2 > c_{\delta} \mid x_0) + \frac{\tau \rho^{j+1}}{1-\rho} V(x_0) < \frac{\delta}{2} + \frac{\delta}{3} < \delta. \end{aligned}$$

Appendix D. Proof of Prop. 3: Shift-invariant square-root kernels

Bochner's theorem (Bochner, 1933; Wendland, 2004, Thm. 6.6) implies that \mathbf{k}_{rt} is a kernel since κ_{rt} is the Fourier transform of a finite Borel measure with Lebesgue density $\sqrt{\widehat{\kappa}}$. Moreover, as $\mathbf{k}_{\text{rt}}(x, \cdot) = \frac{1}{(2\pi)^{d/4}} \mathcal{F}(e^{-i\langle \cdot, x \rangle} \sqrt{\widehat{\kappa}})$ and $e^{-i\langle \cdot, x \rangle} \sqrt{\widehat{\kappa}}$ is integrable and square integrable, the Plancherel-Parseval identity (Wendland, 2004, Proof of Thm. 5.23) implies that

$$\begin{aligned} \int_{\mathbb{R}^d} \mathbf{k}_{\text{rt}}(x, z) \mathbf{k}_{\text{rt}}(y, z) dz & = \int_{\mathbb{R}^d} \frac{1}{(2\pi)^{d/4}} e^{-i\langle \omega, x \rangle} \sqrt{\widehat{\kappa}(\omega)} \frac{1}{(2\pi)^{d/4}} e^{i\langle \omega, y \rangle} \sqrt{\widehat{\kappa}(\omega)} d\omega \\ & = \frac{1}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} e^{-i\langle \omega, x-y \rangle} \widehat{\kappa}(\omega) d\omega = \mathbf{k}_{\star}(x, y) \end{aligned}$$

confirming that \mathbf{k}_{rt} is a square-root kernel of \mathbf{k}_{\star} .

Appendix E. Proof of Thm. 1: MMD guarantee for kernel thinning

By design, KT-SWAP ensures

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

where $\mathcal{S}^{(m,1)}$ denotes the first coreset returned by KT-SPLIT. Next, applying Cor. 5, in particular, the bound (85) yields the desired claim.

Appendix F. Proof of Rem. 4: Finite-time and anytime guarantees

We prove the three claims one by one.

Finite time guarantee For the case with known n , the claim follows simply by noting that

$$\sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \frac{\delta}{n} = \sum_{j=1}^m \frac{2^m}{m} \lfloor \frac{n}{2^m} \rfloor \frac{\delta}{2n} \leq \sum_{j=1}^m \frac{2^m}{m} \frac{n}{2^m} \frac{\delta}{2n} = \frac{\delta}{2}.$$

Any time guarantee When the input size n is not known in advance but is chosen independently of the randomness in kernel thinning, we first note that

$$\sum_{i=1}^{\infty} \frac{1}{(i+1) \log^2(i+1)} \stackrel{(i)}{\leq} 2, \quad \text{and} \quad \sum_{j=1}^m 2^j = 2^{m+1} - 2 \leq 2^{m+1}. \quad (28)$$

where step (i) can be verified using mathematical programming software. Therefore, for any $n \in \mathbb{N}$, with $\delta_i = \frac{m\delta}{2^{m+2}(i+1) \log^2(i+1)}$, we have

$$\begin{aligned} \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \delta_i &\leq \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\infty} \delta_i = \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{\infty} \frac{m\delta}{2^{m+2}(i+1) \log^2(i+1)} \\ &= \frac{\delta}{2^{m+3}} \left(\sum_{j=1}^m 2^j \right) \left(\sum_{i=1}^{\infty} \frac{1}{(i+1) \log^2(i+1)} \right) \\ &\stackrel{(28)}{\leq} \frac{\delta}{2^{m+3}} \cdot 2^{m+1} \cdot 2 \leq \frac{\delta}{2}. \end{aligned}$$

Upper bound on δ^* The probability lower bound in Thm. 1 is $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \delta_i$, which is non-negative only if

$$1 \geq \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \delta_i \geq \frac{2^m}{2} \lfloor \frac{n}{2^m} \rfloor \delta^*,$$

which holds only if $\delta^* \leq \frac{2}{2^m \lfloor n/2^m \rfloor} \leq \frac{6m}{2^m}$ since $m \in [1, \log_2 n]$. The claim follows.

Appendix G. Proof of Cor. 1: MMD rates for kernel thinning

Repeating arguments similar to those deriving (125) in App. N, we find that for the advertised choices of m and for any fixed δ such that $\delta' = \frac{\delta}{2}$ and $\log(1/\delta^*) = \mathcal{O}(\log(n/\delta))$, the

RHS of the bound (9) on MMD from Thm. 1 can be simplified as follows:

$$\begin{aligned}
 & \text{MMD}_{\mathbf{k}_\star}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) \\
 & \leq c \|\mathbf{k}_{\text{rt}}\|_\infty \left(c' \frac{\max(R'_{\mathbf{k}_{\text{rt}},n}, R_{\mathcal{S}_n})^2}{d} \right)^{\frac{d}{4}} d^{\frac{1}{4}} \sqrt{\frac{\log(n/\delta)}{n} \left[\log\left(\frac{8}{\delta}\right) + \log\left(2 + \frac{L_{\mathbf{k}_{\text{rt}}}(R_{\mathbf{k}_{\text{rt}},n} + R_{\mathcal{S}_n})}{\|\mathbf{k}_{\text{rt}}\|_\infty}\right) \right]}, \\
 & = c_{\delta,d} \|\mathbf{k}_{\text{rt}}\|_\infty \left(\max(R'_{\mathbf{k}_{\text{rt}},n}, R_{\mathcal{S}_n}) \right)^{\frac{d}{2}} \sqrt{\frac{\log n}{n}} \sqrt{\log(1 + \max(R_{\mathbf{k}_{\text{rt}},n}, R_{\mathcal{S}_n})) + \log\left(1 + \frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty}\right)},
 \end{aligned} \tag{29}$$

for some universal constants c, c' where to simplify the expressions, we have used the fact that $R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n} \leq R_{\mathcal{S}_n}$ (7). Noting that (29) holds with probability at least $1 - \delta$, Cor. 1 now follows from plugging the assumed growth rate bounds into the estimate (29), and treating $\|\mathbf{k}_{\text{rt}}\|_\infty$ and $\frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty}$ as some constant while n grows.

Appendix H. Proof of Thm. 2: MMD guarantee for square-root L^∞ approximations

Our proof will use the following two lemmas proved in Apps. H.1 and H.2 respectively.

Lemma 11 (Square-root representation of MMD). *For \mathbf{k}_\star satisfying Assump. 1 with square-root kernel $\mathbf{k}_{\text{rt}} \in L^{2,\infty}$ we have, for any distributions μ and ν on \mathbb{R}^d ,*

$$\text{MMD}_{\mathbf{k}_\star}(\mu, \nu) = \sup_{g \in L^2: \|g\|_{L^2} \leq 1} \left| \int g(y) (\mu \mathbf{k}_{\text{rt}}(y) - \nu \mathbf{k}_{\text{rt}}(y)) dy \right|.$$

Lemma 12 (L^∞ bound on L^2 kernel error). *Consider any kernel $\mathbf{k} \in L^{2,\infty}$ satisfying Assump. 1, distributions μ, ν on \mathbb{R}^d , and function $g \in L^2$ with $\|g\|_{L^2} \leq 1$. For any $r, a, b \geq 0$ with $a + b = 1$,*

$$\left| \int g(y) (\mu \mathbf{k}(y) - \nu \mathbf{k}(y)) dy \right| \leq \|\mu \mathbf{k} - \nu \mathbf{k}\|_\infty \text{Vol}^{\frac{1}{2}}(r) + 2\tau_{\mathbf{k}}(ar) + 2\|\mathbf{k}\|_{L^{2,\infty}} \max\{\tau_\mu(br), \tau_\nu(br)\},$$

where $\text{Vol}(r) \triangleq \pi^d / \Gamma(d/2 + 1) r^d$ denotes the volume of the Euclidean ball $\mathcal{B}(0; r)$.

We first note that, by the square-root kernel definition (Def. 5), $\|\mathbf{k}_{\text{rt}}(x, \cdot)\|_{L^2} = \sqrt{\mathbf{k}_\star(x, x)}$ for each $x \in \mathbb{R}^d$. Since \mathbf{k}_\star is bounded, we therefore have $\mathbf{k}_{\text{rt}} \in L^{2,\infty}$ with $\|\mathbf{k}_{\text{rt}}\|_{L^{2,\infty}} = \sqrt{\|\mathbf{k}_\star\|_\infty}$. The result now follows by invoking Thms. 11 and 12 with $\mathbf{k} = \mathbf{k}_{\text{rt}}$.

H.1 Proof of Thm. 11: Square-root representation of MMD

Let $\mathcal{H}_{\mathbf{k}_\star}$ represent the RKHS of \mathbf{k}_\star . By Saitoh (1999, Thms. 1 and 2) and the definition (3) of \mathbf{k}_{rt} , for any $f \in \mathcal{H}_{\mathbf{k}_\star}$, there exists a function $g \in L^2$ such that

$$\|f\|_{\mathbf{k}_\star} = \|g\|_{L^2} \quad \text{and} \quad f(x) = \int g(y) \mathbf{k}_{\text{rt}}(x, y) dy \quad \text{for all } x \in \mathbb{R}^d, \tag{30}$$

and, for any $g \in L^2$, there exists an $f \in \mathcal{H}_{\mathbf{k}_\star}$ such that (30) holds. Note that the integral in (30) is well defined for each x since $g \in L^2$ and $\mathbf{k}_{\text{rt}} \in L^{2,\infty}$. Hence, we have

$$\begin{aligned}
 \text{MMD}_{\mathbf{k}_\star}(\mu, \nu) &= \sup_{f \in \mathcal{H}_{\mathbf{k}_\star}: \|f\|_{\mathbf{k}_\star} \leq 1} |\mu f - \nu f| \\
 &\stackrel{(i)}{=} \sup_{g \in L^2: \|g\|_{L^2} \leq 1} \left| \int \int g(y) \mathbf{k}_{\text{rt}}(y, x) dy d\mu(x) - \int \int g(y) \mathbf{k}_{\text{rt}}(y, x) dy d\nu(x) \right| \\
 &\stackrel{(ii)}{=} \sup_{g \in L^2: \|g\|_{L^2} \leq 1} \left| \int g(y) \mu \mathbf{k}_{\text{rt}}(y) dy - \int g(y) \nu \mathbf{k}_{\text{rt}}(y) dy \right|.
 \end{aligned}$$

where step (i) follows from (30), and we can swap the order of integration to obtain step (ii) using Fubini's theorem along with the following fact justified by Hölder's inequality:

$$\int \int |g(y)\mathbf{k}_{\text{rt}}(y, x)| dy d\tilde{\mu}(x) \leq \|g\|_{L^2} \|\mathbf{k}_{\text{rt}}\|_{L^{2,\infty}} \int d\tilde{\mu}(x) < \infty \text{ for any distribution. } \tilde{\mu} \quad (31)$$

H.2 Proof of Thm. 12: L^∞ bound on L^2 kernel error

Fix any $r \geq 0$, introduce the shorthand $\mathcal{B}(r) = \mathcal{B}(0; r)$, and define the restrictions

$$g_r(x) = g(x)\mathbb{1}_{\mathcal{B}(r)}(x), \quad \mathbf{k}_r(x, z) \triangleq \mathbf{k}(x, z) \cdot \mathbb{1}_{\mathcal{B}(r)}(z), \quad \text{and} \quad \mathbf{k}_r^{(c)} \triangleq \mathbf{k} - \mathbf{k}_r,$$

so that $\mu\mathbf{k} = \mu\mathbf{k}_r + \mu\mathbf{k}_r^{(c)}$. We first note that, by Cauchy-Schwarz, $g_r \in L^1 \cap L^2$ with

$$\|g_r\|_{L^1} \leq \|g_r\|_{L^2} \cdot \sqrt{\text{Vol}(r)} \leq \|g\|_{L^2} \cdot \sqrt{\text{Vol}(r)} \leq \sqrt{\text{Vol}(r)} \quad (32)$$

and that, exactly as in (31), $\int \int |g(y)\mathbf{k}(y, x)| dy d\tilde{\mu}(x) < \infty$ for any distribution $\tilde{\mu}$ so that each of the integrals to follow is well defined. We now apply the triangle inequality and Hölder's inequality to obtain

$$\begin{aligned} \left| \int g(y)(\mu\mathbf{k}(y) - \nu\mathbf{k}(y)) dy \right| &= \left| \int g(y)(\mu\mathbf{k}_r(y) - \nu\mathbf{k}_r(y)) dy + \int g(y)(\mu\mathbf{k}_r^{(c)}(y) - \nu\mathbf{k}_r^{(c)}(y)) dy \right| \\ &\leq \left| \int g_r(y)(\mu\mathbf{k}_r(y) - \nu\mathbf{k}_r(y)) dy \right| + \left| \int_{\|y\|_2 \geq r} g(y)(\mu\mathbf{k}_r^{(c)}(y) - \nu\mathbf{k}_r^{(c)}(y)) dy \right| \\ &\leq \|g_r\|_{L^1} \cdot \|\mu\mathbf{k} - \nu\mathbf{k}\|_\infty + \left| \int_{\|y\|_2 \geq r} g(y)(\mu\mathbf{k}(y) - \nu\mathbf{k}(y)) dy \right|. \end{aligned} \quad (33)$$

Next, we bound the second term in (33). For any $x, y \in \mathbb{R}^d$ with $\|y\|_2 \geq r$ and scalars $a, b \in [0, 1]$ such that $a + b = 1$, either $\|x - y\|_2 \geq ar$ or $\|x\|_2 \geq br$. Hence,

$$\begin{aligned} \left| \int_{\|y\|_2 \geq r} g(y)(\mu\mathbf{k}(y) - \nu\mathbf{k}(y)) dy \right| &= \left| \int_{\|y\|_2 \geq r} g(y) \int_{x \in \mathbb{R}^d} \mathbf{k}(x, y) (d\mu(x) - d\nu(x)) dy \right| \\ &\leq \left| \int_{\|y\|_2 \geq r} \int_{\|x-y\|_2 \geq ar} g(y)\mathbf{k}(x, y) (d\mu(x) - d\nu(x)) dy \right| \\ &\quad + \left| \int_{\|y\|_2 \geq r} \int_{\|x\|_2 \geq br} g(y)\mathbf{k}(x, y) (d\mu(x) - d\nu(x)) dy \right| \\ &=: T_1 + T_2. \end{aligned}$$

Note that both $T_1, T_2 < \infty$ since $g \in L^2$, $\mathbf{k}_{\text{rt}} \in L^{2,\infty}$ and μ, ν are probability measures. We now bound the terms T_1 and T_2 separately in (34a) and (34b) below. These bounds, together with the estimates (32) and (33), yield our claim.

Bounding T_1 Substituting $x - y = z$, we have

$$\begin{aligned} T_1 &\leq \int_{\|x-z\|_2 \geq r} \int_{\|z\|_2 \geq ar} |g(x-z)\mathbf{k}(x, x-z)| |d\mu(x) - d\nu(x)| dz \\ &\leq \int \int_{\|z\|_2 \geq ar} |g(x-z)\mathbf{k}(x, x-z)| dz |d\mu(x) - d\nu(x)| \\ &\stackrel{(i)}{\leq} \int \|g(x-\cdot)\|_{L^2} \sup_{x'} \left(\int_{\|z\|_2 \geq ar} \mathbf{k}^2(x', x'-z) dz \right)^{1/2} |d\mu(x) - d\nu(x)| \\ &\stackrel{(ii)}{=} \int \|g\|_{L^2} \tau_{\mathbf{k}}(ar) |d\mu(x) - d\nu(x)| \stackrel{(iii)}{\leq} 2 \|g\|_{L^2} \tau_{\mathbf{k}}(ar), \end{aligned} \quad (34a)$$

where step (i) follows from Cauchy-Schwarz, step (ii) from the definition (5) of $\tau_{\mathbf{k}}$, and step (iii) from the fact $\int |d\mu(x) - d\nu(x)| \leq 2$.

Bounding T_2 We have

$$\begin{aligned}
 T_2 &\leq \int_{\|x\|_2 \geq br} \left(\int_{\|y\|_2 \geq r} |g(y) \mathbf{k}(x, y)| dy \right) |d\mu(x) - d\nu(x)| \\
 &\leq \int_{\|x\|_2 \geq br} \|g\|_{L^2} \sup_{x'} \|\mathbf{k}(x', \cdot)\|_{L^2} |d\mu(x) - d\nu(x)| \\
 &\stackrel{(iv)}{\leq} 2 \|g\|_{L^2} \cdot \|\mathbf{k}\|_{L^2, \infty} \max\{\tau_\mu(br), \tau_\nu(br)\}.
 \end{aligned} \tag{34b}$$

where step (iv) follows from the definitions (5) and (18) of (τ_μ, τ_ν) and $\|\mathbf{k}\|_{L^2, \infty}$.

Appendix I. Proof of Cor. 2: MMD error from square-root L^∞ error

Let $e'_d \triangleq 2^{\frac{d}{2}} e_d$ for e_d defined in Thm. 2. Notice that the bound (10) for the choice of $a = b = \frac{1}{2}$ can be rewritten as

$$\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{Q}) \leq \inf_r (e'_d r^{\frac{d}{2}} \varepsilon + 2\tilde{\tau}(r)). \tag{35}$$

Then the claims of Cor. 2 follow by optimizing the RHS of (35) over the choice of r depending on the tail decay of $\tilde{\tau}$. Throughout the proofs c, c', ρ denote the (exactly same) constants underlying the assumed tail decay of $\tilde{\tau}$.

Proof for COMPACT part Choosing $r \downarrow c'$, we obtain that

$$\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{Q}) \leq \inf_r (e'_d r^{\frac{d}{2}} \varepsilon + 2c\mathbb{I}(r \leq c')) \leq \varepsilon \cdot e'_d (c')^{\frac{d}{2}}.$$

Proof for SUBGAUSS part Choosing $r = \sqrt{\frac{1}{c'} \log(1 \vee \frac{8cc'}{de'_d \varepsilon})}$, we obtain that

$$\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{Q}) \leq \inf_r (e'_d r^{\frac{d}{2}} \varepsilon + 2ce^{-c'r^2}) \leq \varepsilon \cdot e'_d \left[\left(\frac{\log(1 \vee \frac{8cc'}{de'_d \varepsilon})}{c'} \right)^{\frac{d}{4}} + \frac{d}{4c'} \right].$$

Proof for SUBEXP part Choosing $r = \frac{1}{c'} \log(1 \vee \frac{4cc'}{de'_d \varepsilon})$, we obtain that

$$\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{Q}) \leq \inf_r (e'_d r^{\frac{d}{2}} \varepsilon + 2ce^{-c'r}) \leq \varepsilon \cdot e'_d \left[\left(\frac{\log(1 \vee \frac{4cc'}{de'_d \varepsilon})}{c'} \right)^{\frac{d}{2}} + \frac{d}{2c'} \right].$$

Proof for HEAVYTAIL(ρ) part Choosing $r = (\frac{4c\rho}{de'_d \varepsilon})^{\frac{2}{d+2\rho}}$, we obtain that

$$\text{MMD}_{\mathbf{k}_*}(\mathbb{P}, \mathbb{Q}) \leq \inf_r (e'_d r^{\frac{d}{2}} \varepsilon + 2cr^{-\rho}) \leq (\varepsilon \cdot e'_d)^{\frac{2\rho}{d+2\rho}} \cdot \left(\frac{4c\rho}{d} \right)^{\frac{d}{d+2\rho}} \left(1 + \frac{2\rho}{d} \right).$$

Appendix J. Proof of Thm. 3: Self-balancing Hilbert walk properties

We prove each property from Thm. 3 one by one.

J.1 Property (i): Functional sub-Gaussianity

We prove the functional sub-Gaussianity claim (13) by induction on the iteration $i \in \{0, \dots, n\}$. Our proof uses the following lemma proved in App. J.7, which supplies a convenient decomposition for the self-balancing Hilbert walk iterates.

Lemma 13 (Alternate representation of ψ_i). *For each $i \in [n]$, the iterate ψ_i of the self-balancing Hilbert walk (Alg. 3) satisfies*

$$\langle \psi_i, u \rangle_{\mathcal{H}} = \left\langle \psi_{i-1}, u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\rangle_{\mathcal{H}} + \varepsilon_i \langle f_i, u \rangle_{\mathcal{H}} \quad \text{for all } u \in \mathcal{H} \quad (36)$$

for the random variable $\varepsilon_i \triangleq \mathbb{I}[|\alpha_i| \leq \mathbf{a}_i](\eta_i + \alpha_i/\mathbf{a}_i)$ which satisfies

$$\mathbb{E}[\varepsilon_i | \psi_{i-1}] = 0, \quad \varepsilon_i \in [-2, 2], \quad \text{and} \quad \mathbb{E}[e^{t\varepsilon_i} | \psi_{i-1}] \leq e^{t^2/2} \quad \text{for all } t \in \mathbb{R}. \quad (37)$$

Now we proceed with our induction argument.

Base case For $i = 1$, noting that $\psi_0 = 0$, we have

$$\mathbb{E}[\exp(\langle \psi_1, u \rangle_{\mathcal{H}})] \stackrel{(36)}{=} \mathbb{E}[\exp(\varepsilon_1 \langle f_1, u \rangle_{\mathcal{H}})] \stackrel{(37)}{\leq} \exp\left(\frac{1}{2} \langle f_1, u \rangle_{\mathcal{H}}^2\right) \leq \exp\left(\frac{1}{2} \|f_1\|_{\mathcal{H}}^2 \|u\|_{\mathcal{H}}^2\right),$$

where the last step follows from Cauchy-Schwarz, and thus ψ_1 is sub-Gaussian with parameter $\sigma_1 = \|f_1\|_{\mathcal{H}}$ as desired.

Inductive step Fix any $i \in [n]$ with $i \geq 2$ and assume that the functional sub-Gaussianity claim (13) holds for ψ_{i-1} with σ_{i-1} . We have

$$\begin{aligned} \mathbb{E}[\exp(\langle \psi_i, u \rangle_{\mathcal{H}})] &= \mathbb{E}[\mathbb{E}[\exp(\langle \psi_i, u \rangle_{\mathcal{H}}) | \psi_{i-1}]] \\ &\stackrel{(36)}{=} \mathbb{E}\left[\exp\left(\left\langle \psi_{i-1}, u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\rangle_{\mathcal{H}}\right) \cdot \mathbb{E}[\exp(\varepsilon_i \langle f_i, u \rangle_{\mathcal{H}}) | \psi_{i-1}]\right] \\ &\stackrel{(37)}{\leq} \mathbb{E}\left[\exp\left(\left\langle \psi_{i-1}, u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\rangle_{\mathcal{H}}\right) \cdot \exp\left(\frac{1}{2} \langle f_i, u \rangle_{\mathcal{H}}^2\right)\right] \\ &= \exp\left(\frac{1}{2} \langle f_i, u \rangle_{\mathcal{H}}^2\right) \cdot \mathbb{E}\left[\exp\left(\left\langle \psi_{i-1}, u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\rangle_{\mathcal{H}}\right)\right] \\ &\stackrel{(i)}{\leq} \exp\left(\frac{1}{2} \langle f_i, u \rangle_{\mathcal{H}}^2 + \frac{\sigma_{i-1}^2}{2} \left\| u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\|_{\mathcal{H}}^2\right), \end{aligned} \quad (38)$$

where step (i) follows from the induction hypothesis. Simplifying the exponent in the display (38) using Cauchy-Schwarz and the definition (12) of σ_i , we have

$$\begin{aligned} \frac{1}{2} \langle f_i, u \rangle_{\mathcal{H}}^2 + \frac{\sigma_{i-1}^2}{2} \left\| u - f_i \frac{\langle f_i, u \rangle_{\mathcal{H}}}{\mathbf{a}_i} \right\|_{\mathcal{H}}^2 &= \frac{1}{2} \langle f_i, u \rangle_{\mathcal{H}}^2 + \frac{\sigma_{i-1}^2}{2} \left(\|u\|_{\mathcal{H}}^2 + \langle f_i, u \rangle_{\mathcal{H}}^2 \left(\frac{\|f_i\|_{\mathcal{H}}^2}{\mathbf{a}_i^2} - \frac{2}{\mathbf{a}_i} \right) \right) \\ &= \frac{\sigma_{i-1}^2}{2} \|u\|_{\mathcal{H}}^2 + \langle f_i, u \rangle_{\mathcal{H}}^2 \cdot \left(\frac{1}{2} + \frac{\sigma_{i-1}^2 \|f_i\|_{\mathcal{H}}^2}{2\mathbf{a}_i^2} - \frac{\sigma_{i-1}^2}{\mathbf{a}_i} \right) \\ &\leq \frac{\sigma_{i-1}^2}{2} \|u\|_{\mathcal{H}}^2 + \langle f_i, u \rangle_{\mathcal{H}}^2 \cdot \left(\frac{1}{2} + \frac{\sigma_{i-1}^2 \|f_i\|_{\mathcal{H}}^2}{2\mathbf{a}_i^2} - \frac{\sigma_{i-1}^2}{\mathbf{a}_i} \right)_+ \\ &\leq \frac{\|u\|_{\mathcal{H}}^2}{2} \left(\sigma_{i-1}^2 + \|f_i\|_{\mathcal{H}}^2 \left(1 + \frac{\sigma_{i-1}^2}{\mathbf{a}_i^2} (\|f_i\|_{\mathcal{H}}^2 - 2\mathbf{a}_i) \right)_+ \right) \\ &= \frac{\sigma_i^2}{2} \|u\|_{\mathcal{H}}^2. \end{aligned}$$

J.2 Property (ii): Signed sum representation

Since Alg. 3 adds $\pm f_i$ to ψ_{i-1} whenever $|\alpha_i| = |\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}| \leq \mathbf{a}_i$, by the union bound, it suffices to lower bound the probability of this event by $1 - \delta_i$ for each i . The following lemma establishes this bound using the functional sub-Gaussianity (13) of each ψ_{i-1} .

Lemma 14 (Self-balancing Hilbert walk success probability). *The self-balancing Hilbert walk (Alg. 3) with threshold $\mathbf{a}_i \geq \sigma_{i-1} \|f_i\|_{\mathcal{H}} \sqrt{2 \log(2/\delta_i)}$ for $\delta_i \in (0, 1]$ satisfies*

$$\Pr(\mathcal{E}_i) \geq 1 - \delta_i \quad \text{for} \quad \mathcal{E}_i = \{|\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}| \leq \mathbf{a}_i\}.$$

Proof Instantiate the notation of Thm. 3. The sub-Gaussian Hoeffding inequality (Wainwright, 2019, Prop. 2.5), the functional sub-Gaussianity of ψ_{i-1} (13), and the choice of \mathbf{a}_i imply that

$$\Pr(\mathcal{E}_i^c) = \Pr(|\langle \psi_{i-1}, f_i \rangle_{\mathcal{H}}| > \mathbf{a}_i) \leq 2 \exp(-\mathbf{a}_i^2 / (2\sigma_{i-1}^2 \|f_i\|_{\mathcal{H}}^2)) \leq 2 \exp(-\log(2/\delta_i)) = \delta_i. \quad \blacksquare$$

J.3 Property (iii): Exact halving via symmetrization

Whenever the signed sum representation (14) holds, we have

$$\psi_n = \sum_{i=1}^n \eta_i f_i = \sum_{i=1}^n (\eta_i g_{2i-1} - \eta_i g_{2i}) = \sum_{i=1}^{2n} g_i - 2 \sum_{i \in \mathcal{I}} g_i$$

where the last step follows from the definition of \mathcal{I} .

J.4 Property (iv): Pointwise sub-Gaussianity in RKHS

The reproducing property of the kernel \mathbf{k}_* and the established functional sub-Gaussianity (13) yield

$$\mathbb{E}[\exp(\psi_i(x))] = \mathbb{E}[\exp(\langle \psi_i, \mathbf{k}_*(x, \cdot) \rangle_{\mathcal{H}})] \leq \exp\left(\frac{\sigma_i^2 \|\mathbf{k}_*(x, \cdot)\|_{\mathcal{H}}^2}{2}\right) = \exp\left(\frac{\sigma_i^2 \mathbf{k}_*(x, x)}{2}\right), \quad \forall x \in \mathcal{X}.$$

J.5 Property (v): Sub-Gaussian constant bound

We establish the bound (15) for all $i \in \{0, \dots, n\}$ by induction on the iteration i .

Base case The claim (15) holds for the base case, $i = 0$, since $\sigma_0 = 0$.

Inductive step Fix any $i \in [n]$ and assume that the claim (15) holds for σ_{i-1} .

If either $\|f_i\|_{\mathcal{H}} = 0$ or $\sigma_{i-1}^2 \geq \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \|f_i\|_{\mathcal{H}}^2}$, then $\sigma_i^2 = \sigma_{i-1}^2$ by the definition (12) of σ_i and the assumption that $\frac{\|f_i\|_{\mathcal{H}}^2}{2} \leq \mathbf{a}_i$, completing the inductive step.

If, alternatively, $\sigma_{i-1}^2 < \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \|f_i\|_{\mathcal{H}}^2}$ and $\|f_i\|_{\mathcal{H}} > 0$, then our assumptions $\frac{\|f_i\|_{\mathcal{H}}^2}{1+q} \leq \mathbf{a}_i \leq \frac{\|f_i\|_{\mathcal{H}}^2}{1-q}$ imply that $(\frac{\|f_i\|_{\mathcal{H}}^2}{\mathbf{a}_i} - 1)^2 \leq q^2$. Hence, by the definition (12) of σ_i and the inductive

hypothesis,

$$\begin{aligned}\sigma_i^2 &= \sigma_{i-1}^2 + \|f_i\|_{\mathcal{H}}^2 \left(1 + \frac{\sigma_{i-1}^2}{\mathbf{a}_i^2} (\|f_i\|_{\mathcal{H}}^2 - 2\mathbf{a}_i)\right) \\ &= \|f_i\|_{\mathcal{H}}^2 + \sigma_{i-1}^2 \left(\frac{\|f_i\|_{\mathcal{H}}^2}{\mathbf{a}_i} - 1\right)^2 \\ &\leq (1 - q^2) \frac{\|f_i\|_{\mathcal{H}}^2}{1 - q^2} + q^2 \frac{\max_{j \in [i-1]} \|f_j\|_{\mathcal{H}}^2}{1 - q^2} \leq \frac{\max_{j \in [i]} \|f_j\|_{\mathcal{H}}^2}{1 - q^2},\end{aligned}$$

completing the inductive step.

J.6 Property (vi): Adaptive thresholding

Define $c_1^* = \max(c^*, 1)$ and let $q = \frac{(c_1^*)^2 - 1}{(c_1^*)^2 + 1} \in [0, 1)$ so that

$$\frac{1}{1 - q} = \frac{1 + (c_1^*)^2}{2} \quad \text{and} \quad \frac{1}{1 - q^2} = \frac{(c_1^* + 1/c_1^*)^2}{4} \leq \frac{(c^* + 1/c^*)^2}{4},$$

since $1 = \operatorname{argmin}_{c \geq 0} c + 1/c$. By assumption, $\mathbf{a}_i \geq \|f_i\|_{\mathcal{H}}^2 \geq \frac{\|f_i\|_{\mathcal{H}}^2}{1 + q}$ for all $i \in [n]$.

Now suppose that $\sigma_{i-1}^2 < \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \|f_i\|_{\mathcal{H}}^2}$ and $\|f_i\|_{\mathcal{H}} > 0$. If $\mathbf{a}_i \leq \|f_i\|_{\mathcal{H}}^2$, then $\mathbf{a}_i \leq \frac{\|f_i\|_{\mathcal{H}}^2}{1 - q}$. If, alternatively, $\mathbf{a}_i \leq c_i \sigma_{i-1} \|f_i\|_{\mathcal{H}}$, then

$$\mathbf{a}_i < \frac{1}{2} \|f_i\|_{\mathcal{H}}^2 (1 + c_i^2) \leq \frac{\|f_i\|_{\mathcal{H}}^2}{1 - q}.$$

The conclusion now follows from the sub-Gaussian constant bound (15).

J.7 Proof of Thm. 13: Alternate representation of ψ_i

Alg. 3 and our definition $\varepsilon_i \triangleq \mathbb{I}[|\alpha_i| \leq \mathbf{a}_i](\eta_i + \alpha_i/\mathbf{a}_i)$ give

$$\psi_i = \psi_{i-1} - f_i \alpha_i / \mathbf{a}_i + \mathbb{I}[|\alpha_i| \leq \mathbf{a}_i](f_i \alpha_i / \mathbf{a}_i + \eta_i f_i) = \psi_{i-1} - f_i \frac{\langle f_i, \psi_{i-1} \rangle_{\mathcal{H}}}{\mathbf{a}_i} + \varepsilon_i f_i.$$

Taking an inner product with $u \in \mathcal{H}$ now yields the equality (36). By construction, $\varepsilon_i \in [c_{\min}, c_{\max}] \subseteq [-2, 2]$ for

$$c_{\min} = \max(-2, \min(0, -1 + \alpha_i/\mathbf{a}_i)) \quad \text{and} \quad c_{\max} = \min(2, \max(0, 1 + \alpha_i/\mathbf{a}_i))$$

by construction. Moreover,

$$\begin{aligned}\mathbb{E}[\varepsilon_i \mid \psi_{i-1}, |\alpha_i| > \mathbf{a}_i] &= 0 \\ \mathbb{E}[\varepsilon_i \mid \psi_{i-1}, |\alpha_i| \leq \mathbf{a}_i] &= \left(1 + \frac{\alpha_i}{\mathbf{a}_i}\right) \cdot \frac{1}{2} \left(1 - \frac{\alpha_i}{\mathbf{a}_i}\right) + \left(-1 + \frac{\alpha_i}{\mathbf{a}_i}\right) \cdot \frac{1}{2} \left(1 + \frac{\alpha_i}{\mathbf{a}_i}\right) = 0,\end{aligned}$$

so that $\mathbb{E}[\varepsilon_i \mid \psi_{i-1}] = 0$ as claimed. The conditional sub-Gaussianity claim

$$\mathbb{E}[e^{t\varepsilon_i} \mid \psi_{i-1}] \leq e^{t^2/2} \quad \text{for all } t \in \mathbb{R},$$

now follows from Hoeffding's lemma (Hoeffding, 1963, (4.16)) since ε_i is bounded with $c_{\max} - c_{\min} \leq 2$ and mean-zero conditional on ψ_{i-1} .

Appendix K. Proof of Thm. 4: L^∞ guarantees for kernel halving

We start by showing that after m rounds of kernel halving the output size is $n_{\text{out}} = \lfloor \frac{n}{2^m} \rfloor$ (also see Footnote 3). We prove this by induction. The base case of $m = 1$ can be directly verified. Let $n_j \triangleq |\mathcal{S}^{(j)}|$ denote the output size after j rounds of kernel halving and assume that the hypothesis is true for round j , i.e., $n_j = \lfloor \frac{n}{2^j} \rfloor$. Then for round $j + 1$, we have $n_{j+1} = \lfloor n_j/2 \rfloor = \lfloor \frac{\lfloor n/2^j \rfloor}{2} \rfloor \stackrel{(i)}{=} \lfloor \frac{n}{2^{j+1}} \rfloor$, where step (i) follows from the fact (Graham et al., 1994, Eqn. (3.11)) that $\lfloor \frac{\lfloor \ell \rfloor}{2} \rfloor = \lfloor \frac{\ell}{2} \rfloor$ for any integer ℓ . Our desired claim follows.

Next, define $n_{\text{in}} = 2^m n_{\text{out}} = 2^m \lfloor \frac{n}{2^m} \rfloor$. In Apps. K.1 and K.2 we will show that the respective claims (16) and (17) hold whenever $n = n_{\text{in}}$, that is, whenever 2^m divides n evenly. Now suppose that $n \neq n_{\text{in}}$. Since the output of m rounds of kernel halving depends only on the first n_{in} points, the evenly divisible case of App. K.1 implies that, for part (a),

$$\|\mathbb{P}_{n_{\text{in}}} \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_\infty \leq \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, 1, d, \delta^*, \delta', R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}}) \text{ for } n_{\text{in}} = 2 \lfloor \frac{n}{2} \rfloor = 2n_{\text{out}}, \quad (39)$$

with probability at least $1 - \delta' - \sum_{i=1}^{n_{\text{in}}/2} \delta_i$, and App. K.1 implies that, for part (b),

$$\|\mathbb{P}_{n_{\text{in}}} \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty \leq \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, m, d, \delta^*, \delta', R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}}) \text{ for } n_{\text{in}} = 2^m \lfloor \frac{n}{2^m} \rfloor = 2^m n_{\text{out}}, \quad (40)$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{n_{\text{in}}/2^j} \delta_i$. Next, to recover the bounds (16) and (17), we use the following deterministic inequalities:

$$\|\mathbb{P}_n \mathbf{k} - \mathbb{P}_{n_{\text{in}}} \mathbf{k}\|_\infty \leq \frac{2(n - n_{\text{in}})\|\mathbf{k}\|_\infty}{n} \leq \frac{2(2^m - 1)\|\mathbf{k}\|_\infty}{n} \leq \frac{2\|\mathbf{k}\|_\infty}{n_{\text{out}}}, \quad (41)$$

$$R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}} \leq R_{\mathcal{S}_n, \mathbf{k}, n}, \quad \text{and} \quad (42)$$

$$\begin{aligned} \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, m, d, \delta, \delta', R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}}) + 2 &\stackrel{(i)}{\leq} \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, m, d, \delta, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}) + 2 \\ &\stackrel{(ii)}{\leq} \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}), \end{aligned} \quad (43)$$

where (42) follows directly from the definition (7), and step (i) follows from (42) and the fact that $\mathfrak{M}_{\mathbf{k}}(n, m, d, \delta, \delta', R)$ (8) is non-decreasing in R , and step (ii) follows since $\mathfrak{M}_{\mathbf{k}}(n, m, d, \delta, \delta', R)$ is non-decreasing in n and $n_{\text{in}} \neq n$. For part (a) we conclude, by (39),

$$\begin{aligned} \|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_\infty &\leq \|\mathbb{P}_n \mathbf{k} - \mathbb{P}_{n_{\text{in}}} \mathbf{k}\|_\infty + \|\mathbb{P}_{n_{\text{in}}} \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_\infty \\ &\stackrel{(41)}{\leq} \frac{2\|\mathbf{k}\|_\infty}{n_{\text{out}}} + \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, 1, d, \delta^*, \delta', R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}}) \\ &\stackrel{(43)}{\leq} \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}). \end{aligned}$$

For part (b), we conclude, by (40),

$$\begin{aligned} \|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty &\leq \|\mathbb{P}_n \mathbf{k} - \mathbb{P}_{n_{\text{in}}} \mathbf{k}\|_\infty + \|\mathbb{P}_{n_{\text{in}}} \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty \\ &\stackrel{(41)}{\leq} \frac{2\|\mathbf{k}\|_\infty}{n_{\text{out}}} + \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n_{\text{in}}, 1, d, \delta^*, \delta', R_{\mathcal{S}_{n_{\text{in}}, \mathbf{k}, n_{\text{in}}}}) \\ &\stackrel{(43)}{\leq} \frac{\|\mathbf{k}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}). \end{aligned}$$

K.1 Proof of part (a): Kernel halving yields a 2-thinned L^∞ coreset

As noted earlier, we prove this part assuming n is even. Consider a self-balancing Hilbert walk (Alg. 3) with inputs $(f_i)_{i=1}^{n/2}$ and $(\mathbf{a}_i)_{i=1}^{n/2}$, where $f_i = \mathbf{k}(x_{2i-1}, \cdot) - \mathbf{k}(x_{2i}, \cdot)$ and $\mathbf{a}_i = \max(\|f_i\|_{\mathbf{k}_*}, \sigma_{i-1} \sqrt{2 \log(2/\delta_i)}, \|f_i\|_{\mathbf{k}_*}^2)$, and σ_i was defined in (12). Property (iii) of Thm. 3 implies that for a self-balancing walk with the choices summarized above, the event $\mathcal{E}_{\text{half}} = \{|\alpha_i| \leq \mathbf{a}_i \text{ for all } i \in [n/2]\}$ satisfies

$$\Pr(\mathcal{E}_{\text{half}}) \geq 1 - \sum_{i=1}^{n/2} \delta_i. \quad (44)$$

Consider kernel halving coupled with the above instantiation of self-balancing Hilbert walk. Due to the equivalence with kernel halving on the event $\mathcal{E}_{\text{half}}$, we conclude that the output $\psi_{n/2}$ of self-balancing Hilbert walk matches with that of the kernel halving, and satisfies

$$\frac{1}{n} \psi_{n/2} = \frac{1}{n} \sum_{i=1}^n \mathbf{k}(x_i, \cdot) - \frac{1}{n/2} \sum_{x \in \mathcal{S}^{(1)}} \mathbf{k}(x, \cdot) = \mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}, \quad (45)$$

on the event $\mathcal{E}_{\text{half}}$. Furthermore, on the event $\mathcal{E}_{\text{half}}$, we can also write that the kernel halving coreset satisfies $\mathcal{S}^{(1)} = (x_i)_{i \in \mathcal{I}}$ for $\mathcal{I} = \{2i - \frac{\eta_i - 1}{2} : i \in [n/2]\}$ and η_i defined in Alg. 3. Finally, applying property (vi) of Thm. 3 for $\sigma_{n/2}$ with $c_i = \sqrt{2 \log(2/\delta_i)}$, we obtain that

$$\sigma_{n/2}^2 \leq \frac{\max_{i \leq n/2} \|f_i\|_{\mathbf{k}_*}^2}{4} \cdot 2 \log\left(\frac{2}{\delta^*}\right) \left(1 + \frac{1}{2 \log(2/\delta^*)}\right)^2 \stackrel{(i)}{\leq} 4 \|\mathbf{k}\|_\infty \log\left(\frac{2}{\delta^*}\right) \quad (46)$$

where in step (i), we use the fact that $f_i = \mathbf{k}(x_{2i-1}, \cdot) - \mathbf{k}(x_{2i}, \cdot)$ satisfies

$$\|f_i\|_{\mathbf{k}_*}^2 = \mathbf{k}(x_{2i-1}, x_{2i-1}) + \mathbf{k}(x_{2i}, x_{2i}) - 2\mathbf{k}(x_{2i-1}, x_{2i}) \leq 4\|\mathbf{k}\|_\infty.$$

Next, we split the proof in two parts: **Case (I)** When $R_{\mathcal{S}_n, \mathbf{k}, n} < R_{\mathcal{S}_n}$, and **Case (II)** when $R_{\mathcal{S}_n, \mathbf{k}, n} = R_{\mathcal{S}_n}$, where $R_{\mathcal{S}_n, \mathbf{k}, n}$ was defined in (7). We prove the results for these two cases in Apps. K.1.1 and K.1.2 respectively. In the sequel, we make use of the following tail quantity of the kernel:

$$\bar{\tau}_{\mathbf{k}}(R') \triangleq \sup\{\|\mathbf{k}(x, y)\| : \|x - y\|_2 \geq R'\}. \quad (47)$$

K.1.1 PROOF FOR CASE (I): WHEN $R_{\mathcal{S}_n, \mathbf{k}, n} < R_{\mathcal{S}_n}$

By definition (7), for this case,

$$R_{\mathcal{S}_n, \mathbf{k}, n} = n^{1+\frac{1}{d}} R_{\mathbf{k}, n} + n^{\frac{1}{d}} \frac{\|\mathbf{k}\|_\infty}{L_{\mathbf{k}}}. \quad (48)$$

On the event $\mathcal{E}_{\text{half}}$, the following lemma provides a high probability bound on $\|\psi_{n/2}\|_\infty$ in terms of the kernel parameters, the sub-Gaussianity parameter $\sigma_{n/2}$, and the size of the cover (Wainwright, 2019, Def. 5.1) of a neighborhood of the input points $(x_i)_{i=1}^n$.

Lemma 15 (A direct covering bound on $\|\psi_{n/2}\|_\infty$). *Fix $R \geq r > 0$ and $\delta' > 0$, and suppose $\mathcal{C}^n(r, R)$ is a set of minimum cardinality satisfying*

$$\bigcup_{i=1}^n \mathcal{B}(x_i, R) \subseteq \bigcup_{z \in \mathcal{C}^n(r, R)} \mathcal{B}(z, r). \quad (49)$$

If \mathbf{k} satisfies Assumps. 1 and 2, then, conditional on the event $\mathcal{E}_{\text{half}}$ (45), the event

$$\mathcal{E}_\infty \triangleq \{\|\psi_{n/2}\|_\infty \leq \max(n\bar{\tau}_{\mathbf{k}}(R), nL_{\mathbf{k}}r + \sigma_{n/2} \sqrt{2\|\mathbf{k}\|_\infty \log(2|\mathcal{C}^n(r, R)|/\delta')})\}, \quad (50)$$

occurs with probability at least $1 - \delta'$, i.e., $\Pr(\mathcal{E}_\infty | \mathcal{E}_{\text{half}}) \geq 1 - \delta'$, where $\bar{\tau}_{\mathbf{k}}$ was defined in (47).

Thm. 15 succeeds in controlling $\psi_{n/2}(x)$ for all $x \in \mathbb{R}^d$ since either x lies far from every input point x_i so that each $\mathbf{k}(x_i, x)$ in the expansion (45) is small or x lies near some x_i , in which case $\psi_{n/2}(x)$ is well approximated by $\psi_{n/2}(z)$ for $z \in \mathcal{C}^n(r, R)$. The proof inspired by the covering argument of Phillips and Tai (2020, Lem. 2.1) and using the pointwise sub-Gaussianity property of Thm. 3 over the finite cover \mathcal{C}^n can be found in App. K.3.

Now we put together the pieces to prove Thm. 4.

First, (Wainwright, 2019, Lem. 5.7) implies that $|\mathcal{C}^1(r, R)| \leq (1 + 2R/r)^d$ (i.e., any ball of radius R in \mathbb{R}^d can be covered by $(1 + 2R/r)^d$ balls of radius r). Thus for an arbitrary R , we can conclude that

$$|\mathcal{C}^n(r, R)| \leq n(1 + 2R/r)^d = (n^{1/d} + 2n^{1/d}R/r)^d. \quad (51)$$

Second we fix R and r such that $n\bar{\tau}_{\mathbf{k}}(R) = \|\mathbf{k}\|_\infty$ and $nL_{\mathbf{k}}r = \|\mathbf{k}\|_\infty$, so that $\frac{R}{r} \leftarrow nR_{\mathbf{k},n} \frac{L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty}$ (c.f. (6) and (47)). Substituting these choices of radii in the bound (50) of Thm. 15, we find that conditional to $\mathcal{E}_{\text{half}} \cap \mathcal{E}_\infty$, we have

$$\begin{aligned} \|\psi_{n/2}\|_\infty &\leq \max(n\bar{\tau}_{\mathbf{k}}(R), nL_{\mathbf{k}}r + \sigma_{n/2}\sqrt{2\|\mathbf{k}\|_\infty \log(2|\mathcal{C}^n(r, R)|/\delta')}) \\ &\stackrel{(46)}{\leq} \|\mathbf{k}\|_\infty + 2\sqrt{2}\|\mathbf{k}\|_\infty \sqrt{\log(\frac{2}{\delta^*}) \left[\log(\frac{2}{\delta'}) + d \log\left(n^{\frac{1}{d}} + \frac{2L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty} \cdot n^{1+\frac{1}{d}}R_{\mathbf{k},n}\right) \right]} \\ &\stackrel{(48)}{\leq} \|\mathbf{k}\|_\infty + 2\sqrt{2}\|\mathbf{k}\|_\infty \sqrt{\log(\frac{2}{\delta^*}) \left[\log(\frac{2}{\delta'}) + d \log\left(\frac{2L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty} (R_{\mathbf{k},n} + R_{\mathcal{S}_n, \mathbf{k}, n})\right) \right]} \\ &\stackrel{(8)}{\leq} \|\mathbf{k}\|_\infty \cdot 2\mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}). \end{aligned} \quad (53)$$

Putting (45) and (53) together, we conclude

$$\begin{aligned} \Pr(\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_\infty > \frac{1}{n/2} \mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n})) &\leq \Pr((\mathcal{E}_{\text{half}} \cap \mathcal{E}_\infty)^c) \\ &= \Pr(\mathcal{E}_{\text{half}}^c \cup \mathcal{E}_\infty^c) \\ &= \Pr(\mathcal{E}_{\text{half}}^c) + \Pr(\mathcal{E}_{\text{half}} \cap \mathcal{E}_\infty^c) \\ &= \Pr(\mathcal{E}_{\text{half}}^c) + \Pr(\mathcal{E}_\infty^c | \mathcal{E}_{\text{half}}) \\ &\leq \delta' + \sum_{i=1}^{n/2} \delta_i, \end{aligned}$$

where the last step follows from (44) and Thm. 15. The claim now follows.

K.1.2 PROOF FOR CASE (II): WHEN $R_{\mathcal{S}_n, \mathbf{k}, n} = R_{\mathcal{S}_n}$

In this case, we split the proof for bounding $\|\psi_{n/2}\|_\infty$ into two lemmas. First, we relate the $\|\psi_{n/2}\|_\infty$ in terms of the tail behavior of \mathbf{k} and the supremum of differences for $\psi_{n/2}$ between any pair points on a Euclidean ball (see App. K.4 for the proof):

Lemma 16 (A basic bound on $\|\psi_{n/2}\|_\infty$). *If \mathbf{k} satisfies Assump. 1, then, conditional on the event $\mathcal{E}_{\text{half}}$ (45), for any fixed $R = R' + R_{\mathcal{S}_n}$ with $R' > 0$ and any fixed $\delta' \in (0, 1)$,*

$$\tilde{\mathcal{E}}_\infty = \left\{ \|\psi_{n/2}\|_\infty \leq \max(n\bar{\tau}_{\mathbf{k}}(R'), \sigma_{n/2} \|\mathbf{k}\|_\infty^{\frac{1}{2}} \sqrt{2 \log(\frac{4}{\delta'})}) + \sup_{x, x' \in \mathcal{B}(0, R)} |\psi_{n/2}(x) - \psi_{n/2}(x')| \right\} \quad (54)$$

occurs with probability at least $1 - \frac{\delta'}{2}$, i.e., $\Pr(\tilde{\mathcal{E}}_\infty | \mathcal{E}_{\text{half}}) \geq 1 - \frac{\delta'}{2}$, where $\bar{\tau}_{\mathbf{k}}$ was defined in (47).

Next, to control the supremum term on the RHS of the display (54), we establish a high probability bound in the next lemma. Its proof in App. K.5 proceeds by showing that $\psi_{n/2}$ is an Orlicz process with a suitable metric and then applying standard concentration arguments for such processes.

Lemma 17 (A high probability bound on supremum of $\psi_{n/2}$ differences). *If \mathbf{k} satisfies Assumps. 1 and 2, then, for any fixed $R > 0, \delta' \in (0, 1)$, the event*

$$\mathcal{E}_{\text{sup}} \triangleq \left\{ \sup_{x, x' \in \mathcal{B}(0, R)} |\psi_{n/2}(x) - \psi_{n/2}(x')| \leq 8D_R \left(\sqrt{\log\left(\frac{4}{\delta'}\right)} + 6\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}R}{\|\mathbf{k}\|_{\infty}}\right)} \right) \right\} \quad (55)$$

occurs with probability at least $1 - \frac{\delta'}{2}$, where $D_R \triangleq \sqrt{\frac{32}{3}}\sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}} \min\left(1, \sqrt{\frac{1}{2}L_{\mathbf{k}}R}\right)$.

We now turn to the rest of the proof for Thm. 4. We apply both Thms. 16 and 17 with $R = R_{\mathcal{S}_n} + R_{\mathbf{k}, n}$ (6) and (7). For this R , we have $R' = R_{\mathbf{k}, n}$ in Thm. 16 and hence $n\bar{\tau}_{\mathbf{k}}(R') \leq \|\mathbf{k}\|_{\infty}$. Now, condition on the event $\mathcal{E}_{\text{half}} \cap \tilde{\mathcal{E}}_{\infty} \cap \mathcal{E}_{\text{sup}}$. Then, we have

$$\begin{aligned} & n\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_{\infty} \\ & \stackrel{(45)}{=} \|\psi_{n/2}\|_{\infty} \\ & \stackrel{(54)}{\leq} \max(\|\mathbf{k}\|_{\infty}, \sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}\sqrt{2\log\left(\frac{4}{\delta'}\right)} + \sup_{x, x' \in \mathcal{B}(0, R)} |\psi_{n/2}(x) - \psi_{n/2}(x')|) \\ & \stackrel{(55)}{\leq} \max\left(\|\mathbf{k}\|_{\infty}, \sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}\sqrt{2\log\left(\frac{4}{\delta'}\right)} + \right. \\ & \quad \left. 32\sqrt{\frac{2}{3}}\sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}\left(\sqrt{\log\left(\frac{4}{\delta'}\right)} + 6\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}(R_{\mathcal{S}_n} + R_{\mathbf{k}, n})}{\|\mathbf{k}\|_{\infty}}\right)}\right)\right) \\ & \stackrel{(i)}{\leq} \max(\|\mathbf{k}\|_{\infty}, 32\sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}\left(\sqrt{\log\left(\frac{4}{\delta'}\right)} + 5\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}(R_{\mathcal{S}_n} + R_{\mathbf{k}, n})}{\|\mathbf{k}\|_{\infty}}\right)}\right)) \quad (56) \end{aligned}$$

$$\begin{aligned} & \stackrel{(46)}{\leq} \|\mathbf{k}\|_{\infty} \cdot 64\sqrt{\log\left(\frac{2}{\delta^*}\right)} \left[\sqrt{\log\left(\frac{4}{\delta'}\right)} + 5\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}(R_{\mathcal{S}_n} + R_{\mathbf{k}, n})}{\|\mathbf{k}\|_{\infty}}\right)} \right] \\ & \stackrel{(ii)}{\leq} \|\mathbf{k}\|_{\infty} 2\mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}), \quad (57) \end{aligned}$$

where step (i) follows from the fact that $D_R \leq \sqrt{\frac{32}{3}}\sigma_{n/2}\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}$, and in step (ii) we have used the working assumption for this case, i.e., $R_{\mathcal{S}_n, \mathbf{k}, n} = R_{\mathcal{S}_n}$. As a result,

$$\begin{aligned} \Pr(\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(1)} \mathbf{k}\|_{\infty} > \frac{2}{n}\mathfrak{M}_{\mathbf{k}}(n, 1, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n})) & \leq \Pr((\mathcal{E}_{\text{half}} \cap \tilde{\mathcal{E}}_{\infty} \cap \mathcal{E}_{\text{sup}})^c) \\ & = \Pr(\mathcal{E}_{\text{half}}^c \cup \tilde{\mathcal{E}}_{\infty}^c \cup \mathcal{E}_{\text{sup}}^c) \\ & = \Pr(\mathcal{E}_{\text{half}}^c \cup \mathcal{E}_{\text{sup}}^c) + \Pr((\mathcal{E}_{\text{half}}^c \cup \mathcal{E}_{\text{sup}}^c)^c \cap \tilde{\mathcal{E}}_{\infty}^c) \\ & = \Pr(\mathcal{E}_{\text{half}}^c \cup \mathcal{E}_{\text{sup}}^c) + \Pr(\mathcal{E}_{\text{half}} \cap \mathcal{E}_{\text{sup}} \cap \tilde{\mathcal{E}}_{\infty}^c) \\ & \leq \Pr(\mathcal{E}_{\text{half}}^c \cup \mathcal{E}_{\text{sup}}^c) + \Pr(\mathcal{E}_{\text{half}} \cap \tilde{\mathcal{E}}_{\infty}^c) \\ & \leq \Pr(\mathcal{E}_{\text{half}}^c) + \Pr(\mathcal{E}_{\text{sup}}^c) + \Pr(\tilde{\mathcal{E}}_{\infty}^c | \mathcal{E}_{\text{half}}) \\ & \leq \sum_{i=1}^{n/2} \delta_i + \frac{\delta'}{2} + \frac{\delta'}{2}, \end{aligned}$$

where the last step follows from (44) and Thms. 16 and 17. The desired claim follows.

K.2 Proof of part (b): Repeated kernel halving yields a 2^m -thinned L^∞ coresets

As noted earlier, we prove this part assuming $n/2^m \in \mathbb{N}$. The proof in this section follows by applying the arguments from the previous section, separately for each round and then invoking the sub-Gaussianity of a weighted sum of the output functions from each round.

Note that coresets $\mathcal{S}^{(j-1)}$ is independent of the randomness for the j -th round of kernel halving and thus can be treated as fixed for that round. When running kernel halving with the input $\mathcal{S}^{(j-1)}$, let $f_{i,j}, \mathbf{a}_{i,j}, \psi_{i,j}, \alpha_{i,j}, \eta_{i,j}$ denote the analog of the quantities $f_i, \mathbf{a}_i, \psi_i, \alpha_i, \eta_i$ defined in Alg. 2. Like in part (a), we would couple the j -th round of kernel halving with an instantiation of self-balancing Hilbert walk with inputs $(f_{i,j}, \mathbf{a}_{i,j})_{i=1}^{n/2^j}$ and final output $\psi_{n/2^j,j}$, where $\mathbf{a}_{i,j} = \max(\|f_{i,j}\|_{\mathbf{k}} \sigma_{i-1,j} \sqrt{2 \log(\frac{4m}{2^j \delta_i^*})}, \|f_{i,j}\|_{\mathbf{k}}^2)$ and we define $\sigma_{i,j}$ in a recursive manner as in (12) with $f_{i,j}, \mathbf{a}_{i,j}, \sigma_{i,j}$ taking the role of $f_i, \mathbf{a}_i, \sigma_i$ respectively.

With this set-up, first we apply property (i) of Thm. 3 which implies that given $\mathcal{S}^{(j-1)}$, the function $\psi_{n/2^j,j}$ is $\sigma_{n/2^j,j}$ sub-Gaussian, where

$$\sigma_{n/2^j,j}^2 \leq 4\|\mathbf{k}\|_\infty \log\left(\frac{4m}{2^j \delta_j^*}\right) \quad \text{with} \quad \delta_j^* = \min(\delta_i)_{i=1}^{n/2^j}, \quad (58)$$

using an argument similar to (46) with property (vi) of Thm. 3. Next, we note that for j -th round, the event $\mathcal{E}_{\text{half}}^{(j)} = |\alpha_{i,j}| \leq \mathbf{a}_{i,j}$ for all $i \in [n/2^j]$, satisfies

$$\mathbb{P}(\mathcal{E}_{\text{half}}^{(j)}) \geq 1 - \sum_{i=1}^{n/2^j} \delta_i \frac{2^{j-1}}{m}; \quad (59)$$

and this event serves as the analog the equivalence event $\mathcal{E}_{\text{half}}$ from part (a) for the j -th round of kernel halving. Thus on the event $\mathcal{E}_{\text{half}}^{(j)}$, we can write that the kernel halving coresets $\mathcal{S}^{(j)} = (x_i)_{i \in \mathcal{I}_j}$ for $\mathcal{I}_j = \{2i - \frac{\eta_{i,j}-1}{2} : i \in [n/2^j]\}$ and $\eta_{i,j}$ as defined while running Alg. 3, so that

$$\frac{2^{j-1}}{n} \psi_{n/2^j,j} = \frac{1}{n/2^{j-1}} \sum_{x \in \mathcal{S}^{(j-1)}} \mathbf{k}(x, \cdot) - \frac{1}{n/2^j} \sum_{x \in \mathcal{S}^{(j)}} \mathbf{k}(x, \cdot). \quad (60)$$

Now conditional on the event $\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}$, we conclude that the output of all m kernel halving rounds are equivalent to the output of m different self-balancing Hilbert walks (each with exact two-thinning in each round). Putting the pieces together, we conclude that on the event $\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}$, we have

$$\begin{aligned} \mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k} &= \frac{1}{n} \sum_{x \in \mathcal{S}_n} \mathbf{k}(x, \cdot) - \frac{1}{n/2^m} \sum_{x \in \mathcal{S}^{(m)}} \mathbf{k}(x, \cdot) \\ &\stackrel{(i)}{=} \sum_{j=1}^m \frac{1}{n/2^{j-1}} \sum_{x \in \mathcal{S}^{(j-1)}} \mathbf{k}(x, \cdot) - \frac{1}{n/2^j} \sum_{x \in \mathcal{S}^{(j)}} \mathbf{k}(x, \cdot) \\ &= \sum_{j=1}^m \frac{2^{j-1}}{n} \psi_{n/2^j,j}. \end{aligned} \quad (61)$$

where we abuse notation $\mathcal{S}^{(0)} \triangleq \mathcal{S}_n$ in step (i) for simplicity of expressions.

Next, we use the following basic fact:

Lemma 18 (Sub-Gaussian additivity). *For a sequence of random variables $(Z_j)_{j=1}^m$ such that Z_j is a σ_j sub-Gaussian variable conditional on (Z_1, \dots, Z_{j-1}) , the random variable $Z = \sum_{j=1}^m \theta_j Z_j$ is $(\sum_{j=1}^m \theta_j^2 \sigma_j^2)^{1/2}$ sub-Gaussian.*

Proof We will prove the result for $\mathcal{Z}_s = \sum_{j=1}^s \theta_j Z_j$ by induction on $s \leq m$. The result holds for the base case of $s = 0$ as $\mathcal{Z}_s = 0$ is 0 sub-Gaussian. For the inductive case, suppose the result holds for s . Then we may apply the tower property, our conditional sub-Gaussianity assumption, and our inductive hypothesis in turn to conclude

$$\mathbb{E}[e^{\sum_{j=1}^{s+1} \theta_j Z_j}] = \mathbb{E}[e^{\sum_{j=1}^{s+1} \theta_j Z_j} \mathbb{E}[e^{\theta_{s+1} Z_{s+1}} \mid Z_{1:s}]] \leq \mathbb{E}[e^{\sum_{j=1}^s \theta_j Z_j}] e^{\frac{\theta_{s+1}^2 \sigma_{s+1}^2}{2}} = e^{\frac{\sum_{j=1}^{s+1} \theta_j^2 \sigma_j^2}{2}}.$$

Hence, \mathcal{Z}_{s+1} is $\sqrt{\sum_{j=1}^{s+1} \theta_j^2 \sigma_j^2}$ sub-Gaussian, and the proof is complete. \blacksquare

Applying Thm. 18 to the output sequence $(\psi_{n/2^j, j})_{j=1}^m$ for the m rounds of self-balancing Hilbert walks, we conclude that, the random variable

$$\mathcal{W}_m \triangleq \sum_{j=1}^m \frac{2^{j-1}}{n} \psi_{n/2^j, j} \quad (62)$$

is sub-Gaussian with parameter

$$\sigma_{\mathcal{W}_m} \triangleq \frac{2}{\sqrt{3}} \frac{2^m}{n} \sqrt{\|\mathbf{k}\|_\infty \log\left(\frac{6m}{2^m \delta^*}\right)} \stackrel{(i)}{\geq} \sqrt{\|\mathbf{k}\|_\infty \sum_{j=1}^m \frac{4^j}{n^2} \log\left(\frac{4m}{2^j \delta^*}\right)} \stackrel{(ii)}{\geq} \sqrt{\sum_{j=1}^m \left(\frac{2^{j-1}}{n}\right)^2 \sigma_{n/2^j, j}^2}, \quad (63)$$

conditional to the input \mathcal{S}_n , where step (ii) follows since

$$\|\mathbf{k}\|_\infty \sum_{j=1}^m 4^j \log\left(\frac{4m}{2^j \delta^*}\right) \stackrel{(iii)}{\geq} 4 \|\mathbf{k}\|_\infty \sum_{j=1}^m 4^{j-1} \log\left(\frac{4m}{2^j \delta^*}\right) \stackrel{(58)}{\geq} \sum_{j=1}^m 4^{j-1} \sigma_{n/2^j, j}^2,$$

and step (iii) follows from the fact that $\delta^* = \min(\delta_j^*)_{j=1}^m$ (58). Now to prove step (i), we note that

$$\mathfrak{G}_1 \triangleq \sum_{j=1}^m 4^j = \frac{4}{3}(4^m - 1) \leq \frac{4}{3}4^m, \quad \text{and} \quad \mathfrak{G}_2 \triangleq \sum_{j=1}^m j 4^j = \frac{4m}{3} \cdot 4^m - \frac{\mathfrak{G}_1}{3},$$

which in turn implies that

$$\begin{aligned} \sum_{j=1}^m 4^j \log\left(\frac{4m}{2^j \delta^*}\right) &= \mathfrak{G}_1 \log\left(\frac{4m}{\delta^*}\right) - \log 2 \cdot \mathfrak{G}_2 = \mathfrak{G}_1 \left(\log\left(\frac{4m}{\delta^*}\right) + \frac{\log 2}{3} \right) - \frac{4m \log 2}{3} \cdot 4^m \\ &\leq \mathfrak{G}_1 \cdot \log\left(\frac{6m}{\delta^*}\right) - \frac{4m \log 2}{3} \cdot 4^m \\ &\leq \frac{4}{3}4^m \cdot \log\left(\frac{6m}{\delta^*}\right) - \frac{4}{3}4^m \cdot \log 2^m = \frac{4}{3}4^m \cdot \log\left(\frac{6m}{2^m \delta^*}\right), \end{aligned}$$

thereby establishing step (i).

Next, analogous to the proof of part (a), we split the proof in two parts: **Case (I)** When $R_{\mathcal{S}_n, \mathbf{k}, n} < R_{\mathcal{S}_n}$, in which case, we proceed with a direct covering argument to bound $\|\mathcal{W}_m\|_\infty$ using a lemma analogous to Thm. 15; and **Case (II)** when $R_{\mathcal{S}_n, \mathbf{k}, n} = R_{\mathcal{S}_n}$, in which case, we proceed with a metric-entropy based argument to bound $\|\mathcal{W}_m\|_\infty$ using two lemmas that are analogous Thms. 16 and 17.

K.2.1 PROOF FOR CASE (I): WHEN $R_{\mathcal{S}_n, \mathbf{k}, n} < R_{\mathcal{S}_n}$

Recall that for this case, $R_{\mathcal{S}_n, \mathbf{k}, n}$ is given by (48). The next lemma, a straightforward extension of Thm. 15, provides a high probability control on $\|\mathcal{W}_m\|_\infty$. Its proof is omitted for brevity.

Lemma 19 (A direct covering bound on $\|\mathcal{W}_m\|_\infty$). *Fix $R \geq r > 0$ and $\delta' > 0$, and recall the definition (49) of $\mathcal{C}^n(r, R)$. If \mathbf{k} satisfies Assumps. 1 and 2, then, conditional on the event $\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}$ (60), the event*

$$\mathcal{E}_\infty^{(m)} \triangleq \left\{ \|\mathcal{W}_m\|_\infty \leq \max\left(2^m \bar{\tau}_{\mathbf{k}}(R), 2^m L_{\mathbf{k}} r + \sigma_{\mathcal{W}_m} \sqrt{2\|\mathbf{k}\|_\infty \log(2|\mathcal{C}^n(r, R)|/\delta')}\right) \right\}, \quad (64)$$

occurs with probability at least $1 - \delta'$, i.e., $\Pr(\mathcal{E}_\infty^{(m)} | \cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}) \geq 1 - \delta'$.

Next, we repeat arguments similar to those used earlier around the display (52) and (53). Fix R and r such that $\bar{\tau}_{\mathbf{k}}(R) = \|\mathbf{k}\|_\infty/n$ and $L_{\mathbf{k}} r = \|\mathbf{k}\|_\infty/n$, so that $\frac{R}{r} \leftarrow n R_{\mathbf{k},n} \frac{L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty}$ (c.f. (6) and (47)). Substituting these choices of radii in the bound (64) of Thm. 19, we find that conditional to $\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)} \cap \mathcal{E}_\infty^{(m)}$, we have

$$\begin{aligned} \|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty &\stackrel{(61)}{=} \left\| \sum_{j=1}^m \frac{2^{j-1}}{n} \psi_{n/2^j, j} \right\|_\infty \\ &\stackrel{(62)}{=} \|\mathcal{W}_m\|_\infty \\ &\stackrel{(64)}{\leq} \max\left(2^m \bar{\tau}_{\mathbf{k}}(R), 2^m L_{\mathbf{k}} r + \sigma_{\mathcal{W}_m} \sqrt{2\|\mathbf{k}\|_\infty \log(2|\mathcal{C}^n(r, R)|/\delta')}\right) \\ &\stackrel{(i)}{\leq} \frac{2^m}{n} \|\mathbf{k}\|_\infty \left(1 + 2\sqrt{\frac{2}{3}} \sqrt{\log\left(\frac{6m}{2^m \delta^*}\right)} \left[\log\left(\frac{2}{\delta'}\right) + d \log\left(n^{\frac{1}{d}} + \frac{2L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty} \cdot n^{1+\frac{1}{d}} R_{\mathbf{k},n}\right)\right]\right) \\ &\stackrel{(48)}{\leq} \frac{2^m}{n} \|\mathbf{k}\|_\infty \left(1 + 2\sqrt{\frac{2}{3}} \sqrt{\log\left(\frac{6m}{2^m \delta^*}\right)} \left[\log\left(\frac{2}{\delta'}\right) + d \log\left(\frac{2L_{\mathbf{k}}}{\|\mathbf{k}\|_\infty} (R_{\mathbf{k},n} + R_{\mathcal{S}_n, \mathbf{k}, n})\right)\right]\right) \\ &\stackrel{(8)}{\leq} \|\mathbf{k}\|_\infty \cdot \frac{2^m}{n} \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n}), \end{aligned}$$

where in step (i), we have used the bounds (51) and (63). Putting the pieces together, we conclude

$$\begin{aligned} &\Pr(\|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty > \|\mathbf{k}\|_\infty \cdot \frac{2^m}{n} \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}, n})) \\ &\leq \Pr((\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)} \cap \mathcal{E}_\infty^{(m)})^c) \\ &= \Pr(\cup_{j=1}^m (\mathcal{E}_{\text{half}}^{(j)})^c \cup (\mathcal{E}_\infty^{(m)})^c) \\ &= \Pr(\cup_{j=1}^m (\mathcal{E}_{\text{half}}^{(j)})^c) + \Pr(\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)} \cap (\mathcal{E}_\infty^{(m)})^c) \\ &\leq \sum_{j=1}^m \Pr((\mathcal{E}_{\text{half}}^{(j)})^c) + \Pr((\mathcal{E}_\infty^{(m)})^c | \cap_{j=1}^m (\mathcal{E}_{\text{half}}^{(j)})) \\ &\leq \sum_{j=1}^m \sum_{i=1}^{n/2^j} \delta_i \cdot \frac{2^{j-1}}{m} + \delta', \end{aligned}$$

where the last step follows from (59) and Thm. 19. The claim follows.

K.2.2 PROOF FOR CASE (II): WHEN $R_{\mathcal{S}_n, \mathbf{k}, n} = R_{\mathcal{S}_n}$

In this case, we makes use of two lemmas. Their proofs (omitted for brevity) can be derived essentially by replacing $\psi_{n/2}$ and $\sigma_{n/2}$ with \mathcal{W}_m (62) and $\sigma_{\mathcal{W}_m}$ (63) respectively, and repeating the proof arguments from Thms. 16 and 17.

Lemma 20 (A basic bound on $\|\mathcal{W}_m\|_\infty$). *If \mathbf{k} satisfies Assump. 1, then, conditional on the event $\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}$ (60), for any fixed $R = R' + R_{\mathcal{S}_n}$ with $R' > 0$ and any fixed $\delta' \in (0, 1)$,*

$$\tilde{\mathcal{E}}_\infty^{(m)} = \left\{ \|\mathcal{W}_m\|_\infty \leq \max \left(2^m \bar{\tau}_{\mathbf{k}}(R'), \sigma_{\mathcal{W}_m} \|\mathbf{k}\|_\infty^{\frac{1}{2}} \sqrt{2 \log\left(\frac{4}{\delta'}\right)} + \sup_{x, x' \in \mathcal{B}(0, R)} |\mathcal{W}_m(x) - \mathcal{W}_m(x')| \right) \right\},$$

occurs with probability at least $1 - \frac{\delta'}{2}$, i.e., $\Pr(\tilde{\mathcal{E}}_\infty^{(m)} | \cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}) \geq 1 - \frac{\delta'}{2}$.

Lemma 21 (A high probability bound on supremum of \mathcal{W}_m differences). *If \mathbf{k} satisfies Assumps. 1 and 2, then, for any fixed $R > 0, \delta' \in (0, 1)$, given \mathcal{S}_n , the event*

$$\mathcal{E}_{\text{sup}}^{(m)} \triangleq \left\{ \sup_{x, x' \in \mathcal{B}(0, R)} |\mathcal{W}_m(x) - \mathcal{W}_m(x')| \leq 8D_R \left(\sqrt{\log\left(\frac{4}{\delta'}\right)} + 6\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}} R}{\|\mathbf{k}\|_\infty}\right)} \right) \right\}$$

occurs with probability at least $1 - \delta'/2$, where $D_R \triangleq \sqrt{\frac{32}{3}} \sigma_{\mathcal{W}_m} \|\mathbf{k}\|_\infty^{\frac{1}{2}} \min(1, \sqrt{\frac{1}{2} L_{\mathbf{k}} R})$.

Mimicking the arguments like those in display (56) and (57), with $R = R_{\mathcal{S}_n} + R_{\mathbf{k}, n}$, and $R' = R_{\mathbf{k}, n}$, we find that conditional on the event $\tilde{\mathcal{E}}_\infty^{(m)} \cap \mathcal{E}_{\text{sup}}^{(m)} \cap \cap_{j=1}^m \mathcal{E}_{\text{half}}^{(j)}$,

$$\begin{aligned} \|\mathbb{P}_n \mathbf{k} - \mathbb{Q}_{\text{KH}}^{(m)} \mathbf{k}\|_\infty &\stackrel{(61)}{=} \left\| \sum_{j=1}^m \frac{2^{j-1}}{n} \psi_{n/2^j, j} \right\|_\infty \\ &\stackrel{(62)}{=} \|\mathcal{W}_m\|_\infty \\ &\leq \max\left(\frac{2^m}{n} \|\mathbf{k}\|_\infty, 32\sigma_{\mathcal{W}_m} \|\mathbf{k}\|_\infty^{\frac{1}{2}} \left(\sqrt{\log\left(\frac{4}{\delta'}\right)} + 5\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}(R_{\mathcal{S}_n} + R_{\mathbf{k}, n})}{\|\mathbf{k}\|_\infty}\right)} \right)\right) \\ &\stackrel{(63)}{\leq} \|\mathbf{k}\|_\infty \cdot \frac{2^m}{n} \cdot 37 \sqrt{\log\left(\frac{6m}{2^m \delta^*}\right)} \left[\sqrt{\log\left(\frac{4}{\delta'}\right)} + 5\sqrt{d \log\left(2 + \frac{L_{\mathbf{k}}(R_{\mathcal{S}_n} + R_{\mathbf{k}, n})}{\|\mathbf{k}\|_\infty}\right)} \right] \\ &\stackrel{(8)}{=} \|\mathbf{k}\|_\infty \cdot \frac{2^m}{n} \mathfrak{M}_{\mathbf{k}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_n}), \end{aligned}$$

which does not happen with probability at most

$$\begin{aligned} \Pr\left(\left(\cap_{j=1}^m \mathcal{E}_{\text{half}}^{(m)} \cap \tilde{\mathcal{E}}_\infty^{(m)} \cap \mathcal{E}_{\text{sup}}^{(m)}\right)^c\right) &\leq \sum_{j=1}^m \Pr\left(\left(\mathcal{E}_{\text{half}}^{(j)}\right)^c\right) + \Pr\left(\left(\mathcal{E}_{\text{sup}}^{(m)}\right)^c\right) + \Pr\left(\left(\tilde{\mathcal{E}}_\infty^{(m)}\right)^c | \cap_{j=1}^m \left(\mathcal{E}_{\text{half}}^{(j)}\right)\right) \\ &\leq \sum_{j=1}^m \sum_{i=1}^{\lfloor n/2^j \rfloor} \delta_i \frac{2^{j-1}}{m} + \frac{\delta'}{2} + \frac{\delta'}{2}, \end{aligned}$$

where the last step follows from (59) and Thms. 20 and 21 as claimed. The proof is now complete.

K.3 Proof of Thm. 15: A direct covering bound on $\|\psi_{n/2}\|_\infty$

We claim that conditional on the event $\mathcal{E}_{\text{half}}$ (45), we deterministically have

$$\|\psi_{n/2}\|_\infty \leq \max\left\{ n\bar{\tau}_{\mathbf{k}}(R), nL_{\mathbf{k}}r + \max_{z \in \mathcal{C}^n(r, R)} |\psi_{n/2}(z)| \right\}, \quad (65)$$

and the event

$$\left\{ \max_{z \in \mathcal{C}^n(r, R)} |\psi_{n/2}(z)| \leq \sigma_{n/2} \sqrt{2\|\mathbf{k}\|_\infty \log(2|\mathcal{C}^n(r, R)|/\delta')} \right\} \quad (66)$$

occurs with probability at least $1 - \delta'$. Putting these two claims together yields the lemma. We now prove these two claims separately.

K.3.1 PROOF OF (65)

Note that on the event $\mathcal{E}_{\text{half}}$ (45) we have

$$\psi_{n/2} = n(\mathbb{P}_n \mathbf{k} - \tilde{\mathbb{Q}}_{n/2} \mathbf{k}) = \sum_{i=1}^n \eta_i \mathbf{k}(x_i, \cdot) \quad (67)$$

for some $\eta_i \in \{-1, 1\}$. Now fix any $x \in \mathbb{R}^d$, and introduce the shorthand $\mathcal{C}^n = \mathcal{C}^n(r, R)$. The result follows by considering two cases.

(Case 1) $x \notin \bigcup_{i=1}^n \mathcal{B}(x_i, R)$ In this case, we have $\|x - x_i\|_2 \geq R$ for all $i \in [n]$ and therefore, representation (67) yields that

$$|\psi_{n/2}(x)| = |\sum_{i=1}^n \eta_i \mathbf{k}(x_i, x)| \leq \sum_{i=1}^n |\eta_i| |\mathbf{k}(x_i, x)| = \sum_{i=1}^n |\mathbf{k}(x_i, x)| \leq n\bar{\tau}_{\mathbf{k}}(R),$$

by Cauchy-Schwarz's inequality and the definition (47) of $\bar{\tau}_{\mathbf{k}}$.

(Case 2) $x \in \bigcup_{i=1}^n \mathcal{B}(x_i, R)$ By the definition (49) of our cover \mathcal{C}^n , there exists $z \in \mathcal{C}^n$ such that $\|x - z\|_2 \leq r$. Therefore, on the event $\mathcal{E}_{\text{half}}$, using representation (67), we find that

$$\begin{aligned} |\psi_{n/2}(x)| &\leq |\sum_{i=1}^n \eta_i \mathbf{k}(x_i, x)| = |\sum_{i=1}^n \eta_i (\mathbf{k}(x_i, x) - \mathbf{k}(x_i, z) + \mathbf{k}(x_i, z))| \\ &\leq |\sum_{i=1}^n \eta_i (\mathbf{k}(x_i, x) - \mathbf{k}(x_i, z))| + |\sum_{i=1}^n \eta_i \mathbf{k}(x_i, z)| \\ &\leq \sum_{i=1}^n |\mathbf{k}(x_i, x) - \mathbf{k}(x_i, z)| + |\psi_{n/2}(z)| \\ &\leq nL_{\mathbf{k}}(r) + \sup_{z' \in \mathcal{C}^n} |\psi_{n/2}(z')| \end{aligned}$$

by Cauchy-Schwarz's inequality.

K.3.2 PROOF OF (66)

Introduce the shorthand $\mathcal{C}^n = \mathcal{C}^n(r, R)$. Then applying the union bound, the pointwise sub-Gaussian property (iv) of Thm. 3, and the sub-Gaussian Hoeffding inequality (Wainwright, 2019, Prop. 2.5), we find that

$$\begin{aligned} \Pr(\max_{z \in \mathcal{C}^n} |\psi_{n/2}(z)| > t) &\leq \sum_{z \in \mathcal{C}^n} 2 \exp(-t^2 / (2\sigma_{n/2}^2 \mathbf{k}(z, z))) \\ &\leq 2|\mathcal{C}^n| \exp(-t^2 / (2\sigma_{n/2}^2 \|\mathbf{k}\|_{\infty})) = \delta' \\ \text{for } t &\triangleq \sigma_{n/2} \sqrt{2\|\mathbf{k}\|_{\infty} \log(2|\mathcal{C}^n|/\delta')}, \end{aligned}$$

as claimed.

K.4 Proof of Thm. 16: A basic bound on $\|\psi_{n/2}\|_{\infty}$

For any $R > 0$, we deterministically have

$$\|\psi_{n/2}\|_{\infty} \leq \max(\sup_{x \in \mathcal{B}(0, R)} |\psi_{n/2}(x)|, \sup_{x \in \mathcal{B}^c(0, R)} |\psi_{n/2}(x)|) \quad (68)$$

Since $R = R' + R_{\mathcal{S}_n}$, for any $x \in \mathcal{B}^c(0, R)$, we have $\|x - x_i\|_2 \geq R'$ for all $i \in [n]$. Thus conditional on the event $\mathcal{E}_{\text{half}}$, applying property (ii) from Thm. 3, we find that

$$|\psi_{n/2}(x)| = |\sum_{i=1}^n \eta_i \mathbf{k}(x_i, x)| \leq \sum_{i=1}^n |\eta_i| |\mathbf{k}(x_i, x)| = \sum_{i=1}^n |\mathbf{k}(x_i, x)| \leq n\bar{\tau}_{\mathbf{k}}(R'),$$

by Cauchy-Schwarz's inequality and the definition (47) of $\bar{\tau}_{\mathbf{k}}$.

For the first term on the RHS of (68), we have

$$\begin{aligned} \sup_{x \in \mathcal{B}(0,R)} |\psi_{n/2}(x)| &\leq |\psi_{n/2}(0)| + \sup_{x \in \mathcal{B}(0,R)} |\psi_{n/2}(x) - \psi_{n/2}(0)| \\ &\leq |\psi_{n/2}(0)| + \sup_{x, x' \in \mathcal{B}(0,R)} |\psi_{n/2}(x) - \psi_{n/2}(x')|. \end{aligned}$$

Now, the sub-Gaussianity of $\psi_{n/2}$ (property (iv) of Thm. 3) with $x = 0$, and the sub-Gaussian Hoeffding inequality (Wainwright, 2019, Prop. 2.5) imply that

$$|\psi_{n/2}(0)| \leq \sigma_{n/2} \sqrt{2\mathbf{k}(0,0) \log(4/\delta)} \leq \sigma_{n/2} \sqrt{2\|\mathbf{k}\|_{\infty} \log(4/\delta')},$$

with probability at least $1 - \delta'/2$. Putting the pieces together completes the proof.

K.5 Proof of Thm. 17: A high probability bound on supremum of $\psi_{n/2}$ differences

The proof proceeds by using concentration arguments for Orlicz processes (Wainwright, 2019, Def. 5.5). Given a set $\mathbb{T} \subseteq \mathbb{R}^d$, a random process $\{Z_x, x \in \mathbb{T}\}$ is called an Orlicz Ψ_2 -process with respect to the metric ρ if

$$\|Z_x - Z_{x'}\|_{\Psi_2} \leq \rho(x, x') \quad \text{for all } x, x' \in \mathbb{T},$$

where for any random variable Z , its Orlicz Ψ_2 -norm is defined as

$$\|Z\|_{\Psi_2} = \inf\{\lambda > 0 : \mathbb{E}[\exp(Z^2/\lambda^2)] \leq 2\}.$$

Our next result (see App. K.6 for the proof) establishes that $\psi_{n/2}$ is an Orlicz process with respect to a suitable metric. (For clarity, we use \mathcal{B}_2 to denote the Euclidean ball.)

Lemma 22 ($\psi_{n/2}$ is an Orlicz Ψ_2 -process). *If \mathbf{k} satisfies Assumps. 1 and 2, then, given any fixed $R > 0$, the random process $\{\psi_{n/2}(x), x \in \mathcal{B}_2(0; R)\}$ is an Orlicz Ψ_2 -process with respect to the metric ρ defined in (70), i.e.,*

$$\|\psi_{n/2}(x) - \psi_{n/2}(x')\|_{\Psi_2} \leq \rho(x, x') \quad \text{for all } x, x' \in \mathcal{B}_2(0; R), \quad (69)$$

where the metric ρ is defined as

$$\rho(x, x') = \sqrt{\frac{8}{3}} \cdot \sigma_{n/2} \cdot \min(\sqrt{2L_{\mathbf{k}}\|x - x'\|_2}, 2\sqrt{\|\mathbf{k}\|_{\infty}}). \quad (70)$$

Given Thm. 22, we can invoke high probability bounds for Orlicz processes. For the Orlicz process $\{\psi_{n/2}(x), x \in \mathbb{T}\}$, given any fixed $\delta' \in (0, 1]$, Wainwright (2019, Thm 5.36) implies that

$$\sup_{x, x' \in \mathbb{T}} |\psi_{n/2}(x) - \psi_{n/2}(x')| \leq 8 \left(\mathcal{J}_{\mathbb{T}, \rho}(D) + D \sqrt{\log(4/\delta')} \right), \quad (71)$$

with probability at least $1 - \delta'/2$, where $D \triangleq \sup_{x, x' \in \mathbb{T}} \rho(x, x')$ denotes the ρ -diameter of \mathbb{T} , and the quantity $\mathcal{J}_{\mathbb{T}, \rho}(D)$ is defined as

$$\mathcal{J}_{\mathbb{T}, \rho}(D) \triangleq \int_0^D \sqrt{\log(1 + \mathcal{N}_{\mathbb{T}, \rho}(u))} du. \quad (72)$$

Here $\mathcal{N}_{\mathbb{T},\rho}(u)$ denotes the u -covering number of \mathbb{T} with respect to metric ρ , namely the cardinality of the smallest cover $\mathcal{C}_{\mathbb{T},\rho}(u) \subseteq \mathbb{R}^d$ such that

$$\mathbb{T} \subseteq \cup_{z \in \mathcal{C}_{\mathbb{T},\rho}(u)} \mathcal{B}_\rho(z; u) \quad \text{where} \quad \mathcal{B}_\rho(z; u) \triangleq \{x \in \mathbb{R}^d : \rho(z, x) \leq u\}. \quad (73)$$

To avoid confusion, let \mathcal{B}_2 denote the Euclidean ball (metric induced by $\|\cdot\|_2$). Then in our setting, we have $\mathbb{T} = \mathcal{B}_2(0; R)$ and hence

$$D_R \triangleq \sup_{x, x' \in \mathcal{B}_2(0; R)} \rho(x, x') = \min\left(\sqrt{\frac{\beta_R}{2}}, \sigma_{n/2} \sqrt{\frac{32\|\mathbf{k}\|_\infty}{3}}\right), \quad \text{with } \beta_R = \frac{32}{3}\sigma_{n/2}^2 L_{\mathbf{k}} R. \quad (74)$$

Some algebra establishes that this definition of D_R is identical to that specified in Thm. 17. Next, we derive bounds for $\mathcal{N}_{\mathbb{T},\rho}$ and $\mathcal{J}_{\mathbb{T},\rho}$ (see App. K.7 for the proof).

Lemma 23 (Bounds on $\mathcal{N}_{\mathbb{T},\rho}$ and $\mathcal{J}_{\mathbb{T},\rho}$). *If \mathbf{k} satisfies Assumps. 1 and 2, then, for $\mathbb{T} = \mathcal{B}_2(0; R)$ with $R > 0$ and ρ defined in (70), we have*

$$\mathcal{N}_{\mathbb{T},\rho}(u) \leq (1 + \beta_R/u^2)^d, \quad \text{and} \quad (75)$$

$$\mathcal{J}_{\mathbb{T},\rho}(D_R) \leq \sqrt{d} D_R \left[3 + \sqrt{2 \log(\beta_R/D_R^2)}\right], \quad (76)$$

where D_R and β_R are defined in (74).

Doing some algebra, we find that

$$\frac{\beta_R}{D_R^2} = \max\left(2, \frac{L_{\mathbf{k}} R}{\|\mathbf{k}\|_\infty}\right) \leq 2 + \frac{L_{\mathbf{k}} R}{\|\mathbf{k}\|_\infty}. \quad (77)$$

Moreover, note that $3 + \sqrt{2 \log(2+a)} \leq 6\sqrt{\log(2+a)}$ for all $a \geq 0$, so that the bound (76) can be simplified to

$$\mathcal{J}_{\mathbb{T},\rho}(D_R) \leq 6\sqrt{d} D_R \sqrt{\log\left(2 + \frac{L_{\mathbf{k}} R}{\|\mathbf{k}\|_\infty}\right)}. \quad (78)$$

Substituting the bound (78) in (71) yields that the event \mathcal{E}_{sup} defined in (55) holds with probability at least $1 - \delta'/2$, as claimed.

K.6 Proof of Thm. 22: $\psi_{n/2}$ is an Orlicz Ψ_2 -process

The proof of this lemma follows from the sub-Gaussianity of $\psi_{n/2}$ established in Thm. 3. Introduce the shorthand $Y \triangleq \psi_{n/2}(x) - \psi_{n/2}(x')$. Applying property (i) of Thm. 3 with $u = \mathbf{k}(x, \cdot) - \mathbf{k}(x', \cdot)$ along with the reproducing property of the kernel \mathbf{k} , we find that for any $t \in \mathbb{R}$,

$$\mathbb{E}[\exp(tY)] = \mathbb{E}[\exp(t\langle \psi_{n/2}, \mathbf{k}(x, \cdot) - \mathbf{k}(x', \cdot) \rangle_{\mathbf{k}_*})] \leq \exp\left(\frac{1}{2}t^2 \sigma_{n/2}^2 \|\mathbf{k}(x, \cdot) - \mathbf{k}(x', \cdot)\|_{\mathbf{k}_*}^2\right).$$

That is, the random variable Y is sub-Gaussian with parameter $\sigma_Y \triangleq \sigma_{n/2} \|\mathbf{k}(x, \cdot) - \mathbf{k}(x', \cdot)\|_{\mathbf{k}_*}$.

Next, Wainwright (2019, Thm 2.6 (iv)) yields that $\mathbb{E}[\exp(\frac{3Y^2}{8\sigma_Y^2})] \leq 2$, which in turn implies

that $\|Y\|_{\Psi_2} \leq \sqrt{\frac{8}{3}}\sigma_Y$. Moreover, we have

$$\|\mathbf{k}(x, \cdot) - \mathbf{k}(x', \cdot)\|_{\mathbf{k}_*}^2 = \mathbf{k}(x, x) - \mathbf{k}(x, x') + \mathbf{k}(x', x') - \mathbf{k}(x', x) \leq \min(2L_{\mathbf{k}}\|x - x'\|_2, 4\|\mathbf{k}\|_\infty),$$

using the Lipschitz continuity of \mathbf{k} and the definition of $\|\mathbf{k}\|_\infty$ (4). Putting the pieces together along with the definition (70) of ρ , we find that $\|Y\|_{\Psi_2} \leq \rho(x, x')$ thereby yielding the claim (69).

K.7 Proof of Thm. 23: Bounds on $\mathcal{N}_{\mathbb{T},\rho}$ and $\mathcal{J}_{\mathbb{T},\rho}$

We use the relation of ρ to the Euclidean norm $\|\cdot\|_2$ to establish the bound (75). The definitions (70) and (73) imply that

$$\mathcal{B}_2(z; \frac{u^2}{\alpha}) \subseteq \mathcal{B}_\rho(z; u) \quad \text{for} \quad \alpha \triangleq \frac{16}{3} \sigma_{n/2}^2 L_{\mathbf{k}},$$

since $\rho(x, x') \leq \sqrt{\alpha} \|x - x'\|_2$. Consequently, any u^2/α -cover \mathcal{C} of \mathbb{T} in the Euclidean ($\|\cdot\|_2$) metric automatically yields a u -cover of \mathbb{T} in ρ metric as we note that

$$\mathcal{B}_2(0; R) \subseteq \cup_{z \in \mathcal{C}} \mathcal{B}_2(z; \frac{u^2}{\alpha}) \subseteq \cup_{z \in \mathcal{C}} \mathcal{B}_\rho(z; u).$$

Consequently, the smallest cover $\mathcal{C}_{\mathbb{T},\rho}(u)$ would not be larger than than the smallest cover $\mathcal{C}_{\mathbb{T},\|\cdot\|_2}(u^2/\alpha)$, or equivalently:

$$\mathcal{N}_{\mathbb{T},\rho}(u) \leq \mathcal{N}_{\mathbb{T},\|\cdot\|_2}(\frac{u^2}{\alpha}) \stackrel{(i)}{\leq} (1 + 2R\alpha/u^2)^d,$$

where inequality (i) follows from Wainwright (2019, Lem 5.7) using the fact that $\mathbb{T} = \mathcal{B}_2(0; R)$. Noting that $\beta_R = 2R\alpha$ yields the bound (75).

We now use the bound (75) on the covering number to establish the bound (76).

Proof of bound on $\mathcal{J}_{\mathbb{T},\rho}$ Applying the definition (72) and the bound (75), we find that

$$\begin{aligned} \mathcal{J}_{\mathbb{T},\rho}(D_R) &\leq \int_0^{D_R} \sqrt{\log(1 + (1 + \beta_R/u^2)^d)} du \\ &\stackrel{(i)}{=} \sqrt{\frac{\beta_R}{2}} \int_{\beta_R/D_R^2}^\infty s^{-3/2} \sqrt{\log(1 + (1 + s)^d)} ds, \end{aligned} \quad (79)$$

where step (i) follows from a change of variable $s \leftarrow \beta_R/u^2$. Note that $\log(1 + a) \leq \log a + 1/a$, and $\sqrt{a + b} \leq \sqrt{a} + \sqrt{b}$ for any $a, b \geq 0$. Applying these inequalities, we find that

$$\log(1 + (1 + s)^d) \leq d \log(2 + s) \leq d(\log s + \frac{1}{s} + \frac{1}{1+s}) \leq d(\log s + \frac{2}{s}),$$

for $s \in [\beta_R/D_R^2, \infty) \stackrel{(77)}{\subseteq} [2, \infty)$. Consequently,

$$\int_{\beta_R/D_R^2}^\infty s^{-3/2} \sqrt{\log(1 + (1 + s)^d)} ds \leq \sqrt{d} \int_{\beta_R/D_R^2}^\infty \left(\frac{\sqrt{\log s}}{s^{3/2}} + \frac{\sqrt{2}}{s^2} \right) ds. \quad (80)$$

Next, we note that

$$\int_{\beta_R/D_R^2}^\infty \frac{\sqrt{2}}{s^2} ds = \sqrt{2} \frac{D_R^2}{\beta_R}, \quad \text{and} \quad (81)$$

$$\begin{aligned} \int_{\beta_R/D_R^2}^\infty \frac{\sqrt{\log s}}{s^{3/2}} ds &= -2 \sqrt{\frac{\log s}{s}} - \sqrt{2} \Gamma\left(\frac{1}{2}, \frac{1}{2} \log s\right) \Big|_{s=\beta_R/D_R^2}^{s=\infty} \\ &\stackrel{(i)}{\leq} \frac{2D_R}{\sqrt{\beta_R}} \left(\sqrt{\log(\beta_R/D_R^2)} + \frac{1}{\sqrt{\log(\beta_R/D_R^2)}} \right), \end{aligned} \quad (82)$$

where step (i) follows from the following bound on the incomplete Gamma function:

$$\Gamma\left(\frac{1}{2}, a\right) = \int_a^\infty \frac{1}{\sqrt{t}} e^{-t} dt \leq \frac{1}{\sqrt{a}} \int_a^\infty e^{-t} dt = a^{-\frac{1}{2}} e^{-a} \quad \text{for any } a > 0.$$

Putting the bounds (79) to (82) together, we find that

$$\mathcal{J}_{\mathbb{T},\rho}(D_R) \leq \sqrt{d}D_R \left(\frac{D_R}{\sqrt{\beta_R}} + \sqrt{2} \left(\sqrt{\log(\beta_R/D_R^2)} + 1/\sqrt{\log(\beta_R/D_R^2)} \right) \right).$$

Note that by definition (74) $\frac{D_R}{\sqrt{\beta_R}} \leq \frac{1}{\sqrt{2}}$. Using this observation and the fact that $\frac{1}{\sqrt{2}} + \sqrt{\frac{2}{\log 2}} \leq 3$ yields the claimed bound (76).

Appendix L. Proof of Rem. 8: Example settings for KH L^∞ rates

Note that when the kernel is restricted to a compact domain, we have $R_{\mathbf{k},n} \leq \mathcal{O}(\sigma)$ so that in this case, $\frac{L_{\mathbf{k}}(R_{\mathbf{k},n} \wedge R_{\mathcal{S}_n})}{\|\mathbf{k}\|_\infty} \leq \mathcal{O}(L)$. Next, we establish the bounds on L and $\kappa^\dagger(1/n)$ for each of the kernels, which immediately imply the respective claims both when supported on \mathbb{R}^d and on a compact domain with $\sigma = \sqrt{d}$.

- For a Gaussian kernel, we have $\kappa(r) = e^{-r^2/2}$, which when put together with (98) and (99) implies that $L = \sqrt{2/e}$, and $\kappa^\dagger(1/n) = \sqrt{2 \log n}$.
- For a Matérn kernel, we have $\kappa(r) = c_b r^b K_b(r)$ for $b = \nu - \frac{d}{2}$, for suitable constant c_b , and function K_b ; see Tab. 1. Moreover, (107) and (108) imply that $L = \frac{C_1}{\sqrt{C_2+b}} \leq \frac{C_1}{\sqrt{C_2+1}}$ for universal constants C_1, C_2 , and $\kappa^\dagger(1/n) \leq \mathcal{O}(\max(b \log b, \log n)) = \mathcal{O}(\log n)$.
- For a IMQ kernels, we have $\kappa(r) = \frac{1}{(\gamma^2+r^2)^\nu}$, so that

$$L \leq \sup_r \kappa'(r) = \sup_r \frac{2\nu r}{(\gamma^2+r^2)^{\nu+1}} = \frac{2\nu\gamma}{\gamma^{2(\nu+1)}} \frac{1}{\sqrt{1+2\nu}} \left(\frac{2\nu+1}{2\nu+2} \right)^{\nu+1},$$

$$\text{and } \kappa^\dagger(1/n) = \sqrt{n^{1/\nu} - \gamma^2} \leq n^{1/2\nu}.$$

Appendix M. Proof of Cor. 5: L^∞ and MMD guarantees for KT-SPLIT

First applying Thm. 4(b) with $\mathbf{k} = \mathbf{k}_{\text{rt}}$ yields that

$$\|\mathbb{P}_n \mathbf{k}_{\text{rt}} - \mathbb{Q}_{\text{KT}}^{(m,1)} \mathbf{k}_{\text{rt}}\|_\infty \leq \frac{\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, m, d, \delta^*, \delta', R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n}) \quad (83)$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \delta_i$. Call this event \mathcal{E} .

Let $r' = R'_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n_{\text{out}}} + \varepsilon$ denote a shorthand notation, where $\varepsilon > 0$ is an arbitrary scalar. Applying Thm. 2 with $r = 2r'$ and $a = b = \frac{1}{2}$, we find that on event \mathcal{E} ,

$$\begin{aligned}
 & \text{MMD}_{\mathbf{k}_\star}(\mathcal{S}_n, \mathcal{S}^{(m,1)}) \\
 &= \text{MMD}_{\mathbf{k}_\star}(\mathbb{P}_n, \mathbb{Q}_{\text{KT}}^{(m,1)}) \\
 &\leq \|\mathbb{P}_n \mathbf{k}_{\text{rt}} - \mathbb{Q}_{\text{KT}}^{(m,1)} \mathbf{k}_{\text{rt}}\|_\infty e_d r^{d/2} + 2\tau_{\mathbf{k}_{\text{rt}}}(r') + 2\|\mathbf{k}_\star\|_\infty^{\frac{1}{2}} \cdot \max(\tau_{\mathbb{P}_n}(r'), \tau_{\mathbb{Q}_{\text{KT}}^{(m,1)}}(r')) \\
 &\stackrel{(i)}{=} \|\mathbb{P}_n \mathbf{k}_{\text{rt}} - \mathbb{Q}_{\text{KT}}^{(m,1)} \mathbf{k}_{\text{rt}}\|_\infty \cdot \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (r')^{\frac{d}{2}} + \frac{2\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}}, \\
 &\stackrel{(83)}{\leq} \frac{\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}} \cdot \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, m, d, \delta^\star, \delta', R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n}) \cdot \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (r')^{\frac{d}{2}} + \frac{2\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}} \\
 &= \frac{\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}} \left[2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (r')^{\frac{d}{2}} \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, m, d, \delta^\star, \delta', R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n}) \right] \tag{84}
 \end{aligned}$$

where step (i) uses the following arguments: (a) $e_d r^{d/2} = \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (r')^{\frac{d}{2}}$; (b) $\mathbb{Q}_{\text{KT}}^{(m,1)}$ is supported on a subset of points from \mathcal{S}_n and hence

$$\max\{\tau_{\mathbb{P}_n}(r'), \tau_{\mathbb{Q}_{\text{KT}}^{(m,1)}}(r')\} = \tau_{\mathbb{P}_n}(r') = 0 \quad \text{for any } r' > R_{\mathcal{S}_n};$$

and (c) $\tau_{\mathbf{k}_{\text{rt}}}(r') \leq \frac{\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}}$ for any $r' > R'_{\mathbf{k}_{\text{rt}}, n_{\text{out}}}$. Noting that (84) applies simultaneously for all $\varepsilon > 0$ under event \mathcal{E} , taking the limit $\varepsilon \rightarrow 0$ yields that

$$\text{MMD}_{\mathbf{k}_\star}(\mathcal{S}_n, \mathcal{S}^{(m,1)}) \leq \frac{\|\mathbf{k}_{\text{rt}}\|_\infty}{n_{\text{out}}} \left[2 + \sqrt{\frac{(4\pi)^{d/2}}{\Gamma(\frac{d}{2}+1)}} (R'_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n_{\text{out}}})^{\frac{d}{2}} \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, m, d, \delta^\star, \delta', R_{\mathcal{S}_n, \mathbf{k}_{\text{rt}}, n}) \right] \tag{85}$$

with probability at least $1 - \delta' - \sum_{j=1}^m \frac{2^{j-1}}{m} \sum_{i=1}^{2^{m-j} \lfloor n/2^m \rfloor} \delta_i$, as claimed.

Appendix N. Derivation of Tab. 1: Square-root kernels \mathbf{k}_{rt} for common target kernels \mathbf{k}_\star

In this appendix, we derive the results stated in Tab. 1.

N.1 General proof strategy

Let \mathcal{F} denote the Fourier operator (19). In Tab. 3, we state the continuous κ such that $\mathbf{k}_\star(x, y) = \kappa(x - y)$ in the first column, its Fourier transform (19) $\widehat{\kappa}$ in the second column, the square-root Fourier transform in the third column, and the square-root kernel in the fourth column, given by $\mathbf{k}_{\text{rt}}(x, y) = \frac{1}{(2\pi)^{d/4}} \kappa_{\text{rt}}(x - y)$ with $\kappa_{\text{rt}} = \mathcal{F}(\sqrt{\widehat{\kappa}})$. Prop. 3 along with expressions in Tab. 3 directly establishes the validity of the square-root kernels for the Gaussian and (scaled) B-spline kernels. For completeness, we also illustrate the remaining calculus for the B-spline kernels in App. N.3. We do a similar calculation in App. N.2 for the Matérn kernel for better exposition of the involved expressions.

Expression for $\kappa(z)$	Expression for Fourier transform of $\kappa: \widehat{\kappa}(\omega)$	Square-root Fourier transform: $\sqrt{\widehat{\kappa}(\omega)}$	Expression for $\kappa_{\text{rt}}(z) \triangleq (\frac{1}{2\pi})^{\frac{d}{4}} \mathcal{F}(\sqrt{\widehat{\kappa}})(z)$
$\exp\left(-\frac{\ z\ _2^2}{2\sigma^2}\right)$	$\sigma^d \exp\left(-\frac{\sigma^2\ \omega\ _2^2}{2}\right)$	$\sigma^{\frac{d}{2}} \exp\left(-\frac{\sigma^2\ \omega\ _2^2}{4}\right)$	$\left(\frac{2}{\pi\sigma^2}\right)^{\frac{d}{4}} \exp\left(-\frac{\ z\ _2^2}{\sigma^2}\right)$
$\prod_{j=1}^d \otimes^{2\beta+2} \mathbb{1}_{[-\frac{1}{2}, \frac{1}{2}]}(z_j)$	$\left(\frac{4^{\beta+1}}{\sqrt{2\pi}}\right)^d \prod_{j=1}^d \frac{\sin^{2\beta+2}(\omega_j/2)}{\omega_j^{2\beta+2}}$	$\left(\frac{4^{\beta+1}}{\sqrt{2\pi}}\right)^{\frac{d}{2}} \prod_{j=1}^d \frac{\sin^{\beta+1}(\omega_j/2)}{\omega_j^{\beta+1}}$	$\prod_{j=1}^d \otimes^{\beta+1} \mathbb{1}_{[-\frac{1}{2}, \frac{1}{2}]}(z_j)$

Table 3: Fourier transforms of kernels $\mathbf{k}_*(x, y) = \kappa(x - y)$ and square-root kernels $\mathbf{k}_{\text{rt}}(x, y) = \kappa_{\text{rt}}(x - y)$ from Tab. 1. Here \mathcal{F} denotes the Fourier operator (19), and \otimes^ℓ denotes the convolution operator applied $\ell - 1$ times, with the convention $\otimes^1 f \triangleq f$ and $\otimes^2 f = f \otimes f$ for $(f \otimes g)(x) = \int f(y)g(x - y)dy$. Each Fourier transform is derived from Sriperumbudur et al. (2010, Tab. 2).

N.2 Deriving \mathbf{k}_{rt} for the Matérn kernel

For $\mathbf{k}_* = \text{Matérn}(\nu, \gamma)$ from Tab. 1, we have

$$\mathbf{k}_*(x, y) = \phi_{d, \nu, \gamma} \cdot \Phi_{\nu, \gamma}(x - y) \quad \text{where} \quad \Phi_{\nu, \gamma}(z) \triangleq c_\nu \left(\frac{\|z\|_2}{\gamma}\right)^{\nu - \frac{d}{2}} K_{\nu - \frac{d}{2}}(\gamma \|z\|_2), \quad (86)$$

$$\phi_{d, \nu, \gamma} \triangleq \frac{c_{\nu - d/2} \gamma^{2\nu - d}}{c_\nu} \quad \text{and} \quad c_\nu \triangleq \frac{2^{1-\nu}}{\Gamma(\nu)},$$

and K_a denotes the modified Bessel function of third kind of order a (Wendland, 2004, Def. 5.10). That is, for Matérn kernel with $\mathbf{k}_*(x, y) = \kappa(x - y)$, the function κ is given by $\kappa = \phi_{d, \nu, \gamma} \Phi_{\nu, \gamma}$. Now applying (Wendland, 2004, Thm. is 8.15), we find that

$$\mathcal{F}(\Phi_{\nu, \gamma}) = \frac{1}{(\gamma^2 + \|\omega\|_2^2)^\nu} \implies \mathcal{F}(\sqrt{\mathcal{F}(\Phi_{\nu, \gamma})}) = \frac{1}{(2\pi)^{d/4}} \Phi_{\nu/2, \gamma},$$

where in the last step we also use the facts that $\Phi_{\nu, \gamma}$ is an even function and $\sqrt{\mathcal{F}(\Phi_{\nu, \gamma})} \in L^1$ for all $\nu > d/2$. Thus, by Prop. 3, a valid square-root kernel $\mathbf{k}_{\text{rt}}(x, y) = \kappa_{\text{rt}}(x - y)$ is defined via

$$\kappa_{\text{rt}} = \sqrt{\phi_{d, \nu, \gamma} \frac{1}{(2\pi)^{d/4}} \Phi_{\nu/2, \gamma}}, \implies \mathbf{k}_{\text{rt}} = A_{\nu, \gamma, d} \cdot \text{Matérn}(\nu/2, \gamma), \quad (87)$$

$$\text{where} \quad A_{\nu, \gamma, d} \triangleq \left(\frac{1}{4\pi} \gamma^2\right)^{d/4} \sqrt{\frac{\Gamma(\nu)}{\Gamma(\nu - d/2)}} \cdot \frac{\Gamma((\nu - d)/2)}{\Gamma(\nu/2)}. \quad (88)$$

N.3 Deriving \mathbf{k}_{rt} for the B-Spline kernel

For positive integers β, d , define the constants

$$\mathfrak{B}_\beta \triangleq \frac{1}{(\beta-1)!} \sum_{j=0}^{\lfloor \beta/2 \rfloor} (-1)^j \binom{\beta}{j} \left(\frac{\beta}{2} - j\right)^{\beta-1}, \quad S_{\beta, d} \triangleq \mathfrak{B}_\beta^{-d}, \quad \text{and} \quad \tilde{S}_{\beta, d} \triangleq \frac{\sqrt{S_{2\beta+2, d}}}{S_{\beta+1, d}}. \quad (89)$$

Define the function $\chi_\beta : \mathbb{R}^d \rightarrow \mathbb{R}$ as follows:

$$\chi_\beta(z) \triangleq S_{\beta, d} \prod_{i=1}^d f_\beta(z_i). \quad (90)$$

Then for kernel $\mathbf{k}_\star = \mathbf{B}\text{-spline}(2\beta + 1)$, we have

$$\mathbf{k}_\star(x, y) = \chi_{2\beta+2}(x - y). \quad (91)$$

The second row of Tab. 3 implies that

$$\left(\frac{1}{2\pi}\right)^{\frac{d}{4}} \mathcal{F}\left(\sqrt{\mathcal{F}(\chi_{2\beta+2})}\right) = \frac{\sqrt{S_{2\beta+2,d}}}{S_{\beta+1,d}} \cdot \chi_{\beta+1} \stackrel{(89)}{=} \tilde{S}_{\beta,d} \cdot \chi_{\beta+1},$$

where we also use the fact that χ_β is an even function. Putting the pieces together with Prop. 3, we conclude that

$$\mathbf{k}_{\text{rt}} = \tilde{S}_{\beta,d} \cdot \mathbf{B}\text{-spline}(\beta), \quad (92)$$

(with $\kappa_{\text{rt}} = \tilde{S}_{\beta,d}\chi_{\beta+1}$) is a valid square-root kernel for $\mathbf{k}_\star = \mathbf{B}\text{-spline}(2\beta + 1)$.

Appendix O. Derivation of Tab. 2: Explicit bounds on Thm. 1 quantities for common kernels

We start by collecting some common tools in App. O.1 that we later use for our derivations for the Gaussian, Matérn and B-spline kernels in Apps. O.2 to O.4 respectively. Finally, in App. O.5, we put the pieces together to derive explicit MMD rates, as a function of d, n, δ and kernel parameters for the three kernels, and summarize the rates in Tab. 4. (This table is complementary to the generic results stated in Cor. 1.)

O.1 Common tools for our derivations

We collect some simplified expressions, and techniques that come handy in our derivations to follow.

Simplified bounds on $\mathfrak{M}_{\mathbf{k}_{\text{rt}}}$ From (8), we have

$$\begin{aligned} \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, \tfrac{1}{2} \log_2 n, d, \tfrac{\delta}{n}, \delta', R) &\lesssim \sqrt{\log \frac{n}{\delta} \cdot \left[\log\left(\frac{1}{\delta'}\right) + d \log\left(\frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty} \cdot (R_{\mathbf{k}_{\text{rt}},n} + R)\right) \right]} \\ \implies \mathfrak{M}_{\mathbf{k}_{\text{rt}}}(n, \tfrac{1}{2} \log_2 n, d, \tfrac{\delta}{n}, \delta, R) &\lesssim_\delta \sqrt{d \log n \cdot \log\left(\frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty} \cdot (R_{\mathbf{k}_{\text{rt}},n} + R)\right)}. \end{aligned} \quad (93)$$

Thus, given (93), to get a bound on $\mathfrak{M}_{\mathbf{k}_{\text{rt}}}$, we need to derive bounds on $(L_{\mathbf{k}_{\text{rt}}}, \|\mathbf{k}_{\text{rt}}\|_\infty, R_{\mathbf{k}_{\text{rt}},n})$ for various kernels to obtain the desired guarantees on MMD and L^∞ coresets.

Bounds on Gamma function Our proofs make use of the bounds from Batir (2017, Thm 2.2) on the Gamma function:

$$\Gamma(b+1) \geq (b/e)^b \sqrt{2\pi b} \text{ for any } b \geq 1, \text{ and } \Gamma(b+1) \leq (b/e)^b \sqrt{e^2 b} \text{ for any } b \geq 1.1. \quad (94)$$

General tools for bounding the Lipschitz constant To bound the Lipschitz constant $L_{\mathbf{k}_{\text{rt}}}$, the following two observations come in handy. For a radial kernel $\mathbf{k}_{\text{rt}}(x, y) = \tilde{\kappa}_{\text{rt}}(\|x - y\|_2)$ with $\tilde{\kappa}_{\text{rt}} : \mathbb{R}_+ \rightarrow \mathbb{R}$, we note

$$\begin{aligned} \sup_{\substack{x,y,z: \\ \|y-z\|_2 \leq r}} |\mathbf{k}_{\text{rt}}(x, y) - \mathbf{k}_{\text{rt}}(x, z)| &\leq \sup_{a>0, b \leq r} |\tilde{\kappa}_{\text{rt}}(a) - \tilde{\kappa}_{\text{rt}}(a+b)| \leq \|\tilde{\kappa}'_{\text{rt}}\|_\infty r \\ \implies L_{\mathbf{k}_{\text{rt}}} &\leq \|\tilde{\kappa}'_{\text{rt}}\|_\infty. \end{aligned} \quad (95)$$

For a translation-invariant kernel $\mathbf{k}_{\text{rt}}(x, y) = \kappa_{\text{rt}}(x - y)$ with $\kappa_{\text{rt}} : \mathbb{R}^d \rightarrow \mathbb{R}$, we use the bound

$$\begin{aligned} \sup_{\substack{x, y, z: \\ \|y-z\|_2 \leq r}} |\mathbf{k}_{\text{rt}}(x, y) - \mathbf{k}_{\text{rt}}(x, z)| &= \sup_{z', z'', \|z-z'\|_2 \leq r} |\kappa_{\text{rt}}(z') - \kappa_{\text{rt}}(z'')| \\ &\leq \sup_{z' \in \mathbb{R}^d} \|\nabla \kappa_{\text{rt}}(z')\|_2 r, \\ \implies L_{\mathbf{k}_{\text{rt}}} &\leq \sup_{z' \in \mathbb{R}^d} \|\nabla \kappa_{\text{rt}}(z')\|_2 \stackrel{(i)}{\leq} \sqrt{d} \sup_{z' \in \mathbb{R}^d} \|\nabla \kappa_{\text{rt}}(z')\|_\infty, \end{aligned} \quad (96)$$

where step (i) follows from Cauchy-Schwarz's inequality (and is handy when coordinate wise control on $\nabla \kappa_{\text{rt}}$ is easier to derive). We later apply the inequality (95) for the Gaussian and Matérn kernels, and (96) for the B-spline kernel.

O.2 Proofs for Gaussian kernel

For the kernel $\mathbf{k}_\star = \mathbf{Gauss}(\sigma)$ and its square-root kernel $\mathbf{k}_{\text{rt}} = (\frac{2}{\pi\sigma^2})^{\frac{d}{4}} \mathbf{Gauss}(\frac{\sigma}{\sqrt{2}})$, we claim the following bounds on various quantities

$$\|\mathbf{k}_{\text{rt}}\|_\infty = (\frac{2}{\pi\sigma^2})^{\frac{d}{4}} \quad \text{and} \quad \|\mathbf{k}_\star\|_\infty = 1, \quad (97)$$

$$L_{\mathbf{k}_{\text{rt}}} \leq (\frac{2}{\pi\sigma^2})^{\frac{d}{4}} \frac{\sqrt{2/e}}{\sigma} \implies L_{\mathbf{k}_{\text{rt}}}/\|\mathbf{k}_{\text{rt}}\|_\infty \stackrel{(97)}{=} \frac{1}{\sigma} \sqrt{2/e} \quad (98)$$

$$R_{\mathbf{k}_{\text{rt}}, n} = \sigma \sqrt{\log n}, \quad (99)$$

$$R'_{\mathbf{k}_{\text{rt}}, \sqrt{n}} = \mathcal{O}(\sigma \sqrt{d + \log n}). \quad (100)$$

Substituting these expressions in (93), we find that

$$\mathfrak{M}_{\mathbf{k}_{\text{rt}}}^{\mathbf{Gauss}}(n, \frac{1}{2} \log_2 n, d, \frac{\delta}{2n}, \delta', R) \lesssim \sqrt{\log(\frac{n}{\delta}) [\log \frac{1}{\delta'} + d \log(\sqrt{\log n} + \frac{R}{\sigma})]},$$

as claimed in Tab. 2.

Proof of claims (97) to (100) The claim (97) follows directly from the definition of \mathbf{k}_{rt} . The bound (98) on $L_{\mathbf{k}_{\text{rt}}}$ follows from the fact $\|\frac{d}{dr} e^{-r^2/\sigma^2}\|_\infty = \frac{\sqrt{2/e}}{\sigma}$ and invoking the relation (95). Next, recalling the definition (6) and noting that

$$\sup_{\|x-y\|_2 \geq r} |\mathbf{k}_{\text{rt}}(x, y)| \leq (\frac{2}{\pi\sigma^2})^{\frac{d}{4}} e^{-r^2/\sigma^2}$$

implies the bound (99) on $R_{\mathbf{k}_{\text{rt}}, n}$.

Next, we have

$$\begin{aligned} \tau_{\mathbf{k}_{\text{rt}}}^2(R) &= \int_{\|z\|_2 \geq R} (\frac{2}{\pi\sigma^2})^{\frac{d}{2}} \exp(-2\|z\|_2^2/\sigma^2) dz \\ &= \Pr_{X \sim \mathcal{N}(0, \sigma^2/4 \cdot \mathbf{I}_d)} (\|X\|_2 \geq R) \stackrel{(i)}{\leq} e^{-R^2/\sigma^2}, \end{aligned} \quad (101)$$

for $R \geq \sigma\sqrt{2d}$, where step (i) follows from the standard tail bound for a chi-squared random variable Y with k degree of freedom (Laurent and Massart, 2000, Lem. 1): $\Pr(Y - k > 2\sqrt{kt} + 2t) \leq e^{-t}$, wherein we substitute $k = d$, $Y = \frac{4}{\sigma^2} \|X\|_2^2$, $t = d\alpha$ with $\alpha \geq 2$, and $R^2 = \sigma^2 t$. Using the tail bound (101), the bound (97) on $\|\mathbf{k}_{\text{rt}}\|_\infty$ and the definition (6) of $R'_{\mathbf{k}_{\text{rt}}, \sqrt{n}}$, we find that

$$R'_{\mathbf{k}_{\text{rt}}, \sqrt{n}} = \mathcal{O}(\sigma \max(\sqrt{(\log n - \frac{d}{2} \log(\sigma \sqrt{\frac{\pi}{2}}))_+}, \sqrt{d})) = \mathcal{O}(\sigma \sqrt{d + \log n}).$$

yielding the claim (100).

O.3 Proofs for Matérn kernel

Recall the notations from (86) to (88) for the Matérn kernel related quantities. Let K_b denote the modified Bessel function of the third kind with order b (Wendland, 2004, Def. 5.10). Let $a \triangleq \frac{\nu-d}{2}$. Then, the kernel $\mathbf{k}_\star = \mathbf{Matérn}(\nu, \gamma)$ and its square-root kernel $\mathbf{k}_{\text{rt}} = A_{\nu, \gamma, d} \cdot \mathbf{Matérn}(a, \gamma)$ satisfy

$$\mathbf{k}_\star(x, y) = \tilde{\kappa}_{\nu-d/2}(\gamma \|x - y\|_2), \quad \text{and} \quad (102)$$

$$\mathbf{k}_{\text{rt}}(x, y) = A_{\nu, \gamma, d} \tilde{\kappa}_a(\gamma \|x - y\|_2), \quad (103)$$

$$\text{where } \tilde{\kappa}_b(r) \triangleq c_b r^b K_b(r), \text{ and } c_b = \frac{2^{1-b}}{\Gamma(b)} \text{ for } b > 0. \quad (104)$$

We claim the following bounds on various quantities assuming $a \geq 2.2$, and $d \geq 2$:⁷

$$\|\mathbf{k}_{\text{rt}}\|_\infty = A_{\nu, \gamma, d} \quad \text{and} \quad \|\mathbf{k}_\star\|_\infty = 1, \quad (105)$$

$$A_{\nu, \gamma, d} \leq 5\nu \left(\frac{\gamma^2}{2\pi(a-1)} \right)^{\frac{d}{4}}, \quad (106)$$

$$L_{\mathbf{k}_{\text{rt}}} \leq A_{\nu, \gamma, d} \frac{C_1 \gamma}{\sqrt{a+C_2}} \quad \Longrightarrow \quad \frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty} \stackrel{(105)}{\leq} \frac{C_1 \gamma}{\sqrt{a+C_2}}, \quad (107)$$

$$R_{\mathbf{k}_{\text{rt}}, n} = \frac{1}{\gamma} \mathcal{O}(\max(\log n - a \log(1+a), C_2 a \log(1+a))), \quad (108)$$

$$R'_{\mathbf{k}_{\text{rt}}, \sqrt{n}} = \frac{1}{\gamma} \cdot \mathcal{O}(a + \log n + d \log(\frac{\sqrt{2e\pi}}{\gamma}) + \log(\frac{(\nu-2)^{\nu-\frac{3}{2}}}{(2(a-1))^{2a-1} d^{\frac{d}{2}+1}})) \quad (109)$$

We prove these claims in App. O.3.1 through App. O.3.5. Putting these bounds together with (93) yields that

$$\begin{aligned} \mathfrak{M}_{\mathbf{k}_{\text{rt}}}^{\mathbf{Matérn}}(n, \frac{1}{2} \log_2 n, d, \frac{\delta}{2n}, \delta', R) &\lesssim \begin{cases} \sqrt{\log(\frac{n}{\delta}) \left[\log \frac{1}{\delta} + d \log \left(\frac{1}{\sqrt{1+a}} \cdot (\log n + \gamma R) \right) \right]} & \text{if } a = o(\log n) \\ \sqrt{\log(\frac{n}{\delta}) \left[\log \frac{1}{\delta'} + d \log(\sqrt{a} \log(1+a) + \gamma R) \right]} & \text{if } a = \Omega(\log n) \end{cases} \\ &\lesssim \sqrt{\log(\frac{n}{\delta}) \left[\log \frac{1}{\delta'} + d \log(\log n + B + \gamma R) \right]} \end{aligned}$$

with $B = a \log(1+a)$, as claimed in Tab. 2.

O.3.1 SET-UP FOR PROOFS OF MATÉRN KERNEL QUANTITIES

Before proceeding to the proofs of the claims (105) to (109), we collect some handy inequalities. Applying Wendland (2004, Lem. 5.13, 5.14), we have

$$\tilde{\kappa}_a(\gamma r) \leq \min \left(1, \sqrt{2\pi} c_a (\gamma r)^{a-\frac{1}{2}} e^{-\gamma r + \frac{a^2}{2\gamma r}} \right) \quad \text{for } r > 0. \quad (110)$$

For a large enough r , we also establish the following bound:

$$\tilde{\kappa}_a(\gamma r) \leq \min(1, 4c_a (\gamma r)^{a-1} e^{-\gamma r/2}) \quad \text{for } \gamma r \geq 2(a-1), a \geq 1. \quad (111)$$

7. When $a \in (1, 2.2)$, and $d \in [1, 2)$, analogous bounds with slightly different constants follow from employing the Gamma function upper bound of Batir (2017, Thm 2.3) in place of our upper bound (94). For brevity, we omit these derivations.

Proof of (111) Noting the definition (104) of κ_a , it suffices to show the following bound:

$$|K_a(r)| \leq \frac{4}{r} e^{-r/2} \quad \text{for } r/2 \geq a - 1,$$

where K_a is the modified Bessel function of the third kind. Using Wendland (2004, Def. 5.10), we have

$$K_a(r) = \frac{1}{2} \int_0^\infty e^{-r \cosh t} \cosh(at) dt \stackrel{(i)}{\leq} \int_0^\infty e^{-\frac{r}{2} e^t} e^{at} dt \stackrel{(ii)}{=} \int_{r/2}^\infty e^{-s} \left(\frac{2s}{r}\right)^a \cdot \frac{1}{s} ds \stackrel{(iii)}{\leq} \frac{4}{r} e^{-r/2}$$

where step (i) uses the following inequalities: $\cosh t \triangleq \frac{1}{2}(e^t + e^{-t}) \geq \frac{1}{2}e^t$ and $\cosh(at) \leq e^{at}$ for $a > 0, t > 0$, step (ii) follows from a change of variable $s \leftarrow \frac{r}{2}e^t$, and finally step (iii) uses the following bound on the incomplete Gamma function obtained by substituting $B = 2$ and $A = a$ in Borwein and Chan (2009, Eq. 1.5):

$$\int_r^\infty t^{a-1} e^{-t} dt \leq 2r^{a-1} e^{-r} \quad \text{for } r \geq a - 1.$$

The proof is now complete.

O.3.2 PROOF OF THE BOUND (105) ON $\|\mathbf{k}_{\text{rt}}\|_\infty$ AND $\|\mathbf{k}_\star\|_\infty$

We claim that

$$\|\tilde{\kappa}_b\|_\infty = 1 \quad \text{for all } b > 0. \tag{112}$$

where was defined in (104). Putting the equality (112) with (102), (103), and (110) immediately implies the bounds (105) on the $\|\mathbf{k}_{\text{rt}}\|_\infty$ and $\|\mathbf{k}_\star\|_\infty$. To prove (112), we follow the steps from the proof of Wendland (2004, Lem. 5.14). Using Wendland (2004, Def. 5.10), for $b > 0$ we have

$$K_b(r) = \frac{1}{2} \int_{-\infty}^\infty e^{-r \cosh t} e^{bt} dt = \frac{1}{2} \int_{-\infty}^\infty e^{-\frac{r}{2}(e^t + e^{-t})} e^{bt} dt = k^{-b} \frac{1}{2} \int_0^\infty e^{-r/2(u/k + k/u)} u^{b-1} du$$

where the last step follows by substituting $u = ke^t$. Setting $k = r/2$, we find that

$$r^b K_b(r) = 2^{b-1} \int_0^\infty e^{-u} e^{-r^2/(4u)} u^{b-1} du \leq 2^{b-1} \Gamma(b),$$

where we achieve equality in the last step when we take the limit $r \rightarrow 0$. Noting that $\tilde{\kappa}_b(r) = \frac{2^{1-b}}{\Gamma(b)} r^b K_b(r)$ (104) yields the claim (112).

O.3.3 PROOF OF THE BOUND (106) ON $A_{\nu,\gamma,d}$

Using the definition (88) we have $A_{\nu,\gamma,d} = \left(\frac{1}{4\pi}\gamma^2\right)^{\frac{d}{4}} \cdot A'_{\nu,\gamma,d}$, where

$$\begin{aligned} A'_{\nu,\gamma,d} &= \sqrt{\frac{\Gamma(\nu)}{\Gamma(\nu-d/2)}} \cdot \frac{\Gamma((\nu-d)/2)}{\Gamma(\nu/2)} \stackrel{(94)}{\leq} \sqrt{\frac{e}{\sqrt{2\pi}} \frac{(\frac{\nu-1}{e})^{\nu-1} \sqrt{\nu-1}}{(\frac{\nu-d}{e})^{\nu-d-1} \sqrt{\nu-d-1}}} \cdot \frac{e}{\sqrt{2\pi}} \frac{(\frac{\nu-d-2}{2e})^{\frac{\nu-d-2}{2}} \sqrt{\frac{\nu-d-2}{2}}}{(\frac{\nu-2}{2e})^{\frac{\nu-2}{2}} \sqrt{\frac{\nu-2}{2}}} \\ &\leq \left(\frac{e^2}{2\pi}\right)^{\frac{3}{4}} (2\sqrt{e})^{\frac{d}{2}} \cdot \left(\frac{\nu-1}{\nu-2}\right)^{\frac{\nu-2}{2}} \cdot \sqrt{\frac{\nu-1}{\nu-2}} \cdot (\nu-1)^{\frac{1}{4}} \cdot \left(\frac{\nu-d-2}{\nu-d-1}\right)^{\frac{\nu-d-2}{2}} \cdot \frac{\sqrt{(\nu-d-2)/\sqrt{(\nu-d-1)}}}{(\nu-d-2)^{\frac{1}{2}(\frac{d}{2}-1)}} \\ &\leq \left(\frac{e^2}{2\pi}\right)^{\frac{3}{4}} (2\sqrt{e})^{\frac{d}{2}} \cdot \sqrt{e} \cdot \sqrt{2} \cdot (\nu-1)^{\frac{1}{4}} \cdot (\sqrt{e})^{-\frac{d}{2}+1} \cdot \frac{(\nu-d-2)/(\nu-d-1)^{\frac{1}{4}}}{(\nu-d-2)^{\frac{d}{4}}} \\ &\leq \frac{e^2 \sqrt{2e}}{(2\pi)^{\frac{3}{4}}} (\nu-1)^{\frac{1}{4}} \cdot (\nu-d-1)^{\frac{3}{4}} \left(\frac{\nu-d-2}{4}\right)^{-\frac{d}{4}}. \end{aligned}$$

As a result, we have

$$A_{\nu,\gamma,d} = \left(\frac{1}{4\pi}\gamma^2\right)^{d/4} \cdot A'_{\nu,\gamma,d} \leq 5\nu\left(\frac{\gamma^2}{\pi(\nu-d-2)}\right)^{d/4} = 5\nu\left(\frac{\gamma^2}{2\pi(a-1)}\right)^{d/4},$$

as claimed.

O.3.4 PROOF OF THE BOUND (107) ON $L_{\mathbf{k}_{rt}}$

To derive a bound on the Lipschitz constant $L_{\mathbf{k}_{rt}}$, we bound the derivative of $\tilde{\kappa}_a$. Using $(r^a K_a(r))' = -r^a K_{a-1}(r)$ (Wendland, 2004, Eq. 5.2), we find that

$$(\tilde{\kappa}_a(\gamma r))' = c_a \frac{d}{dr} ((\gamma r)^a K_a(\gamma r)) = -\frac{c_a}{c_{a-1}} \gamma^2 r \tilde{\kappa}_{a-1}(\gamma r).$$

Using (110), we have

$$\frac{c_a}{c_{a-1}} \gamma^2 r \tilde{\kappa}_{a-1}(\gamma r) \leq \frac{c_a \gamma}{c_{a-1}} \min(\gamma r, \sqrt{2\pi} c_{a-1} (\gamma r)^{a-\frac{1}{2}} e^{-\gamma r + \frac{(a-1)^2}{2\gamma r}}) \quad \text{for } r > 0.$$

And (111) implies that

$$\frac{c_a}{c_{a-1}} \gamma^2 r \tilde{\kappa}_{a-1}(\gamma r) \leq \frac{c_a \gamma}{c_{a-1}} \min(\gamma r, 4c_{a-1} (\gamma r)^{a-1} e^{-\gamma r/2}) \quad \text{for } r \geq 2(a-2), a \geq 2.$$

Next, we make use of the following observations: The function $t^b e^{-t/2}$ achieves its maximum at $t = 2b$. Then for any function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $f(t) \leq t$ for all $t \geq 0$ and $f(t) \leq \min(t, C t^b e^{-t/2})$ for $t > t^\dagger$ with $t^\dagger < 2b$, we have

$$\sup_{t \geq 0} f(t) \leq \min(2b, C(2b)^b e^{-b}).$$

As a result, we conclude that

$$\sup_{r \geq 0} \frac{c_a}{c_{a-1}} \gamma^2 r \tilde{\kappa}_{a-1}(\gamma r) \leq \frac{c_a \gamma}{c_{a-1}} \min(2(a-1), 4c_{a-1} (2(a-1))^{a-1} e^{-a+1}). \quad (113)$$

Substituting the value of c_{a-1} and the bound (94), we can bound the second argument inside the min on the RHS of (113) (for $a \geq 3$) as follows:

$$\begin{aligned} 4c_{a-1} (2(a-1))^{a-1} e^{-a+1} &\leq 4 \cdot 2^{2-a} \frac{1}{\sqrt{2\pi(a-2)}} \left(\frac{e}{a-2}\right)^{a-2} 2^{a-1} (a-1)^{a-1} e^{1-a} \\ &\leq \frac{8}{e\sqrt{2\pi}} \left(1 + \frac{1}{a-2}\right)^{a-2} \left(\sqrt{a-2} + \frac{1}{\sqrt{a-2}}\right) \leq \frac{8}{\sqrt{\pi}} \sqrt{a-2} \end{aligned} \quad (114)$$

for $a \geq 3$. When $a \in [2, 3)$, one can directly show that $4c_{a-1} (2(a-1))^{a-1} e^{-a+1} \leq 8/e$. Putting the pieces together, and noting that $\frac{c_a}{c_{a-1}} = \frac{1}{\max(2(a-1), 1)}$, we find that

$$\sup_{r \geq 0} \frac{c_a}{c_{a-1}} \gamma^2 r \tilde{\kappa}_{a-1}(\gamma r) \leq \frac{\gamma}{\max(2(a-1), 1)} \min\left(2(a-1), \frac{8}{e} + \frac{8}{\sqrt{\pi}} \sqrt{a-2}\right) \leq \frac{C_1 \gamma}{\sqrt{a+C_2}}$$

for any $a \geq 2$. And hence

$$L_{\mathbf{k}_{rt}} \leq A_{\nu,\gamma,d} \sup_{r \geq 0} |\kappa'_a(\gamma r)| \leq A_{\nu,\gamma,d} \frac{C_1 \gamma}{\sqrt{a+C_2}} \implies \frac{L_{\mathbf{k}_{rt}}}{\|\mathbf{k}_{rt}\|_\infty} \stackrel{(105)}{\leq} \frac{C_1 \gamma}{\sqrt{a+C_2}},$$

as claimed.

O.3.5 PROOF OF THE BOUND (108) ON $R_{\mathbf{k}_{\text{rt}},n}$

Using arguments similar to those used to obtain (114), we find that

$$\max_{\gamma r \geq 2(a-1)} \tilde{\kappa}_a(\gamma r) \leq 4c_a(2(a-1))^{a-1} e^{-(a-1)} \leq \sqrt{\frac{4}{a-1}} = \sqrt{\frac{8}{\nu-d-2}}.$$

Next, note that

$$\begin{aligned} (\gamma r)^{a-\frac{1}{2}} e^{-\gamma r + \frac{a^2}{2\gamma r}} &\lesssim e^{-\gamma r/2} \quad \text{for } \gamma r \gtrsim a \log(1+a), \\ c_a &= \frac{2^{1-a}}{\Gamma(a)} \stackrel{(94)}{\leq} \left(\frac{2e}{a-1}\right)^{a-1} \frac{1}{\sqrt{2\pi(a-1)}}, \end{aligned} \quad (115)$$

where (115) follows from standard algebra. Thus, we have

$$\begin{aligned} c_a \tilde{\kappa}_a(\gamma r) &\lesssim c_a \exp(-\gamma r/2) \quad \text{for } \gamma r \gtrsim a \log(1+a) \\ \stackrel{(94)}{\implies} R_{\mathbf{k}_{\text{rt}},n} &\lesssim \frac{1}{\gamma} \max(\log n - a \log(1+a), C_1 a \log(1+a)), \end{aligned}$$

as claimed.

 O.3.6 PROOF OF THE BOUND (109) ON $R'_{\mathbf{k}_{\text{rt}},\sqrt{n}}$

Let $V_d = \pi^{d/2}/\Gamma(d/2+1)$ denote the volume of unit Euclidean ball in \mathbb{R}^d . Using (103), we have

$$\begin{aligned} \frac{1}{A_{\nu,\gamma,d}^2} \tau_{\mathbf{k}_{\text{rt}}}^2(R) &= \int_{\|z\|_2 \geq R} \tilde{\kappa}_a^2(\gamma \|z\|_2) dz = dV_d \int_{r>R} r^{d-1} \tilde{\kappa}_a^2(\gamma r) dr \\ &\stackrel{(110)}{\leq} 2\pi c_a^2 dV_d \int_{r>R} r^{d-1} (\gamma r)^{2a-1} \exp(-2\gamma r + \frac{a^2}{\gamma r}) dr \\ &= 2\pi c_a^2 dV_d \gamma^{1-d} \int_{r>R} (\gamma r)^{\nu-2} \exp(-2\gamma r + \frac{a^2}{\gamma r}) dr \end{aligned} \quad (116)$$

where we have also used the fact that $2a+d=\nu$. Next, we note that

$$\exp(-2\gamma r + \frac{a^2}{\gamma r}) \leq \exp(-\frac{3}{2}\gamma r) \quad \text{for } \gamma r \geq \sqrt{2}a \quad (117)$$

For any integer $b > 1$, noticing that $f(t) = \frac{1}{\Gamma(b)} t^{b-1} e^{-t}$ is the density function of an Erlang distribution with shape parameter b and rate parameter 1, and using the expression for its complementary cumulative distribution function, we find that

$$\int_r^\infty t^{b-1} e^{-t} dt = \Gamma(b) \sum_{i=0}^{b-1} \frac{r^i}{i!} e^{-r} \quad (118)$$

Thus, for $R > \frac{\sqrt{2}}{\gamma}a$, we have

$$\begin{aligned} \int_{r>R} (\gamma r)^{\nu-2} \exp(-2\gamma r + \frac{a^2}{\gamma r}) dr &\stackrel{(117)}{\leq} \int_{r>R} (\gamma r)^{\nu-2} \exp(-\frac{3}{2}\gamma r) dr \\ &= \frac{1}{\gamma} \left(\frac{2}{3}\right)^{\nu-1} \int_{3\gamma R/2}^\infty t^{\nu-2} e^{-t} dt \\ &\stackrel{(118)}{=} \frac{1}{\gamma} \left(\frac{2}{3}\right)^{\nu-1} \Gamma(\nu-1) e^{-\frac{3}{2}\gamma R} \sum_{i=0}^{\nu-1} \frac{(3\gamma R/2)^i}{i!} \\ &\leq \frac{1}{\gamma} \Gamma(\nu-1) e^{-\frac{3}{2}\gamma R} e^{\gamma R} \leq \frac{1}{\gamma} \Gamma(\nu-1) e^{-\frac{1}{2}\gamma R}. \end{aligned} \quad (119)$$

Putting the bounds (116) and (119) together for $R \geq \frac{\nu-d}{\sqrt{2}\gamma}$, we find that

$$\tau_{\mathbf{k}_{\text{rt}}}^2(R) \leq A_{\nu,\gamma,d}^2 \cdot \gamma^{-d} \cdot 2\pi \frac{2^{2-2a}}{\Gamma^2(a)} \cdot \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)} \cdot \Gamma(\nu-1) \exp(-\frac{1}{2}\gamma R).$$

Using (94), we have

$$\begin{aligned} 2\pi \frac{2^{2-2a}}{\Gamma^2(a)} \cdot \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)} \cdot \Gamma(\nu-1) &\leq 2\pi \cdot \frac{2^{2-2a} e^{2(a-1)}}{2\pi(a-1)(a-1)^{2(a-1)}} \cdot \frac{\pi^{d/2} (2e)^{d/2}}{\sqrt{\pi d} \cdot d^{d/2}} \cdot e\sqrt{\nu-2} \left(\frac{\nu-2}{e}\right)^{\nu-2} \\ &= \frac{2}{\sqrt{\pi}} e^{2a-2+d/2+1-\nu} (2\pi)^{d/2} (2a-2)^{-(2a-1)} d^{-d/2-1} (\nu-2)^{\nu-3/2} \\ &= \frac{2}{e\sqrt{\pi}} (2e\pi)^{d/2} (\nu-d-2)^{-(\nu-d-1)} d^{-d/2-1} (\nu-2)^{\nu-3/2} \\ &= \frac{2}{e\sqrt{\pi}} (2e\pi)^{d/2} (2(a-1))^{-(2a-1)} d^{-d/2-1} (\nu-2)^{\nu-3/2}. \end{aligned}$$

Putting the pieces together, and solving for $\tau_{\mathbf{k}_{\text{rt}}}(R) \leq \|\mathbf{k}_{\text{rt}}\|_{\infty}/\sqrt{n} \stackrel{(105)}{=} A_{\nu,\gamma,d}/\sqrt{n}$, we find that $R'_{\mathbf{k}_{\text{rt}},\sqrt{n}}$ (6) can be bounded as

$$\begin{aligned} R'_{\mathbf{k}_{\text{rt}},\sqrt{n}} &\leq \frac{2}{\gamma} \cdot \max\left(\frac{a}{\sqrt{2}}, \log n + \log\left(\frac{2}{e\sqrt{\pi}}\right) + d \log\left(\frac{\sqrt{2e\pi}}{\gamma}\right) + \log\left(\frac{(\nu-2)^{\nu-\frac{3}{2}}}{(2(a-1))^{2a-1} d^{\frac{d}{2}+1}}\right)\right) \\ &\lesssim \frac{1}{\gamma} \cdot (a + \log n + d \log\left(\frac{\sqrt{2e\pi}}{\gamma}\right) + \log\left(\frac{(\nu-2)^{\nu-\frac{3}{2}}}{(2(a-1))^{2a-1} d^{\frac{d}{2}+1}}\right)), \end{aligned}$$

as claimed.

O.4 Proofs for B-spline kernel

Recall the notation from (89) to (92). Then, the kernel $\mathbf{k}_{\star} = \mathbf{B}\text{-spline}(2\beta+1)$ and its square-root kernel $\mathbf{k}_{\text{rt}} = \tilde{S}_{\beta,d} \cdot \mathbf{B}\text{-spline}(\beta)$ satisfy

$$\|\mathbf{k}_{\text{rt}}\|_{\infty} = \tilde{S}_{\beta,d} \stackrel{(i)}{=} c_{\beta}^d \quad \text{where} \quad c_{\beta} \triangleq \frac{\mathfrak{B}_{\beta+1}}{\sqrt{\mathfrak{B}_{2\beta+2}}} \begin{cases} \stackrel{(ii)}{=} \frac{2}{\sqrt{3}} & \text{when } \beta = 1 \\ \leq 1 & \text{when } \beta > 1 \end{cases} \quad (120)$$

$$\|\mathbf{k}_{\star}\|_{\infty} = 1, \quad (121)$$

$$\frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_{\infty}} \leq \frac{4}{3}\sqrt{d}, \quad (122)$$

$$R_{\mathbf{k}_{\text{rt}},n} \leq \sqrt{d}(\beta+1)/2, \quad \text{and} \quad R'_{\mathbf{k}_{\text{rt}},\sqrt{n}} \leq \sqrt{d}(\beta+1)/2. \quad (123)$$

While claims (i) and (ii) follow directly from the definitions in the display (89), claim (iii) can be verified numerically, e.g., using SciPy (Virtanen et al., 2020). From numerical verification, we also find that the constant c_{β} in (120) is decreasing with β . See App. O.4.1 for the proofs of the remaining claims. Finally, substituting various quantities from (120), (122), and (123) in (93), we find that

$$\mathfrak{M}_{\mathbf{k}_{\text{rt}}}^{\mathbf{B}\text{-spline}}\left(n, \frac{1}{2} \log_2 n, d, \frac{\delta}{2n}, \delta', R\right) \lesssim \sqrt{\log\left(\frac{n}{\delta}\right) \left(\log\left(\frac{1}{\delta}\right) + d \log(d\beta + \sqrt{d}R)\right)}.$$

O.4.1 PROOFS OF THE BOUNDS ON B-SPLINE KERNEL QUANTITIES

We start with some basic set-up. Consider the (unnormalized) univariate B-splines

$$f_\beta : \mathbb{R} \rightarrow [0, 1] \quad \text{with} \quad f_\beta(a) = \otimes^\beta \mathbb{I}_{[-\frac{1}{2}, \frac{1}{2}]}(a) \stackrel{(i)}{=} \frac{1}{(\beta-1)!} \sum_{j=0}^{\beta} (-1)^j \binom{\beta}{j} (a + \frac{\beta}{2} - j)_+^{\beta-1}$$

where step (i) follows from Schumaker (2007, Eqn 4.46, 4.47, 4.59, p.135, 138). Noting that f_β is an even function with a unique global maxima at 0 (see Schumaker (2007, Ch 4.) for more details), we find that

$$\|f_\beta\|_\infty = f_\beta(0) = \frac{1}{(\beta-1)!} \sum_{j=0}^{\lfloor \beta/2 \rfloor} (-1)^j \binom{\beta}{j} (\frac{\beta}{2} - j)^{\beta-1} \stackrel{(89)}{=} \mathfrak{B}_\beta. \quad (124)$$

Bounds on $\|\mathbf{k}_\star\|_\infty$ and $\|\mathbf{k}_{\text{rt}}\|_\infty$ Recalling (90) and (91), we find that

$$\|\mathbf{k}_\star\|_\infty = \|\chi_{2\beta+2}\|_\infty = S_{2\beta+2,d} \|f_{2\beta+2}\|_\infty^d = S_{2\beta+2,d} \mathfrak{B}_{2\beta+2}^d = 1,$$

thereby establishing (121).

Bounds on $L_{\mathbf{k}_{\text{rt}}}$ We have

$$f'_{\beta+1}(a) = \int f_\beta(b) \frac{d}{da} \mathbb{I}_{[-\frac{1}{2}, \frac{1}{2}]}(a-b) dy = f_\beta(a - \frac{1}{2}) - f_\beta(a + \frac{1}{2})$$

and hence $\|f'_{\beta+1}\|_\infty = \sup_a |f_\beta(a - \frac{1}{2}) - f_\beta(a + \frac{1}{2})| \leq \|f_\beta\|_\infty$ since f_β is non-negative. Putting the pieces together, we have

$$\begin{aligned} \frac{L_{\mathbf{k}_{\text{rt}}}}{\|\mathbf{k}_{\text{rt}}\|_\infty} &= \frac{1}{S_{\beta,d}} L_{\mathbf{k}_{\text{rt}}} \leq \sup_z \|\nabla \chi_{\beta+1}(z)\|_2 \stackrel{(90)}{\leq} \sqrt{d} \cdot S_{\beta+1,d} \|f'_{\beta+1}\|_\infty \cdot \|f_{\beta+1}\|_\infty^{d-1} \\ &\stackrel{(89,124)}{\leq} \sqrt{d} \cdot \mathfrak{B}_{\beta+1}^{-d} \cdot \mathfrak{B}_\beta \cdot \mathfrak{B}_{\beta+1}^{d-1} \\ &= \sqrt{d} \frac{\mathfrak{B}_\beta}{\mathfrak{B}_{\beta+1}} \stackrel{(i)}{\leq} \sqrt{d} \frac{4}{3}, \end{aligned}$$

where step (i) can be verified numerically.

Bounds on $R_{\mathbf{k}_{\text{rt}},n}$, and $R'_{\mathbf{k}_{\text{rt}},\sqrt{n}}$ Using the property of convolution, we find that $f_{\beta+1}(a) = 0$ if $|a| \geq \frac{1}{2}(\beta+1)$. Hence $\kappa_{\text{rt}}(z) = 0$ for $\|z\|_\infty > (\beta+1)/2$ and applying the definitions (6), we find that

$$R_{\mathbf{k}_{\text{rt}},n} \leq \sqrt{d}(\beta+1)/2 \quad \text{and} \quad R'_{\mathbf{k}_{\text{rt}},\sqrt{n}} \leq \sqrt{d}(\beta+1)/2$$

as claimed in (123).

O.5 Explicit MMD rates for common kernels

Putting the quantities from Tab. 2 together with Thm. 1 (with $\delta' = \delta$, and δ treated as a constant) yields the MMD rates summarized in Tab. 4.

For completeness, we illustrate a key simplification that can readily yield the results stated in Tab. 4. Define $\zeta_{\mathbf{k}_{\text{rt}},\mathcal{S}_n}$ as follows:

$$\zeta_{\mathbf{k}_{\text{rt}},\mathcal{S}_n} \triangleq \frac{1}{d} \max(R'_{\mathbf{k}_{\text{rt}},\sqrt{n}}, R_{\mathcal{S}_n})^2,$$

Kernel \mathbf{k}_*	$\text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) \lesssim$
Gaussian(σ)	$C_1^d \cdot \sqrt{\frac{\log n}{n} \cdot [1 + \frac{1}{d}(\log n + (\frac{R_{\mathcal{S}_n}}{\sigma})^2)]^{\frac{d}{2}} \cdot \log(\sqrt{\log n} + \frac{R_{\mathcal{S}_n}}{\sigma})}$
Matérn(ν, γ)	$C_2^d \cdot \sqrt{\frac{\nu^2 \log n}{n} \left[\frac{1}{d(a-1)} (a + \log^2 n + d^2 \log^2(\frac{\sqrt{2e\pi}}{\gamma}) + G_{\nu,d}^2 + \gamma^2 R_{\mathcal{S}_n}^2) \right]^{\frac{d}{2}} \cdot \log(\log n + a + \gamma R_{\mathcal{S}_n})}$
B-spline($2\beta + 1$)	$C_3^d \cdot \sqrt{\frac{\log n}{n} \cdot [\beta^2 + \frac{R_{\mathcal{S}_n}^2}{d}]^{\frac{d}{2}} \cdot \log(\beta + \frac{R_{\mathcal{S}_n}}{\sqrt{d}})}$

Table 4: Explicit kernel thinning MMD guarantees for common kernels. Here, $a \triangleq \frac{1}{2}(\nu - d)$, $G_{\nu,d} \triangleq \log(\frac{(\nu-2)^{\nu-\frac{3}{2}}}{(2(a-1))^{2a-1} d^{\frac{d}{2}+1}})$ and each C_i denotes a universal constant. See App. O.5 for more details on deriving these bounds from Thm. 1.

where $R_{\mathcal{S}_n}$, and $R'_{\mathbf{k}_{\text{rt}}, \sqrt{n}}$ were defined in (6) and (7) respectively. Then applying Thm. 1, and substituting the simplified bound (93) in (9), we find that

$$\begin{aligned} \text{MMD}_{\mathbf{k}_*}(\mathcal{S}_n, \mathcal{S}_{\text{KT}}) &\leq 2 \frac{\|\mathbf{k}_{\text{rt}}\|_{\infty}}{\sqrt{n}} + \frac{\|\mathbf{k}_{\text{rt}}\|_{\infty}}{\sqrt{n}} (B \zeta_{\mathbf{k}_{\text{rt}}, \mathcal{S}_n})^{\frac{d}{4}} \cdot d^{-\frac{1}{4}} \sqrt{\log \frac{n}{\delta} \cdot \left[\log\left(\frac{1}{\delta}\right) + d \log\left(\frac{L_{\mathbf{k}_{\text{rt}}}(R_{\mathbf{k}_{\text{rt}}, n} + R)}{\|\mathbf{k}_{\text{rt}}\|_{\infty}}\right) \right]} \\ &\stackrel{(i)}{\lesssim_{\delta}} 2 \frac{\|\mathbf{k}_{\text{rt}}\|_{\infty}}{\sqrt{n}} + \|\mathbf{k}_{\text{rt}}\|_{\infty} (B \zeta_{\mathbf{k}_{\text{rt}}, \mathcal{S}_n})^{\frac{d}{4}} \cdot d^{\frac{1}{4}} \sqrt{\frac{\log n}{n} \log\left(\frac{L_{\mathbf{k}_{\text{rt}}}(R_{\mathbf{k}_{\text{rt}}, n} + R_{\mathcal{S}_n})}{\|\mathbf{k}_{\text{rt}}\|_{\infty}}\right)}, \end{aligned} \quad (125)$$

with probability at least $1 - 2\delta$, where $B \triangleq 8e\pi$ is a universal constant, and in step (i) we use the following bound: For any $r = \sqrt{\alpha d}$, we have

$$c_d r^{\frac{d}{2}} = \frac{(4\pi)^{\frac{d}{4}}}{(\Gamma(\frac{d}{2}+1))^{\frac{1}{2}}} (\alpha d)^{\frac{d}{4}} \stackrel{(ii)}{\leq} \frac{(8e\pi\alpha)^{\frac{d}{4}}}{(\pi d)^{\frac{1}{4}}} \leq (B\alpha)^{\frac{d}{4}} d^{-\frac{1}{4}} \quad \text{where } B \triangleq 8e\pi,$$

and step (ii) follows (for $d \geq 2$) from the Gamma function bounds (94). Now the results in Tab. 4 follow by simply substituting the various quantities from Tab. 2 in (125).

Appendix P. Supplementary Details for Vignettes of Sec. 7

This section provides supplementary details for the vignettes of Sec. 7.

P.1 Mixture of Gaussians target

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

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

P.2 Empirical decay rates

In Fig. 5, we provide results for regressing the true error rate of $\sqrt{\frac{\log n}{n}}$ and $\frac{\log n}{\sqrt{n}}$ to the input size n (on the log-log scale as in all of mean MMD plots) for various ranges of coreset sizes. Notably, the observed empirical rates are significantly slower than $n^{-\frac{1}{2}}$ for the smallest

coreset range (the same coreset range used in the panels in Fig. 3) but approach $n^{-\frac{1}{2}}$ as the coreset size increases. Hence, the observed empirical rates of Fig. 3 are consistent with $\frac{\log^c n}{\sqrt{n}}$ rates of decay.





Figure 5: Empirical decay rates when true error = $\frac{\log^c n}{\sqrt{n}}$. We display the ordinary least squares fits for regressing the log of true error onto the log of input size n when the true error equals either (a) $\sqrt{\frac{\log n}{n}}$ or (b) $\frac{\log n}{\sqrt{n}}$. As the input range increases, the fitted decay rate tends towards $n^{-\frac{1}{2}}$ despite appearing significantly slower for smaller input ranges.

P.3 MCMC vignette details

For complete details on the targets and sampling algorithms we refer the reader to Riabiz et al. (2021, Sec. 4). When applying standard thinning to any Markov chain output, we adopt the convention of keeping the final sample point. For all experiments, we used only the odd indices of the post burn-in sample points when thinning to form \mathcal{S}_n .

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. For the Hinch experiments, we discard the first 10^6 points as burn-in following Riabiz et al. (2021, App. S5.4). For all n , the parameter σ for the Gaussian kernel is set to the median distance between all pairs of 4^7 points obtained by standard thinning the post-burn-in odd indices. The resulting values for the Goodwin chains were 0.02 for RW; 0.0201 for adaRW; 0.0171 for MALA; and 0.0205 for pMALA. The respective numbers for Lotka-Volterra task were 0.0274 for RW; 0.0283 for adaRW; 0.023 for MALA; and 0.0288 for pMALA. Finally, the numbers for the Hinch task were 8.0676 for RW; 8.3189 for adaRW; 8.621 for MALA; and 8.6649 for pMALA.

Appendix Q. Kernel Thinning with Square-root Dominating Kernels

As alluded to in Sec. 2, it is not necessary to identify an exact square-root kernel \mathbf{k}_{rt} to run kernel thinning. Rather, it suffices to identify any *square-root dominating kernel* $\tilde{\mathbf{k}}_{\text{rt}}$ defined as follows.

Definition 8 (Square-root dominating kernel). *We say a reproducing kernel $\tilde{\mathbf{k}}_{\text{rt}} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a square-root dominating kernel for a reproducing kernel $\mathbf{k}_\star : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ if $\tilde{\mathbf{k}}_{\text{rt}}(x, \cdot)$ is square-integrable for all $x \in \mathbb{R}^d$ and either of the following equivalent conditions hold.*

- (a) *The RKHS of \mathbf{k}_\star belongs to the RKHS of $\tilde{\mathbf{k}}(x, y) \triangleq \int_{\mathbb{R}^d} \tilde{\mathbf{k}}_{\text{rt}}(x, z)\tilde{\mathbf{k}}_{\text{rt}}(y, z)dz$.*
- (b) *The function $\tilde{\mathbf{k}} - c\mathbf{k}_\star$ is positive definite for some $c > 0$.*

Remark 9 (Controlling $\text{MMD}_{\mathbf{k}_\star}$). *A square-root dominating kernel $\tilde{\mathbf{k}}_{\text{rt}}$ is a suitable surrogate for \mathbf{k}_{rt} as, by Zhang and Zhao (2013, Lem. 2.2, Prop. 2.3), $\text{MMD}_{\mathbf{k}_\star}(\mu, \nu) \leq \sqrt{1/c} \cdot \text{MMD}_{\tilde{\mathbf{k}}}(\mu, \nu)$ for c the constant appearing in Def. 8(b) and all distributions μ and ν .*

Def. 8 and Rem. 9 enable us to use convenient dominating surrogates whenever an exact square-root kernel is inconvenient to derive or deploy. For example, our next result, proved in App. Q.1, shows that a standard Matérn kernel is a square-root dominating kernel for every sufficiently-smooth shift-invariant and absolutely integrable \mathbf{k}_\star .

Proposition 4 (Dominating smooth kernels). *If $\mathbf{k}_\star(x, y) = \kappa(x - y)$ and $\kappa \in L^1 \cap C^{2\nu}$ for $\nu > d$, then, for any $\gamma > 0$, the **Matérn** $(\frac{\nu}{2}, \gamma)$ kernel of Tab. 1 is a square-root dominating kernel for \mathbf{k}_\star .*

Checking the square-root dominating condition is also particularly simple for any pair of continuous shift-invariant kernels as the next result, proved in App. Q.2, shows.

Proposition 5 (Dominating shift-invariant kernels). *Suppose that the real-valued kernels $\mathbf{k}_\star(x, y) = \kappa(x - y)$ and $\tilde{\mathbf{k}}_{\text{rt}}(x, y) = \tilde{\kappa}_{\text{rt}}(x - y)$ have respective spectral densities (Def. 6) $\hat{\kappa}$ and $\hat{\kappa}_{\text{rt}}$. If $\hat{\kappa}_{\text{rt}} \in L^2$, then $\tilde{\mathbf{k}}_{\text{rt}}$ is a square-root dominating kernel for \mathbf{k}_\star if and only if*

$$\text{ess sup}_{\omega \in \mathbb{R}^d: \hat{\kappa}(\omega) > 0} \frac{\hat{\kappa}(\omega)}{\hat{\kappa}_{\text{rt}}(\omega)^2} < \infty. \quad (126)$$

In Tab. 5, we use Prop. 5 to derive convenient tailored square-root dominating kernels $\tilde{\mathbf{k}}_{\text{rt}}$ for standard inverse multiquadric kernels, hyperbolic secant (sech) kernels, and Wendland's compactly supported kernels. In each case, we can identify a square-root dominating kernel from the same family.

Name of kernel $\mathbf{k}_*(x, y) = \kappa(x - y)$	Expression for $\kappa(z)$	Name of square-root dominating kernel $\mathbf{k}_{\text{rt}}(x, y) = \tilde{\kappa}_{\text{rt}}(x - y)$
InverseMultiquadric (ν, γ) $\nu > 0, \gamma > 0$	$(\gamma^2 + \ z\ _2^2)^{-\nu}$	InverseMultiquadric (ν', γ') $(\nu', \gamma') \in \mathcal{C}_{\nu, \gamma, d}$; see (131)
Sech (a) $a > 0$	$\prod_{j=1}^d \text{sech}(\sqrt{\frac{\pi}{2}} a z_j)$	Sech $(2a)$
Wendland (s) $s \in \mathbb{N}, s \geq \frac{1}{2}(d + 1)$	$\phi_{d,s}(\ z\ _2)$; see (127)	Wendland (s') $s' \in \mathbb{N}_0, s' \leq \frac{1}{4}(2s - 1 - d)$

Table 5: Square-root dominating kernels $\tilde{\mathbf{k}}_{\text{rt}}$ for common kernels \mathbf{k}_* (see Def. 8). Here \mathbb{N}_0 denotes the non-negative integers. See App. Q.3 for our derivation.

Expressions for Wendland kernels The **Wendland** (s) kernel is a compactly-supported radial kernel $\phi_{d,s}(\|x - y\|_2)$ on \mathbb{R}^d where $\phi_{d,s} : \mathbb{R}_+ \rightarrow \mathbb{R}$ is a truncated minimal-degree polynomial with $2s$ continuous derivatives. We collect here the expressions for $\phi_{d,s}$ for $s = 0, 1, 2$ and refer the readers to Wendland (2004, Ch. 9, Thm. 9.13, Tab. 9.1) for more general s . Let $(r)_+ = \max(0, r)$ and $\ell \triangleq \lfloor d/2 \rfloor + 3$, then we have

$$\begin{aligned} \phi_{d,0}(r) &= (1 - r)_+^{\lfloor d/2 \rfloor + 1}, & \phi_{d,1}(r) &= (1 - r)_+^{\lfloor d/2 \rfloor + 2} [(\lfloor d/2 \rfloor + 2)r + 1] \\ \phi_{d,2}(r) &= (1 - r)_+^{\lfloor d/2 \rfloor + 3} [(\ell^2 + 4\ell + 3)r^2 + (3\ell + 6)r + 3]. \end{aligned} \quad (127)$$

Q.1 Proof of Prop. 4: Dominating smooth kernels

Since $\kappa \in L^1$, $\mathcal{F}(\kappa)$ is bounded by the Babenko-Beckner inequality (Beckner, 1975) and nonnegative by Bochner's theorem (Bochner, 1933; Wendland, 2004, Thm. 6.6). Moreover, since $\kappa \in C^{2\nu}$, Sun (1993, Thm. 4.1) implies that $\int \|\omega\|_2^{2\nu} \mathcal{F}(\kappa)(\omega) d\omega < \infty$. By Wendland (2004, Theorem 8.15), the **Matérn** (ν, γ) kernel $\tilde{\mathbf{k}}(x, y) \propto \Phi_{\nu, \gamma}(x - y)$ for $\Phi_{\nu, \gamma}$ continuous with $\mathcal{F}(\Phi_{\nu, \gamma})(\omega) = (\gamma^2 + \|\omega\|_2^2)^{-\nu}$. Since we have established that

$$\int \frac{\mathcal{F}(\kappa)(\omega)^2}{\mathcal{F}(\Phi_{\nu, \gamma})(\omega)} d\omega = \int (\gamma^2 + \|\omega\|_2^2)^{\nu} \mathcal{F}(\kappa)(\omega)^2 d\omega \leq \|\mathcal{F}(\kappa)\|_{\infty} \int (\gamma^2 + \|\omega\|_2^2)^{\nu} \mathcal{F}(\kappa)(\omega) d\omega < \infty,$$

Wendland (2004, Thm. 10.12) implies that κ belongs to $\mathcal{H}_{\tilde{\mathbf{k}}}$ and hence that $\mathcal{H}_{\mathbf{k}_*} \subseteq \mathcal{H}_{\tilde{\mathbf{k}}}$. Finally, by App. N.1, **Matérn** $(\frac{\nu}{2}, \gamma)$ is a valid square-root dominating kernel for $\tilde{\mathbf{k}}$ and therefore for \mathbf{k}_* .

Q.2 Proof of Prop. 5: Dominating shift-invariant kernels

Def. 6 implies that

$$\mathbf{k}_*(x, y) = \frac{1}{(2\pi)^{d/2}} \int e^{-i\langle \omega, x-y \rangle} \hat{\kappa}(\omega) d\omega \quad \text{and} \quad \tilde{\mathbf{k}}_{\text{rt}}(x, y) = \frac{1}{(2\pi)^{d/2}} \int e^{-i\langle \omega, x-y \rangle} \hat{\kappa}_{\text{rt}}(\omega) d\omega,$$

Moreover, since $\mathbf{k}_{\text{rt}}(x, \cdot) = \mathcal{F}(e^{-i\langle \cdot, x \rangle} \hat{\kappa}_{\text{rt}})$ for $e^{-i\langle \cdot, x \rangle} \hat{\kappa}_{\text{rt}} \in L^1 \cap L^2$, the Plancherel-Parseval identity (Wendland, 2004, Proof of Thm. 5.23) implies that

$$\begin{aligned} \tilde{\mathbf{k}}(x, y) &\triangleq \int_{\mathbb{R}^d} \tilde{\mathbf{k}}_{\text{rt}}(x, z) \tilde{\mathbf{k}}_{\text{rt}}(y, z) dz = \int_{\mathbb{R}^d} e^{-i\langle \omega, x \rangle} \hat{\kappa}_{\text{rt}}(\omega) e^{i\langle \omega, y \rangle} \hat{\kappa}_{\text{rt}}(\omega) d\omega \\ &= \int_{\mathbb{R}^d} e^{-i\langle \omega, x-y \rangle} \hat{\kappa}_{\text{rt}}(\omega)^2 d\omega, \end{aligned}$$

and Bochner's theorem (Bochner, 1933; Wendland, 2004, Thm. 6.6) implies that $\tilde{\mathbf{k}}$ is a kernel. Finally, Prop. 3.1 of Zhang and Zhao (2013) now implies that the RKHS of \mathbf{k}_\star belongs to the RKHS of $\tilde{\mathbf{k}}$ if and only if the ratio condition (126) holds.

Q.3 Derivation of Tab. 5: Square-root dominating kernels $\tilde{\mathbf{k}}_{\text{rt}}$ for common kernels \mathbf{k}_\star

Thanks to Prop. 5, to establish the validity of the square root dominating kernels stated in Tab. 5, it suffices to verify that

$$\widehat{\kappa}_{\text{rt}} \in L^2 \quad \text{and} \quad \widehat{\kappa} \lesssim_{d,\kappa,\tilde{\kappa}_{\text{rt}}} \widehat{\kappa}_{\text{rt}}^2, \quad (128)$$

where the functions κ and $\widehat{\kappa}_{\text{rt}}$ are the spectral densities of κ and $\tilde{\kappa}_{\text{rt}}$.

Inverse multiquadric Consider the positive definite function $\Phi_{\nu,\gamma}$ (86) underlying the Matérn kernel, which is continuous on $\mathbb{R}^d \setminus \{\mathbf{0}\}$. When $\kappa(z) = (\gamma^2 + \|z\|_2^2)^{-\nu}$ and $\tilde{\kappa}_{\text{rt}}(z) = (\gamma'^2 + \|z\|_2^2)^{-\nu'}$, Wendland (2004, Theorem 8.15) implies that

$$\widehat{\kappa}(\omega) = \Phi_{\nu,\gamma}(\omega) \quad \text{and} \quad \Phi_{\nu',\gamma'}^2(\omega) = \widehat{\kappa}_{\text{rt}}(\omega)^2.$$

Let $a(\omega) \asymp_{d,\nu,\gamma} b(\omega)$ denote asymptotic equivalence up to a constant depending on d, ν, γ . Then, by (DLMF, Eqs. 10.25.3 & 10.30.2), we have

$$\Phi_{\nu,\gamma}(\omega) \asymp_{d,\nu,\gamma} \|\omega\|_2^{\nu - \frac{d}{2} - \frac{1}{2}} e^{-\gamma\|\omega\|_2} \quad \text{as} \quad \|\omega\|_2 \rightarrow \infty, \quad (129)$$

$$\Phi_{\nu,\gamma}(\omega) \asymp_{d,\nu,\gamma} \|\omega\|_2^{-(d-2\nu)_+} \quad \text{as} \quad \|\omega\|_2 \rightarrow 0. \quad (130)$$

Hence, $\widehat{\kappa}_{\text{rt}} \in L^2$ whenever $\nu' > \frac{d}{4}$.

Moreover, applying (129), we find that for $\|\omega\|_2 \rightarrow \infty$,

$$\begin{aligned} \frac{\Phi_{\nu,\gamma}(\omega)}{\Phi_{\nu',\gamma'}^2(\omega)} &\asymp_{d,\nu,\gamma,\nu',\gamma'} \|\omega\|_2^{\nu - \frac{d}{2} - \frac{1}{2}} e^{-\gamma\|\omega\|_2} \cdot \|\omega\|_2^{-2\nu' + d + 1} e^{2\gamma'\|\omega\|_2} \\ &= \|\omega\|_2^{\nu + \frac{d}{2} + \frac{1}{2} - 2\nu'} e^{(2\gamma' - \gamma)\|\omega\|_2}. \end{aligned}$$

If $2\gamma' - \gamma < 0$, this expression is bounded for any value of ν' . If $2\gamma' - \gamma = 0$, this expression is bounded when $\nu' \geq \frac{\nu}{2} + \frac{d}{4} + \frac{1}{4}$.

Applying (130), we find that for $\|\omega\|_2 \rightarrow 0$,

$$\frac{\Phi_{\nu,\gamma}(\omega)}{\Phi_{\nu',\gamma'}^2(\omega)} \asymp_{d,\nu,\gamma,\nu',\gamma'} \|\omega\|_2^{-(d-2\nu)_+} \cdot \|\omega\|_2^{2(d-2\nu')_+} = \|\omega\|_2^{2(d-2\nu')_+ - (d-2\nu)_+},$$

If $\nu \geq \frac{d}{2}$, this expression is finite for any value of ν' . If $\nu < \frac{d}{2}$, this expression is finite when $\nu' \leq \frac{\nu}{2} + \frac{d}{4}$. Hence, our condition (128) is verified whenever (ν', γ') belongs to the set

$$\begin{aligned} \mathcal{C}_{\nu,\gamma,d} \triangleq \{(\nu', \gamma') : & (1) \nu' > \frac{d}{4} \quad \text{and} \quad \gamma' \leq \frac{\gamma}{2}, \quad \text{and} \\ & (2) \nu' \leq \frac{d}{4} + \frac{\nu}{2} \quad \text{if} \quad \nu < \frac{d}{2}, \quad \text{and} \\ & (3) \nu' \geq \frac{d}{4} + \frac{\nu}{2} + \frac{1}{4} \quad \text{if} \quad \gamma' = \frac{\gamma}{2}\}. \end{aligned} \quad (131)$$

Sech Define $\kappa_a(z) \triangleq \prod_{j=1}^d \operatorname{sech}(\sqrt{\frac{\pi}{2}} a z_j)$, and suppose $\kappa = \kappa_a$ and $\tilde{\kappa}_{\text{rt}} = \kappa_{2a}$. Huggins and Mackey (2018, Ex. 3.2) yields that

$$\hat{\kappa}(\omega) = \mathcal{F}(\kappa_a)(\omega) = \frac{1}{a^d} \prod_{j=1}^d \operatorname{sech}(\sqrt{\frac{\pi}{2}} \frac{\omega_j}{a}) = a^{-d} \cdot \kappa_{1/a}(\omega), \quad \text{and} \quad (132)$$

$$\hat{\kappa}_{\text{rt}}(\omega) = \mathcal{F}(\tilde{\kappa}_{\text{rt}})(\omega) = (2a)^{-d} \kappa_{1/(2a)}(\omega). \quad (133)$$

Since $\operatorname{sech}^2(b/2) = \frac{4}{e^b + e^{-b} + 2} > \frac{4}{e^b + e^{-b}} = 2 \operatorname{sech}(b)$, we have

$$\kappa_{1/a}(\omega) \leq 2^{-d} \cdot \kappa_{1/(2a)}(\omega)^2. \quad (134)$$

Putting the pieces together, we further have

$$\hat{\kappa}(\omega) \stackrel{(132)}{=} a^{-d} \cdot \kappa_{1/a}(\omega) \stackrel{(134)}{\leq} (2a)^{-d} \cdot \kappa_{1/(2a)}(\omega)^2 \stackrel{(133)}{=} (2a)^d \cdot \hat{\kappa}_{\text{rt}}(\omega)^2.$$

Since $\kappa_{1/(2a)} \in L^2$, we have verified the condition (128).

Wendland When $\kappa(z) = \phi_{d,s}(\|z\|_2)$ and $\tilde{\kappa}_{\text{rt}}(z) = \phi_{d,s'}(\|z\|_2)$, Wendland (2004, Thm. 10.35) implies that

$$\hat{\kappa}(\omega) \lesssim_{d,s,s'} \frac{1}{(1+\|\omega\|_2^2)^{d+2s+1}} \leq \frac{1}{(1+\|\omega\|_2^2)^{2d+4s'+2}} \lesssim_{d,s,s'} \hat{\kappa}_{\text{rt}}(\omega)^2 \lesssim_{d,s,s'} \frac{1}{(1+\|\omega\|_2^2)^{2d+4s'+2}},$$

with $s' \in \mathbb{N}_0$, $s' \leq \frac{(2s-1-d)}{4}$, thereby establishing the condition (128).

Appendix R. Online Vector Balancing in Euclidean Space

Using Thm. 3, we recover the online vector balancing result of Alweiss et al. (2021, Thm. 1.2) with improved constants and a less conservative setting of the thresholds \mathbf{a}_i . Note that we capture the usual obliviousness assumption by treating the sequence $(f_i)_{i=1}^n$ as a fixed, deterministic input to Alg. 3.

Corollary 7 (Online vector balancing in Euclidean space). *If $\mathcal{H} = \mathbb{R}^d$ equipped with the Euclidean dot product, each $\|f_i\|_2 \leq 1$, and each $\mathbf{a}_i = \frac{1}{2} + \log(4n/\delta)$, then, with probability at least $1 - \delta$, the self-balancing Hilbert walk (Alg. 3) returns ψ_n satisfying*

$$\|\psi_n\|_\infty = \|\sum_{i=1}^n \eta_i f_i\|_\infty \leq \sqrt{2 \log(4d/\delta) \log(4n/\delta)}.$$

Proof Instantiate the notation of Thm. 3. With our choice of \mathbf{a}_i , the signed sum representation property (ii) implies that $\psi_n = \sum_{i=1}^n \eta_i f_i$ with probability at least $1 - \delta/2$. Moreover, the union bound, the functional sub-Gaussianity property (i), and the sub-Gaussian Hoeffding inequality (Wainwright, 2019, Prop. 2.5) now imply

$$\Pr(\|\psi_n\|_\infty > t \mid \mathcal{F}_n) \leq \sum_{j=1}^d \mathbb{P}(|\langle \psi_n, e_j \rangle| > t \mid \mathcal{F}_n) \leq 2d \exp(-t^2/(2\sigma_n^2)) = \delta/2$$

for $t \triangleq \sigma_n \sqrt{2 \log(4d/\delta)}$. We claim that

$$\sigma_n^2 \leq \alpha_n \triangleq \max_{j \leq n} \max(r_j, \|f_j\|_{\mathcal{H}}^2) \quad \text{with} \quad r_j \triangleq \frac{\mathbf{a}_j^2}{2\mathbf{a}_j - \|f_j\|_{\mathcal{H}}^2}. \quad (135)$$

Assuming this claim as given at the moment, we finish the proof. Since $\|f_j\|_{\mathcal{H}} = \|f_j\|_2 \leq 1$ for each j and $\frac{\frac{1}{2} + \log(4/\delta)}{\log(4/\delta)}$ is increasing in $\delta \in (0, 1]$,

$$\begin{aligned} \sigma_n^2 &\leq 1 \vee \max_{j \leq n} \frac{\alpha_j^2}{2\alpha_j - \|f_j\|_{\mathcal{H}}^2} \leq 1 \vee \max_{j \leq n} \frac{\alpha_j^2}{2\alpha_j - 1} = \log(4n/\delta) \frac{(\frac{1}{2} + \log(4n/\delta))^2}{2(\log(4n/\delta))^2} \\ &\leq \log(4n/\delta) \frac{(\frac{1}{2} + \log(4))^2}{2(\log(4))^2} \leq \log(4n/\delta). \end{aligned}$$

The advertised result now follows from the union bound.

We will now establish our earlier claim (135) using induction. For the base case, we have $\sigma_1^2 = \|f_1\|_{\mathcal{H}}^2 \leq \alpha_1$ from the definition of σ_1^2 (12).

Now, for the inductive case, assume that $\sigma_{i-1}^2 \leq \alpha_{i-1}$ for some $i > 1$, and define the function $g(x) = x + \|f_i\|_{\mathcal{H}}^2(1 - x/r_i)_+$ for $x \geq 0$. Using the definition of σ_i^2 (12), we find

$$\sigma_i^2 = \sigma_{i-1}^2 + \|f_i\|_{\mathcal{H}}^2 \left(1 + \frac{\sigma_{i-1}^2}{\alpha_i^2} (\|f_i\|_{\mathcal{H}}^2 - 2\alpha_i)\right)_+ = \sigma_{i-1}^2 + \|f_i\|_{\mathcal{H}}^2(1 - \sigma_{i-1}^2/r_i)_+ = g(\sigma_{i-1}^2).$$

Since $\alpha_{i-1} \leq \alpha_i$, our inductive assumption implies $\sigma_{i-1}^2 \leq \alpha_i$ and hence

$$\sigma_i^2 = g(\sigma_{i-1}^2) \leq \max_{x \in [0, \alpha_i]} g(x) = \max(\max_{x \in [0, r_i]} g(x), \max_{x \in (r_i, \alpha_i]} g(x)).$$

Next, we observe that g is continuous and that the derivative of g is $1 - \|f_i\|_{\mathcal{H}}^2/r_i$ for $x < r_i$ and 1 for $x > r_i$. Consequently, whenever $\|f_i\|_{\mathcal{H}}^2 < r_i$, the function g is increasing so that $\max_{x \in [0, \alpha_i]} g(x) \leq g(\alpha_i)$. On the other hand, when $\|f_i\|_{\mathcal{H}}^2 \geq r_i$, the function g is non-increasing in $[0, r_i]$ and increasing on (r_i, ∞) so that $\max_{x \in [0, r_i]} g(x) \leq g(0)$ and $\max_{x \in (r_i, \alpha_i]} g(x) \leq g(\alpha_i)$. Moreover, $g(\alpha_i) = \alpha_i$ since $\alpha_i \geq r_i$. Therefore, we always have

$$\sigma_i^2 \leq \max_{x \in [0, \alpha_i]} g(x) \leq \max(g(0), g(\alpha_i)) = \max(\|f_i\|_{\mathcal{H}}^2, \alpha_i) = \alpha_i.$$

Since $i > 1$ was arbitrary, the desired claim (135) follows by induction. \blacksquare

Appendix S. L^∞ Coresets of Phillips and Tai (2020) and Tai (2020)

Here we provide more details on the L^∞ coreset construction of Phillips and Tai (2020) and Tai (2020) discussed in Sec. 8.3. Given input points $(x_i)_{i=1}^n$, the Phillips-Tai (PT) construction forms the matrix $K = (\mathbf{k}_*(x_i, x_j))_{i,j=1}^n$ of pairwise kernel evaluations, finds a matrix square-root $V \in \mathbb{R}^{n \times n}$ satisfying $K = VV^\top$, and augments V with a row of ones (to encourage near-halving in the next step). Then, the PT construction runs the Gram-Schmidt (GS) walk of Bansal et al. (2018) to identify approximately half of the columns of V as coreset members and *rebalances* the coreset until exactly half of the input points belong to the coreset. The GS walk and rebalancing steps are recursively repeated $\Omega(\log(n))$ times to obtain an $(n^{\frac{1}{2}}, \mathcal{O}_p(\sqrt{dn}^{-\frac{1}{2}} \sqrt{\log n}))$ - L^∞ coreset. The low-dimensional Gaussian kernel construction of Tai (2020) first partitions the input points into balls of radius $2\sqrt{\log n}$ and then applies the PT construction separately to each ball. The result is an $(n^{\frac{1}{2}}, \mathcal{O}_p(2^d n^{-\frac{1}{2}} \sqrt{\log(d \log n)}))$ - L^∞ coreset with an additional superexponential $\Omega(d^{5d})$ running time dependence.

Acknowledgments

RD acknowledges support by National Science Foundation under Grant No. DMS2023528 for the Foundations of Data Science Institute (FODSI). The authors thank François-Xavier Briol, Lucas Janson, Lingxiao Li, Chris Oates, Art Owen, and the anonymous reviewers for their valuable feedback on this work and Jeffrey Rosenthal for helpful discussions surrounding geometric ergodicity. Part of this work was done when RD was interning at Microsoft Research New England.

References

- Ryan Alweiss, Yang P Liu, and Mehtaab Sawhney. Discrepancy minimization via a self-balancing walk. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 14–20, 2021.
- 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.
- Francis Bach, Simon Lacoste-Julien, and Guillaume Obozinski. On the equivalence between herding and conditional gradient algorithms. In *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML’12*, page 1355–1362, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.
- Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt walk: A cure for the Banaszczyk blues. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 587–597, 2018.
- Rémi Bardenet and Adrien Hardy. Monte Carlo with determinantal point processes. *The Annals of Applied Probability*, 30(1):368–417, 2020.
- 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>.
- William Beckner. Inequalities in Fourier analysis. *Annals of Mathematics*, pages 159–182, 1975.
- Ayoub Belhadji, Rémi Bardenet, and Pierre Chainais. Kernel quadrature with DPPs. In *Advances in Neural Information Processing Systems*, volume 32, pages 12927–12937. Curran Associates, Inc., 2019.
- Ayoub Belhadji, Rémi Bardenet, and Pierre Chainais. Kernel interpolation with continuous volume sampling. In *International Conference on Machine Learning*, pages 725–735. PMLR, 2020.
- Salomon Bochner. Monotone funktionen, stieltjessche integrale und harmonische analyse. *Mathematische Annalen*, 108(1):378–410, 1933.

- Sergiy V Borodachov, Douglas P Hardin, and Edward B Saff. Low complexity methods for discretizing manifolds via Riesz energy minimization. *Foundations of Computational Mathematics*, 14(6):1173–1208, 2014.
- Jonathan M Borwein and O-Yeat Chan. Uniform bounds for the complementary incomplete gamma function. *Mathematical Inequalities and Applications*, 12:115–121, 2009.
- François-Xavier Briol, Chris J. Oates, Mark Girolami, and Michael A. Osborne. Frank-Wolfe Bayesian Quadrature: Probabilistic integration with theoretical guarantees. In *Advances in Neural Information Processing Systems*, pages 1162–1170, 2015.
- Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov chain Monte Carlo*. CRC press, 2011.
- V. V. Buldygin and Yu. V. Kozachenko. Sub-gaussian random variables. *Ukrainian Mathematical Journal*, 32(6):483–489, 1980. doi: 10.1007/BF01087176. URL <https://doi.org/10.1007/BF01087176>.
- Trevor Campbell and Tamara Broderick. Automated scalable Bayesian inference via Hilbert coresets. *The Journal of Machine Learning Research*, 20(1):551–588, 2019.
- W WL Chen and MM Skriganov. Explicit constructions in the classical mean squares problem in irregularities of point distribution. *Journal für die Reine und Angewandte Mathematik*, (545):67–95, 2002.
- W. Y. Chen, L. Mackey, J. Gorham, F-X. Briol, and C. J. Oates. Stein points. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- 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.
- Stefano De Marchi, Robert Schaback, and Holger Wendland. Near-optimal data-independent point locations for radial basis function interpolation. *Advances in Computational Mathematics*, 23(3):317–330, 2005.
- DLMF. *NIST Digital Library of Mathematical Functions*. <http://dlmf.nist.gov/>, Release 1.1.1 of 2021-03-15. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds.
- Randal Douc, Eric Moulines, Pierre Priouret, and Philippe Soulier. *Markov chains*. Springer, 2018.
- Rick Durrett. *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 5 edition, 2019. doi: 10.1017/9781108591034.

- Raaz Dwivedi and Lester Mackey. Generalized kernel thinning. In *International Conference on Learning Representations*, 2022.
- 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.
- Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010.
- M. A. Gallegos-Herrada, D. Ledvinka, and J. S. Rosenthal. Equivalences of geometric ergodicity of markov chains, 2023.
- Damien Garreau, Wittawat Jitkrittum, and Motonobu Kanagawa. Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*, 2017.
- 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.
- Ronald L. Graham, Donald Ervin Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, Reading, MA, Second edition, 1994. URL <https://www.csie.ntu.edu.tw/~r97002/temp/Concrete%20Mathematics%20e.pdf>.
- 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.
- Karthik S Gurumoorthy, Amit Dhurandhar, Guillermo Cecchi, and Charu Aggarwal. Efficient data representation by selecting prototypes with importance weights. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 260–269. IEEE, 2019.
- 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.
- Antoine Havet, Matthieu Lerasle, Eric Moulines, and Elodie Vernet. A quantitative McDiarmid’s inequality for geometrically ergodic Markov chains. *Electronic Communications in Probability*, 25(none):1 – 11, 2020. doi: 10.1214/20-ECP286. URL <https://doi.org/10.1214/20-ECP286>.
- 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.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Jonathan Huggins and Lester Mackey. Random feature Stein discrepancies. In *Advances in Neural Information Processing Systems*, pages 1899–1909, 2018.
- Ferenc Huszár and David Kristjansson Duvenaud. Optimally-weighted herding is bayesian quadrature. *ArXiv*, abs/1408.2049, 2012.
- V Roshan Joseph, Tirthankar Dasgupta, Rui Tuo, and CF Jeff Wu. Sequential exploration of complex surfaces using minimum energy designs. *Technometrics*, 57(1):64–74, 2015.
- V Roshan Joseph, Dianpeng Wang, Li Gu, Shiji Lyu, and Rui Tuo. Deterministic sampling of expensive posteriors using minimum energy designs. *Technometrics*, 2019.
- Sarang Joshi, Raj Varma Kommaraji, Jeff M Phillips, and Suresh Venkatasubramanian. Comparing distributions and shapes using the kernel distance. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 47–56, 2011.
- Zohar Karnin and Edo Liberty. Discrepancy, coresets, and sketches in machine learning. In *Conference on Learning Theory*, pages 1975–1993. PMLR, 2019.
- Toni Karvonen, Chris J Oates, and Simo Sarkka. A Bayes-Sard cubature method. In *Advances in Neural Information Processing Systems*, pages 5882–5893, 2018.
- Toni Karvonen, Motonobu Kanagawa, and Simo Särkkä. On the positivity and magnitudes of Bayesian quadrature weights. *Statistics and Computing*, 29(6):1317–1333, 2019.
- Rajiv Khanna and Michael W Mahoney. On linear convergence of weighted kernel herding. *arXiv preprint arXiv:1907.08410*, 2019.
- Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. Examples are not enough, learn to criticize! Criticism for interpretability. *Advances in neural information processing systems*, 29, 2016.
- Vladimir Koltchinskii and Evarist Giné. Random matrix approximation of spectra of integral operators. *Bernoulli*, pages 113–167, 2000.
- Simon Lacoste-Julien, Fredrik Lindsten, and Francis Bach. Sequential kernel herding: Frank-Wolfe optimization for particle filtering. In *Artificial Intelligence and Statistics*, pages 544–552. PMLR, 2015.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Qiang Liu and Jason Lee. Black-box importance sampling. In *Artificial Intelligence and Statistics*, pages 952–961. PMLR, 2017.

- 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.
- Sean P Meyn and Richard L Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.
- Ha Quang Minh. Some properties of gaussian reproducing kernel hilbert spaces and their implications for function approximation and learning theory. *Constructive Approximation*, 32(2):307–338, 2010.
- Arunava Mukherjea. A remark on tonelli’s theorem on integration in product spaces. *Pacific Journal of Mathematics*, 42(1):177–185, 1972.
- 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.
- Chris J Oates, Mark Girolami, and Nicolas Chopin. Control functionals for Monte Carlo integration. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):695–718, 2017.
- Chris J Oates, Jon Cockayne, François-Xavier Briol, and Mark Girolami. Convergence rates for a class of estimators based on Stein’s method. *Bernoulli*, 25(2):1141–1159, 2019.
- Anthony O’Hagan. Bayes–Hermite quadrature. *Journal of statistical planning and inference*, 29(3):245–260, 1991.
- 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. ε -samples for kernels. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1622–1632. SIAM, 2013.
- Jeff M Phillips and Wai Ming Tai. Improved coresets for kernel density estimates. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2718–2727. SIAM, 2018.
- Jeff M Phillips and Wai Ming Tai. Near-optimal coresets of kernel density estimates. *Discrete & Computational Geometry*, 63(4):867–887, 2020.
- Alireza Rezaei and Shayan Oveis Gharan. A polynomial time mcmc method for sampling from continuous determinantal point processes. In *International Conference on Machine Learning*, pages 5438–5447. PMLR, 2019.

- 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, 2020. URL <https://doi.org/10.7910/DVN/MDKNWM>. Accessed on Mar 23, 2021.
- Marina Riabiz, Wilson Chen, Jon Cockayne, Pawel Swietach, Steven A Niederer, Lester Mackey, and Chris Oates. Optimal thinning of MCMC output. *To appear: Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2021.
- Gareth O Roberts and Jeffrey S Rosenthal. General state space Markov chains and MCMC algorithms. *Probability surveys*, 1:20–71, 2004.
- Gareth O Roberts and Richard L Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- Saburo Saitoh. Applications of the general theory of reproducing kernels. In *Reproducing Kernels and Their Applications*, pages 165–188. Springer, 1999.
- Gabriele Santin and Bernard Haasdonk. Convergence rate of the data-independent p-greedy algorithm in kernel-based approximation. *Dolomites Research Notes on Approximation*, 10(Special Issue), 2017.
- Larry Schumaker. *Spline functions: basic theory*. Cambridge University Press, 2007.
- Abhishek Shetty, Raaz Dwivedi, and Lester Mackey. Distribution compression in near-linear time. In *International Conference on Learning Representations*, 2022.
- Joel Spencer. Balancing games. *Journal of Combinatorial Theory, Series B*, 23(1):68–74, 1977.
- Bharath K Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert RG Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11(Apr):1517–1561, 2010.
- Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- Ingo Steinwart and Clint Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35(3):363–417, 2012.
- Marina Strocchi, 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.
- Xingping Sun. Conditionally positive definite functions and their application to multivariate interpolations. *Journal of approximation theory*, 74(2):159–180, 1993.

- Wai Ming Tai. New nearly-optimal coresets for kernel density estimation. *arXiv preprint arXiv:2007.08031*, 2020.
- 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.
- Paxton Turner, Jingbo Liu, and Philippe Rigollet. A statistical perspective on coresets density estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 2512–2520. PMLR, 2021.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, and SciPy 1.0 others. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 2020. doi: <https://doi.org/10.1038/s41592-019-0686-2>.
- 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. The covering number in learning theory. *Journal of Complexity*, 18(3): 739–767, 2002.