
Debiased Distribution Compression

Lingxiao Li¹ Raaz Dwivedi² Lester Mackey³

Abstract

Modern compression methods can summarize a target distribution \mathbb{P} more succinctly than i.i.d. sampling but require access to a low-bias input sequence like a Markov chain converging quickly to \mathbb{P} . We introduce a new suite of compression methods suitable for compression with biased input sequences. Given n points targeting the wrong distribution and quadratic time, Stein Kernel Thinning (SKT) returns \sqrt{n} equal-weighted points with $\tilde{O}(n^{-1/2})$ maximum mean discrepancy (MMD) to \mathbb{P} . For larger-scale compression tasks, Low-rank SKT achieves the same feat in sub-quadratic time using an adaptive low-rank debiasing procedure that may be of independent interest. For downstream tasks that support simplex or constant-preserving weights, Stein Recombination and Stein Cholesky achieve even greater parsimony, matching the guarantees of SKT with as few as poly-log(n) weighted points. Underlying these advances are new guarantees for the quality of simplex-weighted coresets, the spectral decay of kernel matrices, and the covering numbers of Stein kernel Hilbert spaces. In our experiments, our techniques provide succinct and accurate posterior summaries while overcoming biases due to burn-in, approximate Markov chain Monte Carlo, and tempering.

1. Introduction

Distribution compression is the problem of summarizing a target probability distribution \mathbb{P} with a small set of representative points. Such compact summaries are particularly valuable for tasks that incur substantial downstream computation costs per summary point, like organ and tissue

¹MIT CSAIL ²Cornell Tech ³Microsoft Research New England. Correspondence to: Lingxiao Li <lingxiao@mit.edu>, Raaz Dwivedi <dwivedi@cornell.edu>, Lester Mackey <lmackey@microsoft.com>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

modeling in which each simulation consumes thousands of CPU hours (Niederer et al., 2011).

Remarkably, modern compression methods can summarize a distribution more succinctly than i.i.d. sampling. For example, kernel thinning (KT) (Dwivedi and Mackey, 2021; 2022), Compress++ (Shetty et al., 2022), recombination (Hayakawa et al., 2023), and randomly pivoted Cholesky (Epperly and Moreno, 2024) all provide $\tilde{O}(1/m)$ approximation error using m points, a significant improvement over the $\Omega(1/\sqrt{m})$ approximation provided by i.i.d. sampling from \mathbb{P} . However, each of these constructions relies on access to an accurate input sequence, like an i.i.d. sample from \mathbb{P} or a Markov chain converging quickly to \mathbb{P} .

Much more commonly, one only has access to n *biased* sample points approximating a distribution $\mathbb{Q} \neq \mathbb{P}$. Such biases are a common occurrence in Markov chain Monte Carlo (MCMC)-based inference due to tempering (where one targets a less peaked and more dispersed distribution to achieve faster convergence, Gramacy et al., 2010), burn-in (where the initial state of a Markov chain biases the distribution of chain iterates, Cowles and Carlin, 1996), or approximate MCMC (where one runs a cheaper approximate Markov chain to avoid the prohibitive costs of an exact MCMC algorithm, e.g., Ahn et al., 2012). The Stein thinning (ST) method of Riabiz et al. (2022) was developed to provide accurate compression even when the input sample sequence provides a poor approximation to the target. ST operates by greedily thinning the input sample to minimize the maximum mean discrepancy (MMD, Gretton et al., 2012) to \mathbb{P} . However, ST is only known to provide an $O(1/\sqrt{m})$ approximation to \mathbb{P} ; this guarantee is no better than that of i.i.d. sampling and a far cry from the $\tilde{O}(1/m)$ error achieved with unbiased coreset constructions.

In this work, we address this deficit by developing new, efficient coreset constructions that provably yield better-than-i.i.d. error even when the input sample is biased. For \mathbb{P} on \mathbb{R}^d , our primary contributions are fourfold and summarized in Tab. 1. First, for the task of equal-weighted compression, we introduce *Stein Kernel Thinning* (SKT, Alg. 1), a strategy that combines the greedy bias correction properties of ST with the unbiased compression of KT to produce \sqrt{n} summary points with error $\tilde{O}(n^{-1/2})$ in $O(n^2)$ time. In contrast, ST would require $\Omega(n)$ points to guaran-

Table 1: **Methods for debiased distribution compression.** For each method, we report the smallest coreset size m and running time, up to logarithmic factors, sufficient to guarantee $\tilde{O}(n^{-1/2})$ $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}$ to \mathbb{P} given a LOGGROWTH kernel $\mathbf{k}_{\mathbb{P}}$ and n slow-growing input points $\mathcal{S}_n = (x_i)_{i=1}^n$ from a fast-mixing Markov chain targeting \mathbb{Q} with tails no lighter than \mathbb{P} (see Thm. 1 and Def. 3). For generic slow-growing \mathcal{S}_n , identical guarantees hold for excess $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}$ (2) relative to the best simplex reweighting of \mathcal{S}_n .

Method	Compression Type	Coreset Size m	Runtime	Source
Stein Thinning (Riabiz et al., 2022)	equal-weighted	n	$d_{\mathbf{k}_{\mathbb{P}}} n^2$	App. D.1
Stein Kernel Thinning	Greedy (Alg. 1)	equal-weighted	$d_{\mathbf{k}_{\mathbb{P}}} n^2$	Thm. 3
	Low-rank (Alg. 3)		$d_{\mathbf{k}_{\mathbb{P}}} n^{1.5}$	Thm. 5
Stein Recombination	Greedy (Alg. 5)	simplex-weighted	$d_{\mathbf{k}_{\mathbb{P}}} n^2$	Thm. 6
	Low-rank		$d_{\mathbf{k}_{\mathbb{P}}} n + n^{1.5}$	
Stein Cholesky	Greedy (Alg. 7)	constant-preserving	$d_{\mathbf{k}_{\mathbb{P}}} n^2$	Thm. 7
	Low-rank		$d_{\mathbf{k}_{\mathbb{P}}} n + n^{1.5}$	

tee this error. Second, for larger-scale compression problems, we propose *Low-rank SKT* (Alg. 3), a strategy that combines the scalable summarization of Compress++ with a new low-rank debiasing procedure (Alg. 2) to match the SKT guarantees in sub-quadratic $o(n^2)$ time.

Third, for the task of simplex-weighted compression, in which summary points are accompanied by weights in the simplex, we propose greedy and low-rank *Stein Recombination* (Alg. 5) constructions that match the guarantees of SKT with as few as $\text{poly-log}(n)$ points. Finally, for the task of constant-preserving compression, in which summary points are accompanied by real-valued weights summing to 1, we introduce greedy and low-rank *Stein Cholesky* (Alg. 7) constructions that again match the guarantees of SKT using as few as $\text{poly-log}(n)$ points.

Underlying these advances are new guarantees for the quality of simplex-weighted coresets (Thms. 1 and 2), the spectral decay of kernel matrices (Cor. B.1), and the covering numbers of Stein kernel Hilbert spaces (Prop. 1) that may be of independent interest. In Sec. 5, we employ our new procedures to produce compact summaries of complex target distributions given input points biased by burn-in, approximate MCMC, or tempering.

Notation We assume Borel-measurable sets and functions and define $[n] \triangleq \{1, \dots, n\}$, $\Delta_{n-1} \triangleq \{w \in \mathbb{R}^n : w \geq 0, \mathbf{1}^\top w = 1\}$, $\|x\|_0 \triangleq |\{i : x_i \neq 0\}|$, and $\|x\|_p^p \triangleq \sum_i |x_i|^p$ for $x \in \mathbb{R}^d$ and $p \geq 1$. For $x \in \mathbb{R}^d$, δ_x denotes the delta measure at x . We let $\mathcal{H}_{\mathbf{k}}$ denote the reproducing kernel Hilbert space (RKHS) of a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ (Aronszajn, 1950) and $\|f\|_{\mathbf{k}}$ denote the RKHS norm of $f \in \mathcal{H}_{\mathbf{k}}$. For a measure μ and separately μ -integrable \mathbf{k} and f , we write $\mu f \triangleq \int f(x) d\mu(x)$ and $\mu \mathbf{k}(x) \triangleq \int \mathbf{k}(x, y) d\mu(y)$. The divergence of a differentiable matrix-valued function A is $(\nabla_x \cdot A(x))_j = \sum_i \partial_{x_i} A_{ij}(x)$. For random variables

$(X_n)_{n \in \mathbb{N}}$, we say $X_n = O(f(n, \delta))$ holds with probability $\geq 1 - \delta$ if $\Pr(X_n \leq C f(n, \delta)) \geq 1 - \delta$ for a constant C independent of (n, δ) and all n sufficiently large. When using this notation, we view all algorithm parameters except δ as functions of n . For $A \in \mathbb{R}^{n \times n}$ and $v \in \mathbb{R}^n$, $\text{diag}(A)$ and $\text{diag}(v)$ are $n \times n$ diagonal matrices with A_{ii} and v_i respectively as the i -th diagonal entry.

2. Debiased Distribution Compression

Throughout, we aim to summarize a fixed target distribution \mathbb{P} on \mathbb{R}^d using a sequence $\mathcal{S}_n \triangleq (x_i)_{i=1}^n$ of potentially biased candidate points in \mathbb{R}^d .¹ Correcting for unknown biases in \mathcal{S}_n requires some auxiliary knowledge of \mathbb{P} . For us, this knowledge comes in the form of a kernel function $\mathbf{k}_{\mathbb{P}}$ with known expectation under \mathbb{P} . Without loss of generality, we can take this kernel mean to be identically zero.²

Assumption 1 (Mean-zero kernel). *For some $p \geq 1/2$, $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)^p] < \infty$ and $\mathbb{P} \mathbf{k}_{\mathbb{P}} \equiv 0$.*

Given a target compression size m , our goal is to output an weight vector $w \in \mathbb{R}^n$ with $\|w\|_0 \leq m$, $\mathbf{1}_n^\top w = 1$, and $o(m^{-1/2})$ (better-than-i.i.d.) maximum mean discrepancy (MMD) to \mathbb{P} :

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\sum_{i=1}^n w_i \delta_{x_i}, \mathbb{P}) \triangleq \sqrt{\sum_{i,j=1}^n w_i w_j \mathbf{k}_{\mathbb{P}}(x_i, x_j)}.$$

We consider three standard compression tasks with $\|w\|_0 \leq m$. In *equal-weighted compression* one selects m possibly repeated points from \mathcal{S}_n and assigns each a weight of $\frac{1}{m}$; because of repeats, the induced weight vector

¹Our coreset constructions will in fact apply to any sample space, but our analysis will focus on \mathbb{R}^d .

²For $\mathbb{P} \mathbf{k}_{\mathbb{P}} \neq 0$, the kernel $\mathbf{k}_{\mathbb{P}'}(x, y) = \mathbf{k}_{\mathbb{P}}(x, y) - \mathbb{P} \mathbf{k}_{\mathbb{P}}(x) - \mathbb{P} \mathbf{k}_{\mathbb{P}}(y) + \mathbb{P} \mathbb{P} \mathbf{k}_{\mathbb{P}}$ satisfies $\mathbb{P} \mathbf{k}_{\mathbb{P}'} \equiv 0$ and $\text{MMD}_{\mathbf{k}_{\mathbb{P}'}} = \text{MMD}_{\mathbf{k}_{\mathbb{P}}}$.

over \mathcal{S}_n satisfies $w \in \Delta_{n-1} \cap (\frac{\mathbb{N}_0}{m})^n$. In *simplex-weighted compression* we allow any $w \in \Delta_{n-1}$, and in *constant-preserving compression* we simply enforce $\mathbf{1}_n^\top w = 1$. All three coresets types exactly preserve constants and hence satisfy standard coherence constraints.

When making big-O statements, we will treat \mathcal{S}_n as the prefix of an infinite sequence $\mathcal{S}_\infty \triangleq (x_i)_{i \in \mathbb{N}}$. We also write $\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n[\mathbf{J}], \mathcal{S}_n[\mathbf{J}]) \triangleq [\mathbf{k}_{\mathbb{P}}(x_i, x_j)]_{i,j \in \mathbf{J}}$ for the principal kernel submatrix with indices $\mathbf{J} \subseteq [n]$.

2.1. Kernel assumptions

Many practical *Stein kernel* constructions are available for generating mean-zero kernels for a target \mathbb{P} (Chwialkowski et al., 2016; Liu et al., 2016; Gorham and Mackey, 2017; Gorham et al., 2019; Barp et al., 2019; Yang et al., 2018; Afzali and Muthukumarana, 2023). We will use the most prominent of these Stein kernels as a running example:

Definition 1 (Stein kernel). *Given a differentiable base kernel \mathbf{k} and a symmetric positive semidefinite matrix M , the Stein kernel $\mathbf{k}_p : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ for \mathbb{P} with positive differentiable Lebesgue density p is defined as*

$$\mathbf{k}_p(x, y) \triangleq \frac{1}{p(x)p(y)} \nabla_x \cdot \nabla_y \cdot (p(x)M\mathbf{k}(x, y)p(y)).$$

While our algorithms apply to any mean zero kernel, including the aforementioned Stein kernels, the Sobolev kernels used in quasi-Monte Carlo (Kuo, 2003), and the popular centered Gaussian kernel (Chen et al., 2010; Lacoste-Julien et al., 2015), our guarantees adapt to the underlying smoothness of the kernel. Our next definition and assumption make this precise.

Definition 2 (Covering number). *For a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ with $\mathcal{B}_{\mathbf{k}} \triangleq \{f \in \mathcal{H}_{\mathbf{k}} : \|f\|_{\mathbf{k}} \leq 1\}$, a set $A \subset \mathbb{R}^d$, and $\varepsilon > 0$, the covering number $\mathcal{N}_{\mathbf{k}}(A, \varepsilon)$ is the minimum cardinality of all sets $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}}$ satisfying*

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

Assumption (α, β)-kernel. *For some $\mathfrak{C}_d > 0$, all $r > 0$ and $\varepsilon \in (0, 1)$, and $\mathcal{B}_2(r) \triangleq \{x \in \mathbb{R}^d : \|x\|_2 \leq r\}$, a kernel \mathbf{k} is either POLYGROWTH(α, β), i.e.,*

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \varepsilon) \leq \mathfrak{C}_d (1/\varepsilon)^\alpha (r+1)^\beta,$$

with $\alpha < 2$ or LOGGROWTH(α, β), i.e.,

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \varepsilon) \leq \mathfrak{C}_d \log(e/\varepsilon)^\alpha (r+1)^\beta.$$

In Cor. B.1 we show that the eigenvalues of kernel matrices with POLYGROWTH and LOGGROWTH kernels have polynomial and exponential decay respectively. Dwivedi and Mackey (2022, Prop. 2) showed that all sufficiently differentiable kernels satisfy the POLYGROWTH condition and

that bounded radially analytic kernels are LOGGROWTH. Our next result, proved in App. B.2, shows that a Stein kernel \mathbf{k}_p can inherit the growth properties of its base kernel even if \mathbf{k}_p is itself unbounded and non-smooth.

Proposition 1 (Stein kernel growth rates). *A Stein kernel \mathbf{k}_p with $\sup_{\|x\|_2 \leq r} \|\nabla \log p(x)\|_2 = O(r^{d_\ell})$ for $d_\ell \geq 0$ is*

- (a) LOGGROWTH($d+1, 2d+\delta$) for any $\delta > 0$ if the base kernel \mathbf{k} is radially analytic (Def. B.3) and
- (b) POLYGROWTH($\frac{d}{s-1}, (1+\frac{d_\ell}{s})d$) if the base kernel \mathbf{k} is s -times continuously differentiable (Def. B.2) for $s > 1$.

Notably, the popular Gaussian (Ex. B.1) and inverse multiquadric (Ex. B.2) base kernels satisfy the LOGGROWTH preconditions, while Matérn, B-spline, sinc, sech, and Wendland’s compactly supported kernels satisfy the POLYGROWTH precondition (Dwivedi and Mackey, 2022, Prop. 3). To our knowledge, Prop. 1 and Cor. B.1 provide the first covering number bounds and eigenvalue decay rates for the typically unbounded Stein kernels \mathbf{k}_p .

2.2. Input point desiderata

Our primary desideratum for the input points is that they can be debiased into an accurate estimate of \mathbb{P} . Indeed, our high-level strategy for debiased compression is to first use $\mathbf{k}_{\mathbb{P}}$ to debias the input points into a more accurate approximation of \mathbb{P} and then compress that approximation into a more succinct representation. Fortunately, even when the input \mathcal{S}_n targets a distribution $\mathbb{Q} \neq \mathbb{P}$, effective debiasing is often achievable via simplex reweighting, i.e., by solving the convex optimization problem

$$w_{\text{OPT}} \in \arg \min_{w \in \Delta_{n-1}} \sum_{i,j=1}^n w_i w_j \mathbf{k}_{\mathbb{P}}(x_i, x_j) \quad (1)$$

with $\text{MMD}_{\text{OPT}} \triangleq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\sum_{i=1}^n w_{\text{OPT}i} \delta_{x_i}, \mathbb{P})$.

For example, Hodgkinson et al. (2020, Thm. 1b) showed that simplex reweighting can correct for biases due to off-target i.i.d. or MCMC sampling. Our next result (proved in App. C.2) significantly relaxes their conditions.

Theorem 1 (Debiasing to i.i.d. quality via simplex reweighting). *Consider a kernel $\mathbf{k}_{\mathbb{P}}$ satisfying Assum. 1 with $\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ separable, and suppose $(x_i)_{i=1}^\infty$ are the iterates of a homogeneous ϕ -irreducible geometrically ergodic Markov chain (Gallegos-Herrada et al., 2023, Thm. 1) with stationary distribution \mathbb{Q} and initial distribution absolutely continuous with respect to \mathbb{P} . If $\mathbb{E}_{x \sim \mathbb{P}}[\frac{d\mathbb{P}}{d\mathbb{Q}}(x)^{2q-1} \mathbf{k}_{\mathbb{P}}(x, x)^q] < \infty$ for some $q > 1$ then $\text{MMD}_{\text{OPT}} = O(n^{-1/2})$ in probability.*

Remark 1. $\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ is separable whenever $\mathbf{k}_{\mathbb{P}}$ is continuous (Steinwart and Christmann, 2008, Lem. 4.33).

Since n points sampled i.i.d. from \mathbb{P} have $\Theta(n^{-1/2})$ root mean squared MMD (see Prop. C.1), Thm. 1 shows that

a debiased off-target sample can be as accurate as a direct sample from \mathbb{P} . Moreover, Thm. 1 applies to many practical examples. The simplest example of a geometrically ergodic chain is i.i.d. sampling from \mathbb{Q} , but geometric ergodicity has also been established for a variety of popular Markov chains including random walk Metropolis (Roberts and Tweedie, 1996, Thm. 3.2), independent Metropolis-Hastings (Atchadé and Perron, 2007, Thm. 2.2), the unadjusted Langevin algorithm (Durmus and Éric Moulines, 2017, Prop. 8), the Metropolis-adjusted Langevin algorithm (Durmus and Moulines, 2022, Thm. 1), Hamiltonian Monte Carlo (Durmus et al., 2020, Thm. 10 and Thm. 11), stochastic gradient Langevin dynamics (Li et al., 2023, Thm. 2.1), and the Gibbs sampler (Johnson, 2009). Moreover, for \mathbb{Q} absolutely continuous with respect to \mathbb{P} , the importance weight $\frac{d\mathbb{P}}{d\mathbb{Q}}$ is typically bounded or slowly growing when the tails of \mathbb{Q} are not much lighter than those of \mathbb{P} .

Remarkably, under more stringent conditions, Thm. 2 (proved in App. C.3) shows that simplex reweighting can decrease MMD to \mathbb{P} at an even-faster-than-i.i.d. rate.

Theorem 2 (Better-than-i.i.d. debiasing via simplex reweighting). *Consider a kernel $k_{\mathbb{P}}$ satisfying Assum. 1 with $p = 2$ and points $(x_i)_{i=1}^{\infty}$ drawn i.i.d. from a distribution \mathbb{Q} with $\frac{d\mathbb{P}}{d\mathbb{Q}}$ bounded. If $\mathbb{E}[k_{\mathbb{P}}(x_1, x_1)^q] < \infty$ for some $q > 3$, then $\mathbb{E}[\text{MMD}_{\text{OPT}}^2] = o(n^{-1})$.*

The work of Liu and Lee (2017, Thm. 3.3) also established $o(n^{-1/2})$ MMD error for simplex reweighting but only under a uniformly bounded eigenfunctions assumption that is often violated (Minh, 2010, Thm. 1, Zhou, 2002, Ex. 1) and difficult to verify (Steinwart and Scovel, 2012).

We highlight that our remaining results make no particular assumption about the input points but rather upper bound the excess MMD

$$\Delta \text{MMD}_{k_{\mathbb{P}}}(w) \triangleq \text{MMD}_{k_{\mathbb{P}}}(\sum_{i \in [n]} w_i \delta_{x_i}, \mathbb{P}) - \text{MMD}_{\text{OPT}} \quad (2)$$

of a candidate weighting w in terms of the input point radius $R_n \triangleq \max_{i \in [n]} \|x_i\|_2 \vee 1$ and kernel radius $\|k_{\mathbb{P}}\|_n \triangleq \max_{i \in [n]} k_{\mathbb{P}}(x_i, x_i)$. While these results apply to any input points, we consider the following running example of slow-growing input points throughout the paper.

Definition 3 (Slow-growing input points). *We say \mathcal{S}_n is γ -slow-growing if $R_n = O((\log n)^\gamma)$ for some $\gamma \geq 0$ and $\|k_{\mathbb{P}}\|_n = \tilde{O}(1)$.*

Notably, \mathcal{S}_n is 1-slow-growing with probability 1 when $k_{\mathbb{P}}(x, x)$ is polynomially bounded by $\|x\|_2$ and the input points are drawn from a homogeneous ϕ -irreducible geometrically ergodic Markov chain with a sub-exponential target \mathbb{Q} , i.e., $\mathbb{E}[e^{c\|x\|_2}] < \infty$ for some $c > 0$ (Dwivedi and Mackey, 2021, Prop. 2). For a Stein kernel k_p (Def. 1),

by Prop. B.3, $k_p(x, x)$ is polynomially bounded by $\|x\|_2$ if $k(x, x)$, $\|\nabla_x \nabla_y k(x, x)\|_2$, and $\|\nabla \log p(x)\|_2$ are all polynomially bounded by $\|x\|_2$. Moreover, $\|\nabla \log p(x)\|_2$ is automatically polynomially bounded by $\|x\|_2$ when $\nabla \log p$ is Lipschitz or, more generally, pseudo-Lipschitz (Erdogdu et al., 2018, Eq. (2.5)).

2.3. Debiased compression via Stein Kernel Thinning

Off-the-shelf solvers based on mirror descent and Frank Wolfe can solve the convex debiasing program (1) in $O(n^3)$ time by generating weights with $O(n^{-1/2} \|k_{\mathbb{P}}\|_n)$ excess MMD (Liu and Lee, 2017). We instead employ a more efficient, greedy debiasing strategy based on Stein thinning (ST). After n rounds, ST outputs an equal-weighted coreset of size n with $O(n^{-1/2} \|k_{\mathbb{P}}\|_n)$ excess MMD (Riabiz et al., 2022, Thm. 1). Moreover, while the original implementation of Riabiz et al. (2022) has cubic runtime, our implementation (Alg. D.1) based on sufficient statistics improves the runtime to $O(n^2 d_{k_{\mathbb{P}}})$ where $d_{k_{\mathbb{P}}}$ denotes the runtime of a single kernel evaluation.³

The equal-weighted output of ST serves as the perfect input for the kernel thinning (KT) algorithm which compresses an equal-weighted sample of size n into a coreset of any target size $m \leq n$ in $O(n^2 d_{k_{\mathbb{P}}})$ time. We adapt the target KT algorithm slightly to target MMD error to \mathbb{P} and to include a baseline ST coreset of size m in the KT-SWAP step (see Alg. D.3). Combining the two routines we obtain Stein Kernel Thinning (SKT), our first solution for equal-weighted debiased distribution compression:

Algorithm 1 Stein Kernel Thinning (SKT)

Input: mean-zero kernel $k_{\mathbb{P}}$, points \mathcal{S}_n , output size m , KT failure probability δ
 $n' \leftarrow m 2^{\lceil \log_2 \frac{m}{m} \rceil}$
 $w \leftarrow \text{SteinThinning}(k_{\mathbb{P}}, \mathcal{S}_n, n')$
 $w_{\text{SKT}} \leftarrow \text{KernelThinning}(k_{\mathbb{P}}, \mathcal{S}_n, n', w, m, \delta)$
Return: $w_{\text{SKT}} \in \Delta_{n-1} \cap \binom{[n]}{m}^n \triangleright \text{hence } \|w_{\text{SKT}}\|_0 \leq m$

Our next result, proved in App. D.3, shows that SKT yields better-than-i.i.d. excess MMD whenever the radii (R_n and $\|k_{\mathbb{P}}\|_n$) and kernel covering number exhibit slow growth.

Theorem 3 (MMD guarantee for SKT). *Given a kernel $k_{\mathbb{P}}$ satisfying Assums. 1 and (α, β) -kernel, Stein Kernel Thinning (Alg. 1) outputs w_{SKT} in $O(n^2 d_{k_{\mathbb{P}}})$ time satisfying*

$$\Delta \text{MMD}_{k_{\mathbb{P}}}(w_{\text{SKT}}) = O\left(\frac{\sqrt{\|k_{\mathbb{P}}\|_n \ell_\delta \cdot \log n \cdot R_n^\beta G_m^\alpha}}{\min(m, \sqrt{n})}\right)$$

with probability at least $1 - \delta$, where $\ell_\delta \triangleq \log^2(\frac{e}{\delta})$ and

$$G_m \triangleq \begin{cases} 1 + \log m & \text{LOGGROWTH}(\alpha, \beta), \\ m & \text{POLYGROWTH}(\alpha, \beta). \end{cases}$$

³Often, $d_{k_{\mathbb{P}}} = \Theta(d)$ as in the case of Stein kernels (App. I.1).

Example 1. Under the assumptions of Thm. 3 with γ -slow-growing input points (Def. 3), LOGGROWTH $\mathbf{k}_{\mathbb{P}}$, and a coresets size $m \leq \sqrt{n}$, SKT with high probability delivers $\tilde{O}(m^{-1})$ excess MMD, a nearly minimax optimal rate for equal-weighted coresets (Phillips and Tai, 2020, Thm. 3.1) and a significant improvement over the $\Omega(m^{-1/2})$ error rate of i.i.d. sampling.

Remark 2. When $m < \sqrt{n}$, we can uniformly subsample or, in the case of MCMC inputs, *standard thin* (i.e., keep only every $\frac{n}{m^2}$ -th point of) the input sequence down to size m^2 before running SKT to reduce runtime from $O(n^2)$ to $O(m^4)$ while incurring only $O(m^{-1})$ excess error. A similar remark applies to LSKT algorithm introduced in Sec. 3.

3. Accelerated Debiased Compression

To enable larger-scale debiased compression, we next introduce a sub-quadratic-time version of SKT built via a new low-rank debiasing scheme and the near-linear-time compression algorithm of Shetty et al. (2022).

3.1. Fast bias correction via low-rank approximation

At a high level, our approach to accelerated debiasing involves four components. First, we form a rank- r approximation FF^\top of the kernel matrix $K = \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)$ in $O(nrd_{\mathbf{k}_{\mathbb{P}}} + nr^2)$ time using a weighted extension (WeightedRPCholesky, Alg. F.1) of the randomly pivoted Cholesky algorithm of Chen et al. (2022, Alg. 2.1). Second, we correct the diagonal to form $K' = FF^\top + \text{diag}(K - FF^\top)$. Third, we solve the reweighting problem (1) with K' substituted for K using T iterations of accelerated entropic mirror descent (AMD, Wang et al., 2023, Alg. 14 with $\phi(w) = \sum_i w_i \log w_i$). The acceleration ensures $O(1/T^2)$ suboptimality after T iterations, and each iteration takes only $O(nr)$ time thanks to the low-rank plus diagonal approximation. Finally, we repeat this three-step procedure Q times, each time using the weights outputted by the prior round to update the low-rank approximation \hat{K} . On these subsequent adaptive rounds, WeightedRPCholesky approximates the leading subspace of a *weighted* kernel matrix $\text{diag}(\sqrt{\tilde{w}})K \text{diag}(\sqrt{\tilde{w}})$ for $\tilde{w} \in \Delta_{n-1}$ before undoing the row and column reweighting. Since each round's weights are closer to optimal, this adaptive updating has the effect of upweighting more relevant subspaces for subsequent debiasing. For added sparsity, we prune the weights outputted by the prior round using stratified residual resampling (Resample, Alg. E.3, Douc and Cappé, 2005). Our complete Low-rank Debiasing (LD) scheme, summarized in Alg. 2, enjoys $o(n^2)$ runtime whenever $r = o(n^{1/2})$, $T = O(n^{1/2})$, and $Q = O(1)$.

Moreover, our next result, proved in App. F.1, shows that LD provides i.i.d.-level precision whenever $T \geq \sqrt{n}$, $Q =$

Algorithm 2 Low-rank Debiasing (LD)

Input: mean-zero kernel $\mathbf{k}_{\mathbb{P}}$, points $\mathcal{S}_n = (x_i)_{i=1}^n$, rank r , AMD steps T , adaptive rounds Q
 $w^{(0)} \leftarrow (\frac{1}{n}, \dots, \frac{1}{n}) \in \mathbb{R}^n$
for $q = 1$ **to** Q **do**
 $\tilde{w} \leftarrow \text{Resample}(w^{(q-1)}, n)$
 $I, F \leftarrow \text{WeightedRPCholesky}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, \tilde{w}, r)$
 $K' \leftarrow FF^\top + \text{diag}(\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)) - \text{diag}(FF^\top)$
 $w^{(q)} \leftarrow \text{AMD}(K', T, \tilde{w}, \text{AGG} = \mathbb{1}_{q>1})$
if $(w^{(q)})^\top K' w^{(q)} > \tilde{w}^\top K' \tilde{w}$ **then** $w^{(q)} \leftarrow \tilde{w}$
end for
Return: $w_{\text{LD}} \leftarrow w^{(Q)} \in \Delta_{n-1}$

$O(1)$, and r grows appropriately with the input radius and kernel covering number.

Assumption (α, β)-params. *The kernel $\mathbf{k}_{\mathbb{P}}$ satisfies Assums. 1 and (α, β)-kernel, the output size and rank $m, r \geq (\frac{\mathfrak{c}_d R_n^\beta + 1}{\sqrt{\log 2}} + 2\sqrt{\log 2})^2$, the AMD step count $T \geq \sqrt{n}$, and the adaptive round count $Q = O(1)$.*⁴

Theorem 4 (Debiasing guarantee for LD). *Under Assum. (α, β)-params, Low-rank Debiasing (Alg. 2) takes $O((d_{\mathbf{k}_{\mathbb{P}}} + r + T)nr)$ time to output w_{LD} satisfying*

$$\Delta \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{LD}}) = O\left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n \max(\log n, 1/\delta)}{n}} + \sqrt{\frac{nH_{n,r}}{\delta}}\right)$$

with probability at least $1 - \delta$, for any $\delta \in (0, 1)$ and $H_{n,r}$ defined in (46) that satisfies

$$H_{n,r} = \begin{cases} O\left(\sqrt{r} \left(\frac{R_n^{2\beta}}{r}\right)^{\frac{1}{\alpha}}\right) & \text{POLYGROWTH}(\alpha, \beta), \\ O\left(\sqrt{r} \exp\left(-\left(\frac{0.83\sqrt{r} - 2.39}{\mathfrak{c}_d R_n^\beta}\right)^{\frac{1}{\alpha}}\right)\right) & \text{LOGGROWTH}(\alpha, \beta). \end{cases}$$

Example 2. Under the assumptions of Thm. 4 with γ -slow-growing input points (Def. 3), LOGGROWTH $\mathbf{k}_{\mathbb{P}}$, $T = \Theta(\sqrt{n})$, and $r = (\log n)^{2(\alpha+\beta\gamma)+\epsilon}$ for any $\epsilon > 0$, LD delivers $\tilde{O}(n^{-1/2})$ excess MMD with high probability in $\tilde{O}(n^{1.5})$ time.

3.2. Fast debiased compression via Low-rank Stein KT

To achieve debiased compression in sub-quadratic time, we next propose Low-rank SKT (Alg. 3). LSKT debiases the input using LD, converts the LD output into an equal-weighted coresets using Resample, and finally combines KT with the divide-and-conquer Compress++ framework (Shetty et al., 2022) to compress n equal-weighted points into \sqrt{n} in near-linear time.

Our next result (proved in App. F) shows that LSKT can provide better-than-i.i.d. excess MMD in $o(n^2)$ time.

⁴To unify the presentation of our results, Assum. (α, β)-params constrains all common algorithm input parameters with the understanding that the conditions are enforced only when the input is relevant to a given algorithm.

Algorithm 3 Low-rank Stein Kernel Thinning (LSKT)

Input: mean-zero kernel $k_{\mathbb{P}}$, points $\mathcal{S}_n = (x_i)_{i=1}^n$, rank r , AGM steps T , adaptive rounds Q , oversampling parameter g , failure prob. δ

$w \leftarrow \text{Low-rankDebiasing}(k_{\mathbb{P}}, \mathcal{S}_n, r, T, Q)$
 $n' \leftarrow 4^{\lceil \log_4 n \rceil}$, $m \leftarrow \sqrt{n'} \triangleright \text{output size } \sqrt{n} \leq m < 2\sqrt{n}$
 $w \leftarrow \text{Resample}(w, n')$
 $w_{\text{LSKT}} \leftarrow \text{KT-Compress++}(k_{\mathbb{P}}, \mathcal{S}_n, n', w, g, \frac{\delta}{3})$

Return: $w_{\text{LSKT}} \in \Delta_{n-1} \cap (\frac{N_0}{m})^n \triangleright \text{hence } \|w_{\text{LSKT}}\|_0 \leq m$

Theorem 5 (MMD guarantee for LSKT). *Under Assum. (α, β) -params, Low-rank SKT (Alg. 3) with $g \in [\log_2 \log(n+1) + 3.1, \log_4(\sqrt{n}/\log n)]$ and $\delta \in (0, 1)$ outputs w_{LSKT} in $O((d_{k_{\mathbb{P}}} + r + T)nr + d_{k_{\mathbb{P}}}n^{1.5})$ time satisfying, with probability at least $1 - \delta$,*

$$\begin{aligned} & \Delta_{\text{MMD}}_{k_{\mathbb{P}}}(w_{\text{LSKT}}) \\ &= O\left(\sqrt{\frac{\|k_{\mathbb{P}}\|_n \max(1/\delta, \ell_{\delta}(\log n)n^{\gamma\beta}G_{\sqrt{n}}^{\alpha})}{n}} + \sqrt{\frac{nH_{n,r}}{\delta}}\right), \end{aligned}$$

for $G_m, H_{n,r}$ as in Thms. 3 and 5.

Example 3. Under the assumptions of Thm. 5 with γ -slow-growing input points (Def. 3), LOGGROWTH $k_{\mathbb{P}}, T = \Theta(\sqrt{n})$, and $r = (\log n)^{2(\alpha+\beta\gamma)+\epsilon}$ for any $\epsilon > 0$, LSKT delivers, with high probability, $\tilde{O}(n^{-1/2})$ excess MMD in $\tilde{O}(n^{1.5})$ time using a coreset of size $m \in [\sqrt{n}, 2\sqrt{n})$.

4. Weighted Debiased Compression

The prior sections developed debiased equal-weighted coresets with better-than-i.i.d. compression guarantees. Equal-weighted coresets are typically chosen for their compatibility with unweighted downstream tasks and easy visualization. For downstream tasks that support weights, we next match the equal-weighted guarantees with significantly smaller simplex-weighted or constant-preserving coresets.

4.1. Simplex-weighted coresets via Stein Recombination

Simplex-weighted coresets automatically preserve the constraints of convex integrands, support straightforward direct sampling, and, when compared with schemes involving negative weights, offer improved numerical stability in the presence of integral evaluations errors (Karvonen et al., 2019). Inspired by the coreset constructions of Hayakawa et al. (2022; 2023), we first introduce a simplex-weighted compression algorithm, **RecombinationThinning** (RT, Alg. 4), suitable for summarizing a debiased input sequence. To produce a coreset given input weights $w \in \Delta_{n-1}$, RT first prunes small weights using **Resample** and then uses **WeightedRPCholesky** to identify $m-1$ test vectors that capture most of the variability in the weighted

Algorithm 4 Recombination Thinning (RT)

Input: mean-zero kernel $k_{\mathbb{P}}$, points $\mathcal{S}_n = (x_i)_{i=1}^n$, weights $w \in \Delta_{n-1}$, output size m

$\tilde{w} \leftarrow \text{Resample}(w, n)$
 $I, F \leftarrow \text{WeightedRPCholesky}(k_{\mathbb{P}}, \mathcal{S}_n, \tilde{w}, m-1)$
 $w' \leftarrow \text{Recombination}([F, \mathbf{1}_n]^{\top}, \tilde{w}) \triangleright [F, \mathbf{1}_n]^{\top} \in \mathbb{R}^{m \times n}$
 $\triangleright F^{\top} \tilde{w} = F^{\top} w', w' \in \Delta_{n-1}$, and $\|w'\|_0 \leq m$
 $w'' \leftarrow \text{KT-Swap-LS}(k_{\mathbb{P}}, \mathcal{S}_n, w', \text{SPLX}); J \leftarrow \{i: w''_i > 0\}$
 $w''[J] \leftarrow \text{argmin}_{w' \in \Delta_{|J|-1}} w'^{\top} k_{\mathbb{P}}(\mathcal{S}_n[J], \mathcal{S}_n[J])w' \triangleright \text{use any } O(|J|^3) \text{ quadratic programming solver}$

Return: $w_{\text{RT}} \leftarrow w'' \in \Delta_{n-1}$ with $\|w_{\text{RT}}\|_0 \leq m$

kernel matrix. Next, **Recombination** (Alg. G.1) (Tchernychova, 2016, Alg. 1) identifies a sparse simplex vector w' with $\|w'\|_0 \leq m$ that exactly matches the inner product of its input with each of the test vectors. Then, we run **KT-Swap-LS** (Alg. G.2), a new, line-search version of KT-SWAP (Dwivedi and Mackey, 2021, Alg. 1b) that greedily improves MMD to \mathbb{P} while maintaining both the sparsity and simplex constraint of its input. Finally, we optimize the weights of the remaining support points using any cubic-time quadratic programming solver.

In Prop. G.1 we show that RT runs in time $O((d_{k_{\mathbb{P}}} + m)nm + m^3 \log n)$ and nearly preserves the MMD of its input whenever m grows appropriately with the kernel covering number. Combining RT with **SteinThinning** or **Low-rankDebiasing** in Alg. 5, we obtain Stein Recombination (SR) and Low-rank SR (LSR), our approaches to debiased simplex-weighted compression. Remarkably, SR and LSR can match the MMD error rates established for SKT and LSKT using substantially fewer coreset points, as our next result (proved in App. G.2) shows.

Algorithm 5 (Low-rank) Stein Recombination (SR / LSR)

Input: mean-zero kernel $k_{\mathbb{P}}$, points \mathcal{S}_n , output size m , rank r , AGM steps T , adaptive rounds Q

$w \leftarrow \begin{cases} \text{Low-rankDebiasing}(k_{\mathbb{P}}, \mathcal{S}_n, r, T, Q) & \text{if low-rank} \\ \text{SteinThinning}(k_{\mathbb{P}}, \mathcal{S}_n) & \text{otherwise} \end{cases}$

$w_{\text{SR}} \leftarrow \text{RecombinationThinning}(k_{\mathbb{P}}, \mathcal{S}_n, w, m)$

Return: $w_{\text{SR}} \in \Delta_{n-1}$ with $\|w_{\text{SR}}\|_0 \leq m$

Theorem 6 (MMD guarantee for SR/LSR). *Under Assum. (α, β) -params, Stein Recombination (Alg. 5) takes $O(d_{k_{\mathbb{P}}}n^2 + (d_{k_{\mathbb{P}}} + m)nm + m^3 \log n)$ to output w_{SR} , and Low-rank SR takes $O((d_{k_{\mathbb{P}}} + r + T)nr + (d_{k_{\mathbb{P}}} + m)nm + m^3 \log n)$ time to output w_{LSR} . Moreover, for any $\delta \in (0, 1)$ and $H_{n,r}$ as in Thm. 4, each of the following bounds holds (sepa-*

rately) with probability at least $1 - \delta$:

$$\Delta\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{SR}}) = O\left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n(\log n \vee \frac{1}{\delta})}{n} + \frac{nH_{n,m}}{\delta}}\right) \text{ and}$$

$$\Delta\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{LSR}}) = O\left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n(\log n \vee \frac{1}{\delta})}{n} + \frac{n(H_{n,m} + H_{n,r})}{\delta}}\right).$$

Example 4. Instantiate the assumptions of Thm. 6 with γ -slow-growing input points (Def. 3), LOGGROWTH $\mathbf{k}_{\mathbb{P}}$, and a heavily compressed coreset size $m = (\log n)^{2(\alpha+\beta\gamma)+\epsilon}$ for any $\epsilon > 0$. Then SR delivers $\tilde{O}(n^{-1/2})$ excess MMD with high probability in $O(n^2)$ time, and LSR with $r = m$ and $T = \Theta(\sqrt{n})$ achieves the same in $\tilde{O}(n^{1.5})$ time.

4.2. Constant-preserving coresets via Stein Cholesky

For applications supporting negative weights, constant-preserving coresets offer the possibility of higher accuracy at the cost of potential numerical instability and poorer robustness to errors in integral evaluation (Karvonen et al., 2019). We introduce a constant-preserving compression algorithm, CholeskyThinning (CT, Alg. 6), suitable for summarizing a debiased input sequence. CT first applies WeightedRPCholesky to a constant-regularized kernel $\mathbf{k}_{\mathbb{P}}(x, y) + c$ to select an initial coreset and then uses a combination of KT-Swap-LS and closed-form optimal constant-preserving reweighting to greedily refine the support and weights. The regularized kernel ensures that the coreset output by WeightedRPCholesky is of high quality when paired with the best constant-preserving weights, and our CT standalone analysis (Prop. H.1) improves upon the runtime and error guarantees of RT. In Alg. 7, we combine CT with SteinThinning or Low-rankDebiasing to obtain Stein Cholesky (SC) and Low-rank SC (LSC), our approaches to debiased constant-preserving compression. Our MMD guarantees for SC and LSC (proved in App. H.2) improve upon the rates of Thm. 6.

Algorithm 6 Cholesky Thinning (CT)

Input: mean-zero kernel $\mathbf{k}_{\mathbb{P}}$, points $\mathcal{S}_n = (x_i)_{i=1}^n$, weights $w \in \Delta_{n-1}$, output size m
 $c \leftarrow \text{AVERAGE}(\text{Largest } m \text{ entries of } (\mathbf{k}_{\mathbb{P}}(x_i, x_i))_{i=1}^n)$
 $\mathbf{I}, F \leftarrow \text{WeightedRPCholesky}(\mathbf{k}_{\mathbb{P}} + c, \mathcal{S}_n, w, m); w \leftarrow \mathbf{0}_n$
 $w[\mathbf{I}] \leftarrow \text{argmin}_{w' \in \mathbb{R}^{|\mathbf{I}|}; \sum_i w'_i = 1} w'^{\top} \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n[\mathbf{I}], \mathcal{S}_n[\mathbf{I}]) w'$
 $w \leftarrow \text{KT-Swap-LS}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, w, \text{CP}); \mathbf{I} \leftarrow \{i : w_i \neq 0\}$
 $w[\mathbf{I}] \leftarrow \text{argmin}_{w' \in \mathbb{R}^{|\mathbf{I}|}; \sum_i w'_i = 1} w'^{\top} \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n[\mathbf{I}], \mathcal{S}_n[\mathbf{I}]) w'$
Return: $w_{\text{CT}} \leftarrow w \in \mathbb{R}^n$ with $\|w_{\text{CT}}\|_0 \leq m, \mathbf{1}_n^{\top} w_{\text{CT}} = 1$

Theorem 7 (MMD guarantee for SC / LSC). *Under Assum. (α, β) -params, Stein Cholesky (Alg. 7) takes $O(d_{\mathbf{k}_{\mathbb{P}}} n^2 + (d_{\mathbf{k}_{\mathbb{P}}} + m)nm + m^3)$ time to output w_{SC} , and Low-rank SC takes $O((d_{\mathbf{k}_{\mathbb{P}}} + r + T)nr + (d_{\mathbf{k}_{\mathbb{P}}} + m)nm + m^3)$ time to output w_{LSC} . Moreover, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, each of the following bounds hold:*

Algorithm 7 (Low-rank) Stein Cholesky (SC / LSC)

Input: mean-zero kernel $\mathbf{k}_{\mathbb{P}}$, points \mathcal{S}_n , output size m , rank r , AGM steps T , adaptive rounds Q
 $w \leftarrow \begin{cases} \text{Low-rankDebiasing}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, r, T, Q) & \text{if low-rank} \\ \text{SteinThinning}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n) & \text{otherwise} \end{cases}$
 $w_{\text{SC}} \leftarrow \text{CholeskyThinning}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, w, m)$
Return: $w_{\text{SC}} \in \mathbb{R}^n$ with $\|w_{\text{SC}}\|_0 \leq m$ and $\mathbf{1}_n^{\top} w_{\text{SC}} = 1$

$$\Delta\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{SC}}) = 2\text{MMD}_{\text{OPT}}$$

$$+ O\left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n \log n}{\delta n} + \frac{H_{n,m'}}{\delta}}\right) \text{ and}$$

$$\Delta\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{LSC}}) = 2\text{MMD}_{\text{OPT}}$$

$$+ O\left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n(\log n \vee 1/\delta)}{\delta n} + \frac{H_{n,m'}}{\delta} + \frac{nH_{n,r}}{\delta^2}}\right)$$

for $H_{n,r}$ as in Thm. 4 and $m' \triangleq m + \log 2 - 2\sqrt{m \log 2} + 1$.

Example 5. Instantiate the assumptions of Thm. 7 with γ -slow-growing input points (Def. 3), LOGGROWTH $\mathbf{k}_{\mathbb{P}}$, and a heavily compressed coreset size $m = (\log n)^{2(\alpha+\beta\gamma)+\epsilon}$ for any $\epsilon > 0$. Then SC delivers $\tilde{O}(n^{-1/2})$ excess MMD with high probability in $O(n^2)$ time, and LSC with $r = m$ and $T = \Theta(\sqrt{n})$ achieves the same in $\tilde{O}(n^{1.5})$ time.

Remark 3. While we present our results for a target precision of $1/\sqrt{n}$, a coarser target precision of $1/\sqrt{n_0}$ for $n_0 < n$ can be achieved more quickly by random subsampling/standard thinning the input sequence down to size n_0 before running SR, LSR, SC, or LSC.

5. Experiments

We next evaluate the practical utility of our procedures when faced with three common sources of bias: (1) burn-in, (2) approximate MCMC, and (3) tempering. In all experiments, we use a Stein kernel $\mathbf{k}_{\mathbb{P}}$ with an inverse multiquadric (IMQ) base kernel $\mathbf{k}(x, y) = (1 + \|x - y\|_M^2 / \sigma^2)^{-1/2}$ for σ equal to the median pairwise $\|\cdot\|_M$ distance amongst 1000 points standard thinned from the input. To vary output MMD precision, we first standard thin the input to size $n_0 \in \{2^{10}, 2^{12}, 2^{14}, 2^{16}, 2^{18}, 2^{20}\}$ before applying any method, as discussed in Rems. 2 and 3. For low-rank or weighted coreset methods, we show results for $m = r = n^{\tau}$. When comparing weighted coresets, we optimally reweight every coreset. We report the median over 5 independent runs for all error metrics. We implement our algorithms in JAX (Bradbury et al., 2018) and refer the reader to App. I for additional experiment details (including runtime comparison in Tab. I.1). Our open-source code is available as part of the GoodPoints Python library at <https://github.com/microsoft/goodpoints>.

Correcting for burn-in The initial iterates of a Markov chain are biased by its starting point and need not accu-

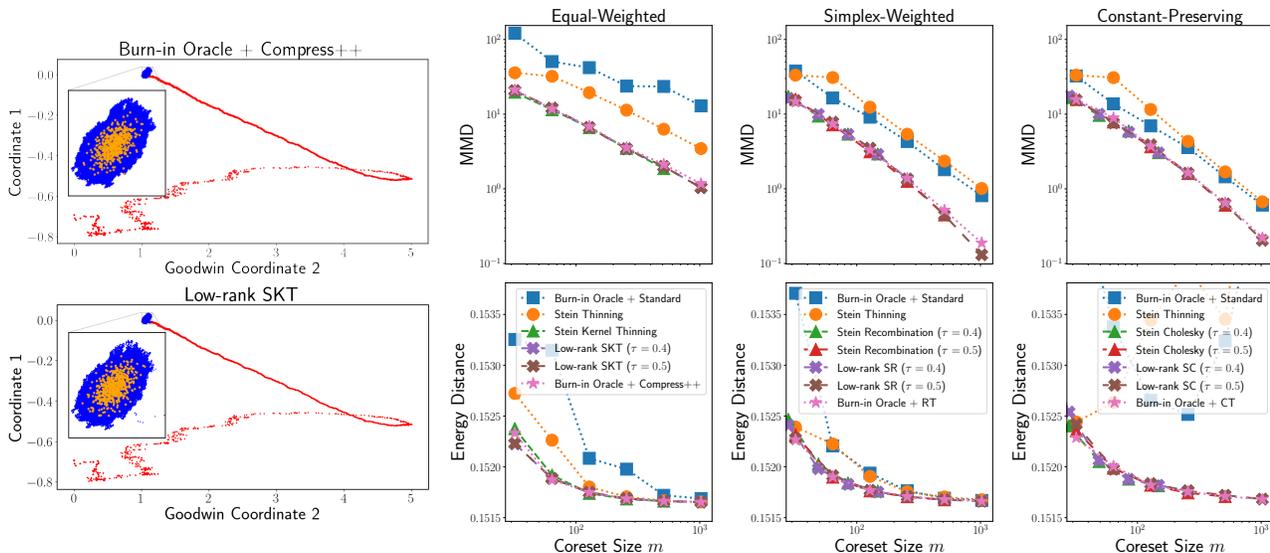


Figure 1: **Correcting for burn-in.** *Left:* Before selecting coresets (orange), the burn-in oracle uses 6 independent Markov chains to discard burn-in (red) while **LSKT** identifies the same high-density region (blue) with 1 chain. *Right:* Using only one chain, our methods consistently outperform the Stein and standard thinning baselines and match the 6-chain oracle.

rately reflect the target distribution \mathbb{P} . Classical burn-in corrections use convergence diagnostics to detect and discard these iterates but typically require running multiple independent Markov chains (Cowles et al., 1999). Alternatively, our proposed debiased compression methods can be used to correct for burn-in given just a single chain.

We test this claim using an experimental setup from Ribiz et al. (2022, Sec. 4.1) and the 6-chain “burn-in oracle” diagnostic of Vats and Knudson (2021).⁵ We aim to compress a posterior \mathbb{P} over the parameters in the Goodwin model of oscillatory enzymatic control ($d = 4$) using $n = 2 \times 10^6$ points from a preconditioned Metropolis-adjusted Langevin algorithm (P-MALA) chain. We repeat this experiment with three alternative MCMC algorithms in App. I.3. Our primary metric is $\text{MMD}_{k_{\mathbb{P}}}$ to \mathbb{P} with $M = I$, but, for external validation, we also measure the energy distance (Ribiz et al., 2022, Eq. 11) to an auxiliary MCMC chain of length n . Trajectory plots of the first two coordinates (Fig. 1, left) highlight the substantial burn-in period for the Goodwin chain and the ability of **LSKT** to mimic the 6-chain burn-in oracle using only a single chain. In Fig. 1 (right), for both the MMD metric and the auxiliary energy distance, our proposed methods consistently outperform Stein thinning and match the quality of 6-chain burn-in removal paired with unbiased compression. The spike

⁵The Vats and Knudson diagnostic is an improvement on the popular Gelman and Rubin (1992) diagnostic (GRD). The GRD was designed to diagnose MCMC convergence rather than burn-in but is commonly used to identify and discard burn-in points automatically through packages like `coda` (Plummer et al., 2006).

in baseline energy distance for the constant-preserving task can be attributed to the selection of overly large weight values due to poor matrix conditioning; the simplex-weighted task does not suffer from this issue due to its regularizing nonnegativity constraint.

Correcting for approximate MCMC In posterior inference, MCMC algorithms typically require iterating over every datapoint to draw each new sample point. When datasets are large, approximating MCMC using datapoint mini-batches can reduce sampling time at the cost of persistent bias and an unknown stationary distribution that prohibits debiasing via importance sampling. Our proposed methods can correct for these biases during compression by computing full-dataset scores on a small subset of n_0 standard thinned points. To evaluate this protocol, we compress a Bayesian logistic regression posterior conditioned on the Forest Covtype dataset ($d = 54$) using $n = 2^{24}$ approximate MCMC points from the stochastic gradient Fisher scoring sampler (Ahn et al., 2012) with batch size 32. Following Wang et al. (2024), we set $M = -\nabla^2 \log p(x_{\text{mode}})$ at the sample mode x_{mode} and use 2^{20} surrogate ground truth points from the No U-turn Sampler (Hoffman and Gelman, 2014) to evaluate energy distance. We find that our proposals improve upon standard thinning and Stein thinning for each compression task, not just in the optimized MMD metric (Fig. 2, top) but also in the auxiliary energy distance (Fig. 2, middle) and when measuring integration error for the mean (Fig. I.4).

Correcting for tempering Tempering, targeting a less-peaked and more dispersed distribution \mathbb{Q} , is a popular

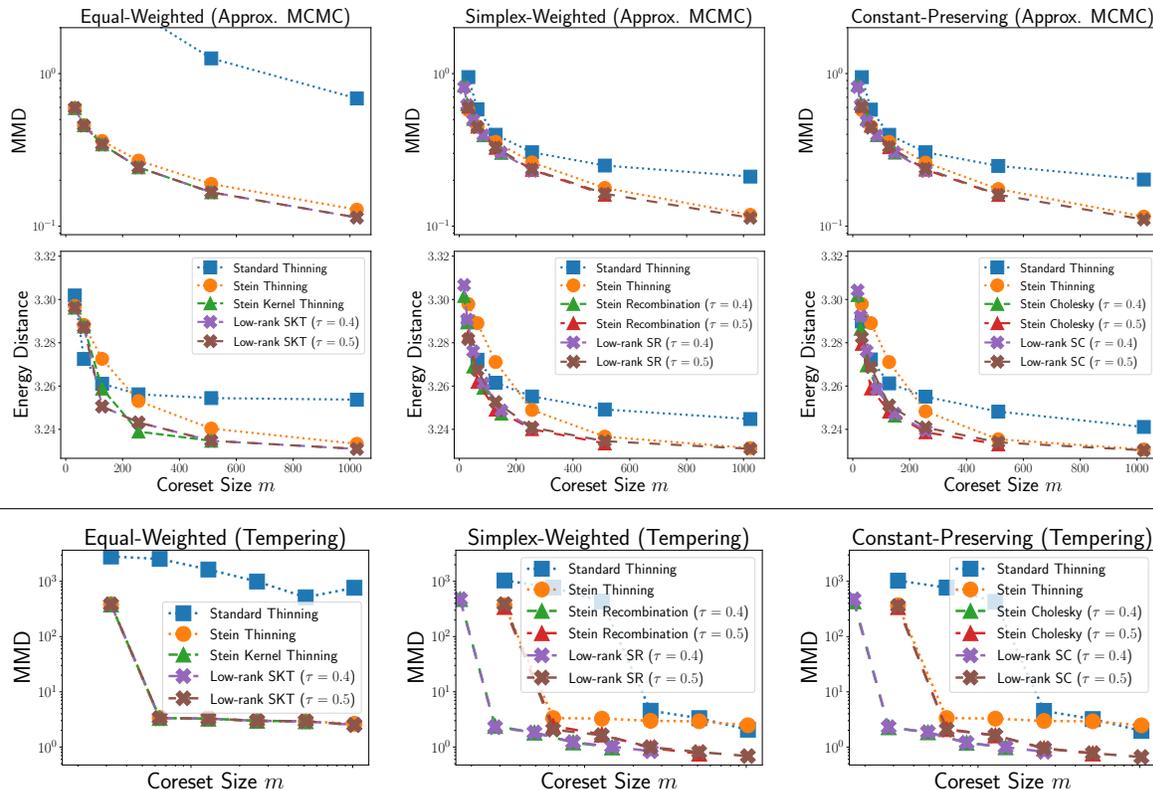


Figure 2: **Correcting for approximate MCMC (top) and tempering (bottom).** For posterior inference over the parameters of Bayesian logistic regression ($d = 54$, top) and a cardiac calcium signaling model ($d = 38$, bottom), our concise coreset constructions correct for approximate MCMC and tempering biases without need for explicit importance sampling.

technique to improve the speed of MCMC convergence. One can correct for the sample bias using importance sampling, but this requires knowledge of the tempered density and can introduce substantial variance (Gramacy et al., 2010). Alternatively, one can use constructions of this work to correct for tempering during compression; this requires no importance weighting and no knowledge of \mathbb{Q} . To test this proposal, we compress the cardiac calcium signaling model posterior ($d = 38$) of Riabiz et al. (2022, Sec. 4.3) with $M = I$ and $n = 3 \times 10^6$ tempered points from a Gaussian random walk Metropolis-Hastings chain. As discussed by Riabiz et al., compression is essential in this setting as the ultimate aim is to propagate posterior uncertainty through a human heart simulator, a feat which requires over 1000 CPU hours for each summary point retained. Our methods perform on par with Stein thinning for equal-weighted compression and yield substantial gains over Stein (and standard) thinning for the two weighted compression tasks.

6. Conclusions and Future Work

We have introduced and analyzed a suite of new procedures for compressing a biased input sequence into an accurate summary of a target distribution. For equal-weighted compression, Stein kernel thinning delivers \sqrt{n} points with $\tilde{O}(n^{-1/2})$ MMD in $O(n^2)$ time, and low-rank SKT can improve this running time to $\tilde{O}(n^{3/2})$. For simplex-weighted and constant-preserving compression, Stein recombination and Stein Cholesky provide enhanced parsimony, matching these guarantees with as few as poly-log(n) points. Recent work has identified some limitations of score-based discrepancies, like Stein kernel MMDs, and developed modified objectives that are more sensitive to the relative density of isolated modes (Liu et al., 2023; Bénard et al., 2024). A valuable next step would be to extend our constructions to provide compression guarantees for these modified discrepancy measures. Other opportunities for future work include marrying the better-than-i.i.d. guarantees of this work with the non-myopic compression of Teymur et al. (2021), the control-variate compression of Chopin and Ducrocq (2021), and the online compression of Koppel et al. (2024).

Acknowledgments

We thank Marina Riabiz for making the Markov chain data used in Sec. 5 available and Jeffrey Rosenthal for helpful discussions concerning the geometric ergodicity of Markov chains. Lingxiao Li acknowledges the generous support of Army Research Office grants W911NF2010168 and W911NF2110293, of Air Force Office of Scientific Research award FA9550-19-1-031, of National Science Foundation grant CHS-1955697, from the CSAIL Systems that Learn program, from the MIT-IBM Watson AI Laboratory, from the Toyota-CSAIL Joint Research Center, from a gift from Adobe Systems, and from a Google Research Scholar award.

Impact Statement

This paper presents work with the aim of advancing the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- E. Afzali and S. Muthukumarana. Gradient-free kernel conditional Stein discrepancy goodness of fit testing. *Machine Learning with Applications*, 12:100463, 2023. ISSN 2666-8270. doi: <https://doi.org/10.1016/j.mlwa.2023.100463>. URL <https://www.sciencedirect.com/science/article/pii/S2666827023000166>. (Cited on page 3.)
- S. Ahn, A. Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient Fisher scoring. In J. Langford and J. Pineau, editors, *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, ICML '12, pages 1591–1598, New York, NY, USA, July 2012. Omnipress. ISBN 978-1-4503-1285-1. (Cited on pages 1, 8, and 55.)
- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404, 1950. (Cited on page 2.)
- Y. F. Atchadé and F. Perron. On the geometric ergodicity of metropolis-hastings algorithms. *Statistics*, 41(1):77–84, 2007. (Cited on page 4.)
- A. Barp, F.-X. Briol, A. Duncan, M. Girolami, and L. Mackey. Minimum stein discrepancy estimators. *Advances in Neural Information Processing Systems*, 32, 2019. (Cited on page 3.)
- A. Barp, C.-J. Simon-Gabriel, M. Girolami, and L. Mackey. Targeted separation and convergence with kernel discrepancies. *arXiv preprint arXiv:2209.12835*, 2022. (Cited on page 19.)
- C. Bénard, B. Staber, and S. Da Veiga. Kernel Stein discrepancy thinning: A theoretical perspective of pathologies and a practical fix with regularization. *Advances in Neural Information Processing Systems*, 36, 2024. (Cited on page 9.)
- P. Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013. (Cited on page 26.)
- J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>. (Cited on pages 7 and 54.)
- R. C. Bradley. Basic Properties of Strong Mixing Conditions. A Survey and Some Open Questions. *Probability Surveys*, 2(none):107 – 144, 2005. doi: 10.1214/154957805100000104. URL <https://doi.org/10.1214/154957805100000104>. (Cited on page 25.)
- C. Carmeli, E. De Vito, and A. Toigo. Vector valued reproducing kernel Hilbert spaces of integrable functions and Mercer theorem. *Analysis and Applications*, 4(04): 377–408, 2006. (Cited on pages 18 and 19.)
- T. K. Chandra. De la vallée poussin’s theorem, uniform integrability, tightness and moments. *Statistics & Probability Letters*, 107:136–141, 2015. (Cited on page 26.)
- Y. Chen, M. Welling, and A. Smola. Super-samples from kernel herding. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 109–116, 2010. (Cited on page 3.)
- Y. Chen, E. N. Epperly, J. A. Tropp, and R. J. Webber. Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations. *arXiv preprint arXiv:2207.06503*, 2022. (Cited on pages 5, 40, 42, and 51.)
- N. Chopin and G. Ducrocq. Fast compression of MCMC output. *Entropy*, 23(8):1017, 2021. (Cited on page 9.)
- K. Chwialkowski, H. Strathmann, and A. Gretton. A kernel test of goodness of fit. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2606–2615, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/chwialkowski16.html>. (Cited on page 3.)
- M. K. Cowles and B. P. Carlin. Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904, 1996. (Cited on page 1.)

- M. K. Cowles, G. O. Roberts, and J. S. Rosenthal. Possible biases induced by mcmc convergence diagnostics. *Journal of Statistical Computation and Simulation*, 64 (1):87–104, 1999. (Cited on page 8.)
- A. Dax. Low-rank positive approximants of symmetric matrices. *Advances in Linear Algebra & Matrix Theory*, 4 (03):172, 2014. (Cited on page 51.)
- R. Douc and O. Cappé. Comparison of resampling schemes for particle filtering. In *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005.*, pages 64–69. IEEE, 2005. (Cited on pages 5 and 37.)
- R. Douc, E. Moulines, P. Priouret, and P. Soulier. *Markov chains*, volume 1. Springer, 2018. (Cited on page 25.)
- A. Durmus and É. Moulines. On the geometric convergence for MALA under verifiable conditions. *arXiv preprint arXiv:2201.01951*, 2022. (Cited on page 4.)
- A. Durmus, É. Moulines, and E. Saksman. Irreducibility and geometric ergodicity of Hamiltonian Monte Carlo. *The Annals of Statistics*, 48(6):3545–3564, 2020. (Cited on page 4.)
- A. Durmus and Éric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017. ISSN 10505164. URL <http://www.jstor.org/stable/26361410>. (Cited on page 2.)
- R. Dwivedi and L. Mackey. Kernel thinning. *arXiv preprint arXiv:2105.05842*, 2021. (Cited on pages 1, 4, 6, and 34.)
- R. Dwivedi and L. Mackey. Generalized kernel thinning. In *International Conference on Learning Representations*, 2022. (Cited on pages 1, 3, 21, 34, and 35.)
- E. Epperly and E. Moreno. Kernel quadrature with randomly pivoted Cholesky. *Advances in Neural Information Processing Systems*, 36, 2024. (Cited on pages 1 and 51.)
- M. A. Erdogdu, L. Mackey, and O. Shamir. Global non-convex optimization with discretized diffusions. *Advances in Neural Information Processing Systems*, 31, 2018. (Cited on page 4.)
- M. A. Gallegos-Herrada, D. Ledvinka, and J. S. Rosenthal. Equivalences of geometric ergodicity of Markov chains. *Journal of Theoretical Probability*, pages 1–27, 2023. (Cited on pages 3 and 25.)
- A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical science*, 7 (4):457–472, 1992. (Cited on page 8.)
- B. Ghoghogh, A. Ghodsi, F. Karray, and M. Crowley. KKT conditions, first-order and second-order optimization, and distributed optimization: tutorial and survey. *arXiv preprint arXiv:2110.01858*, 2021. (Cited on page 52.)
- J. Gorham and L. Mackey. Measuring sample quality with kernels. In *International Conference on Machine Learning*, pages 1292–1301. PMLR, 2017. (Cited on page 3.)
- J. Gorham, A. B. Duncan, S. J. Vollmer, and L. Mackey. Measuring sample quality with diffusions. *The Annals of Applied Probability*, 29(5):2884–2928, 2019. (Cited on page 3.)
- R. Gramacy, R. Samworth, and R. King. Importance tempering. *Statistics and Computing*, 20:1–7, 2010. (Cited on pages 1 and 9.)
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. (Cited on page 1.)
- S. Hayakawa, H. Oberhauser, and T. Lyons. Positively weighted kernel quadrature via subsampling. *Advances in Neural Information Processing Systems*, 35:6886–6900, 2022. (Cited on page 6.)
- S. Hayakawa, H. Oberhauser, and T. Lyons. Sampling-based Nyström approximation and kernel quadrature. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023. (Cited on pages 1 and 6.)
- L. Hodgkinson, R. Salomone, and F. Roosta. The reproducing Stein kernel approach for post-hoc corrected sampling. *arXiv preprint arXiv:2001.09266*, 2020. (Cited on page 3.)
- M. D. Hoffman and A. Gelman. The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014. (Cited on pages 8 and 55.)
- A. A. Johnson. *Geometric ergodicity of Gibbs samplers*. university of minnesota, 2009. (Cited on page 4.)
- T. Karvonen, M. Kanagawa, and S. Särkkä. On the positivity and magnitudes of Bayesian quadrature weights. *Statistics and Computing*, 29:1317–1333, 2019. (Cited on pages 6 and 7.)
- A. Koppel, J. Eappen, S. Bhatt, C. Hawkins, and S. Ganesh. Online MCMC thinning with kernelized Stein discrepancy. *SIAM J. Math. Data Sci.*, 6(1):51–75, 2024. doi: 10.1137/22M1510108. URL <https://doi.org/10.1137/22m1510108>. (Cited on page 9.)

- F. Kuo. Component-by-component constructions achieve the optimal rate of convergence for multivariate integration in weighted Korobov and Sobolev spaces. *Journal of Complexity*, 19(3):301–320, 2003. ISSN 0885-064X. doi: [https://doi.org/10.1016/S0885-064X\(03\)00006-2](https://doi.org/10.1016/S0885-064X(03)00006-2). URL <https://www.sciencedirect.com/science/article/pii/S0885064X03000062>. Oberwolfach Special Issue. (Cited on page 3.)
- S. Lacoste-Julien, F. Lindsten, and F. Bach. Sequential kernel herding: Frank-Wolfe optimization for particle filtering. In *Artificial Intelligence and Statistics*, pages 544–552. PMLR, 2015. (Cited on page 3.)
- L. Li, J.-G. Liu, and Y. Wang. Geometric ergodicity of SGLD via reflection coupling. *arXiv preprint arXiv:2301.06769*, 2023. (Cited on page 4.)
- Q. Liu and J. Lee. Black-box importance sampling. In *Artificial Intelligence and Statistics*, pages 952–961. PMLR, 2017. (Cited on pages 4, 26, 28, and 55.)
- Q. Liu, J. Lee, and M. Jordan. A kernelized Stein discrepancy for goodness-of-fit tests. In *International conference on machine learning*, pages 276–284. PMLR, 2016. (Cited on page 3.)
- X. Liu, A. B. Duncan, and A. Gandy. Using perturbation to improve goodness-of-fit tests based on kernelized Stein discrepancy. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 21527–21547. PMLR, 23–29 Jul 2023. (Cited on page 9.)
- F. Merlevède, M. Peligrad, and S. Utev. Sharp conditions for the CLT of linear processes in a Hilbert space. *Journal of Theoretical Probability*, 10(3):681–693, 1997. (Cited on page 25.)
- S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012. (Cited on pages 25 and 26.)
- H. Q. Minh. Some properties of Gaussian reproducing kernel Hilbert spaces and their implications for function approximation and learning theory. *Constructive Approximation*, 32:307–338, 2010. (Cited on page 4.)
- S. Niederer, L. Mitchell, N. Smith, and G. Plank. Simulating human cardiac electrophysiology on clinical time-scales. *Frontiers in physiology*, 2:14, 2011. (Cited on page 1.)
- V. I. Paulsen and M. Raghupathi. *An introduction to the theory of reproducing kernel Hilbert spaces*, volume 152. Cambridge university press, 2016. (Cited on pages 15, 23, and 25.)
- D. Phan, N. Pradhan, and M. Jankowiak. Composable effects for flexible and accelerated probabilistic programming in numpyro. *arXiv preprint arXiv:1912.11554*, 2019. (Cited on page 55.)
- J. M. Phillips and W. M. Tai. Near-optimal coresets of kernel density estimates. *Discrete & Computational Geometry*, 63:867–887, 2020. (Cited on page 5.)
- I. Pinelis. Exact lower and upper bounds on the incomplete gamma function. *arXiv preprint arXiv:2005.06384*, 2020. (Cited on page 43.)
- Y. Pitcan. A note on concentration inequalities for U-statistics. *arXiv preprint arXiv:1712.06160*, 2017. (Cited on page 31.)
- M. Plummer, N. Best, K. Cowles, and K. Vines. Coda: Convergence diagnosis and output analysis for mcmc. *R News*, 6(1):7–11, 2006. URL <https://journal.r-project.org/archive/>. (Cited on page 8.)
- Q. Qin. Geometric ergodicity of trans-dimensional Markov chain Monte Carlo algorithms. *arXiv preprint arXiv:2308.00139*, 2023. (Cited on page 25.)
- M. Riabiz, W. Y. Chen, J. Cockayne, P. Swietach, S. A. Niederer, L. Mackey, and C. J. Oates. Replication Data for: Optimal Thinning of MCMC Output, 2020. URL <https://doi.org/10.7910/DVN/MDKNWM>. Accessed on Mar 23, 2021. (Cited on page 55.)
- M. Riabiz, W. Y. Chen, J. Cockayne, P. Swietach, S. A. Niederer, L. Mackey, and C. J. Oates. Optimal thinning of MCMC output. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(4):1059–1081, 2022. (Cited on pages 1, 2, 4, 8, 9, 33, 36, and 54.)
- G. O. Roberts and R. L. Tweedie. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika*, 83(1):95–110, 1996. (Cited on page 4.)
- A. Shetty, R. Dwivedi, and L. Mackey. Distribution compression in near-linear time. In *International Conference on Learning Representations*, 2022. (Cited on pages 1, 5, 34, 40, 45, 46, 47, and 54.)
- I. Steinwart and A. Christmann. *Support vector machines*. Springer Science & Business Media, 2008. (Cited on pages 3, 15, 16, 17, 19, 27, and 29.)
- I. Steinwart and C. Scovel. Mercer’s theorem on general domains: On the interaction between measures, kernels, and RKHSs. *Constructive Approximation*, 35:363–417, 2012. (Cited on pages 4, 14, and 15.)

- H.-W. Sun and D.-X. Zhou. Reproducing kernel Hilbert spaces associated with analytic translation-invariant Mercer kernels. *Journal of Fourier Analysis and Applications*, 14(1):89–101, 2008. (Cited on pages 21 and 23.)
- M. Tchernychova. *Caratheodory cubature measures*. PhD thesis, University of Oxford, 2016. (Cited on pages 6, 48, and 49.)
- O. Teymur, J. Gorham, M. Riabiz, and C. Oates. Optimal quantisation of probability measures using maximum mean discrepancy. In *International Conference on Artificial Intelligence and Statistics*, pages 1027–1035. PMLR, 2021. (Cited on page 9.)
- D. Vats and C. Knudson. Revisiting the Gelman–Rubin diagnostic. *Statistical Science*, 36(4):518–529, 2021. (Cited on page 8.)
- M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019. (Cited on pages 23, 29, and 31.)
- C. Wang, Y. Chen, H. Kanagawa, and C. J. Oates. Stein II-importance sampling. *Advances in Neural Information Processing Systems*, 36, 2024. (Cited on page 8.)
- J.-K. Wang, J. Abernethy, and K. Y. Levy. No-regret dynamics in the Fenchel game: A unified framework for algorithmic convex optimization. *Mathematical Programming*, pages 1–66, 2023. (Cited on pages 5, 40, 41, and 43.)
- J. Wellner. *Weak convergence and empirical processes: With applications to statistics*. Springer Science & Business Media, 2013. (Cited on page 25.)
- J. Yang, Q. Liu, V. Rao, and J. Neville. Goodness-of-fit testing for discrete distributions via Stein discrepancy. In *International Conference on Machine Learning*, pages 5561–5570. PMLR, 2018. (Cited on page 3.)
- F. Zhang. *The Schur complement and its applications*, volume 4. Springer Science & Business Media, 2006. (Cited on page 51.)
- D.-X. Zhou. The covering number in learning theory. *Journal of Complexity*, 18(3):739–767, 2002. (Cited on page 4.)

Appendix Contents

A. Appendix Notation

For the point sequence $\mathcal{S}_n = (x_i)_{i \in [n]}$, we use $\mathbb{S}_n \triangleq \frac{1}{n} \sum_{i \in [n]} \delta_{x_i}$ to denote the empirical distribution. For a weight vector $w \in \mathbb{R}^n$, we define the support $\text{supp}(w) \triangleq \{i \in [n] : w_i \neq 0\}$ and the signed measure $\mathbb{S}_n^w \triangleq \sum_{i \in [n]} w_i \delta_{x_i}$. For a matrix $K \in \mathbb{R}^{n \times n}$ and $w \in \Delta_{n-1}$, we define the weighted matrix $K^w \triangleq \text{diag}(\sqrt{w})K \text{diag}(\sqrt{w})$. For positive semidefinite (PSD) matrices (A, B) , we use $A \succeq B$ (resp. $A \preceq B$) to mean $A - B$ (resp. $B - A$) is PSD. For a symmetric PSD (SPSD) matrix M , we let $M^{1/2}$ denote a symmetric matrix square root satisfying $M = M^{1/2}M^{1/2}$. For $A \in \mathbb{R}^{n \times m}$, we denote $\|A\|_p \triangleq \sup_{x \neq 0} \frac{\|Ax\|_p}{\|x\|_p}$. We will use $\mathbb{1}_E$ to denote the indicator function for an event E .

Notation used only in a specific section will be introduced therein.

B. Spectral Analysis of Kernel Matrices

The goal of this section is to develop spectral bounds for kernel matrices.

From covering numbers of kernels to eigenvalues of kernel matrices In App. B.1, we transfer the bounds on covering numbers from the definition of POLYGROWTH or LOGGROWTH kernels to bounds on the eigenvalues of the kernel matrices. This sets the theoretical foundation for the algorithms in later sections as their error guarantees rely on the fast decay of eigenvalues of kernel matrices.

Covering numbers and eigenvalues of Stein kernels In App. B.2, we show that Stein kernels are POLYGROWTH (resp. LOGGROWTH) provided that their base kernels are differentiable (resp. radially analytic). Putting the results together from App. B.1, we obtain spectral bounds for a wide range of Stein kernels.

Notation For a normed space E , we use $\|\cdot\|_E$ to denote its norm, $\mathcal{B}_E(p, r) \triangleq \{x \in E : \|x - p\|_E \leq r\}$ to denote the closed ball of radius r centered at p in E with the shorthand $\mathcal{B}_E(r) \triangleq \mathcal{B}_E(0, r)$ and $\mathcal{B}_E \triangleq \mathcal{B}_E(1)$. When E is an RKHS with kernel \mathbf{k} , for brevity we use \mathbf{k} in place of E in the subscript. Let $\mathfrak{F}(\mathcal{X}, \mathcal{Y})$ denote the space of functions from \mathcal{X} to \mathcal{Y} , and $\mathfrak{B}(E, F)$ denote the space of bounded linear functions between normed spaces E, F . For a set A , we use $\ell_\infty(A)$ to denote the space of bounded \mathbb{R} -valued functions on A equipped with the sup-norm $\|f\|_{\infty, A} \triangleq \sup_{x \in A} |f(x)|$. We use $E \hookrightarrow F$ to denote the inclusion map. We use $\lambda_\ell(T)$ to denote the ℓ -th largest eigenvalue of an operator T .

B.1. From covering numbers of kernels to eigenvalues of kernel matrices

We first introduce the general Mercer representation theorem from [Steinwart and Scovel \(2012\)](#), which shows the existence of a discrete spectrum of the integral operator associated with a continuous square-integrable kernel. The theorem also provides a series expansion of the kernel, i.e., the Mercer representation, in terms of the eigenvalues and eigenfunctions.

Lemma B.1 (General Mercer representation ([Steinwart and Scovel, 2012](#))). *Consider a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ that is jointly continuous in both inputs and a probability measure μ such that $\int \mathbf{k}(x, x) d\mu(x) < \infty$. Then the following holds.*

- (a) *The inclusion $\mathcal{H}_{\mathbf{k}} \hookrightarrow \mathcal{L}^2(\mu)$ is a compact operator, i.e., $\mathcal{B}_{\mathbf{k}}$ is a compact subset of $\mathcal{L}^2(\mu)$. In particular, this inclusion is continuous.*
- (b) *The Hilbert-space adjoint of the inclusion $\mathcal{H}_{\mathbf{k}} \hookrightarrow \mathcal{L}^2(\mu)$ is the compact operator $S_{\mathbf{k}, \mu} : \mathcal{L}^2(\mu) \rightarrow \mathcal{H}_{\mathbf{k}}$ defined as*

$$S_{\mathbf{k}, \mu} f \triangleq \int \mathbf{k}(\cdot, x) f(x) d\mu(x). \quad (3)$$

We also have $S_{\mathbf{k}, \mu}^ \triangleq \mathcal{H}_{\mathbf{k}} \hookrightarrow \mathcal{L}^2(\mu)$. Hence the operator*

$$T_{\mathbf{k}, \mu} \triangleq S_{\mathbf{k}, \mu}^* S_{\mathbf{k}, \mu} : \mathcal{L}^2(\mu) \rightarrow \mathcal{L}^2(\mu) \quad (4)$$

is also compact.

(c) There exist $\{\lambda_\ell\}_{\ell=1}^\infty$ with $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ and $\{\phi_\ell\}_{\ell=1}^\infty \subset \mathcal{H}_{\mathbf{k}}$ such that $\{\phi_\ell\}_{\ell=1}^\infty$ is an orthonormal system in $\mathcal{L}^2(\mu)$ and $\{\lambda_\ell\}_{\ell=1}^\infty$ (resp. $\{\phi_\ell\}_{\ell=1}^\infty$) consists of the eigenvalues (resp. eigenfunctions) of $T_{\mathbf{k},\mu}$ with eigendecomposition, for $f \in \mathcal{L}^2(\mu)$,

$$T_{\mathbf{k},\mu}f = \sum_{\ell=1}^\infty \lambda_\ell \langle f, \phi_\ell \rangle_{\mathcal{L}^2(\mu)} \phi_\ell$$

with convergence in $\mathcal{L}^2(\mu)$.

(d) We have the following series expansion

$$\mathbf{k}(x, x') = \sum_{\ell=1}^\infty \lambda_\ell \phi_\ell(x) \phi_\ell(x'), \quad (5)$$

where the series convergence is absolute and uniform in x, x' on all $A \times A \subset \text{supp } \mu \times \text{supp } \mu$.

Proof of Lem. B.1. Part (a) and (b) follow respectively from Steinwart and Scovel (2012, Lem. 2.3 and 2.2). Part (c) follows from part (a) and Steinwart and Scovel (2012, Lem. 2.12). Finally, part (d) follows from Steinwart and Scovel (2012, Cor. 3.5). \square

We will use the following lemma regarding the restriction of covering numbers.

Lemma B.2 (Covering number is preserved in restriction). *For a kernel $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and a set $A \subset \mathbb{R}^d$, we have $\mathcal{N}_{\mathbf{k}}(A, \epsilon) = \mathcal{N}_{\mathbf{k}|_A}(A, \epsilon)$, for $\mathbf{k}|_A$, the restricted kernel of \mathbf{k} to A (Paulsen and Raghupathi, 2016, Sec. 5.4).*

Proof of Lem. B.2. It suffices to show that a $(\mathbf{k}, A, \epsilon)$ cover can be converted to a cover of $(\mathbf{k}|_A, A, \epsilon)$ of the same cardinality and vice versa.

Let $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}|_A}$ be a $(\mathbf{k}|_A, A, \epsilon)$ cover. For any $f \in \mathcal{C}$, we have $\|f\|_{\mathbf{k}|_A} = \inf \left\{ \|\tilde{f}\|_{\mathbf{k}} : \tilde{f} \in \mathcal{H}_{\mathbf{k}}, \tilde{f}|_A = f \right\} \leq 1$ (Paulsen and Raghupathi, 2016, Corollary 5.8). Moreover, the infimum is attained by some $\tilde{f} \in \mathcal{H}_{\mathbf{k}}$ such that $\|\tilde{f}\|_{\mathbf{k}} = \|f\|_{\mathbf{k}|_A} \leq 1$ and $\tilde{f}|_A = f$. Now form $\tilde{\mathcal{C}} = \{\tilde{f} : f \in \mathcal{C}\}$. For any $\tilde{h} \in \mathcal{B}_{\mathbf{k}}$, there exists $f \in \mathcal{C}$ such that

$$\|\tilde{h}|_A - f\|_{\infty, A} \leq \epsilon \implies \|\tilde{h} - \tilde{f}\|_{\infty, A} \leq \epsilon,$$

so $\tilde{\mathcal{C}}$ is a $(\mathbf{k}|_A, A, \epsilon)$ cover.

For the other direction, let $\tilde{\mathcal{C}} \subset \mathcal{B}_{\mathbf{k}}$ be a $(\mathbf{k}, A, \epsilon)$ cover. Define $\mathcal{C} = \{\tilde{f}|_A : \tilde{f} \in \tilde{\mathcal{C}}\} \subset \mathcal{H}_{\mathbf{k}|_A}$. Since $\|\tilde{f}|_A\|_{\mathbf{k}|_A} \leq \|\tilde{f}\|_{\mathbf{k}}$, we have $\mathcal{C} \subset \mathcal{B}_{\mathbf{k}|_A}$. For any $h \in \mathcal{B}_{\mathbf{k}|_A}$, again by Paulsen and Raghupathi (2016, Corollary 5.8), there exists $\tilde{h} \in \mathcal{H}_{\mathbf{k}}$ such that $\|\tilde{h}\|_{\mathbf{k}} = \|h\|_{\mathbf{k}|_A} \leq 1$, so there exists $\tilde{f} \in \tilde{\mathcal{C}}$ such that

$$\|\tilde{h} - \tilde{f}\|_{\infty, A} \leq \epsilon \implies \|h - \tilde{f}|_A\|_{\infty, A} \leq \epsilon,$$

Hence \mathcal{C} is a $(\mathbf{k}|_A, A, \epsilon)$ cover. \square

The goal for the rest of this section is to transfer the bounds of the covering number in the definition of a POLYGROWTH or LOGGROWTH kernel from Assum. (α, β) -kernel to bounds on entropy numbers (Steinwart and Christmann, 2008, Def. 6.20) that are closely related to eigenvalues of the integral operator (4).

Definition B.1 (Entropy number of a bounded linear map). *For a bounded linear operator $S : E \rightarrow F$ between normed spaces E, F , for $\ell \in \mathbb{N}$, the ℓ -th entropy number of S is defined as*

$$e_\ell(S) \triangleq \inf \left\{ \epsilon > 0 : \exists s_1, \dots, s_{2^\ell-1} \in S(\mathcal{B}_E) \text{ such that } S(\mathcal{B}_E) \subset \bigcup_{i=1}^{2^\ell-1} \mathcal{B}_F(s_i, \epsilon) \right\}.$$

The following lemma shows the relation between covering numbers and entropy numbers.

Lemma B.3 (Relation between covering number and entropy number). *Suppose a kernel \mathbf{k} is jointly continuous and $A \subset \mathbb{R}^d$ is bounded. Then for any $\epsilon > 0$,*

$$e_{\lceil \log_2 \mathcal{N}_{\mathbf{k}}(A, \epsilon) \rceil + 1}(\mathcal{H}_{\mathbf{k}|_A} \hookrightarrow \ell_\infty(A)) \leq \epsilon.$$

Proof of Lem. B.3. First, the assumption implies $\mathbf{k}|_A$ is a bounded kernel, so by [Steinwart and Christmann \(2008, Lemma 4.23\)](#), the inclusion $\mathcal{H}_{\mathbf{k}|_A} \hookrightarrow \ell_\infty(A)$ is continuous. By the definition of $\mathcal{N}_{\mathbf{k}|_A}(A, \epsilon)$, by adding arbitrary elements into the cover if necessary, there exists a $(\mathbf{k}|_A, A, \epsilon)$ cover of $\mathcal{B}_{\mathbf{k}|_A}$ of cardinality $2^{\lceil \log_2(\mathcal{N}_{\mathbf{k}|_A}(A, \epsilon)) \rceil} \geq \mathcal{N}_{\mathbf{k}|_A}(A, \epsilon)$. Hence

$$e_{\lceil \log_2 \mathcal{N}_{\mathbf{k}|_A}(A, \epsilon) \rceil + 1}(\mathcal{H}_{\mathbf{k}|_A} \hookrightarrow \ell_\infty(A)) \leq \epsilon.$$

The claim follows since $\mathcal{N}_{\mathbf{k}|_A}(A, \epsilon) = \mathcal{N}_{\mathbf{k}}(A, \epsilon)$ by [Lem. B.2](#). \square

Proposition B.1 (ℓ_∞ -entropy number bound for POLYGROWTH or LOGGROWTH \mathbf{k}). *Suppose a kernel \mathbf{k} satisfies Assum. (α, β) -kernel. Let $\mathfrak{C}_d > 0$ denote the constant that appears in the Assum. (α, β) -kernel. Define*

$$L_{\mathbf{k}}(r) \triangleq \frac{\mathfrak{C}_d}{\log 2} r^\beta. \quad (6)$$

Then for any $r > 0$ and $\ell \in \mathbb{N}$ that satisfies $\ell > L_{\mathbf{k}}(r+1) + 1$, we have

$$e_\ell(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(r)}} \hookrightarrow \ell_\infty(\mathcal{B}_2(r))) \leq \begin{cases} \left(\frac{L_{\mathbf{k}}(r+1)}{\ell-1} \right)^{\frac{1}{\alpha}} & \text{if } \mathbf{k} \text{ is POLYGROWTH}(\alpha, \beta), \text{ and} \\ \exp\left(1 - \left(\frac{\ell-1}{L_{\mathbf{k}}(r+1)} \right)^{\frac{1}{\alpha}}\right) & \text{if } \mathbf{k} \text{ is LOGGROWTH}(\alpha, \beta). \end{cases}$$

Proof of Prop. B.1. By [Lem. B.3](#) and the fact that e_ℓ is monotonically decreasing in ℓ by definition, if $\ell \geq \log_2 \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) + 1$ for some $\epsilon > 0$, then

$$e_\ell(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(r)}} \hookrightarrow \ell_\infty(\mathcal{B}_2(r))) \leq e_{\lceil \log_2 \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) \rceil + 1}(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(r)}} \hookrightarrow \ell_\infty(\mathcal{B}_2(r))) \leq \epsilon. \quad (7)$$

For the POLYGROWTH case, by its definition, the condition $\ell \geq \log_2 \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) + 1$ is met if $\epsilon \in (0, 1)$ and

$$\ell \geq \frac{\mathfrak{C}_d}{\log 2} (1/\epsilon)^\alpha (r+1)^\beta + 1 \iff \epsilon \leq \left(\frac{L_{\mathbf{k}}(r+1)}{\ell-1} \right)^{\frac{1}{\alpha}}.$$

Hence (7) holds with $\epsilon = \left(\frac{L_{\mathbf{k}}(r+1)}{\ell-1} \right)^{\frac{1}{\alpha}}$, as long as $\epsilon \in (0, 1)$, so ℓ needs to satisfy

$$1 > \left(\frac{L_{\mathbf{k}}(r+1)}{\ell-1} \right)^{\frac{1}{\alpha}} \iff \ell > L_{\mathbf{k}}(r+1) + 1.$$

Similarly, for the LOGGROWTH case, the condition $\ell \geq \log_2 \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) + 1$ is met if $\epsilon \in (0, 1)$ and

$$\ell \geq \frac{\mathfrak{C}_d}{\log 2} (\log(1/\epsilon) + 1)^\alpha (r+1)^\beta + 1 \iff \epsilon \leq \exp\left(1 - \left(\frac{\ell-1}{L_{\mathbf{k}}(r+1)} \right)^{\frac{1}{\alpha}}\right).$$

Hence (7) holds with $\epsilon = \exp\left(1 - \left(\frac{\ell-1}{L_{\mathbf{k}}(r+1)} \right)^{\frac{1}{\alpha}}\right)$, as long as $\epsilon \in (0, 1)$, so ℓ needs to satisfy

$$1 > \exp\left(1 - \left(\frac{\ell-1}{L_{\mathbf{k}}(r+1)} \right)^{\frac{1}{\alpha}}\right) \iff \ell > L_{\mathbf{k}}(r+1) + 1.$$

\square

Next, we show that we can transfer bounds on entropy numbers to obtain bounds for the eigenvalues of kernel matrices, which will become handy when we develop sub-quadratic-time algorithms in [Sec. 3](#). We rely on the following lemma, which summarizes the relevant facts from [Steinwart and Christmann \(2008, Appendix A\)](#).

Lemma B.4 (Eigenvalue is bounded by entropy number). *Let \mathbf{k} be a jointly continuous kernel and \mathbb{P} be a distribution such that $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}(x, x)] < \infty$, and recall that $\lambda_\ell(\cdot)$ denotes the ℓ -th largest eigenvalue of a linear operator. Then, for all $\ell \in \mathbb{N}$,*

$$\lambda_\ell(T_{\mathbf{k}, \mathbb{P}}) \leq 4e_\ell^2(\mathcal{H}_{\mathbf{k}} \hookrightarrow \mathcal{L}^2(\mathbb{P})).$$

Proof of Lem. B.4. For any bounded linear operator $S : \mathcal{H}_1 \rightarrow \mathcal{H}_2$ between Hilbert spaces \mathcal{H}_1 and \mathcal{H}_2 , we have $a_\ell(S) \leq 2e_\ell(S)$, where a_ℓ is the ℓ -th approximation number defined in [Steinwart and Christmann \(2008, \(A.29\)\)](#). Recall the operator $S_{\mathbf{k}, \mathbb{P}}^* = \mathcal{H}_{\mathbf{k}} \hookrightarrow \mathcal{L}^2(\mathbb{P})$ from [\(3\)](#), which is compact (in particular bounded) by [Lem. B.1\(a\)](#). Thus

$$s_\ell(S_{\mathbf{k}, \mathbb{P}}^*) = a_\ell(S_{\mathbf{k}, \mathbb{P}}^*) \leq 2e_\ell(S_{\mathbf{k}, \mathbb{P}}^*),$$

where the first equality follows from the paragraph below [Steinwart and Christmann \(2008, \(A.29\)\)](#) and s_ℓ is ℓ -th singular number of an operator ([Steinwart and Christmann, 2008, \(A.25\)](#)). Then using the identities mentioned under [Steinwart and Christmann \(2008, \(A.25\)\)](#) and [Steinwart and Christmann \(2008, \(A.27\)\)](#) and that all operators involved are compact by [Lem. B.1\(b\)](#), we have

$$\lambda_\ell(T_{\mathbf{k}, \mathbb{P}}) = \lambda_\ell(S_{\mathbf{k}, \mathbb{P}}^* S_{\mathbf{k}, \mathbb{P}}) = s_\ell(S_{\mathbf{k}, \mathbb{P}}^* S_{\mathbf{k}, \mathbb{P}}) = s_\ell^2(S_{\mathbf{k}, \mathbb{P}}^*) \leq 4e_\ell^2(S_{\mathbf{k}, \mathbb{P}}^*).$$

□

The previous lemma allows us to bound eigenvalues of kernel matrices by ℓ_∞ -entropy numbers.

Proposition B.2 (Eigenvalue of kernel matrix is bounded by ℓ_∞ -entropy number). *Let \mathbf{k} be a jointly continuous kernel. Define $K \triangleq \mathbf{k}(\mathcal{S}_n, \mathcal{S}_n)$ for the sequence of points $\mathcal{S}_n = (x_1, \dots, x_n) \subset \mathbb{R}^d$. For any $w \in \Delta_{n-1}$, recall the notation $\mathbb{S}_n^w = \sum_{i \in [n]} w_i \delta_{x_i}$, $K^w = \text{diag}(\sqrt{w})K \text{diag}(\sqrt{w})$, and $R_n = 1 + \sup_{i \in [n]} \|x_i\|_2$. Then for all $\ell \in \mathbb{N}$,*

$$\lambda_\ell(K^w) \stackrel{(i)}{=} \lambda_\ell(T_{\mathbf{k}, \mathbb{S}_n^w}) \stackrel{(ii)}{\leq} 4e_\ell^2(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(R_n-1)}} \hookrightarrow \ell_\infty(\mathcal{B}_2(R_n-1))). \quad (8)$$

Proof of Prop. B.2. Without loss of generality, we assume $w_i > 0$ for all $i \in [n]$, since otherwise, we can consider a smaller set of points by removing the ones with zero weights.

Proof of equality (i) from display (8) Note that $\mathcal{L}^2(\mathbb{S}_n^w)$ is isometric to \mathbb{R}^n . Let $K \triangleq \mathbf{k}(\mathcal{S}_n, \mathcal{S}_n)$ denote the kernel matrix. The action of $T_{\mathbf{k}, \mathbb{S}_n^w}$ is given by, for $i \in [n]$,

$$T_{\mathbf{k}, \mathbb{S}_n^w} f(x_i) = \sum_{j \in [n]} w_j \mathbf{k}(x_i, x_j) f(x_j),$$

so in matrix form, $T_{\mathbf{k}, \mathbb{S}_n^w} f = K \text{diag}(w) f$, and hence $T_{\mathbf{k}, \mathbb{S}_n^w} = K \text{diag}(w)$. If λ_ℓ is an eigenvalue of $K \text{diag}(w)$ with eigenvector v_ℓ , then

$$\begin{aligned} K \text{diag}(w) v_\ell = \lambda_\ell v_\ell &\iff \text{diag}(\sqrt{w}) K \text{diag}(w) v_\ell = \lambda_\ell \text{diag}(\sqrt{w}) v_\ell \\ &\iff \text{diag}(\sqrt{w}) K \text{diag}(\sqrt{w}) (\text{diag}(\sqrt{w}) v_\ell) = \lambda_\ell \text{diag}(\sqrt{w}) v_\ell, \end{aligned}$$

where we used $w_i > 0$ for all $i \in [n]$. Hence the eigenspectrum of $T_{\mathbf{k}, \mathbb{S}_n^w}$ agrees with that of $\text{diag}(\sqrt{w}) K \text{diag}(\sqrt{w})$.

Proof of bound (ii) from display (8) By [Lem. B.4](#), we have $\lambda_\ell(T_{\mathbf{k}, \mathbb{S}_n^w}) \leq 4e_\ell^2(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(R_n-1)}} \hookrightarrow \mathcal{L}^2(\mathbb{S}_n^w))$. Finally, using [Def. B.1](#), we have $e_\ell(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(R_n-1)}} \hookrightarrow \mathcal{L}^2(\mathbb{S}_n^w)) \leq e_\ell(\mathcal{H}_{\mathbf{k}|_{\mathcal{B}_2(R_n-1)}} \hookrightarrow \ell_\infty(\mathcal{B}_2(R_n-1)))$ because \mathbb{S}_n^w is supported in $\mathcal{B}_2(R_n-1)$ and the fact that $\|\cdot\|_{\mathcal{L}^2(\mathbb{P})} \leq \|\cdot\|_\infty$ for any \mathbb{P} . □

Combining the tools developed so far, we have the following corollary for bounding the eigenvalues of POLYGROWTH and LOGGROWTH kernel matrices.

Corollary B.1 (Eigenvalue bound for POLYGROWTH or LOGGROWTH kernel matrix). *Suppose a kernel \mathbf{k} satisfies Assum. (α, β) -kernel. Let $\mathcal{S}_n = (x_1, \dots, x_n) \subset \mathbb{R}^d$ be a sequence of points. For any $w \in \Delta_{n-1}$, using the notation $L_{\mathbf{k}}$ from [\(6\)](#), for any $\ell > L_{\mathbf{k}}(R_n) + 1$, we have*

$$\lambda_\ell(K^w) \leq \begin{cases} 4 \left(\frac{L_{\mathbf{k}}(R_n)}{\ell-1} \right)^{\frac{2}{\alpha}} & \text{POLYGROWTH}(\alpha, \beta) \quad \text{and} \\ 4 \exp \left(2 - 2 \left(\frac{\ell-1}{L_{\mathbf{k}}(R_n)} \right)^{\frac{1}{\alpha}} \right) & \text{LOGGROWTH}(\alpha, \beta). \end{cases}$$

Proof of Cor. B.1. The claim follows by applying [Prop. B.2](#) and [Prop. B.1](#). □

B.2. Covering numbers and eigenvalues of Stein kernels

The goal of this section is to show that a Stein kernel \mathbf{k}_p satisfies Assum. (α, β) -kernel provided that the base kernel is sufficiently smooth and to derive the parameters α, β for POLYGROWTH and LOGGROWTH cases. Such results when put together with results from App. B.1 imply spectral decay rates for Stein kernel matrices.

For a Stein kernel \mathbf{k}_p with preconditioning matrix M , we define

$$\mathfrak{S}_p(r) \triangleq \max \left(1, \sup_{\|x\|_2 \leq r} \|M^{1/2} \nabla \log p(x)\|_2 \right). \quad (9)$$

We start by noting a useful alternative expression for a Stein kernel where we only need access to the density via the score $\nabla \log p$.

Proposition B.3 (Alternative expression for Stein kernel). *The Stein kernel \mathbf{k}_p has the following alternative form:*

$$\begin{aligned} \mathbf{k}_p(x, y) = & \langle \nabla \log p(x), M \nabla \log p(y) \rangle \mathbf{k}(x, y) + \langle \nabla \log p(x), M \nabla_y \mathbf{k}(x, y) \rangle + \\ & \langle \nabla \log p(y), M \nabla_x \mathbf{k}(x, y) \rangle + \text{tr}(M \nabla_x \nabla_y \mathbf{k}(x, y)), \end{aligned} \quad (10)$$

where $\nabla_x \nabla_y \mathbf{k}(x, y)$ denotes the $d \times d$ matrix $(\partial_{x_i} \partial_{y_j} \mathbf{k}(x, y))_{i, j \in [d]}$.

Proof of Prop. B.3. We compute

$$\begin{aligned} (\nabla_x \cdot (p(x) M \mathbf{k}(x, y) p(y)))_j &= \sum_{i \in [d]} M_{ij} (\partial_{x_i} p(x) \mathbf{k}(x, y) p(y) + p(x) \partial_{x_i} \mathbf{k}(x, y) p(y)). \\ \nabla_y \cdot \nabla_x \cdot (p(x) M \mathbf{k}(x, y) p(y)) &= \sum_{i, j \in [d]} M_{ij} (\partial_{x_i} p(x) \partial_{y_j} p(y) \mathbf{k}(x, y) + \partial_{x_i} p(x) \partial_{y_j} \mathbf{k}(x, y) p(y)) \\ &+ \sum_{i, j \in [d]} M_{ij} (p(x) \partial_{y_j} p(y) \partial_{x_i} \mathbf{k}(x, y) + p(x) \partial_{x_i} \partial_{y_j} \mathbf{k}(x, y) p(y)). \end{aligned}$$

The four terms in the last equation correspond to the four terms in (10). \square

In what follows, we will make use of a matrix-valued kernel $\mathbf{K} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ which generates an RKHS $\mathcal{H}_{\mathbf{K}}$ of vector-valued functions. See Carmeli et al. (2006) for an introduction to vector-valued RKHS theory.

Our next goal is to build a Hilbert-space isometry between the direct sum Hilbert space $\mathcal{H}_{\mathbf{k}}^{\oplus d}$ and $\mathcal{H}_{\mathbf{k}_p}$ to represent functions in $\mathcal{H}_{\mathbf{k}_p}$ using functions from $\mathcal{H}_{\mathbf{k}}$.

Lemma B.5 (Preconditioned matrix-valued RKHS from a scalar kernel). *Let $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be kernel and $\mathcal{H}_{\mathbf{k}}$ be the corresponding RKHS. Let $M \in \mathbb{R}^{d \times d}$ be an SPSD matrix. Consider the map $\iota : \mathcal{H}_{\mathbf{k}}^{\oplus d} \rightarrow \mathfrak{F}(\mathbb{R}^d, \mathbb{R}^d)$ defined by $(f_1, \dots, f_d) \mapsto [x \mapsto M^{1/2}(f_1(x), \dots, f_d(x))]$, where $\mathcal{H}_{\mathbf{k}}^{\oplus d}$ is the direct-sum Hilbert space of d copies of $\mathcal{H}_{\mathbf{k}}$. Then ι is a Hilbert-space isometry onto a vector-valued RKHS $\mathcal{H}_{\mathbf{K}}$ with matrix-valued reproducing kernel given by $\mathbf{K}(x, y) = \mathbf{k}(x, y)M$.*

Proof of Lem. B.5. Define $\gamma : \mathbb{R}^d \rightarrow \mathfrak{F}(\mathbb{R}^d, \mathcal{H}_{\mathbf{k}}^{\oplus d})$ via

$$\gamma(x)(\alpha) \triangleq \mathbf{k}(x, \cdot) M^{1/2} \alpha.$$

We have

$$\|\gamma(x)(\alpha)\|_{\mathcal{H}_{\mathbf{k}}^{\oplus d}} \leq \|\mathbf{k}(x, \cdot)\|_{\mathbf{k}} \|M^{1/2}\|_2 \|\alpha\|_2,$$

so $\gamma(x)$ is bounded. Since $\gamma(x)$ is also linear, we have $\gamma(x) \in \mathfrak{B}(\mathbb{R}^d, \mathcal{H}_{\mathbf{k}}^{\oplus d})$. Let $\gamma(x)^* : \mathcal{H}_{\mathbf{k}}^{\oplus d} \rightarrow \mathbb{R}^d$ denote the Hilbert-space adjoint of $\gamma(x)$. Then for any $(f_1, \dots, f_d) \in \mathcal{H}_{\mathbf{k}}^{\oplus d}$, $\alpha \in \mathbb{R}^d$, we have

$$\begin{aligned} \langle \gamma(x)^*(f_1, \dots, f_d), \alpha \rangle &= \langle (f_1, \dots, f_d), \gamma(x)(\alpha) \rangle_{\mathcal{H}_{\mathbf{k}}^{\oplus d}} \\ &= \langle (f_1, \dots, f_d), \mathbf{k}(x, \cdot) M^{1/2} \alpha \rangle_{\mathcal{H}_{\mathbf{k}}^{\oplus d}} \\ &= \langle (f_1(x), \dots, f_d(x)), M^{1/2} \alpha \rangle \\ &= \langle M^{1/2}(f_1(x), \dots, f_d(x)), \alpha \rangle. \end{aligned}$$

Hence $\gamma(x)^*(f_1, \dots, f_d) = M^{1/2}(f_1(x), \dots, f_d(x))$, so $\iota(f_1, \dots, f_d)(x) = \gamma(x)^*(f_1, \dots, f_d)$. By Carmeli et al. (2006, Proposition 2.4), we see that ι is a partial isometry from $\mathcal{H}_{\mathbf{k}}^{\oplus d}$ onto a vector-valued RKHS space with reproducing kernel $\mathbf{K}(x, y) = \gamma(x)^*\gamma(y) : \mathbb{R}^d \rightarrow \mathbb{R}^d$. For $\alpha \in \mathbb{R}^d$, previous calculation implies

$$\gamma(x)^*\gamma(y)(\alpha) = \gamma(x)^*(\mathbf{k}(y, \cdot)M^{1/2}\alpha) = M^{1/2}\mathbf{k}(y, x)M^{1/2}\alpha = \mathbf{k}(x, y)M.$$

□

Lemma B.6 (Stein operator is an isometry). *Consider a Stein kernel \mathbf{k}_p with base kernel \mathbf{k} and preconditioning matrix M . Then, the Stein operator \mathcal{S}_p defined by $\mathcal{S}_p(v) \triangleq \frac{1}{p}\nabla \cdot (pv)$ is an isometry from $\mathcal{H}_{\mathbf{K}}$ with $\mathbf{K} \triangleq \mathbf{k}M$ to $\mathcal{H}_{\mathbf{k}_p}$.*

Proof. This follows from Barp et al. (2022, Theorem 2.6) applied to \mathbf{K} . □

The previous two lemmas show that $\mathcal{S}_p \circ \iota$ is a Hilbert space isometry from $\mathcal{H}_{\mathbf{k}}^{\oplus d}$ to $\mathcal{H}_{\mathbf{k}_p}$. Note that $\mathcal{S}_p(v) = \langle \nabla \log p, h \rangle + \nabla \cdot h$. Hence, we immediately have

$$\mathcal{H}_{\mathbf{k}_p} = \left\{ \langle \nabla \log p, M^{1/2}f \rangle + \nabla \cdot (M^{1/2}f) : f = (f_1, \dots, f_d) \in \mathcal{H}_{\mathbf{k}}^{\oplus d} \right\}. \quad (11)$$

We next build a divergence RKHS which is one of the summands used to form $\mathcal{H}_{\mathbf{k}_p}$.

Lemma B.7 (Divergence RKHS). *Let $\mathbf{k} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuously differentiable kernel. Let M be an SPSD matrix. Define $\nabla^{\otimes 2} \cdot (M\mathbf{k}) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ via*

$$\nabla^{\otimes 2} \cdot (M\mathbf{k})(x, y) \triangleq \nabla_y \cdot \nabla_x \cdot (M\mathbf{k}(x, y)) = \text{tr}(M\nabla_x \nabla_y \mathbf{k}(x, y)). \quad (12)$$

Then $\nabla^{\otimes 2} \cdot (M\mathbf{k})$ is a kernel, and its RKHS $\mathcal{H}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ has the following explicit form

$$\mathcal{H}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})} = \nabla \cdot \mathcal{H}_{\mathbf{K}} = \left\{ \nabla \cdot (M^{1/2}f) : f = (f_1, \dots, f_d) \in \mathcal{H}_{\mathbf{k}}^{\oplus d} \right\}, \quad (13)$$

where $\mathbf{K} = M\mathbf{k}$. Moreover, $\nabla \cdot : \mathcal{H}_{\mathbf{K}} \rightarrow \mathcal{H}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ is an isometry.

Proof of Lem. B.7. First of all, by Steinwart and Christmann (2008, Corollary 4.36), every $f \in \mathcal{H}_{\mathbf{k}}$ is continuously differentiable, so $\partial_{x_i} f$ exists. By Lem. B.5, $\nabla \cdot \mathcal{H}_{\mathbf{K}}$ is well-defined and the right equality in (13) holds.

Define $\gamma : \mathbb{R}^d \rightarrow \mathfrak{F}(\mathbb{R}, \mathcal{H}_{\mathbf{K}})$ via

$$\gamma(x)(t) \triangleq t \sum_{i=1}^d \partial_{x_i} \mathbf{K}(x, \cdot) e_i,$$

where $e_i \in \mathbb{R}^d$ is the i th standard basis vector in \mathbb{R}^d ; by Barp et al. (2022, Lemma C.8) we have $\partial_{x_i} \mathbf{K}(x, \cdot) e_i \in \mathcal{H}_{\mathbf{K}}$. Note that

$$\|\gamma(x)(t)\|_{\mathbf{K}} = |t| \left\| \sum_{i=1}^d \partial_{x_i} \mathbf{K}(x, \cdot) e_i \right\|_{\mathbf{K}},$$

so $\gamma(x) \in \mathfrak{B}(\mathbb{R}, \mathcal{H}_{\mathbf{K}})$. The adjoint $\gamma(x)^* \in \mathfrak{B}(\mathcal{H}_{\mathbf{K}}, \mathbb{R})$ must satisfy, for any $h \in \mathcal{H}_{\mathbf{K}}$,

$$t\gamma(x)^*h = \langle h, \gamma(x)(t) \rangle_{\mathbf{K}} = \left\langle h, t \sum_{i=1}^d \partial_{x_i} \mathbf{K}(x, \cdot) e_i \right\rangle_{\mathbf{K}} = t\nabla \cdot h,$$

where we used the fact (Barp et al., 2022, Lemma C.8) that, for $c \in \mathbb{R}^d$, $h \in \mathcal{H}_{\mathbf{K}}$, $\langle \partial_{x_i} \mathbf{K}(x, \cdot) c, h \rangle = c^\top \partial_{x_i} h(x)$. So we find $\gamma(x)^*(h) = \nabla \cdot h(x)$. By Carmeli et al. (2006, Proposition 2.4), the map $A : \mathcal{H}_{\mathbf{K}} \rightarrow \mathfrak{F}(\mathbb{R}^d, \mathbb{R})$ defined by $A(h)(x) = \gamma^*(x)(h) = \nabla \cdot h(x)$, i.e., $A = \nabla \cdot$, is a partial isometry from $\mathcal{H}_{\mathbf{K}}$ to an RKHS $\mathcal{H}_{\nabla \cdot \mathbf{K}}$ with reproducing kernel

$$\gamma(x)^*\gamma(y) = \nabla \cdot \left(\sum_{i=1}^d \partial_{x_i} \mathbf{K}(x, \cdot) e_i \right) (y) = \nabla_y \cdot \nabla_x \cdot \mathbf{K}(x, y) = \nabla^{\otimes 2} \cdot (M\mathbf{k})(x, y).$$

□

The following lemma shows that we can project a covering of $\mathcal{B}_{\mathbf{k}}$ consisting of arbitrary functions to a covering using functions only in $\mathcal{B}_{\mathbf{k}}$ while inflating the covering radius by at most 2.

Lemma B.8 (Projection of coverings into RKHS balls). *Let \mathbf{k} be a kernel, $A \subset \mathbb{R}^d$ be a set, and $\epsilon > 0$. Let \mathcal{C} be a set of functions such that for any $f \in \mathcal{B}_{\mathbf{k}}$, there exists $g \in \mathcal{C}$ such that $\|f - g\|_{\infty, A} \leq \epsilon$. Then*

$$\mathcal{N}_{\mathbf{k}}(A, 2\epsilon) \leq |\mathcal{C}|.$$

Proof. We will build a $(\mathbf{k}, A, 2\epsilon)$ covering \mathcal{C}' as follows. For any $h \in \mathcal{C}$, if there exists $h' \in \mathcal{B}_{\mathbf{k}}$ with $\|h' - h\|_{\infty, A} \leq \epsilon$, then we include h' in \mathcal{C}' . By construction, $|\mathcal{C}'| \leq |\mathcal{C}|$. Then, for any $f \in \mathcal{B}_{\mathbf{k}}$, by assumption, there exists $g \in \mathcal{C}$ such that $\|f - g\|_{\infty, A} \leq \epsilon$. By construction, there exists $g' \in \mathcal{C}'$ such that $\|g' - g\|_{\infty, A} \leq \epsilon$. Thus

$$\|f - g'\|_{\infty, A} \leq \|f - g\|_{\infty, A} + \|g - g'\|_{\infty, A} \leq 2\epsilon.$$

Hence \mathcal{C}' is a $(\mathbf{k}, A, 2\epsilon)$ covering. \square

We are now ready to bound the covering numbers of \mathbf{k}_p by those of \mathbf{k} and $\nabla^{\otimes 2} \cdot (M\mathbf{k})$. Our key insight towards this end is that any element in $\mathcal{H}_{\mathbf{k}_p}$ can be decomposed as a sum of functions originated from $\mathcal{H}_{\mathbf{k}}$ and a function from the divergence RKHS $\mathcal{H}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$.

Lemma B.9 (Upper bounding covering number of Stein kernel with that of its base kernel). *Let \mathbf{k}_p be a Stein kernel with density p and preconditioning matrix M . For any $A \subset \mathbb{R}^d$, $\epsilon_1, \epsilon_2 > 0$,*

$$\mathcal{N}_{\mathbf{k}_p}(A, \epsilon) \leq \mathcal{N}_{\mathbf{k}}(A, \epsilon_1)^d \mathcal{N}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(A, \epsilon_2),$$

for $\epsilon = 2(\sqrt{d}\epsilon_1 \sup_{x \in A} \|M^{1/2} \nabla \log p(x)\| + \epsilon_2)$.

Proof of Lem. B.9. Let $\mathcal{C}_{\mathbf{k}}$ be a $(\mathbf{k}, A, \epsilon_1)$ covering and $\mathcal{C}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ be a $(\nabla^{\otimes 2} \cdot (M\mathbf{k}), A, \epsilon_2)$ covering. Define $b \triangleq M^{1/2} \nabla \log p$. Form

$$\mathcal{C} \triangleq \left\{ \langle b, \tilde{f} \rangle + \tilde{g} : \tilde{f} = (\tilde{f}_1, \dots, \tilde{f}_d) \in (\mathcal{C}_{\mathbf{k}})^d, \tilde{g} \in \mathcal{C}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})} \right\} \subset \mathfrak{F}(\mathbb{R}^d, \mathbb{R}).$$

Then $|\mathcal{C}| \leq |\mathcal{C}_{\mathbf{k}}|^d |\mathcal{C}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}|$. Let $\mathbf{K} \triangleq \mathbf{k}M$. For any $h \in \mathcal{B}_{\mathbf{k}_p}$, by (11), we can find $f = (f_1, \dots, f_d) \in \mathcal{H}_{\mathbf{k}}^{\oplus d}$ with $f_i \in \mathcal{H}_{\mathbf{k}}$ such that

$$h = \mathcal{S}_p \circ \iota(f) = \langle \nabla \log p, M^{1/2} f \rangle + \nabla \cdot (M^{1/2} f) = \langle b, f \rangle + \nabla \cdot (M^{1/2} f).$$

Since ι and \mathcal{S}_p are isometries, we have $f \in \mathcal{B}_{\mathcal{H}_{\mathbf{k}}^{\oplus d}}$. Since, for each i ,

$$\|f_i\|_{\mathbf{k}} \leq \sqrt{\sum_{j=1}^d \|f_j\|_{\mathbf{k}}^2} = \|f\|_{\mathcal{H}_{\mathbf{k}}^{\oplus d}} \leq 1,$$

we have $f_i \in \mathcal{B}_{\mathbf{k}}$. By Lem. B.7, $\nabla \cdot : \mathcal{H}_{\mathbf{K}} \rightarrow \mathcal{H}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ is also an isometry, so $\nabla \cdot (M^{1/2} f) \in \mathcal{B}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$. Thus there exist $\tilde{f}_i \in \mathcal{C}_{\mathbf{k}}$ for each i and $\tilde{g} \in \mathcal{C}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ such that

$$\|f_i - \tilde{f}_i\|_{\infty, A} \leq \epsilon_1, \quad \|\nabla \cdot (M^{1/2} f) - \tilde{g}\|_{\infty, A} \leq \epsilon_2.$$

Let

$$\tilde{h}(x) \triangleq \langle b, \tilde{f} \rangle + \tilde{g} \in \mathcal{C}.$$

Then for $x \in A$,

$$\begin{aligned} |h(x) - \tilde{h}(x)| &= \left| \langle b(x), f(x) - \tilde{f}(x) \rangle + \nabla \cdot (M^{1/2} f(x)) - \tilde{g}(x) \right| \\ &\leq \|b(x)\| \sqrt{\sum_{i=1}^d (f_i(x) - \tilde{f}_i(x))^2} + |\nabla \cdot (M^{1/2} f(x)) - \tilde{g}(x)| \\ &\leq \sqrt{d}\epsilon_1 \|b(x)\| + \epsilon_2. \end{aligned}$$

Hence

$$\|h - \tilde{h}\|_{\infty, A} \leq \sqrt{d}\epsilon_1 \sup_{x \in A} \|b(x)\| + \epsilon_2 \triangleq \epsilon_3.$$

Note that \mathcal{C} that we constructed is not necessarily contained in $\mathcal{B}_{\mathbf{k}_p}$. By Lem. B.8, we can get a $(\mathbf{k}_p, A, 2\epsilon_3)$ covering and we are done. \square

Corollary B.2 (Log-covering number bound for Stein kernel). *Let \mathbf{k}_p be a Stein kernel and $A \subset \mathbb{R}^d$. For any $r > 0$, $\epsilon > 0$,*

$$\log \mathcal{N}_{\mathbf{k}_p}(A, \epsilon) \leq d \log \mathcal{N}_{\mathbf{k}} \left(A, \frac{\epsilon}{4\sqrt{d}\mathfrak{S}_p(r)} \right) + \log \mathcal{N}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})} \left(A, \frac{\epsilon}{4} \right),$$

where \mathfrak{S}_p is defined in (9).

Proof. This is direct from Lem. B.9 with $\epsilon_1 = \frac{\epsilon}{4\sqrt{d}\mathfrak{S}_p(r)}$, $\epsilon_2 = \frac{\epsilon}{4}$. \square

B.2.1. CASE OF DIFFERENTIABLE BASE KERNEL

Definition B.2 (*s*-times continuously differentiable kernel). *A kernel \mathbf{k} is *s*-times continuously differentiable for $s \in \mathbb{N}$ if all partial derivatives $\partial^{\alpha, \alpha} \mathbf{k}$ exist and are continuous for all multi-indices $\alpha \in \mathbb{N}_0^d$ with $|\alpha| \leq s$.*

Proposition B.4 (Covering number bound for \mathbf{k}_p with differentiable base kernel). *Suppose \mathbf{k}_p is a Stein kernel with an *s*-times continuously differentiable base kernel \mathbf{k} for $s \geq 2$. Then there exist a constant $\mathfrak{C}_d > 0$ depending only on (s, d, \mathbf{k}, M) such that for any $r > 0$, $\epsilon \in (0, 1)$,*

$$\log \mathcal{N}_{\mathbf{k}_p}(\mathcal{B}_2(r), \epsilon) \leq \mathfrak{C}_d r^d \mathfrak{S}_p^{d/s}(r) (1/\epsilon)^{d/(s-1)}.$$

Proof of Prop. B.4. Since \mathbf{k} is *s*-times continuously differentiable, the divergence kernel $\nabla^{\otimes 2} \cdot (M\mathbf{k})$ is $(s-1)$ -times continuously differentiable. By Dwivedi and Mackey (2022, Proposition 2(b)), there exists constants c_1, c_2 depending only on (s, d, \mathbf{k}, M) such that, for any $r > 0$, $\epsilon_1, \epsilon_2 > 0$,

$$\begin{aligned} \log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon_1) &\leq c_1 r^d (1/\epsilon)^{d/s}, \\ \log \mathcal{N}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(\mathcal{B}_2(r), \epsilon_2) &\leq c_2 r^d (1/\epsilon)^{d/(s-1)}. \end{aligned}$$

By Cor. B.2 with $A = \mathcal{B}_2(r)$, we have, for any $r > 0$ and $\epsilon \in (0, 1)$,

$$\begin{aligned} \log \mathcal{N}_{\mathbf{k}_p}(\mathcal{B}_2(r), \epsilon) &\leq c_1 d r^d (4\sqrt{d}\mathfrak{S}_p(r))^{d/s} (1/\epsilon)^{d/\epsilon} + c_2 r^d (4/\epsilon)^{d/(s-1)} \\ &\leq \mathfrak{C}_d r^d \mathfrak{S}_p^{d/s}(r) (1/\epsilon)^{d/(s-1)} \end{aligned}$$

for some $\mathfrak{C}_d > 0$ depending only on (s, d, \mathbf{k}, M) . \square

B.2.2. CASE OF RADIALLY ANALYTIC BASE KERNEL

For a symmetric positive definite $M \in \mathbb{R}^{d \times d}$, we define, for $x \in \mathbb{R}^d$,

$$\|x\|_M \triangleq \sqrt{x^\top M^{-1} x}.$$

Definition B.3 (Radially analytic kernel). *A kernel \mathbf{k} is radially analytic if \mathbf{k} satisfies $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2)$ for a symmetric positive definite matrix $M \in \mathbb{R}^{d \times d}$ and a function $\kappa : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ real-analytic everywhere with convergence radius $R_\kappa > 0$ such that there exists a constant $C_\kappa > 0$ for which*

$$\left| \frac{1}{j!} \kappa_+^{(j)}(0) \right| \leq C_\kappa (2/R_\kappa)^j, \text{ for all } j \in \mathbb{N}_0, \quad (14)$$

where $\kappa_+^{(j)}$ indicates the *j*-th right-sided derivative of κ .

Example B.1 (Gaussian kernel). Consider the Gaussian kernel $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2)$ with $\kappa(t) = e^{-\frac{t}{2\sigma^2}}$ where $\sigma > 0$. Note the exponential function is real-analytic everywhere, and so is κ . Since $\kappa(t) = \sum_{j=0}^{\infty} \frac{(-t/2\sigma^2)^j}{j!}$, we find $\frac{1}{j!} \kappa^{(j)}(0) = \frac{(-1)^j}{j(2\sigma^2)^j}$. Hence (14) holds with $C_\kappa = 1$ and $R_\kappa = 2 \inf_{j \geq 0} (j(2\sigma^2)^j)^{1/j} = 4\sigma^2$.

Example B.2 (IMQ kernel). Consider the inverse multiquadric kernel $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2)$ with $\kappa(t) = (c^2 + t)^{-\beta}$ where $c, \beta > 0$. By Sun and Zhou (2008, Example 3), κ is real-analytic everywhere with $C_\kappa = c^{-2\beta} (2\beta + 1)^{\beta+1}$ and $R_\kappa = c^2$.

A simple calculation yields the following lemma.

Proposition B.5 (Expression for \mathbf{k}_p with a radially analytic base kernel). *Suppose a Stein kernel \mathbf{k}_p has a symmetric positive definite preconditioning matrix and a base kernel $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2)$ where κ is twice-differentiable. Then*

$$\begin{aligned} \mathbf{k}_p(x, y) &= \langle \nabla \log p(x), M \nabla \log p(y) \rangle \kappa(\|x - y\|_M^2) - \\ &2\kappa'_+(\|x - y\|_M^2) \langle x - y, \nabla \log p(x) - \nabla \log p(y) \rangle - \\ &4\kappa''_+(\|x - y\|_M^2) \|x - y\|_M^2 - 2d\kappa'_+(\|x - y\|_M^2). \end{aligned} \quad (15)$$

In particular,

$$\mathbf{k}_p(x, x) = \|M^{1/2} \nabla \log p(x)\|_2^2 \kappa(0) - 2d\kappa'_+(0).$$

Proof of Prop. B.5. From $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2) = \kappa((x - y)^\top M^{-1}(x - y))$, we compute, using (12),

$$\begin{aligned} \nabla_y \mathbf{k}(x, y) &= -2\kappa'_+(\|x - y\|_M^2) M^{-1}(x - y) \\ \nabla_x \nabla_y \mathbf{k}(x, y) &= -2\kappa'_+(\|x - y\|_M^2) M^{-1} - 4\kappa''_+(\|x - y\|_M^2) M^{-1}(x - y)((x - y)M^{-1})^\top \\ \nabla^{\otimes 2} \cdot (M\mathbf{k})(x, y) &= \text{tr}(M \nabla_x \nabla_y \mathbf{k}(x, y)) = -4\kappa''_+(\|x - y\|_M^2) \|x - y\|_M^2 - 2d\kappa'_+(\|x - y\|_M^2). \end{aligned} \quad (16)$$

The form (15) then follows from applying Prop. B.3. \square

We next show that the divergence kernel $\nabla^{\otimes 2} \cdot (M\mathbf{k})$ is radially analytic given that \mathbf{k} is.

Lemma B.10 (Convergence radius of divergence kernel). *Let \mathbf{k} be a radially analytic kernel with the corresponding real-analytic function κ , convergence radius R_κ with constant C_κ , and a symmetric positive definite matrix M . Then*

$$\nabla^{\otimes 2} \cdot (M\mathbf{k})(x, y) = \kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(\|x - y\|_M^2),$$

where $\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is the real-analytic function defined as

$$\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(t) \triangleq -4\kappa''_+(t)t - 2d\kappa'_+(t).$$

Moreover, $\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ has convergence radius with constant

$$R_{\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}} = \frac{R_\kappa}{4d+8}, \quad C_{\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}} = 4dC_\kappa/R_\kappa.$$

Proof of Lem. B.10. The first statement regarding the form of $\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ directly follows from (16). Next, iterative differentiation yields, for $j \in \mathbb{N}_0$,

$$\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}^{(j)}(t) = -(2d + 4j)\kappa_+^{(j+1)}(t) - 4\kappa_+^{(j+2)}(t)t.$$

Thus

$$\begin{aligned} \left| \frac{1}{j!} \kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}^{(j)}(0) \right| &= \frac{2d+4j}{j!} \kappa_+^{(j+1)}(0) \\ &= \frac{(2d+4j)(j+1)}{(j+1)!} \kappa_+^{(j+1)}(0) \\ &\leq (2d + 4j)(j + 1)C_\kappa(2/R_\kappa)^{j+1}. \end{aligned} \quad (17)$$

For $j \geq 1$,

$$\begin{aligned} \left| \frac{1}{j!} \kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}^{(j)}(0) \right| &\leq (2C_\kappa/R_\kappa) \left(((2d + 4j)(j + 1))^{1/j} 2/R_\kappa \right)^j \\ &\leq (2C_\kappa/R_\kappa) \left((2(2d + 4) \cdot 2/R_\kappa)^j \right). \end{aligned}$$

where we used the fact that $((2d + 4j)(j + 1))^{1/j}$ is decreasing in j . For $j = 0$, (17) is just $2dC_\kappa \cdot 2/R_\kappa$. Hence $\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}$ is analytic with $C_{\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}} = 4dC_\kappa/R_\kappa$ and $R_{\kappa_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}} = \frac{R_\kappa}{4d+8}$. \square

We will use the following lemma repeatedly.

Lemma B.11 (Covering number of radially analytic kernel with M -metric). *Let \mathbf{k}_0 be a radially analytic kernel with $\mathbf{k}_0(x, y) = \kappa(\|x - y\|_2^2)$. For any symmetric positive definite $M \in \mathbb{R}^{d \times d}$, consider the radially analytic kernel $\mathbf{k}(x, y) \triangleq \kappa(\|x - y\|_M^2)$. Then for any $A \subset \mathbb{R}^d$ and $\epsilon > 0$, we have*

$$\mathcal{N}_{\mathbf{k}}(M^{-1/2}(A), \epsilon) = \mathcal{N}_{\mathbf{k}_0}(A, \epsilon).$$

In particular, for any $r > 0$,

$$\mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) \leq \mathcal{N}_{\mathbf{k}_0}(\mathcal{B}_2(r\|M^{1/2}\|_2), \epsilon).$$

Proof. Note that $\mathbf{k}(x, y) = \mathbf{k}_0(M^{-1/2}x, M^{-1/2}y)$. By [Paulsen and Raghupathi \(2016, Theorem 5.7\)](#), $\mathcal{H}_{\mathbf{k}} = \{f \circ M^{-1/2} : f \in \mathcal{H}_{\mathbf{k}_0}\}$, and moreover $\mathcal{B}_{\mathbf{k}} = \{f \circ M^{-1} : f \in \mathcal{B}_{\mathbf{k}_0}\}$. Let \mathcal{C}_0 be a $(\mathbf{k}_0, A, \epsilon)$ covering. Form $\mathcal{C} = \{h \circ M^{-1/2} : h \in \mathcal{C}_0\} \subset \mathcal{B}_{\mathbf{k}}$. For any element $f \circ M^{-1/2} \in \mathcal{B}_{\mathbf{k}}$ where $f \in \mathcal{B}_{\mathbf{k}_0}$, there exists $h \in \mathcal{C}_0$ such that $\|f - h\|_{\infty, A} \leq \epsilon$. Thus

$$\|f \circ M^{-1/2} - h \circ M^{-1/2}\|_{\infty, M^{-1/2}(A)} = \|f - h\|_{\infty, A} \leq \epsilon.$$

Thus $\mathcal{N}_{\mathbf{k}}(M^{-1/2}(A), \epsilon) \leq \mathcal{N}_{\mathbf{k}_0}(A, \epsilon)$. By considering M^{-1} in place of M , we get our desired equality.

For the second statement, by letting $A = M^{1/2}\mathcal{B}_2(r)$, we have

$$\mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) = \mathcal{N}_{\mathbf{k}_0}(M^{1/2}\mathcal{B}_2(r), \epsilon) \leq \mathcal{N}_{\mathbf{k}_0}(\mathcal{B}_2(r\|M^{1/2}\|_2), \epsilon),$$

where we use the fact that $M^{1/2}\mathcal{B}_2(r) \subset \mathcal{B}_2(r\|M^{1/2}\|_2)$. □

In the next lemma, we rephrase the result from [Sun and Zhou \(2008, Theorem 2\)](#) for bounding the covering number of a radially analytic kernel.

Lemma B.12 (Covering number bound for radially analytic kernel). *Let \mathbf{k} be a radially analytic kernel with $\mathbf{k}(x, y) = \kappa(\|x - y\|_2^2)$. Then, there exist a polynomial $P(r)$ of degree $2d$ and a constant C depending only on (κ, d) such that for any $r > 0$, $\epsilon \in (0, 1/2)$,*

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) \leq P(r)(\log(1/\epsilon) + C)^{d+1}.$$

Proof of Lem. B.12. Let R_κ, C_κ denote the constants for κ as in (14). By and [Sun and Zhou \(2008, Theorem 2\)](#) with $R = 1$, $D = 2r$, and Lem. B.2, for $\epsilon \in (0, 1/2)$, we have

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon) \leq N_2(\mathcal{B}_2(r), r^\dagger/2) \left(4 \log(1/\epsilon) + 2 + 4 \log(16\sqrt{C_\kappa} + 1)\right)^{d+1},$$

where $r^\dagger = \min(\frac{\sqrt{R_\kappa}}{2d}, \sqrt{R_\kappa + (2r)^2} - 2r)$, and $N_2(\mathcal{B}_2(r), r^\dagger/2)$ is the covering number of $\mathcal{B}_2(r)$ as a subset of \mathbb{R}^d , which can be further bounded by ([Wainwright, 2019, \(5.8\)](#))

$$N_2(\mathcal{B}_2(r), r^\dagger/2) \leq \left(1 + \frac{4r}{r^\dagger}\right)^d.$$

If $r^\dagger = \sqrt{R_\kappa + (2r)^2} - 2r$, then $\frac{r}{r^\dagger} = \frac{r}{\sqrt{R_\kappa + (2r)^2} - 2r} = \frac{r(\sqrt{R_\kappa + (2r)^2} + 2r)}{R_\kappa} \leq \frac{r(\sqrt{R_\kappa} + 4r)}{R_\kappa}$ which is a quadratic polynomial in r . Hence for a constant $C > 0$ and a polynomial $P(r)$ of degree $2d$ that depend only on (κ, d) , we have the claim. □

Proposition B.6 (Covering number bound for \mathbf{k}_p with radially analytic base kernel). *Suppose \mathbf{k}_p is a Stein kernel with a preconditioning matrix M and a radially analytic base kernel \mathbf{k} based on a real-analytic function κ . Then there exist a constant $C > 0$ and a polynomial $P(r)$ of degree $2d$ depending only on (κ, d, M) such that for any $r > 0$, $\epsilon \in (0, 1)$,*

$$\log \mathcal{N}_{\mathbf{k}_p}(\mathcal{B}_2(r), \epsilon) \leq \left(\log \frac{\mathfrak{S}_p(r)}{\epsilon} + C\right)^{d+1} P(r). \quad (18)$$

Proof of Prop. B.6. Recall $\mathbf{k}(x, y) = \kappa(\|x - y\|_M^2)$. Consider $\mathbf{k}_0(x, y) \triangleq \kappa(\|x - y\|_2^2)$. For $\epsilon_1 \in (0, 1/2)$, by Lem. B.11, we have

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r/\|M^{1/2}\|_2), \epsilon_1) \leq \log \mathcal{N}_{\mathbf{k}_0}(\mathcal{B}_2(r), \epsilon_1/2).$$

Thus by Lem. B.12, there exists a polynomial $P_{\mathbf{k}}(r)$ of degree $2d$ and a constant $C_{\mathbf{k}}$ depending only on (κ, d, M) such that

$$\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(r), \epsilon_1) \leq P_{\mathbf{k}}(r)(\log(1/\epsilon_1) + C_{\mathbf{k}})^{d+1}$$

Similarly, for $\epsilon_2 \in (0, 1/2)$, by Lem. B.10 and Lem. B.12, we have, for a constant $C_{\nabla^{\otimes 2} \cdot (M\mathbf{k})} > 0$ and a polynomial $P_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(r)$ of degree $2d$ that depend only on (κ, d, M) ,

$$\log \mathcal{N}_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(\mathcal{B}_2(r), \epsilon_2) \leq P_{\nabla^{\otimes 2} \cdot (M\mathbf{k})}(r)(\log(1/\epsilon_2) + C_{\nabla^{\otimes 2} \cdot (M\mathbf{k})})^{d+1}.$$

For a given $\epsilon \in (0, 1)$, let $\epsilon_1 = \frac{\epsilon}{4\sqrt{d}\mathfrak{S}_p(r)}$ and $\epsilon_2 = \frac{\epsilon}{4}$. Then since $\mathfrak{S}_p \geq 1$, we have $\epsilon_1, \epsilon_2 \in (0, 1/2)$. By Cor. B.2 with $A = \mathcal{B}_2(r)$, we obtain, for a constants $C > 0$ and a polynomial $P(r)$ of degree $2d$ that depend only on (κ, d, M) ,

$$\log \mathcal{N}_{\mathbf{k}_p}(\mathcal{B}_2(r), \epsilon) \leq P(r)(\log(1/\epsilon) + \log \mathfrak{S}_p(r) + C)^{d+1}.$$

Hence (18) is shown. \square

When $\log \mathfrak{S}_p(r)$ grows polynomially in r , we apply Prop. B.6 to immediately obtain the following.

Corollary B.3. *Under the assumption of Prop. B.6, suppose $\mathfrak{S}_p(r) = O(\text{poly}(r))$. Then for any $\delta > 0$, there exists $\mathfrak{C}_d > 0$ such that*

$$\log \mathcal{N}_{\mathbf{k}_p}(\mathcal{B}_2(r), \epsilon) \leq \mathfrak{C}_d \log(e/\epsilon)^{d+1} (r+1)^{2d+\delta}.$$

Proof of Cor. B.3. This immediately follows from Prop. B.6 by using $\delta > 0$ to absorb the $\log \mathfrak{S}_p(r) = O(r^\delta)$ term. \square

B.2.3. PROOF OF PROP. 1: STEIN KERNEL GROWTH RATES

This follows from Prop. B.4 and Cor. B.3, and by noticing that if $\sup_{\|x\|_2 \leq r} \|\nabla \log p(x)\|_2$ is bounded by a degree d_ℓ polynomial, then so is

$$\mathfrak{S}_p(r) = \sup_{\|x\|_2 \leq r} \|M^{1/2} \nabla \log p(x)\|_2 \leq \|M^{1/2}\|_2 \sup_{\|x\|_2 \leq r} \|\nabla \log p(x)\|_2.$$

\square

C. Analysis of Debiasing Benchmarks

We start with a result on the MMD quality of i.i.d. sample points in App. C.1, followed by the proofs of our i.i.d.-like and better-than-i.i.d. guarantees for debiasing from Thms. 1 and 2 in Apps. C.2 and C.3 respectively.

C.1. MMD of unbiased i.i.d. sample points

We start by showing that sequence of n points sampled i.i.d. from \mathbb{P} achieves $\Theta(n^{-1})$ squared $\text{MMD}_{\mathbf{k}_\mathbb{P}}$ to \mathbb{P} in expectation.

Proposition C.1 (MMD of unbiased i.i.d. sample points). *Let $\mathbf{k}_\mathbb{P}$ be a kernel satisfying Assum. 1 with $\mathfrak{p} \geq 1$. Let $\mathcal{S}_n = (x_i)_{i \in [n]}$ be n i.i.d. samples from \mathbb{P} . Then*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_\mathbb{P}}(\mathcal{S}_n, \mathbb{P})^2] = \frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_\mathbb{P}(x, x)]}{n}.$$

Proof of Prop. C.1. We compute

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_\mathbb{P}}(\mathcal{S}_n, \mathbb{P})^2] = \mathbb{E}[\sum_{i, j \in [n]} \frac{1}{n^2} \mathbf{k}_\mathbb{P}(x_i, x_j)] = \frac{1}{n^2} \sum_{i, j \in [n]} \mathbb{E}[\mathbf{k}_\mathbb{P}(x_i, x_j)] = \frac{1}{n} \mathbb{E}[\mathbf{k}_\mathbb{P}(x_1, x_1)],$$

where we used the fact that $\mathbf{k}_\mathbb{P}$ is mean-zero with respect to \mathbb{P} and the independence of x_i, x_j for $i \neq j$. \square

C.2. Proof of Thm. 1: Debiasing to i.i.d. quality via simplex reweighting

We make use of the self-normalized importance sampling weights $w_j^{\text{SNIS}} = \frac{d\mathbb{P}}{d\mathbb{Q}}(x_j) / \sum_{i \in [n]} \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i)$ for $j \in [n]$ in our proofs. Notice that $(w_1^{\text{SNIS}}, \dots, w_n^{\text{SNIS}})^\top \in \Delta_{n-1}$ and hence

$$\text{MMD}_{\text{OPT}} \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_i^{\text{SNIS}} \delta_{x_i}, \mathbb{P}) = \frac{\|\sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot)\|_{\mathbf{k}_{\mathbb{P}}}}{\sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i)} = \frac{\|\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot)\|_{\mathbf{k}_{\mathbb{P}}}}{\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i)}.$$

Introduce the bounded in probability notation $X_n = O_p(g_n)$ to mean $\Pr(|X_n/g_n| > C_\epsilon) \leq \epsilon$ for all $n \geq N_\epsilon$ for any $\epsilon > 0$. Then we claim that under the conditions assumed in Thm. 1,

$$\|\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot)\|_{\mathbf{k}_{\mathbb{P}}} = O_p(n^{-\frac{1}{2}}) \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \rightarrow 1 \quad \text{almost surely,} \quad (19)$$

so that by Slutsky's theorem (Wellner, 2013, Ex. 1.4.7), we have $\text{MMD}_{\text{OPT}} = O_p(n^{-\frac{1}{2}})$ as desired. We prove the claims in (19) in two main steps: (a) first, we construct a weighted RKHS and then (b) establish a central limit theorem (CLT) that allows us to conclude both claims from (19)

Constructing a weighted and separable RKHS Define the kernel $\mathbf{k}_{\mathbb{Q}}(x, y) \triangleq \frac{d\mathbb{P}}{d\mathbb{Q}}(x) \mathbf{k}_{\mathbb{P}}(x, y) \frac{d\mathbb{P}}{d\mathbb{Q}}(y)$ with Hilbert space $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}} = \frac{d\mathbb{P}}{d\mathbb{Q}} \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ and the elements $\xi_i \triangleq \mathbf{k}_{\mathbb{Q}}(x_i, \cdot) = \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot) \frac{d\mathbb{P}}{d\mathbb{Q}}(\cdot) \in \mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$ for each $i \in \mathbb{N}$. By Paulsen and Raghupathi (2016, Prop. 5.20), any element in $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$ is represented by $\frac{d\mathbb{P}}{d\mathbb{Q}} f$ for some $f \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ and moreover, $f \mapsto \frac{d\mathbb{P}}{d\mathbb{Q}} f$ preserves inner product between the two RKHSs, i.e., $\langle f, g \rangle_{\mathbf{k}_{\mathbb{P}}} = \langle \frac{d\mathbb{P}}{d\mathbb{Q}} f, \frac{d\mathbb{P}}{d\mathbb{Q}} g \rangle_{\mathbf{k}_{\mathbb{Q}}}$ for $f, g \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$, which in turn implies $\|f\|_{\mathbf{k}_{\mathbb{P}}} = \|\frac{d\mathbb{P}}{d\mathbb{Q}} f\|_{\mathbf{k}_{\mathbb{Q}}}$. As a result, we also have that

$$\|\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot)\|_{\mathbf{k}_{\mathbb{P}}} = \|\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot) \frac{d\mathbb{P}}{d\mathbb{Q}}(\cdot)\|_{\mathbf{k}_{\mathbb{Q}}} = \|\frac{1}{n} \sum_{i=1}^n \xi_i\|_{\mathbf{k}_{\mathbb{Q}}}. \quad (20)$$

Since $\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ is separable, there exists a dense countable subset $(f_n)_{n \in \mathbb{N}} \subset \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$. For any $\frac{d\mathbb{P}}{d\mathbb{Q}} f \in \mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$, there exists $\{n_k\}_{k \in \mathbb{N}}$ such that $\lim_{k \rightarrow \infty} \|f_{n_k} - f\|_{\mathbf{k}_{\mathbb{P}}} = 0$. Since $\|\frac{d\mathbb{P}}{d\mathbb{Q}} f_{n_k} - \frac{d\mathbb{P}}{d\mathbb{Q}} f\|_{\mathbf{k}_{\mathbb{Q}}} = \|\frac{d\mathbb{P}}{d\mathbb{Q}}(f_{n_k} - f)\|_{\mathbf{k}_{\mathbb{Q}}} = \|f_{n_k} - f\|_{\mathbf{k}_{\mathbb{P}}}$ due to inner-product preservation, we thus have $\lim_{k \rightarrow \infty} \|\frac{d\mathbb{P}}{d\mathbb{Q}} f_{n_k} - \frac{d\mathbb{P}}{d\mathbb{Q}} f\|_{\mathbf{k}_{\mathbb{Q}}} = \lim_{k \rightarrow \infty} \|f_{n_k} - f\|_{\mathbf{k}_{\mathbb{P}}} = 0$, so $(\frac{d\mathbb{P}}{d\mathbb{Q}} f_n)_{n \in \mathbb{N}}$ is dense in $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$, showing that $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$ is separable.

Harris recurrence of the chain $(x_i)_{i \in \mathbb{N}}$ Let μ_1 denote the distribution of x_1 . Since $\mathcal{S}_\infty = (x_i)_{i=1}^\infty$ is a homogeneous ϕ -irreducible geometrically ergodic Markov chain with stationary distribution \mathbb{Q} , it is also positive (Meyn and Tweedie, 2012, Ch. 10) by definition and aperiodic by Douc et al. (2018, Lem. 9.3.9). Moreover, since \mathcal{S}_∞ is ϕ -irreducible, aperiodic, and geometrically ergodic in the sense of Gallegos-Herrada et al. (2023, Thm. 1) and μ_1 is absolutely continuous with respect to \mathbb{P} , we will assume, without loss of generality, that \mathcal{S}_∞ is Harris recurrent (Meyn and Tweedie, 2012, Ch. 9), since, by Qin (2023, Lem. 9), \mathcal{S}_∞ is equal to a geometrically ergodic Harris chain with probability 1.

CLT for $\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i$ We now show that $\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i$ converges to a Gaussian random element taking values in $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$. We separate the proof in two parts: first when the initial distribution $\mu_1 = \mathbb{Q}$ and next when $\mu_1 \neq \mathbb{Q}$.

Case 1: $\mu_1 = \mathbb{Q}$ When $\mu_1 = \mathbb{Q}$, \mathcal{S}_∞ is a strictly stationary chain. By Bradley (2005, Thm. 3.7 and (1.11)), since \mathcal{S}_∞ is geometrically ergodic, its strong mixing coefficients $(\tilde{\alpha}_i)_{i \in \mathbb{N}}$ satisfy $\tilde{\alpha}_i \leq C\rho^i$ for some $C > 0$ and $\rho \in [0, 1)$ and all $i \in \mathbb{N}$. Since each ξ_i is a measurable function of x_i , the strong mixing coefficients $(\alpha_i)_{i \in \mathbb{N}}$ of $(\xi_i)_{i \in \mathbb{N}}$ satisfy $\alpha_i \leq \tilde{\alpha}_i \leq C\rho^i$ for each $i \in \mathbb{N}$. Consequently, $\sum_{i \in \mathbb{N}} i^{2/\delta} \alpha_i < \infty$ for $\delta = 2q - 2 > 0$. Note that we also have

$$\mathbb{E}_{z \sim \mathbb{Q}}[\|\mathbf{k}_{\mathbb{Q}}(z, \cdot)\|_{\mathbf{k}_{\mathbb{Q}}}^{2+\delta}] = \mathbb{E}_{z \sim \mathbb{Q}}[\mathbf{k}_{\mathbb{Q}}(z, z)^q] = \mathbb{E}_{z \sim \mathbb{Q}}[\frac{d\mathbb{P}}{d\mathbb{Q}}(z)^{2q} \mathbf{k}_{\mathbb{P}}(z, z)^q] = \mathbb{E}_{x \sim \mathbb{P}}[\frac{d\mathbb{P}}{d\mathbb{Q}}(x)^{2q-1} \mathbf{k}_{\mathbb{P}}(x, x)^q] < \infty,$$

$\mathbb{E}_{x_i \sim \mathbb{Q}}[\xi_i] = \mathbb{E}_{x_i \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x_i, \cdot)] = 0$ and that $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$ is separable. Since \mathcal{S}_∞ is a strictly stationary chain, we conclude that $(\xi_i)_{i \in \mathbb{N}}$ is a strictly stationary centered sequence of $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$ -valued random variables satisfying the conditions needed to invoke Merlevède et al. (1997, Cor. 1), and hence $\sum_{i=1}^n \xi_i / \sqrt{n}$ converges in distribution to a Gaussian random element taking values in $\mathcal{H}_{\mathbf{k}_{\mathbb{Q}}}$.

Case 2: $\mu_1 \neq \mathbb{Q}$ Since \mathcal{S}_∞ is positive Harris and $\sum_{i=1}^n \xi_i / \sqrt{n}$ satisfies a CLT for the initial distribution $\mu_1 = \mathbb{Q}$, Meyn and Tweedie (2012, Prop. 17.1.6) implies that $\sum_{i=1}^n \xi_i / \sqrt{n}$ also satisfies the same CLT for any initial distribution μ_1 .

Putting the pieces together for (19) Since, for any initial distribution for x_1 , the sequence $(\sum_{i=1}^n \xi_i / \sqrt{n})_{n \in \mathbb{N}}$ converges in distribution and that $\mathcal{H}_{\mathbf{k}_Q}$ is separable and (by virtue of being a Hilbert space) complete, Prokhorov's theorem (Billingsley, 2013, Thm. 5.2) implies that the sequence is also tight, i.e., $\|\sum_{i=1}^n \xi_i\|_{\mathbf{k}_Q} / \sqrt{n} = O_p(1)$. Consequently,

$$\|\frac{1}{n} \sum_{i=1}^n \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i) \mathbf{k}_{\mathbb{P}}(x_i, \cdot)\|_{\mathbf{k}_{\mathbb{P}}} \stackrel{(20)}{=} \|\frac{1}{n} \sum_{i=1}^n \xi_i\|_{\mathbf{k}_Q} = \frac{1}{\sqrt{n}} \cdot \|\frac{\sum_{i=1}^n \xi_i}{\sqrt{n}}\|_{\mathbf{k}_Q} = O_p(n^{-\frac{1}{2}}),$$

as desired for the first claim in (19). Moreover, the strong law of large numbers for positive Harris chains (Meyn and Tweedie, 2012, Thm. 17.0.1(i)) implies that $\frac{1}{n} \sum_{i \in [n]} \frac{d\mathbb{P}}{d\mathbb{Q}}(x_i)$ converges almost surely to $\mathbb{E}_{z \sim \mathbb{Q}}[\frac{d\mathbb{P}}{d\mathbb{Q}}(z)] = 1$ as desired for the second claim in (19). \square

C.3. Proof of Thm. 2: Better-than-i.i.d. debiasing via simplex reweighting

We start with Thm. C.1, proved in App. C.4, that bounds MMD_{OPT} in terms of the eigenvalues of the integral operator of the kernel $\mathbf{k}_{\mathbb{P}}$. Our proof makes use of a weight construction from Liu and Lee (2017, Theorem 3.2), but is a non-trivial generalization of their proof as we no longer assume uniform bounds on the eigenfunctions, and instead leverage truncated variations of Bernstein's inequality (Lems. C.2 and C.3) to establish suitable concentration bounds.

Theorem C.1 (Debiasing via i.i.d. simplex reweighting). *Consider a kernel $\mathbf{k}_{\mathbb{P}}$ satisfying Assum. 1 with $\mathbf{p} = 2$. Let $(\lambda_\ell)_{\ell=1}^\infty$ be the decreasing sequence of eigenvalues of $T_{\mathbf{k}_{\mathbb{P}}, \mathbb{P}}$ defined in (4). Let \mathcal{S}_n be a sequence of $n \in 2\mathbb{N}$ i.i.d. random variables with law \mathbb{Q} such that \mathbb{P} is absolutely continuous with respect to \mathbb{Q} and $\|\frac{d\mathbb{P}}{d\mathbb{Q}}\|_\infty \leq M$ for some $M > 0$. Furthermore, assume there exist constants $\delta_n, B_n > 0$ such that $\Pr(\|\mathbf{k}_{\mathbb{P}}\|_n > B_n) < \delta_n$. Then for all $L \in \mathbb{N}$ such that $\lambda_L > 0$, we have*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathcal{S}_n^{\text{wOPT}}, \mathbb{P})] \leq \frac{8M}{n} \left(\frac{2M}{n} \frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)]}{\lambda_L} + \sum_{\ell > L} \lambda_\ell \right) + \epsilon_n \mathbb{E}[\mathbf{k}_{\mathbb{P}}^2(x_1, x_1)], \quad (21)$$

where

$$\epsilon_n^2 \triangleq n \exp\left(\frac{-3n}{16MB_n/\lambda_L}\right) + 2 \exp\left(\frac{-n}{16M^2}\right) + 2 \exp\left(-\frac{n}{64M^2(\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)] + B_n/12)/\lambda_L}\right) + \delta_n. \quad (22)$$

Given Thm. C.1, Thm. 2 follows, i.e., we have $\mathbb{E}[\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathcal{S}_n^{\text{wOPT}}, \mathbb{P})] = o(n^{-1})$, as long as we can show (i) $\mathbb{E}[\mathbf{k}_{\mathbb{P}}^2(x_1, x_1)] < \infty$, which in turn holds when $\mathfrak{q} > 3$ as assumed in Thm. 2, and (ii) find sequences $(B_n)_{n=1}^\infty$, $(\delta_n)_{n=1}^\infty$, and $(L_n)_{n=1}^\infty$ such that $\Pr(\|\mathbf{k}_{\mathbb{P}}\|_n > B_n) < \delta_n$ for all n and the following conditions are met:

- (a) $\frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)]}{\lambda_{L_n}} = o(n)$;
- (b) $\frac{B_n}{\lambda_{L_n}} = O(n^\beta)$, for some $\beta < 1$;
- (c) $\sum_{\ell > L_n} \lambda_\ell = o(1)$;
- (d) $\delta_n = o(n^{-2})$.

We now proceed to establish these conditions under the assumptions of Thm. 2.

Condition (d) By the de La Vallée Poussin Theorem (Chandra, 2015, Thm. 1.3) applied to the \mathbb{Q} -integrable function $x \mapsto \mathbf{k}_{\mathbb{P}}(x, x)^\mathfrak{q}$ (which is a uniformly integrable family with one function), there exists a convex increasing function G such that $\lim_{t \rightarrow \infty} \frac{G(t)}{t} = \infty$ and $\mathbb{E}[G(\mathbf{k}_{\mathbb{P}}(x_1, x_1)^\mathfrak{q})] < \infty$. Thus,

$$\begin{aligned} \Pr(\mathbf{k}_{\mathbb{P}}(x_1, x_1) > n^{3/\mathfrak{q}}) &= \Pr(\mathbf{k}_{\mathbb{P}}(x_1, x_1)^\mathfrak{q} > n^3) = \Pr(G(\mathbf{k}_{\mathbb{P}}(x_1, x_1)^\mathfrak{q}) > G(n^3)) \\ &\leq \frac{\mathbb{E}[G(\mathbf{k}_{\mathbb{P}}(x_1, x_1)^\mathfrak{q})]}{G(n^3)} = o(n^{-3}), \end{aligned}$$

where the last step uses $\lim_{t \rightarrow \infty} \frac{G(t)}{t} = \infty$. Hence by the union bound,

$$\Pr(\|\mathbf{k}_{\mathbb{P}}\|_n > n^{3/\mathfrak{q}}) = \Pr(\max_{i \in [n]} \mathbf{k}_{\mathbb{P}}(x_i, x_i) > n^{3/\mathfrak{q}}) \leq n \Pr(\mathbf{k}_{\mathbb{P}}(x_1, x_1) > n^{3/\mathfrak{q}}) = o(n^{-2}).$$

Hence if we set $B_n = n^\tau$ for $\tau \triangleq 3/\mathfrak{q} < 1$, there exists $(\delta_n)_{n=1}^\infty$ such that $\delta_n = o(n^{-2})$. This fulfills (d) and that $\Pr(\|\mathbf{k}_{\mathbb{P}}\|_n > B_n) < \delta_n$.

To prove remaining conditions, without loss of generality, we assume that $\lambda_\ell > 0$ for all $\ell \in \mathbb{N}$, since otherwise we can choose L_n to be, for all n , the largest ℓ such that $\lambda_\ell > 0$. Then $\sum_{\ell > L_n} \lambda_{L_n} = 0$ and all other conditions are met.

Condition (c) If $L_n \rightarrow \infty$, then (c) is fulfilled since $\sum_\ell \lambda_\ell < \infty$, which follows from Lem. B.1(d) and that

$$\sum_\ell \lambda_\ell = \sum_{\ell=1}^{\infty} \lambda_\ell \mathbb{E}_{x \sim \mathbb{P}}[\phi_\ell(x)^2] = \sum_{\ell=1}^{\infty} \lambda_\ell \mathbb{E}_{x \sim \mathbb{P}}[\phi_\ell(x)^2] = \mathbb{E}_{x \sim \mathbb{P}}[\sum_{\ell=1}^{\infty} \lambda_\ell \phi_\ell(x)^2] = \mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)] < \infty.$$

Conditions (a) and (b) Note that the condition (a) is subsumed by (b) since $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)] < \infty$. It remains to choose $(L_n)_{n=1}^{\infty}$ to satisfy (b) such that $\lim_{n \rightarrow \infty} L_n = \infty$. Define $L_n \triangleq \max\{\ell \in \mathbb{N} : \lambda_\ell \geq n^{-\frac{\tau-1}{2}}\}$. Then L_n is well-defined for $n \geq (\frac{1}{\lambda_1})^{\frac{2}{1-\tau}}$, since for such n we have $\lambda_1 \geq n^{\frac{\tau-1}{2}}$. Hence for $n \geq (\frac{1}{\lambda_1})^{\frac{2}{1-\tau}}$, we have

$$\frac{B_n}{\lambda_{L_n}} \leq \frac{n^\tau}{n^{\frac{\tau-1}{2}}} = n^{\frac{\tau+1}{2}},$$

so (b) is satisfied with $\beta = \frac{\tau+1}{2} < 1$. Since $\tau < 1$, L_n is non-decreasing in n and $n^{-\frac{\tau-1}{2}}$ decreases to 0. Since each $\lambda_\ell > 0$, we therefore have $\lim_{n \rightarrow \infty} L_n = \infty$. \square

C.4. Proof of Thm. C.1: Debiasing via i.i.d. simplex reweighting

We will slowly build up towards proving Thm. C.1. First notice $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)] < \infty$ implies $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)] < \infty$, so Lem. B.1 holds. Fix any $L \in \mathbb{N}$ satisfying $\lambda_L > 0$. Since n is even, we can define $\mathcal{D}_0 \triangleq [n/2]$ and $\mathcal{D}_1 \triangleq [n] \setminus \mathcal{D}_0$. We will use $\mathcal{S}_{\mathcal{D}_0}$ and $\mathcal{S}_{\mathcal{D}_1}$ to denote the subsets of \mathcal{S}_n with indices in \mathcal{D}_0 and \mathcal{D}_1 respectively. Let $(\phi_\ell)_{\ell=1}^{\infty} \subset \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ be eigenfunctions corresponding to the eigenvalues $(\lambda_\ell)_{\ell=1}^{\infty}$ by Lem. B.1(c), so that $(\phi_\ell)_{\ell=1}^{\infty}$ is an orthonormal system of $\mathcal{L}^2(\mathbb{P})$.

We start with a useful lemma.

Lemma C.1 ($\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ consists of mean-zero functions). *Let $\mathbf{k}_{\mathbb{P}}$ be a kernel satisfying Assum. 1. Then for any $f \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$, we have $\mathbb{P}f = 0$.*

Proof. Fix $f \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$. By Steinwart and Christmann (2008, Thm 4.26), f is \mathbb{P} integrable. Consider the linear operator I that maps $f \mapsto \mathbb{P}f$. Since

$$|I(f)| = |\mathbb{P}f| \leq \mathbb{P}|f| = \int |\langle f, \mathbf{k}_{\mathbb{P}}(x, \cdot) \rangle_{\mathbf{k}_{\mathbb{P}}} |d\mathbb{P} \leq \int \|f\|_{\mathbf{k}_{\mathbb{P}}} \sqrt{\mathbf{k}_{\mathbb{P}}(x, x)} d\mathbb{P} = \|f\|_{\mathbf{k}_{\mathbb{P}}} \mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)^{\frac{1}{2}}].$$

Hence I is a continuous linear operator, so by the Riez representation theorem (Steinwart and Christmann, 2008, Thm. A.5.12), there exists $g \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ such that $I(h) = \langle h, g \rangle_{\mathbf{k}_{\mathbb{P}}}$ for any $h \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$.

By Steinwart and Christmann (2008, Thm. 4.21), the set

$$H_{\text{pre}} \triangleq \left\{ \sum_{i=1}^n \alpha_i \mathbf{k}_{\mathbb{P}}(\cdot, x_i) : n \in \mathbb{N}, (\alpha_i)_{i \in [n]} \subset \mathbb{R}, (x_i)_{i \in [n]} \subset \mathbb{R}^d \right\}$$

is dense in $\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$. Note that H_{pre} consists of mean zero functions under \mathbb{P} by linearity. So there exists f_n converging to f in $\mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$ where each f_n has $\mathbb{P}f_n = I(f_n) = \langle f_n, g \rangle_{\mathbf{k}_{\mathbb{P}}} = 0$. Since

$$\lim_{n \rightarrow \infty} |\langle f, g \rangle_{\mathbf{k}_{\mathbb{P}}} - \langle f_n, g \rangle_{\mathbf{k}_{\mathbb{P}}}| = \lim_{n \rightarrow \infty} |\langle f - f_n, g \rangle_{\mathbf{k}_{\mathbb{P}}}| \leq \lim_{n \rightarrow \infty} \|f - f_n\|_{\mathbf{k}_{\mathbb{P}}} \|g\|_{\mathbf{k}_{\mathbb{P}}} = 0,$$

we have $\mathbb{P}f = \langle f, g \rangle_{\mathbf{k}_{\mathbb{P}}} = 0$. \square

In particular, the assumption $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)] < \infty$ of Thm. C.1 implies $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)^{\frac{1}{2}}] < \infty$, so Lem. C.1 holds.

Step 1. Build control variate weights

Fix any $L \geq 1$ and $h \in \mathcal{H}_{\mathbf{k}_{\mathbb{P}}}$, and let $\hat{h}_{\mathcal{D}_0}$ denote the eigen-expansion truncated approximation of h based on \mathcal{D}_0 ,

$$\hat{h}_{\mathcal{D}_0}(x) \triangleq \sum_{\ell=1}^L \hat{\beta}_{\ell,0} \phi_\ell(x) \quad \text{for} \quad \hat{\beta}_{\ell,0} \triangleq \frac{2}{n} \sum_{i \in \mathcal{D}_0} h(x_i) \phi_\ell(x_i) \xi(x_i).$$

Then

$$\mathbb{E}[\hat{\beta}_{\ell,0}] = \mathbb{E} \left[\frac{2}{n} \sum_{i \in \mathcal{D}_0} h(x_i) \phi_\ell(x_i) \xi(x_i) \right] = \langle h, \phi_\ell \rangle_{\mathcal{L}^2(\mathbb{P})}. \quad (23)$$

Next, define the control variate

$$\hat{Z}_0[h] = \frac{2}{n} \sum_{i \in \mathcal{D}_1} \left(\xi(x_i)(h(x_i) - \hat{h}_{\mathcal{D}_0}(x_i)) \right). \quad (24)$$

which satisfies

$$\mathbb{E}[\hat{Z}_0[h]] = \mathbb{E}_{x \sim \mathbb{P}} \left[h(x) - \sum_{\ell=1}^L \mathbb{E}[\hat{\beta}_{\ell,0}] \phi_\ell(x) \right] = 0, \quad (25)$$

since functions in \mathcal{H}_{k_p} have mean 0 with respect to \mathbb{P} (Lem. C.1). Similarly, we define $\hat{Z}_1[h]$ by swapping \mathcal{D}_0 and \mathcal{D}_1 . Then we form $\hat{Z}[h] \triangleq \frac{\hat{Z}_0[h] + \hat{Z}_1[h]}{2}$. We can rewrite $\hat{Z}[h]$ as a quadrature rule over \mathcal{S}_n (Liu and Lee, 2017, Lemma B.6)

$$\hat{Z}[h] = \sum_{i \in [n]} w_i h(x_i), \quad (26)$$

where w_i is defined as (whose randomness depends on the randomness in \mathcal{S}_n)

$$w_i \triangleq \begin{cases} \frac{1}{n} \xi(x_i) - \frac{2}{n^2} \sum_{j \in \mathcal{D}_1} \xi(x_i) \xi(x_j) \langle \Phi_L(x_i), \Phi_L(x_j) \rangle, \forall i \in \mathcal{D}_0, \\ \frac{1}{n} \xi(x_i) - \frac{2}{n^2} \sum_{j \in \mathcal{D}_0} \xi(x_i) \xi(x_j) \langle \Phi_L(x_i), \Phi_L(x_j) \rangle, \forall i \in \mathcal{D}_1, \end{cases} \quad (27)$$

and $\Phi_L(x) \triangleq (\phi_1(x), \dots, \phi_L(x))$.

Step 2. Show $\mathbb{E}[\text{MMD}_{k_p}^2(\mathbb{S}_n^w, \mathbb{P})] = o(n^{-1})$

We first bound the variance of the control variate $\hat{Z}_0[h]$ for $h = \phi_{\ell'}$ for $\ell' \in \mathbb{N}$. Let us fix $\ell' \in \mathbb{N}$. From (24), we compute

$$\begin{aligned} \mathbb{E}[\hat{Z}_0[h]^2] &= \frac{4}{n^2} \mathbb{E} \left[\left(\sum_{i \in \mathcal{D}_1} \xi(x_i)(h(x_i) - \hat{h}_{\mathcal{D}_0}(x_i)) \right)^2 \right] = \frac{4}{n^2} \mathbb{E} \left[\sum_{i \in \mathcal{D}_1} \xi(x_i)^2 (h(x_i) - \hat{h}_{\mathcal{D}_0}(x_i))^2 \right] \\ &= \frac{2}{n} \mathbb{E}[\mathbb{E}_{x \sim \mathbb{Q}}[\xi(x)^2 (h(x) - \hat{h}_{\mathcal{D}_0}(x))^2 | \mathcal{S}_{\mathcal{D}_0}]] \\ &= \frac{2}{n} \mathbb{E}[\mathbb{E}_{x \sim \mathbb{P}}[\xi(x)(h(x) - \hat{h}_{\mathcal{D}_0}(x))^2 | \mathcal{S}_{\mathcal{D}_0}]] \\ &\leq \frac{2M}{n} \mathbb{E}[\mathbb{E}_{x \sim \mathbb{P}}[(h(x) - \hat{h}_{\mathcal{D}_0}(x))^2 | \mathcal{S}_{\mathcal{D}_0}]], \end{aligned}$$

where in the second equality, the cross terms are zero due to the independence of points x_i and the equality (25). By the definition of $\hat{h}_{\mathcal{D}_0}$, we compute

$$\begin{aligned} \mathbb{E}_{x \sim \mathbb{P}}[(h(x) - \hat{h}_{\mathcal{D}_0}(x))^2 | \mathcal{S}_{\mathcal{D}_0}] &= \mathbb{E}_{x \sim \mathbb{P}} \left[\left(\phi_{\ell'}(x) - \sum_{\ell \leq L} \hat{\beta}_{\ell,0} \phi_\ell(x) \right)^2 \middle| \mathcal{S}_{\mathcal{D}_0} \right] \\ &= \mathbb{E}_{x \sim \mathbb{P}} \left[\phi_{\ell'}^2(x) + \sum_{\ell \leq L} \hat{\beta}_{\ell,0}^2 \phi_\ell^2(x) - 2\phi_{\ell'}(x) \sum_{\ell \leq L} \hat{\beta}_{\ell,0} \phi_\ell(x) \middle| \mathcal{S}_{\mathcal{D}_0} \right] \\ &= 1 + \sum_{\ell \leq L} \hat{\beta}_{\ell,0}^2 - 2 \sum_{\ell \leq L} \hat{\beta}_{\ell,0} \mathbb{1}_{\ell'=\ell} \\ &= 1 + \sum_{\ell \leq L} \hat{\beta}_{\ell,0}^2 - 2\hat{\beta}_{\ell',0} \mathbb{1}_{\ell' \leq L}, \end{aligned}$$

where we use the fact that $(\phi_\ell)_{\ell=1}^\infty$ is an orthonormal system in $\mathcal{L}^2(\mathbb{P})$. By (23) with $h = \phi_{\ell'}$, we have $\mathbb{E}[\hat{\beta}_{\ell',0}] = 1$. On the other hand, we can bound, again using the orthonormality of $(\phi_\ell)_{\ell=1}^\infty$,

$$\mathbb{E}[\hat{\beta}_{\ell,0}^2] = \mathbb{E} \left[\left(\frac{2}{n} \sum_{i \in \mathcal{D}_0} \phi_\ell(x_i) \phi_{\ell'}(x_i) \xi(x_i) \right)^2 \right] = \frac{4}{n^2} \mathbb{E} \left[\sum_{i \in \mathcal{D}_0} (\phi_\ell(x_i) \phi_{\ell'}(x_i) \xi(x_i))^2 \right] \leq \frac{2M}{n} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_\ell(x) \phi_{\ell'}(x))^2].$$

Thus for all $\ell' \in \mathbb{N}$,

$$\mathbb{E}[\hat{Z}_0[\phi_{\ell'}]^2] \leq \frac{2M}{n} \left(1 + \frac{2M}{n} \sum_{\ell \leq L} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_\ell(x) \phi_{\ell'}(x))^2] - 2\mathbb{1}_{\ell' \leq L} \right) \leq \frac{2M}{n} \left(\frac{2M}{n} \sum_{\ell \leq L} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_\ell(x) \phi_{\ell'}(x))^2] + \mathbb{1}_{\ell' > L} \right).$$

Since $\hat{Z}[h] = \frac{\hat{Z}_0[h] + \hat{Z}_1[h]}{2}$ and $(\frac{a+b}{2})^2 \leq \frac{a^2+b^2}{2}$ for $a, b \in \mathbb{R}$, and, by symmetry, $\mathbb{E}[\hat{Z}_0[h]^2] = \mathbb{E}[\hat{Z}_1[h]^2]$, we have

$$\mathbb{E}[\hat{Z}[\phi_{\ell'}]^2] \leq \frac{2M}{n} \left(\frac{2M}{n} \sum_{\ell \leq L} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_\ell(x) \phi_{\ell'}(x))^2] + \mathbb{1}_{\ell' > L} \right). \quad (28)$$

Now we have

$$\begin{aligned}
 \mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^w, \mathbb{P})] &= \mathbb{E}\left[\sum_{i,j \in [n]} w_i w_j \mathbf{k}_P(x_i, x_j)\right] = \mathbb{E}\left[\sum_{i,j \in [n]} w_i w_j \sum_{\ell'=1}^{\infty} \lambda_{\ell'} \phi_{\ell'}(x_i) \phi_{\ell'}(x_j)\right] \\
 &= \mathbb{E}\left[\sum_{\ell'=1}^{\infty} \sum_{i,j \in [n]} w_i w_j \lambda_{\ell'} \phi_{\ell'}(x_i) \phi_{\ell'}(x_j)\right] \\
 &= \mathbb{E}\left[\sum_{\ell'=1}^{\infty} \lambda_{\ell'} \left(\sum_{i \in [n]} w_i \phi_{\ell'}(x_i)\right)^2\right] \\
 &= \sum_{\ell'=1}^{\infty} \lambda_{\ell'} \mathbb{E}\left[\left(\sum_{i \in [n]} w_i \phi_{\ell'}(x_i)\right)^2\right] = \sum_{\ell'=1}^{\infty} \lambda_{\ell'} \mathbb{E}[\hat{Z}[\phi_{\ell'}]^2],
 \end{aligned}$$

where the second and third equalities are due to the absolute convergence of the Mercer series (Lem. B.1(d)), the fourth equality follows from Tonelli's theorem (Steinwart and Christmann, 2008, Thm. A.3.10), and the last step is due to (26). Plugging in (28), we have

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^w, \mathbb{P})] \leq \frac{2M}{n} \left(\frac{2M}{n} \sum_{\ell'=1}^{\infty} \sum_{\ell \leq L} \lambda_{\ell'} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_{\ell}(x) \phi_{\ell'}(x))^2] + \sum_{\ell > L} \lambda_{\ell} \right).$$

Since the eigenvalues are nonnegative and non-increasing, we can write, by (5),

$$\mathbf{k}_P^2(x, x) = \left(\sum_{\ell=1}^{\infty} \lambda_{\ell} \phi_{\ell}(x)\right)^2 \geq \sum_{\ell'=1}^{\infty} \sum_{\ell \leq L} \lambda_{\ell'} \lambda_{\ell} (\phi_{\ell}(x) \phi_{\ell'}(x))^2 \geq \lambda_L \sum_{\ell'=1}^{\infty} \sum_{\ell \leq L} \lambda_{\ell'} (\phi_{\ell}(x) \phi_{\ell'}(x))^2.$$

Thus by Tonelli's theorem (Steinwart and Christmann, 2008, Thm. A.3.10),

$$\sum_{\ell'=1}^{\infty} \sum_{\ell \leq L} \lambda_{\ell'} \mathbb{E}_{x \sim \mathbb{P}}[(\phi_{\ell}(x) \phi_{\ell'}(x))^2] = \mathbb{E}_{x \sim \mathbb{P}} \left[\sum_{\ell'=1}^{\infty} \sum_{\ell \leq L} \lambda_{\ell'} (\phi_{\ell}(x) \phi_{\ell'}(x))^2 \right] \leq \frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_P^2(x, x)]}{\lambda_L}.$$

Finally, we have

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^w, \mathbb{P})] \leq \frac{2M}{n} \left(\frac{2M}{n} \frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_P^2(x, x)]}{\lambda_L} + \sum_{\ell > L} \lambda_{\ell} \right). \quad (29)$$

Step 3. Meet the non-negative constraint

We now show that the weights (27) are nonnegative and sum close to one with high probability. For $i \in \mathcal{D}_0$, we have

$$w_i = \frac{1}{n} \xi(x_i) (1 - T_i) \quad \text{for} \quad T_i \triangleq \frac{2}{n} \sum_{j \in \mathcal{D}_1} \xi(x_j) \langle \Phi_L(x_i), \Phi_L(x_j) \rangle.$$

Our first goal is to derive an upper bound for $\Pr(\min_{i \in \mathcal{D}_0} w_i < 0)$. Define the event

$$E_n \triangleq \{\|\mathbf{k}_P\|_n \leq B_n\}, \quad (30)$$

so $\Pr(E_n^c) < \delta_n$ by the assumption on $\|\mathbf{k}_P\|_n$. Then

$$\Pr(\min_{i \in [n]} w_i < 0, E_n) = \Pr(\max_{i \in [n]} T_i > 1, E_n) \leq n \Pr(T_1 \mathbb{1}_{E_n} > 1), \quad (31)$$

where we applied the union bound and used the fact that $T_i \mathbb{1}_{E_n}$ has the same law for different i . To further bound $\Pr(T_1 \mathbb{1}_{E_n} > 1)$, we will use the following lemma.

Lemma C.2 (Truncated Bernstein inequality). *Let X_1, \dots, X_n be i.i.d. random variables with $\mathbb{E}[X_1] = 0$ and $\mathbb{E}[X_1^2] < \infty$. For any $B > 0, t > 0$,*

$$\Pr\left(\frac{1}{n} \sum_{i \in [n]} X_i \mathbb{1}_{X_i \leq B} > t\right) \leq \exp\left(\frac{-nt^2}{2(\mathbb{E}[X_1^2] + \frac{Bt}{3})}\right).$$

Proof of Lem. C.2. Fix any $B > 0$ and $t > 0$ and define, for each $i \in [n]$, $Y_i \triangleq X_i \mathbb{1}_{X_i \leq B}$. Then $Y_i \leq B$,

$$\begin{aligned}
 \mathbb{E}[Y_i] &= \mathbb{E}[X_i \mathbb{1}_{X_i \leq B}] \leq \mathbb{E}[X_i \mathbb{1}_{X_i \leq B}] + \mathbb{E}[X_i \mathbb{1}_{X_i > B}] = \mathbb{E}[X_i] = 0, \quad \text{and} \\
 \mathbb{E}[Y_i^2] &= \mathbb{E}[X_i^2 \mathbb{1}_{X_i \leq B}] \leq \mathbb{E}[X_i^2] = \mathbb{E}[X_1^2].
 \end{aligned}$$

Now we can invoke the non-positivity of $\mathbb{E}[Y_i]$, the one-sided Bernstein inequality (Wainwright, 2019, Prop. 2.14), and the relation $\mathbb{E}[Y_i^2] \leq \mathbb{E}[X_1^2]$ to conclude that

$$\Pr\left(\frac{1}{n} \sum_{i \in [n]} Y_i > t\right) \leq \Pr\left(\frac{1}{n} \sum_{i \in [n]} (Y_i - \mathbb{E}[Y_i]) > t\right) \leq \exp\left(\frac{-nt^2}{2\left(\frac{1}{n} \sum_{i \in [n]} \mathbb{E}[Y_i^2] + \frac{Bt}{3}\right)}\right) \leq \exp\left(\frac{-nt^2}{2(\mathbb{E}[X_1^2] + \frac{Bt}{3})}\right).$$

□

For $j \in \mathcal{D}_1$, define $X_j \triangleq \xi(x_j) \langle \Phi_L(x_1), \Phi_L(x_j) \rangle$ and note that

$$\begin{aligned} \mathbb{E}[X_j | x_1] &= \mathbb{E}_{x \sim \mathbb{Q}}[\xi(x) \langle \Phi_L(x_1), \Phi_L(x) \rangle | x_1] = \mathbb{E}_{x \sim \mathbb{P}}[\langle \Phi_L(x_1), \Phi_L(x) \rangle | x_1] = 0 \\ \mathbb{E}[X_j^2 | x_1] &= \mathbb{E}[\xi(x_j)^2 \langle \Phi_L(x_1), \Phi_L(x_j) \rangle^2 | x_1] \leq M \mathbb{E}_{x \sim \mathbb{P}}[\langle \Phi_L(x_1), \Phi_L(x) \rangle^2 | x_1] \\ &= M \mathbb{E}_{x \sim \mathbb{P}} \left[\sum_{\ell, \ell' \leq L} \phi_\ell(x_1) \phi_{\ell'}(x_1) \phi_\ell(x) \phi_{\ell'}(x) \mid x_1 \right] \\ &= M \sum_{\ell, \ell' \leq L} \phi_\ell(x_1) \phi_{\ell'}(x_1) \mathbb{E}_{x \sim \mathbb{P}}[\phi_\ell(x) \phi_{\ell'}(x)] \\ &= M \|\Phi_L(x_1)\|_2^2. \end{aligned}$$

Since $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$, for any $x \in \mathbb{R}^d$, we can bound $\|\Phi_L(x)\|_2^2$ via

$$\|\Phi_L(x)\|_2^2 = \sum_{\ell \leq L} \phi_\ell(x)^2 \leq \frac{\sum_{\ell \leq L} \lambda_\ell \phi_\ell(x)^2}{\lambda_L} \leq \frac{\sum_{\ell=1}^{\infty} \lambda_\ell \phi_\ell(x)^2}{\lambda_L} = \frac{\mathbf{k}_{\mathbb{P}}(x, x)}{\lambda_L}, \quad (32)$$

where we applied Lem. B.1(d) for the last equality. Thus

$$|X_j| \leq M \|\Phi_L(x_1)\|_2 \|\Phi_L(x_j)\|_2 \leq M \|\Phi_L(x_1)\|_2 \sqrt{\frac{\mathbf{k}_{\mathbb{P}}(x_j, x_j)}{\lambda_L}},$$

so if we let $B \triangleq \sqrt{\frac{B_n}{\lambda_L}} M \|\Phi_L(x_1)\|_2$, then

$$E_n = \left\{ \sup_{i \in [n]} \mathbf{k}_{\mathbb{P}}(x_i, x_i) \leq B_n \right\} \subset \bigcap_{j \in \mathcal{D}_1} \{|X_j| \leq B\}.$$

Since $T_1 = \frac{2}{n} \sum_{j \in \mathcal{D}_1} X_j$, we have inclusions of events

$$\{T_1 \mathbb{1}_{E_n} > 1\} = \{T_1 > 1\} \cap E_n \subset \left\{ \frac{2}{n} \sum_{j \in \mathcal{D}_1} X_j \mathbb{1}_{|X_j| \leq B} > 1 \right\}.$$

Thus Lem. C.2 with $t = 1$ and conditioned on x_1 implies

$$\begin{aligned} \Pr(T_1 \mathbb{1}_{E_n} > 1 | x_1) &\leq \Pr\left(\frac{2}{n} \sum_{j \in \mathcal{D}_1} X_j \mathbb{1}_{|X_j| \leq B} > 1 \mid x_1\right) \\ &\leq \exp\left(\frac{-n}{4(M\|\Phi_L(x_1)\|_2^2 + \sqrt{\frac{B_n}{\lambda_L}} M \|\Phi_L(x_1)\|_2 / 3)}\right). \end{aligned}$$

On event $\{\mathbf{k}_{\mathbb{P}}(x_1, x_1) \leq B_n\}$, by (32), we have

$$\|\Phi_L(x_1)\|_2 \leq \sqrt{\frac{B_n}{\lambda_L}}.$$

Hence

$$\Pr(T_1 \mathbb{1}_{E_n} > 1 | x_1) \mathbb{1}_{\mathbf{k}_{\mathbb{P}}(x_1, x_1) \leq B_n} \leq \exp\left(\frac{-n}{\frac{16}{3} M \frac{B_n}{\lambda_L}}\right).$$

On the other hand, $\{\mathbf{k}_{\mathbb{P}}(x_1, x_1) > B_n\} \notin E_n$, so

$$\Pr(T_1 \mathbb{1}_{E_n} > 1 | x_1) \mathbb{1}_{\mathbf{k}_{\mathbb{P}}(x_1, x_1) > B_n} = 0$$

Thus

$$\Pr(T_1 \mathbb{1}_{E_n} > 1) = \mathbb{E}[\Pr(T_1 \mathbb{1}_{E_n} > 1 | x_1)] \leq \exp\left(\frac{-n}{\frac{16}{3} M \frac{B_n}{\lambda_L}}\right).$$

Combining the last inequality with (31), we have:

$$\Pr(\min_{i \in [n]} w_i < 0, E_n) \leq n \exp\left(\frac{-n}{\frac{16}{3} M \frac{B_n}{\lambda_L}}\right). \quad (33)$$

Step 4. Meet the sum-to-one constraint

Let

$$S \triangleq \sum_{i \in \mathcal{D}_0} w_i = \sum_{i \in \mathcal{D}_0} \frac{1}{n} \xi(x_i) \left(1 - \frac{2}{n} \sum_{j \in \mathcal{D}_1} \xi(x_j) \langle \Phi_L(x_i), \Phi_L(x_j) \rangle \right).$$

We now derive a bound for $\Pr(S < 1/2 - t/2)$ for $t \in (0, 1)$. Let

$$S_1 \triangleq \frac{1}{n} \sum_{i \in \mathcal{D}_0} \xi(x_i), \quad S_2 \triangleq -\frac{2}{n^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} \xi(x_i) \xi(x_j) \langle \Phi_L(x_i), \Phi_L(x_j) \rangle,$$

so $S = S_1 + S_2$. Note that $\mathbb{E}[S_1] = 1/2$ and $\mathbb{E}[S_2] = 0$ since \mathcal{D}_0 and \mathcal{D}_1 are disjoint. Let E_n be the same event defined as in (30). For $t_1 \in (0, t/2)$ to be determined later and $t_2 \triangleq t/2 - t_1$, we have, by the union bound

$$\Pr(S < 1/2 - t/2, E_n) \leq \Pr(S_1 < 1/2 - t_1, E_n) + \Pr(S_2 < -t_2, E_n).$$

By Hoeffding's inequality and the assumption $\xi(x) \leq M$, we have

$$\Pr(S_1 < 1/2 - t_1, E_n) \leq \Pr\left(\frac{2}{n} \sum_{i \in \mathcal{D}_0} \frac{\xi(x_i)}{2} - 1/2 < -t_1\right) \leq \exp\left(\frac{-2(n/2)t_1^2}{(M/2)^2}\right) = \exp\left(\frac{-4nt_1^2}{M^2}\right). \quad (34)$$

To give a concentration bound for $\Pr(S_2 < -t_2, E_n)$, we will use the following lemma.

Lemma C.3 (U-statistic Bernstein's inequality). *Let $h : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a function bounded above by $b > 0$. Assume $n \in 2\mathbb{N}$ and let x_1, \dots, x_n be i.i.d. random variables taking values in \mathcal{X} . Denote $m_h \triangleq \mathbb{E}[h(x_1, x_2)]$ and $\sigma_h^2 \triangleq \text{Var}[h(x_1, x_2)]$. Let $\mathcal{D}_0 = \lfloor n/2 \rfloor$ and $\mathcal{D}_1 = \lfloor n \rfloor \setminus \lfloor n/2 \rfloor$. Define*

$$U \triangleq \frac{1}{(n/2)^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} h(x_i, x_j).$$

Then

$$\Pr(U - m_h > t) \leq \exp\left(\frac{-nt^2}{4(\sigma_h^2 + \frac{bt}{3})}\right).$$

Proof of Lem. C.3. We adapt the proof from Pitcan (2017, Section 3) as follows. Let $k \triangleq n/2$. Define $V : \mathcal{X}^n \rightarrow \mathbb{R}$ as

$$V(x_1, \dots, x_n) \triangleq \frac{1}{k} \sum_{i \in [k]} h(x_i, x_{i+k}).$$

Then note that

$$U = \frac{1}{k!} \sum_{\sigma \in \text{perm}(k)} V_\sigma, \\ V_\sigma \triangleq V(x_{\sigma_1}, \dots, x_{\sigma_k}),$$

where $\text{perm}(k)$ is the set of all permutations of $[k]$; this is because every $h(x_i, x_j)$ term for $i \in \mathcal{D}_0, j \in \mathcal{D}_1$ will appear in the summation an equal number of times. For a fixed $\sigma \in \text{perm}(k)$, the random variable $V(x_{\sigma_1}, \dots, x_{\sigma_k}, x_{k+1}, \dots, x_n)$ is a sum of k i.i.d. terms $h(x_{\sigma_i}, x_{i+k})$. Denote $V = V(x_1, \dots, x_n)$. For any $s > 0$, we have, by independence,

$$\begin{aligned} \mathbb{E}[e^{s(V-m_h)}] &= \mathbb{E}\left[\exp\left(\frac{s}{k} \sum_{i \in [k]} (h(x_i, x_{i+k}) - m_h)\right)\right] \\ &= \left(\mathbb{E}\left[\exp\left(\frac{s}{k} (h(x_1, x_2) - m_h)\right)\right]\right)^k \end{aligned}$$

By the one-sided Bernstein's lemma Wainwright (2019, Prop. 2.14) applied to $\frac{h(x_1, x_2)}{k}$ which is upper bounded by $\frac{b}{k}$ with variance $\frac{\sigma_h^2}{k^2}$, we have

$$\mathbb{E}\left[\exp\left(s \frac{h(x_1, x_2) - m_h}{k}\right)\right] \leq \exp\left(\frac{s^2 \sigma_h^2 / 2}{k(k - \frac{bs}{3})}\right),$$

for $s \in [0, 3k/b)$. Next, by Markov's inequality and Jensen's inequality,

$$\begin{aligned} \Pr(U - m_h > t) &= \Pr(e^{s(U-m_h)} > e^{st}) \leq \mathbb{E}[e^{s(U-m_h)}] e^{-st} \\ &= \mathbb{E}\left[\exp\left(\frac{1}{(n/2)!} \sum_{\sigma \in \text{perm}(n/2)} s(V_\sigma - m_h)\right)\right] e^{-st} \\ &\leq \mathbb{E}\left[\frac{1}{(n/2)!} \sum_{\sigma \in \text{perm}(n/2)} \exp(s(V_\sigma - m_h))\right] e^{-st} \\ &= \mathbb{E}[e^{s(V-m_h)}] e^{-st}. \end{aligned}$$

Therefore,

$$\Pr(U - m_h > t) \leq \exp\left(\frac{s^2 \sigma_h^2}{2(k - \frac{bs}{3})} - st\right).$$

Now, we get the desired bound if we pick $s = \frac{k^2 t}{k\sigma_h^2 + \frac{k\epsilon b}{3}} \in [0, 3k/b)$ and simplify. \square

Let

$$\begin{aligned} h(x, x') &\triangleq \xi(x)\xi(x')\langle \Phi_L(x), \Phi_L(x') \rangle \\ \bar{h}(x, x') &\triangleq h(x, x') \mathbb{1}_{h(x, x') \leq M^2 \frac{B_n}{\lambda_L}}. \end{aligned}$$

Then

$$\begin{aligned} \Pr(S_2 < -t_2, E_n) &= \Pr\left(\frac{1}{(n/2)^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} h(x_i, x_j) > 2t_2, E_n\right) \\ &\leq \Pr\left(\frac{1}{(n/2)^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} \bar{h}(x_i, x_j) > 2t_2\right), \end{aligned} \quad (35)$$

where the last inequality used the fact that, for $i \in \mathcal{D}_0, j \in \mathcal{D}_1$,

$$E_n \subset \{\max(\mathbf{k}_{\mathbb{P}}(x_i, x_i), \mathbf{k}_{\mathbb{P}}(x_j, x_j)) \leq B_n\} \subset \left\{h(x_i, x_j) \leq M^2 \frac{B_n}{\lambda_L}\right\},$$

using (32). We further compute

$$\begin{aligned} m_{\bar{h}} &= \mathbb{E}[\bar{h}(x_1, x_2)] \leq \mathbb{E}[h(x_1, x_2)] = \mathbb{E}[\xi(x_1)\xi(x_2)\langle \Phi_L(x_1), \Phi_L(x_2) \rangle] \\ &= \sum_{\ell \leq L} \mathbb{E}[\xi(x_1)\xi(x_2)\phi_{\ell}(x_1)\phi_{\ell}(x_2)] \\ &= \sum_{\ell \leq L} (\mathbb{E}_{x \sim \mathbb{P}}[\phi_{\ell}(x)])^2 = 0, \end{aligned}$$

and

$$\begin{aligned} \sigma_{\bar{h}}^2 &= \text{Var}[\bar{h}(x_1, x_2)] \leq \mathbb{E}[\bar{h}(x_1, x_2)^2] \leq \mathbb{E}[h(x_1, x_2)^2] \\ &= \mathbb{E}\left[\left(\xi(x_1)\xi(x_2)\langle \Phi_L(x_1), \Phi_L(x_2) \rangle\right)^2\right] \\ &\leq M^2 \mathbb{E}_{(x, x') \sim \mathbb{P} \times \mathbb{P}}[\langle \Phi_L(x), \Phi_L(x') \rangle^2] \\ &= M^2 \mathbb{E}_{(x, x') \sim \mathbb{P} \times \mathbb{P}}\left[\sum_{\ell, \ell' \leq L} \phi_{\ell}(x)\phi_{\ell'}(x)\phi_{\ell}(x')\phi_{\ell'}(x')\right] \\ &= M^2 \sum_{\ell, \ell' \leq L} (\mathbb{E}[\phi_{\ell}(x)\phi_{\ell'}(x)])^2 = LM^2. \end{aligned}$$

Since $\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}(x, x)] = \sum_{\ell} \lambda_{\ell} \geq L\lambda_L$, we have $L \leq \frac{\|\mathbf{k}_{\mathbb{P}}\|_{\mathcal{L}^2(\mathbb{P})}^2}{\lambda_L}$, so that $\sigma_{\bar{h}}^2 \leq \frac{M^2 \|\mathbf{k}_{\mathbb{P}}\|_{\mathcal{L}^2(\mathbb{P})}^2}{\lambda_L}$. Applying Lem. C.3 to \bar{h} , which is bounded by $M^2 \frac{B_n}{\lambda_L}$ and using the fact that $m_{\bar{h}} \leq 0$, we have

$$\begin{aligned} \Pr\left(\frac{1}{(n/2)^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} \bar{h}(x_i, x_j) > 2t_2\right) &\leq \Pr\left(\frac{1}{(n/2)^2} \sum_{i \in \mathcal{D}_0} \sum_{j \in \mathcal{D}_1} \bar{h}(x_i, x_j) - m_{\bar{h}} > 2t_2\right) \\ &\leq \exp\left(\frac{-n(2t_2)^2}{4\left(\frac{M^2 \|\mathbf{k}_{\mathbb{P}}\|_{\mathcal{L}^2(\mathbb{P})}^2}{\lambda_L} + 2M^2 \frac{B_n}{\lambda_L} t_2/3\right)}\right). \end{aligned} \quad (36)$$

Thus combining (34), (35), (36), we get

$$\Pr(S < 1/2 - t/2, E_n) \leq \exp\left(\frac{-4nt_1^2}{M^2}\right) + \exp\left(\frac{-nt_2^2}{\left(\frac{M^2 \|\mathbf{k}_{\mathbb{P}}\|_{\mathcal{L}^2(\mathbb{P})}^2}{\lambda_L} + 2M^2 \frac{B_n}{\lambda_L} t_2/3\right)}\right).$$

Finally, by symmetry and the union bound, for $t \in (0, 1)$, $t \in (0, t/2)$ and $t_2 = t/2 - t_1$, we have

$$\begin{aligned} \Pr\left(\sum_{i \in [n]} w_i < 1 - t, E_n\right) &\leq \Pr\left(\sum_{i \in \mathcal{D}_0} w_i < 1/2 - t/2, E_n\right) + \Pr\left(\sum_{i \in \mathcal{D}_1} w_i < 1/2 - t/2, E_n\right) \\ &= 2 \Pr(S < 1/2 - t/2, E_n) \\ &\leq 2 \left(\exp\left(\frac{-4nt_1^2}{M^2}\right) + \exp\left(\frac{-nt_2^2}{\left(\frac{M^2 \|\mathbf{k}_{\mathbb{P}}\|_{\mathcal{L}^2(\mathbb{P})}^2}{\lambda_L} + 2M^2 \frac{B_n}{\lambda_L} t_2/3\right)}\right) \right). \end{aligned} \quad (37)$$

Step 5. Putting it all together

Define the event

$$F_n = \left\{ \min_{i \in [n]} w_i \geq 0, \sum_{i \in [n]} w_i \geq \frac{1}{2} \right\}.$$

Then, by the union bound,

$$\Pr(F_n^c) \leq \Pr(\min_{i \in [n]} w_i < 0, E_n) + \Pr\left(\sum_{i \in [n]} w_i < \frac{1}{2}, E_n\right) + \Pr(E_n^c).$$

Applying (33) and (37) to bound the last expression with $t = 1/2$, $t_1 = t_2 = 1/8$, we have $\Pr(F_n^c) \leq \epsilon_n^2$ for ϵ_n defined in (22). On the event F_n , if we define $w^+ \in \Delta_{n-1}$ via

$$w_i^+ \triangleq \frac{w_i}{\sum_{i \in [n]} w_i},$$

then $w_i^+ = \alpha w_i$ for $i \in [n]$ and $\alpha \triangleq \frac{1}{\sum_{i \in [n]} w_i} \leq 2$. Let $\tilde{w} \in \Delta_{n-1}$ be the weight defined by $\tilde{w}_1 = 1$ and $\tilde{w}_i = 0$ for $i > 1$.

Since w_{OPT} is the best simplex weight, we have $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) \leq \min(\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^+}, \mathbb{P}), \text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{\tilde{w}}, \mathbb{P}))$. Hence

$$\begin{aligned} \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P})] &= \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) \mathbb{1}_{F_n}] + \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) \mathbb{1}_{F_n^c}] \\ &\leq \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^+}, \mathbb{P}) \mathbb{1}_{F_n}] + \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{\tilde{w}}, \mathbb{P}) \mathbb{1}_{F_n^c}]. \end{aligned}$$

For the first term, we have the bound

$$\begin{aligned} \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^+}, \mathbb{P}) \mathbb{1}_{F_n}] &= \mathbb{E} \left[\sum_{i,j \in [n]} w_i^+ w_j^+ \mathbf{k}_{\mathbb{P}}(x_i, x_j) \mathbb{1}_{F_n} \right] = \mathbb{E} \left[\alpha^2 \sum_{i,j \in [n]} w_i w_j \mathbf{k}_{\mathbb{P}}(x_i, x_j) \mathbb{1}_{F_n} \right] \\ &\leq 4 \mathbb{E} \left[\sum_{i,j \in [n]} w_i w_j \mathbf{k}_{\mathbb{P}}(x_i, x_j) \right] \leq \frac{8M}{n} \left(\frac{2M}{n} \frac{\mathbb{E}_{x \sim \mathbb{P}}[\mathbf{k}_{\mathbb{P}}^2(x, x)]}{\lambda_L} + \sum_{\ell > L} \lambda_{\ell} \right), \end{aligned}$$

where we applied (29) for the last inequality. For the second term, by the Cauchy-Schwartz inequality,

$$\begin{aligned} \mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{\tilde{w}}, \mathbb{P}) \mathbb{1}_{F_n^c}] &\leq \sqrt{\Pr(F_n^c)} \sqrt{\mathbb{E} \left[\left(\sum_{i,j \in [n]} \mathbf{k}_{\mathbb{P}}(x_i, x_j) \tilde{w}_i \tilde{w}_j \right)^2 \right]} \\ &\leq \sqrt{\Pr(F_n^c)} \sqrt{\mathbb{E}[\mathbf{k}_{\mathbb{P}}(x_1, x_1)^2]}. \end{aligned}$$

Putting everything together we obtain (21). \square

D. Stein Kernel Thinning

In this section, we detail our Stein thinning implementation in App. D.1, our kernel thinning implementation and analysis in App. D.2, and our proof of Thm. 3 in App. D.3.

D.1. Stein Thinning with sufficient statistics

For an input point set of size n , the original implementation of Stein Thinning of Riabiz et al. (2022) takes $O(nm^2)$ time to output a coreset of size m . In Alg. D.1, we show that this runtime can be improved to $O(nm)$ using sufficient statistics. The idea is to maintain a vector $g \in \mathbb{R}^n$ such that $g = 2\mathbf{k}_{\mathbb{P}}(\mathbb{S}_n, \mathbb{S}_n)w$ where w is the weight representing the current coreset.

Algorithm D.1 SteinThinning (ST) with sufficient statistics

Input: kernel $\mathbf{k}_{\mathbb{P}}$ with zero-mean under \mathbb{P} , input points $\mathcal{S}_n = (x_i)_{i \in [n]}$, output size m
 $w \leftarrow \mathbf{0} \in \mathbb{R}^n$
 $j \leftarrow \operatorname{argmin}_{i \in [n]} \mathbf{k}_{\mathbb{P}}(x_i, x_i)$
 $w_j \leftarrow 1$
 $g \leftarrow 2\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, x_j) \triangleright$ maintain sufficient statistics $g = 2\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)w$
for $t = 1$ **to** $m - 1$ **do**
 $j \leftarrow \operatorname{argmin}_{i \in [n]} \{tg_i + \mathbf{k}_{\mathbb{P}}(x_i, x_i)\}$
 $w \leftarrow \frac{t}{t+1}w + \frac{1}{t+1}e_j$
 $g \leftarrow \frac{t}{t+1}g + \frac{2}{t+1}\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, x_j)$
end for
Return: w

Algorithm D.2 KernelThinning (KT) (adapted from Dwivedi and Mackey (2022, Alg. 1))

Input: kernel $\mathbf{k}_{\mathbb{P}}$ with zero-mean under \mathbb{P} , input points $\mathcal{S}_n = (x_i)_{i \in [n]}$, multiplicity n' with $\log_2 \frac{n'}{m} \in \mathbb{N}$, weight $w \in \Delta_{n-1} \cap (\frac{\mathbb{N}_0}{n'})^n$, output size m with $\frac{n'}{m} \in 2^{\mathbb{N}}$, failure probability δ
 $\mathbf{S} \leftarrow$ index sequence where $k \in [n]$ appears $n'w_k$ times
 $t \leftarrow \log_2 \frac{n'}{m} \in \mathbb{N}$
 $(\mathbf{I}^{(\ell)})_{\ell \in [2^t]} \leftarrow \text{KT-SPLIT}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n[\mathbf{S}], t, \delta/n') \triangleright$ KT-SPLIT is from Dwivedi and Mackey (2022, Algorithm 1a) and we set $\delta_i = \delta$ for all i
 $\mathbf{I}^{(\ell)} \leftarrow \mathbf{S}[\mathbf{I}^{(\ell)}]$ for each $\ell \in [2^t]$
 $\mathbf{I} \leftarrow \text{KT-Swap}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, (\mathbf{I}^{(\ell)})_{\ell \in [2^t]})$
 $w_{\text{KT}} \leftarrow$ simplex weight corresponding to $\mathbf{I} \triangleright w_i = \frac{\text{number of occurrences of } i \text{ in } \mathbf{I}}{|\mathbf{I}|}$
Return: $w_{\text{KT}} \in \Delta_{n-1} \cap (\frac{\mathbb{N}_0}{m})^n \triangleright$ Hence $\|w_{\text{KT}}\|_0 \leq m$

D.2. Kernel Thinning targeting \mathbb{P}

Our KernelThinning implementation is detailed in Alg. D.2. Since we are able to directly compute $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^w, \mathbb{P})$, we use KT-Swap (Alg. D.3) in place of the standard KT-SWAP subroutine (Dwivedi and Mackey, 2022, Algorithm 1b) to choose candidate points to swap in so as to greedily minimize $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^w, \mathbb{P})$. To facilitate our subsequent SKT analysis, we restate the guarantees of KT-SPLIT (Dwivedi and Mackey, 2022, Theorem 2) in the sub-Gaussian format of (Shetty et al., 2022, Definition 3).

Lemma D.1 (Sub-Gaussian guarantee for KT-SPLIT). *Let \mathcal{S}_n be a sequence of n points and \mathbf{k} a kernel. For any $\delta \in (0, 1)$ and $m \in \mathbb{N}$ such that $\log_2 \frac{n}{m} \in \mathbb{N}$, consider the KT-SPLIT algorithm (Dwivedi and Mackey, 2022, Algorithm 1a) with $\mathbf{k}_{\text{split}} = \mathbf{k}$, thinning parameter $t = \log_2 \frac{n}{m}$, and $\delta_i = \frac{\delta}{n}$ to compress \mathcal{S}_n to 2^t coresets $\{\mathcal{S}_{\text{out}}^{(i)}\}_{i \in [2^t]}$ where each $\mathcal{S}_{\text{out}}^{(i)}$ has m points. Denote the signed measure $\phi^{(i)} \triangleq \frac{1}{n} \sum_{x \in \mathcal{S}_n} \delta_x - \frac{2^t}{n} \sum_{x \in \mathcal{S}_{\text{out}}^{(i)}} \delta_x$. Then for each $i \in [2^t]$, on an event $\mathcal{E}_{\text{equi}}^{(i)}$ with $\mathbb{P}(\mathcal{E}_{\text{equi}}^{(i)}) \geq 1 - \frac{\delta}{2}$, $\phi^{(i)} = \tilde{\phi}^{(i)}$ for a random signed measure $\tilde{\phi}^{(i)}$ ⁶ such that, for any $\delta' \in (0, 1)$,*

$$\Pr \left(\left\| \tilde{\phi}^{(i)} \mathbf{k} \right\|_{\mathcal{H}_{\mathbf{k}}} \geq a_{n,m} + v_{n,m} \sqrt{\log \left(\frac{1}{\delta'} \right)} \right) \leq \delta',$$

where

$$a_{n,m} \triangleq \frac{1}{m} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log \left(\frac{6(\log_2 \frac{n}{m})m}{\delta} \right) \log(4\mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1}))} \right),$$

$$v_{n,m} \triangleq \frac{1}{m} \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log \left(\frac{6(\log_2 \frac{n}{m})m}{\delta} \right)}.$$

⁶This is the signed measure returned by repeated applications of self-balancing Hilbert walk (SBHW) (Dwivedi and Mackey, 2021, Algorithm 3). Although SBHW returns an element of $\mathcal{H}_{\mathbf{k}}$, by tracing the algorithm, the returned output is equivalent to a signed measure via the correspondence $\sum_{i \in [n]} c_i \mathbf{k}(x_i, \cdot) \Leftrightarrow \sum_{i \in [n]} c_i \delta_{x_i}$. The usage of signed measures is consistent with Shetty et al. (2022).

Algorithm D.3 KT-Swap (modified Dwivedi and Mackey (2022, Alg. 1b) to minimize MMD to \mathbb{P})

Input: kernel $\mathbf{k}_{\mathbb{P}}$ with zero-mean under \mathbb{P} , input points $\mathcal{S}_n = (x_i)_{i \in [n]}$, candidate coresets indices $(\mathbf{I}^{(\ell)})_{\ell \in [L]}$
 $m \leftarrow |\mathbf{I}^{(0)}| \triangleright$ all candidate coresets are of the same size
 $\mathbf{I} \leftarrow \mathbf{I}^{(\ell^*)}$ for $\ell^* \in \operatorname{argmin}_{\ell \in [L]} \operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathcal{S}_n[\mathbf{I}^{(\ell)}], \mathbb{P}) \triangleright$ select the best KT-SPLIT coreset
 $\mathbf{I}_{\text{ST}} \leftarrow$ index sequence of SteinThinning($\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, m$) \triangleright add Stein thinning baseline
 $\mathbf{C} = \{\mathbf{I}, \mathbf{I}_{\text{ST}}\} \triangleright$ shortlisted candidates
for $\mathbf{I} \in \mathbf{C}$ **do**
 $g \leftarrow \mathbf{0} \in \mathbb{R}^n \triangleright$ maintain sufficient statistics $g = \sum_{j \in [m]} \mathbf{k}_{\mathbb{P}}(x_{\mathbf{I}_j}, \mathcal{S}_n)$
 $\mathbf{K}\text{diag} \leftarrow (\mathbf{k}_{\mathbb{P}}(x_i, x_i))_{i \in [n]}$
 for $j = 1$ **to** m **do**
 $g \leftarrow g + \mathbf{k}_{\mathbb{P}}(x_{\mathbf{I}_j}, \mathcal{S}_n)$
 end for
 for $j = 1$ **to** m **do**
 $\Delta = 2(g - \mathbf{k}_{\mathbb{P}}(x_{\mathbf{I}_j}, \mathcal{S}_n)) + \mathbf{K}\text{diag} \triangleright$ this is the change in $\operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathcal{S}_n[\mathbf{I}], \mathbb{P})$ if we were to replace \mathbf{I}_j
 $k \leftarrow \operatorname{argmin}_{i \in [n]} \Delta_i$
 $g = g - \mathbf{k}_{\mathbb{P}}(x_{\mathbf{I}_j}, \mathcal{S}_n) + \mathbf{k}_{\mathbb{P}}(x_k, \mathcal{S}_n)$
 $\mathbf{I}_j \leftarrow k$
 end for
end for
Return: $\mathbf{I} = \operatorname{argmin}_{\mathbf{I} \in \mathbf{C}} \operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathcal{S}_n[\mathbf{I}], \mathbb{P})$

Proof of Lem. D.1. Fix $i \in [2^t]$, $\delta \in (0, 1)$ and $n, m \in \mathbb{N}$ such that $t = \log_2 \frac{n}{m} \in \mathbb{N}$. Define $\phi \triangleq \phi^{(i)}$. By the proof of Dwivedi and Mackey (2022, Thms. 1 and 2), there exists an event $\mathcal{E}_{\text{equi}}$ with $\Pr(\mathcal{E}_{\text{equi}}^c) \leq \frac{\delta}{2}$ such that, on this event, $\phi = \tilde{\phi}$ where $\tilde{\phi}$ is a signed measure such that, for any $\delta' \in (0, 1)$, with probability at least $1 - \delta'$,

$$\left\| \tilde{\phi} \mathbf{k} \right\|_{\mathcal{H}_{\mathbf{k}}} \leq \inf_{\epsilon \in (0, 1), A: \mathcal{S}_n \subset A} 2\epsilon + \frac{2^t}{n} \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6tn}{2^t \delta}\right) \left[\log \frac{4}{\delta'} + \log \mathcal{N}_{\mathbf{k}}(A, \epsilon) \right]}.$$

Note that on $\mathcal{E}_{\text{equi}}$, $\left\| \tilde{\phi} \mathbf{k} \right\|_{\mathcal{H}_{\mathbf{k}}} = \|\phi \mathbf{k}\|_{\mathcal{H}_{\mathbf{k}}}$. We choose $A = \mathcal{B}_2(R_n)$ and $\epsilon = \frac{2^t}{n} = m^{-1}$, so that, with probability at least $1 - \delta'$, using the fact that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$,

$$\begin{aligned} \left\| \tilde{\phi} \mathbf{k} \right\|_{\mathcal{H}_{\mathbf{k}}} &\leq \frac{2^{t+1}}{n} + \frac{2^t}{n} \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6tn}{2^t \delta}\right) \left[\log \frac{4}{\delta'} + \log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1}) \right]} \\ &\leq \frac{2^{t+1}}{n} + \frac{2^t}{n} \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6tn}{2^t \delta}\right) \left[\sqrt{\log \frac{1}{\delta'}} + \sqrt{\log 4 \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1})} \right]} \\ &\leq a_{n,m} + v_{n,m} \sqrt{\log\left(\frac{1}{\delta'}\right)}, \end{aligned} \tag{38}$$

for $a_{n,m}, v_{n,m}$ in Lem. D.1. □

Corollary D.1 (MMD guarantee for KT-SPLIT). *Let \mathcal{S}_{∞} be an infinite sequence of points in \mathbb{R}^d and \mathbf{k} a kernel. For any $\delta \in (0, 1)$ and $n, m \in \mathbb{N}$ such that $\log_2 \frac{n}{m} \in \mathbb{N}$, consider the KT-SPLIT algorithm (Dwivedi and Mackey, 2022, Algorithm 1a) with parameters $\mathbf{k}_{\text{split}} = \mathbf{k}$ and $\delta_i = \frac{\delta}{n}$ to compress \mathcal{S}_n to 2^t coresets $\{\mathcal{S}_{\text{out}}^{(i)}\}_{i \in [2^t]}$ where $t = \log_2 \frac{n}{m}$, each with m points. Then for any $i \in [2^t]$, with probability at least $1 - \delta$,*

$$\operatorname{MMD}_{\mathbf{k}}(\mathcal{S}_n, \mathcal{S}_{\text{out}}^{(i)}) \leq \frac{1}{m} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6(\log_2 \frac{n}{m})m}{\delta}\right) \left(\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1}) + \log \frac{8}{\delta} \right)} \right). \tag{39}$$

Proof of Cor. D.1. Fix $i \in [2^t]$. By taking $\delta' = \frac{\delta}{2}$ in (38), we obtain (39). This occurs with probability

$$\begin{aligned}
 & \Pr\left(\text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}}^{(i)}) < a_{n,m} + v_{n,m} \sqrt{\log\left(\frac{1}{\delta'}\right)}\right) \\
 &= 1 - \Pr\left(\text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}}^{(i)}) \geq a_{n,m} + v_{n,m} \sqrt{\log\left(\frac{1}{\delta'}\right)}\right) \\
 &\geq 1 - \Pr\left(\mathcal{E}_{\text{equi}}^{(i)}, \text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}}^{(i)}) \geq a_{n,m} + v_{n,m} \sqrt{\log\left(\frac{1}{\delta'}\right)}\right) - \Pr\left(\mathcal{E}_{\text{equi}}^{(i)c}\right) \\
 &\geq 1 - \Pr\left(\left\|\tilde{\phi}^{(i)} \mathbf{k}\right\|_{\mathcal{H}_{\mathbf{k}}} \geq a_{n,m} + v_{n,m} \sqrt{\log\left(\frac{1}{\delta'}\right)}\right) - \Pr\left(\mathcal{E}_{\text{equi}}^{(i)c}\right) \\
 &\geq 1 - \frac{\delta}{2} - \frac{\delta}{2} = 1 - \delta.
 \end{aligned}$$

□

D.3. Proof of Thm. 3: MMD guarantee for SKT

Thm. 3 follows directly from Assum. (α, β) -kernel and the following result for a kernel with generic covering number.

Theorem D.1. *Let $\mathbf{k}_{\mathbb{P}}$ be a kernel satisfying Assum. 1. Let \mathcal{S}_{∞} be an infinite sequence of points. Then for a prefix sequence \mathcal{S}_n of n points, $m \in [n]$, and $n' \triangleq m2^{\lceil \log_2 \frac{n}{m} \rceil}$, SKT outputs w_{SKT} in $O(n^2 d_{\mathbf{k}_{\mathbb{P}}})$ time that satisfies, with probability at least $1 - \delta$,*

$$\Delta \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(w_{\text{SKT}}) \leq \sqrt{\left(\frac{1 + \log n'}{n'}\right) \|\mathbf{k}_{\mathbb{P}}\|_n} + \frac{1}{m} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6(\log_2 \frac{n'}{m})m}{\delta}\right) (\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1}) + \log \frac{8}{\delta})}\right).$$

Proof of Thm. D.1. The runtime of SKT comes from the fact that all of SteinThinning (with output size n), KT-SPLIT, and KT-Swap take $O(d_{\mathbf{k}_{\mathbb{P}}} n^2)$ time.

By Riabiz et al. (2022, Theorem 1), SteinThinning (which is a deterministic algorithm) from n points to n' points has the following guarantee

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^\dagger}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \left(\frac{1 + \log n'}{n'}\right) \|\mathbf{k}_{\mathbb{P}}\|_n,$$

where we denote the output weight of SteinThinning as w^\dagger . Using $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$, we have

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^\dagger}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{\left(\frac{1 + \log n'}{n'}\right) \|\mathbf{k}_{\mathbb{P}}\|_n}.$$

Fix $\delta \in (0, 1)$. By Cor. D.1 with $\mathbf{k} = \mathbf{k}_{\mathbb{P}}$ and $\mathfrak{t} = \log_2 \frac{n'}{m}$, with probability at least $1 - \delta$, we have, for any $i \in [2^t]$,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^\dagger}, \mathbb{S}_{\text{out}}^{(i)}) \leq \frac{1}{m} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6(\log_2 \frac{n'}{m})m}{\delta}\right) (\log \mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), m^{-1}) + \log \frac{8}{\delta})}\right),$$

where $\mathbb{S}_{\text{out}}^{(i)}$ is the i -th coreset output by KT-SPLIT. Since KT-Swap can only decrease the MMD to \mathbb{P} , we have, by the triangle inequality of $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}$,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w_{\text{SKT}}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_{\text{out}}^{(1)}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_{\text{out}}^{(1)}, \mathbb{S}_n^{w^\dagger}) + \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^\dagger}, \mathbb{P}),$$

which gives the desired bound. □

Thm. 3 now follows from Thm. D.1, the kernel growth definitions in Assum. (α, β) -kernel, $n \leq n' \leq 2n$, and that $\log_2(\frac{n'}{m})m \leq n'$. □

E. Resampling of Simplex Weights

Integral to many of our algorithms is a resampling procedure that turns a simplex-weighted point set of size n into an equal-weighted point set of size m while incurring at most $O(1/\sqrt{m})$ MMD error. The motivation for wanting an equal-weighted point set is two-fold: First, in [LSKT](#), we need to provide an equal-weighted point set to [KT-Compress++](#), but the output of [LD](#) is a simplex weight. Secondly, we can exploit the fact that non-zero weights are bounded away from zero in equal-weighted point sets to provide a tighter analysis of [WeightedRPCholesky](#). While i.i.d. resampling also achieves the $O(1/\sqrt{m})$ goal, we choose [Resample](#) (Alg. E.3), a stratified residual resampling algorithm ([Douc and Cappé, 2005](#), Sec. 3.2, 3.3). In this section, we derive an MMD bound for [Resample](#) and show that it is better in expectation than using i.i.d. resampling or residual resampling alone.

Let D_w^{inv} be the inverse of the cumulative distribution function of the multinomial distribution with weight w , i.e.,

$$D_w^{\text{inv}}(u) \triangleq \min \left\{ i \in [n] : u \leq \sum_{j=1}^i w_j \right\}.$$

Algorithm E.1 i.i.d. resampling

Input: Weights $w \in \Delta_{n-1}$, output size m

$w' \leftarrow \mathbf{0} \in \mathbb{R}^n$

for $j = 1$ **to** m **do**

 Draw $U_j \sim \text{Uniform}([0, 1])$

$I_j \leftarrow D_w^{\text{inv}}(U_j)$

$w'_{I_j} \leftarrow w'_{I_j} + \frac{1}{m}$

end for

Return: $w' \in \Delta_{n-1} \cap \left(\frac{\mathbb{N}_0}{m}\right)^n$

Algorithm E.2 Residual resampling

Input: Weights $w \in \Delta_{n-1}$, output size m

$w'_i \leftarrow \frac{\lfloor mw_i \rfloor}{m}, \forall i \in [n]$

$r \leftarrow m - \sum_{i \in [n]} \lfloor mw_i \rfloor \in \mathbb{N}$

$\eta_i \leftarrow \frac{mw_i - \lfloor mw_i \rfloor}{r}, \forall i \in [n]$

for $j = 1$ **to** r **do**

 Draw $U_j \sim \text{Uniform}([0, 1])$

$I_j \leftarrow D_\eta^{\text{inv}}(U_j)$

$w'_{I_j} \leftarrow w'_{I_j} + \frac{1}{m}$

end for

Return: $w' \in \Delta_{n-1} \cap \left(\frac{\mathbb{N}_0}{m}\right)^n$

Algorithm E.3 Stratified residual resampling (Resample)

Input: Weights $w \in \Delta_{n-1}$, output size m

$w'_i \leftarrow \frac{\lfloor mw_i \rfloor}{m}, \forall i \in [n]$

$r \leftarrow m - \sum_{i \in [n]} \lfloor mw_i \rfloor \in \mathbb{N}$

$\eta_i \leftarrow \frac{mw_i - \lfloor mw_i \rfloor}{r}, \forall i \in [n]$

for $j = 1$ **to** r **do**

 Draw $U_j \sim \text{Uniform}\left(\left[\frac{j}{r}, \frac{j+1}{r}\right]\right)$

$I_j \leftarrow D_\eta^{\text{inv}}(U_j)$

$w'_{I_j} \leftarrow w'_{I_j} + \frac{1}{m}$

end for

Return: $w' \in \Delta_{n-1} \cap \left(\frac{\mathbb{N}_0}{m}\right)^n$

Proposition E.1 (MMD guarantee of resampling algorithms). *Consider any kernel \mathbf{k} , points $\mathcal{S}_n = (x_1, \dots, x_n) \subset \mathbb{R}^d$, and a weight vector $w \in \Delta_{n-1}$.*

(a) *Using the notation from Alg. E.1, let X, X' be independent random variables with law \mathbb{S}_n^w . Then, the output weight vector $w^{\text{i.i.d.}} \triangleq w'$ of Alg. E.1 satisfies*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w^{\text{i.i.d.}}}, \mathbb{S}_n^w)] = \frac{\mathbb{E}\mathbf{k}(X, X) - \mathbb{E}\mathbf{k}(X, X')}{m}. \quad (40)$$

(b) *Using the notation from Alg. E.2, let R, R' be independent random variables with law \mathbb{S}_n^η . Then, the output weight vector $w^{\text{resid}} \triangleq w'$ of Alg. E.2 satisfies*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w^{\text{resid}}}, \mathbb{S}_n^w)] = \frac{r(\mathbb{E}\mathbf{k}(R, R) - \mathbb{E}\mathbf{k}(R, R'))}{m^2}. \quad (41)$$

(c) *Using the notation from Alg. E.3, let $R_j \triangleq x_{\mathbb{I}_j}$ and R'_j be an independent copy of R_j . Let R be an independent random variable with law \mathbb{S}_n^η . Then, the output weight vector $w^{\text{sr}} \triangleq w'$ of Alg. E.3 satisfies*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^w)] = \frac{r\mathbb{E}\mathbf{k}(R, R) - \sum_{j \in [r]} \mathbb{E}\mathbf{k}(R_j, R'_j)}{m^2}. \quad (42)$$

Proof of Prop. E.1(a). Let $X_i \triangleq x_{\mathbb{I}_i}$. As random signed measures, we have

$$\mathbb{S}_n^{w'} - \mathbb{S}_n^w = \frac{1}{m} \sum_{i \in [m]} \delta_{X_i} - \sum_{i \in [n]} w_i \delta_{x_i}.$$

Hence

$$\begin{aligned} \text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w) &= ((\mathbb{S}_n^{w'} - \mathbb{S}_n^w) \times (\mathbb{S}_n^{w'} - \mathbb{S}_n^w))_{\mathbf{k}} \\ &= \frac{1}{m^2} \sum_{i, i' \in [m]} \mathbf{k}(X_i, X_{i'}) - \frac{2}{m} \sum_{i \in [m], i' \in [n]} w_{i'} \mathbf{k}(X_i, x_{i'}) + \sum_{i, i' \in [n]} w_i w_{i'} \mathbf{k}(x_i, x_{i'}). \end{aligned}$$

Since each X_i is distributed to \mathbb{S}_n^w and X_i and $X_{i'}$ are independent for $i \neq i'$, taking expectation, we have

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w)] = \frac{1}{m} \mathbb{E}\mathbf{k}(X, X) + \frac{m-1}{m} \mathbb{E}\mathbf{k}(X, X') - 2\mathbb{E}\mathbf{k}(X, X') + \mathbb{E}\mathbf{k}(X, X').$$

This gives the bound (40). \square

Proof of Prop. E.1(b). Let $R_j \triangleq x_{\mathbb{I}_j}$. As random signed measures, we have

$$\begin{aligned} \mathbb{S}_n^{w'} - \mathbb{S}_n^w &= \left(\sum_{i \in [n]} \frac{\lfloor mw_i \rfloor}{m} \delta_{x_i} + \frac{1}{m} \sum_{j \in [r]} \delta_{R_j} \right) - \sum_{i \in [n]} w_i \delta_{x_i} \\ &= \frac{1}{m} \sum_{j \in [r]} \delta_{R_j} - \sum_{i \in [n]} \left(w_i - \frac{\lfloor mw_i \rfloor}{m} \right) \delta_{x_i} \\ &= \frac{1}{m} \sum_{j \in [r]} \delta_{R_j} - \frac{r}{m} \sum_{i \in [n]} \eta_i \delta_{x_i}. \end{aligned}$$

Hence

$$\begin{aligned} \text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w) &= ((\mathbb{S}_n^{w'} - \mathbb{S}_n^w) \times (\mathbb{S}_n^{w'} - \mathbb{S}_n^w))_{\mathbf{k}} \\ &= \frac{1}{m^2} \sum_{j, j' \in [r]} \mathbf{k}(R_j, R_{j'}) - \frac{2r}{m^2} \sum_{j \in [r], i \in [n]} \eta_i \mathbf{k}(R_j, x_i) + \frac{r^2}{m^2} \sum_{i, i' \in [n]} \eta_i \eta_{i'} \mathbf{k}(x_i, x_{i'}). \end{aligned} \quad (43)$$

Since each R_j is distributed to \mathbb{S}_n^η and R_j and $R_{j'}$ are independent for $j \neq j'$, taking expectation, we have

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w)] = \frac{r}{m^2} \mathbb{E}\mathbf{k}(R, R) + \frac{r(r-1)}{m^2} \mathbb{E}\mathbf{k}(R, R') - \frac{2r^2}{m^2} \mathbb{E}\mathbf{k}(R, R') + \frac{r^2}{m^2} \mathbb{E}\mathbf{k}(R, R').$$

This gives the bound (41). \square

Proof of Prop. E.1(c). We repeat the same steps from the previous part of the proof to get (43). In the case of (c), R_j 's are not identically distributed so the analysis is different. Let R' be an independent copy of R . Taking expectation of (43), we have

$$m^2 \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w)] = \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R_j) + \sum_{j \in [r]} \sum_{j' \in [r] \setminus \{j\}} \mathbb{E} \mathbf{k}(R_j, R_{j'}) - 2r \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R) + r^2 \mathbb{E} \mathbf{k}(R, R').$$

Note

$$\sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R_j) = \sum_{j \in [r]} r \int_{[\frac{j}{r}, \frac{j+1}{r})} \mathbf{k}(x_{D_{\eta}^{\text{inv}}(u)}, x_{D_{\eta}^{\text{inv}}(u)}) du = r \int_0^1 \mathbf{k}(x_{D_{\eta}^{\text{inv}}(u)}, x_{D_{\eta}^{\text{inv}}(u)}) du = r \mathbb{E} \mathbf{k}(R, R),$$

where we used the fact that $x_{D_{\eta}^{\text{inv}}(U)} \stackrel{D}{=} R$ for $U \sim \text{Uniform}([0, 1])$. Similarly, we deduce

$$\begin{aligned} \sum_{j \in [r]} \sum_{j' \in [r] \setminus \{j\}} \mathbb{E} \mathbf{k}(R_j, R_{j'}) &= \sum_{j \in [r]} \left(\sum_{j' \in [r]} \mathbb{E} \mathbf{k}(R_j, R_{j'}) - \mathbb{E} \mathbf{k}(R_j, R_j) \right) \\ &= \sum_{j \in [r]} (r \mathbb{E} \mathbf{k}(R_j, R') - \mathbb{E} \mathbf{k}(R_j, R_j)) \\ &= r^2 \mathbb{E} \mathbf{k}(R, R') - \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R_j), \end{aligned}$$

and also

$$\sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R) = r \mathbb{E} \mathbf{k}(R, R').$$

Combining terms, we get

$$\begin{aligned} m^2 \mathbb{E} \text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w'}, \mathbb{S}_n^w) &= r \mathbb{E} \mathbf{k}(R, R) + r^2 \mathbb{E} \mathbf{k}(R, R') - \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R_j) - 2r^2 \mathbb{E} \mathbf{k}(R, R') + r^2 \mathbb{E} \mathbf{k}(R, R') \\ &= r \mathbb{E} \mathbf{k}(R, R) - \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R_j), \end{aligned}$$

which yields the desired bound (42). \square

The next proposition shows that stratifying the residuals always improves upon using i.i.d. sampling or residual resampling alone. We need the following convexity lemma.

Lemma E.1 (Convexity of squared MMD). *Let \mathbf{k} be a kernel. Let $\mathcal{S}_n = (x_1, \dots, x_n)$ be an arbitrary set of points. The function $E_{\mathbf{k}} : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by*

$$E_{\mathbf{k}}(w) \triangleq \|\mathbb{S}_n^w \mathbf{k}\|_{\mathcal{H}_{\mathbf{k}}}^2 = \sum_{i, j \in [n]} w_i w_j \mathbf{k}(x_i, x_j)$$

is convex on \mathbb{R}^n .

Proof of Lem. E.1. Since \mathbf{k} is a kernel, the Hessian $\nabla^2 E_{\mathbf{k}} = 2\mathbf{k}(\mathcal{S}_n, \mathcal{S}_n)$ is PSD, and hence $E_{\mathbf{k}}$ is convex. \square

Proposition E.2 (Stratified residual resampling improves MMD). *Under the assumptions of Prop. E.1, we have*

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{i.i.d.}}, \mathbb{S}_n^w)] \geq \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{resid}}, \mathbb{S}_n^w)] \geq \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{sr}}, \mathbb{S}_n^w)].$$

Proof of Prop. E.2. Let $K \triangleq \mathbf{k}(\mathcal{S}_n, \mathcal{S}_n)$. To show the first inequality, note that since $\eta = \frac{mw - \lfloor mw \rfloor}{r}$, by Prop. E.1,

$$\begin{aligned} \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{resid}}, \mathbb{S}_n^w)] &= \frac{r(\mathbb{E} \mathbf{k}(R, R) - \mathbb{E} \mathbf{k}(R, R'))}{m^2} \\ &= \frac{r(\sum_{i \in [n]} K_{ii} \eta_i - \sum_{i, j \in [n]} K_{ij} \eta_i \eta_j)}{m^2} \\ &= \frac{1}{m} \left(\sum_{i \in [n]} K_{ii} \left(w_i - \frac{\lfloor mw_i \rfloor}{m} \right) - \frac{m}{r} \left(w - \frac{\lfloor mw \rfloor}{m} \right)^\top K \left(w - \frac{\lfloor mw \rfloor}{m} \right) \right). \end{aligned}$$

Hence

$$\begin{aligned} &\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{i.i.d.}}, \mathbb{S}_n^w)] - \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{w, \text{resid}}, \mathbb{S}_n^w)] \\ &= \frac{1}{m} \left(\sum_{i \in [n]} K_{ii} \frac{\lfloor mw_i \rfloor}{m} + \frac{m}{r} \left(w - \frac{\lfloor mw \rfloor}{m} \right)^\top K \left(w - \frac{\lfloor mw \rfloor}{m} \right) - w^\top K w \right) \\ &= \frac{1}{m} \left((1 - \theta) \sum_{i \in [n]} K_{ii} \xi_i + \theta \eta^\top K \eta - w^\top K w \right), \end{aligned}$$

where we let $\xi \triangleq \frac{m}{m-r} \lfloor \frac{mw}{m} \rfloor$ and $\theta \triangleq \frac{r}{m}$. Note that $w = \theta\eta + (1 - \theta)\xi$. By Lem. E.1 and Jensen's inequality, we have

$$w^\top K w = E_{\mathbf{k}}(w) \leq \theta E_{\mathbf{k}}(\eta) + (1 - \theta) E_{\mathbf{k}}(\xi) = \theta \eta^\top K \eta + (1 - \theta) \xi^\top K \xi \leq \theta \eta^\top K \eta + (1 - \theta) \sum_{i \in [n]} K_{ii} \xi_i,$$

where the last inequality follows from Prop. E.1(a) with $w = \xi$ and the fact that MMD is nonnegative. Hence we have shown

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{i.i.d.}}, \mathbb{S}_n^w)] - \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{resid}}, \mathbb{S}_n^w)] \geq 0,$$

as desired.

For the second inequality, by Prop. E.1, we compute

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{resid}}, \mathbb{S}_n^w)] - \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{sr}}, \mathbb{S}_n^w)] = \frac{r}{m^2} \left(\frac{1}{r} \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R'_j) - \mathbb{E} \mathbf{k}(R, R') \right).$$

Note that

$$\mathbb{E} \mathbf{k}(R, R') = \int_{[0,1]} \int_{[0,1]} k(x_{D_\eta^{\text{inv}}(u)}, x_{D_\eta^{\text{inv}}(v)}) du dv = E_{\mathbf{k}} \left((D_\eta^{\text{inv}})_{\#} \text{Uniform}[0, 1] \right),$$

where we used $T_{\#} \mu$ to denote the pushforward measure of μ by T . Similarly,

$$\begin{aligned} \frac{1}{r} \sum_{j \in [r]} \mathbb{E} \mathbf{k}(R_j, R'_j) &= \frac{1}{r} \sum_{j \in [r]} \int_{[\frac{j}{r}, \frac{j+1}{r}]} \int_{[\frac{j}{r}, \frac{j+1}{r}]} k(x_{D_\eta^{\text{inv}}(u)}, x_{D_\eta^{\text{inv}}(v)}) du dv \\ &= \frac{1}{r} \sum_{j \in [r]} E_{\mathbf{k}} \left((D_\eta^{\text{inv}})_{\#} \text{Uniform} \left[\frac{j}{r}, \frac{j+1}{r} \right] \right) \\ &\leq E_{\mathbf{k}} \left((D_\eta^{\text{inv}})_{\#} \text{Uniform}[0, 1] \right) = \mathbb{E} \mathbf{k}(R, R'), \end{aligned}$$

where in the last inequality we applied Jensen's inequality since $E_{\mathbf{k}}$ is convex by Lem. E.1. Hence we have shown

$$\mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{resid}}, \mathbb{S}_n^w)] - \mathbb{E}[\text{MMD}_{\mathbf{k}}^2(\mathbb{S}_n^{\text{sr}}, \mathbb{S}_n^w)] \geq 0$$

and the proof is complete. \square

F. Accelerated Debiased Compression

In this section, we provide supplementary algorithmic details and deferred analyses for **LSKT** (Alg. 3). In **WeightedRPC-holesky** (Alg. F.1), we provide details for the weighted extension of [Chen et al. \(2022, Alg. 2.1\)](#) that is used extensively in our algorithms. The details of **AMD** ([Wang et al., 2023, Alg. 14](#)) are provided in Alg. F.2. In App. F.1, we give the proof of Thm. 4 for the MMD error guarantee of **LD** (Alg. 2). In App. F.2, we provide details on **KT-Compress++** modified from **Compress++** ([Shetty et al., 2022](#)) to minimize MMD to \mathbb{P} . Finally, Thm. 5 is proved in App. F.3.

Algorithm F.1 Weighted Randomly Pivoted Cholesky (WeightedRPCholesky) (extension of [Chen et al. \(2022, Alg. 2.1\)](#))

Input: kernel \mathbf{k} , points $\mathbb{S}_n = (x_i)_{i=1}^n$, simplex weights $w \in \Delta_{n-1}$, rank r

$\tilde{\mathbf{k}}(i, j) \triangleq \mathbf{k}(x_i, x_j) \sqrt{w_i} \sqrt{w_j} \triangleright$ reweighted kernel matrix function

$F \leftarrow \mathbf{0}_{n \times r}$, $\mathbf{S} \leftarrow \{\}$, $d \leftarrow (\tilde{\mathbf{k}}(i, i))_{i \in [n]}$

for $i = 1$ **to** r **do**

Sample $s \sim d / \sum_{j \in [n]} d_j$

$\mathbf{S} \leftarrow \mathbf{S} \cup \{s\}$

$g \leftarrow \tilde{\mathbf{k}}(:, s) - F(:, 1 : i - 1) F(s, 1 : i - 1)^\top$

$F(:, i) \leftarrow g / \sqrt{g_s}$

$d \leftarrow d - F(:, i)^2 \triangleright F(:, i)^2$ denotes a vector with entry-wise squared values of $F(:, i)$

$d \leftarrow \max(d, 0) \triangleright$ numerical stability fix, helpful in practice

end for

$F \leftarrow \text{diag}((1/\sqrt{w_i})_{i \in [n]}) F \triangleright$ undo weighting; treat $1/\sqrt{w_i} = 0$ if $w_i = 0$

Return: $\mathbf{S} \subset [n]$ with $|\mathbf{S}| = r$ and $F \in \mathbb{R}^{n \times r}$

Algorithm F.2 Accelerated Entropic Mirror Descent (AMD) (modification of Wang et al. (2023, Alg. 14))

Input: kernel matrix $K \in \mathbb{R}^{n \times n}$, number of steps T , initial weight $w_0 \in \Delta_{n-1}$, aggressive flag AGG

 $\eta \leftarrow \frac{1}{8w_0^\top \text{diag}(K)}$ if AGG else $\frac{1}{8 \max_{i \in [n]} K_{ii}}$
 $v_0 \leftarrow w_0$
for $t = 1$ **to** T **do**
 $\beta_t \leftarrow \frac{2}{t+1}$
 $z_t \leftarrow (1 - \beta_t)w_{t-1} + \beta_t v_{t-1}$
 $g \leftarrow 2t\eta K z_t \triangleright$ this is $\gamma_t \nabla f(z_t)$ in Wang et al. (2023, Alg. 14) for $f(w) = w^\top K w$
 $v_t \leftarrow v_{t-1} \cdot \exp(-g) \triangleright$ component-wise exponentiation and multiplication

 $v_t \leftarrow v_t / \|v_t\|_1 \triangleright v_t = \arg\min_{w \in \Delta_{n-1}} \langle g, w \rangle + D_{v_{t-1}}^\phi(w)$ for $\phi(w) = \sum_{i \in [n]} w_i \log w_i$
 $w_t \leftarrow (1 - \beta_t)w_{t-1} + \beta_t v_t$
end for
Return: $w_T \in \Delta_{n-1}$

F.1. Proof of Thm. 4: Debiasing guarantee for LD

 We start with a useful lemma that bounds $w^\top (K - \hat{K})w$ by $\text{tr}(K - \hat{K})$ for any simplex weights w .

Lemma F.1. For any PSD matrix $A \in \mathbb{R}^{n \times n}$ and $w \in \Delta_{n-1}$, we have

$$w^\top A w \leq \text{tr}(A^w) \leq \max_{i \in [n]} A_{ii} \leq \lambda_1(A),$$

 where $\lambda_1(A)$ denotes the largest eigenvalue of A .

Proof of Lem. F.1. Note that

$$w^\top A w = \sqrt{w}^\top \text{diag}(\sqrt{w}) A \text{diag}(\sqrt{w}) \sqrt{w} = \sqrt{w}^\top A^w \sqrt{w}.$$

 The condition that $w \in \Delta_{n-1}$ implies $\|\sqrt{w}\|_2 = 1$, so that

$$\sqrt{w}^\top A^w \sqrt{w} \leq \lambda_1(A^w) \leq \text{tr}(A^w).$$

 To see $\text{tr}(A^w) \leq \max_{i \in [n]} A_{ii}$, note that $\text{tr}(A^w) = \sum_{i \in [n]} A_{ii} w_i \leq \max_{i \in [n]} A_{ii}$ since $w \in \Delta_{n-1}$.

 Since $\lambda_1(A) = \sup_{x: \|x\|_2=1} x^\top A x$, if we let $i^* \triangleq \arg\min_{i \in [n]} A_{ii}$, then the simplex weight with 1 on the i^* -th entry has two-norm 1, so we see that $\max_{i \in [n]} A_{ii} \leq \lambda_1(A)$. \square

 Our next lemma bounds the suboptimality of surrogate optimization of a low-rank plus diagonal approximation of K .

Lemma F.2 (Suboptimality of surrogate optimization). Let $\mathbf{k}_{\mathbb{P}}$ be a kernel satisfying Assum. 1. Let $\mathcal{S}_n = (x_1, \dots, x_n) \subset \mathbb{R}^d$ be a sequence of points. Define $K \triangleq \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n) \in \mathbb{R}^{n \times n}$. Suppose $\hat{K} \in \mathbb{R}^{n \times n}$ is another PSD matrix such that $K \succeq \hat{K}$. Define $D \triangleq \text{diag}(K - \hat{K})$, the diagonal part of $K - \hat{K}$, and form $K' \triangleq \hat{K} + D$. Let $w' \in \arg\min_{w \in \Delta_{n-1}} w'^\top K' w'$. Then for any $w \in \Delta_{n-1}$,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^w, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \text{tr}\left((K - \hat{K})^w\right) + \max_{i \in [n]} (K - \hat{K})_{ii} + (w^\top K' w - w'^\top K' w'). \quad (44)$$

Proof of Lem. F.2. Since $K = K' + (K - \hat{K}) - D$ by construction, we have

$$\begin{aligned} w^\top K w &= w^\top K' w + w^\top (K - \hat{K}) w - w^\top D w \\ &\leq w^\top K' w + w^\top (K - \hat{K}) w \\ &= (w^\top K' w - w'^\top K' w') + w'^\top K' w' + w^\top (K - \hat{K}) w \\ &\leq (w^\top K' w - w'^\top K' w') + w'^\top K' w' + \text{tr}\left((K - \hat{K})^w\right), \end{aligned}$$

where we used the fact that $D \succeq 0$ and Lem. F.1. Next, by the definition of w' , we have

$$\begin{aligned} w'^{\top} K' w' &\leq (w_{\text{OPT}})^{\top} K' w_{\text{OPT}} = (w_{\text{OPT}})^{\top} (K' - K) w_{\text{OPT}} + (w_{\text{OPT}})^{\top} K w_{\text{OPT}} \\ &= (w_{\text{OPT}})^{\top} (D - (K - \widehat{K})) w_{\text{OPT}} + (w_{\text{OPT}})^{\top} K w_{\text{OPT}} \\ &\leq (w_{\text{OPT}})^{\top} D w_{\text{OPT}} + (w_{\text{OPT}})^{\top} K w_{\text{OPT}} \\ &\leq \max_{i \in [n]} (K - \widehat{K})_{ii} + (w_{\text{OPT}})^{\top} K w_{\text{OPT}}, \end{aligned}$$

where we used the fact $K \succeq \widehat{K}$ in the penultimate step and Lem. F.1 in the last step. Hence we have shown our claim. \square

Lem. F.2 shows that to control $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^w, \mathbb{P})$, it suffices to separately control the approximation error in terms of $\text{tr}(K - \widehat{K})$ and the optimization error $(w^{\top} K' w - w'^{\top} K' w')$. The next result establishes that using [WeightedRPCholesky](#), we can obtain polynomial and exponential decay bounds for $\text{tr}(K - \widehat{K})$ in expectation depending on the kernel growth of $\mathbf{k}_{\mathbb{P}}$.

Proposition F.1 (Approximation error of [WeightedRPCholesky](#)). *Let \mathbf{k} be a kernel satisfying Assum. (α, β) -kernel. Let \mathcal{S}_{∞} be an infinite sequence of points in \mathbb{R}^d . For any $w \in \Delta_{n-1}$, let F be the low-rank approximation factor output by [WeightedRPCholesky](#)($\mathbf{k}, \mathcal{S}_n, w, r$). Define $K \triangleq \mathbf{k}(\mathcal{S}_n, \mathcal{S}_n)$. If $r \geq (\frac{c_d R_n^{\beta} + 1}{\sqrt{\log 2}} + \sqrt{\log 2})^2 - \frac{1}{\log 2}$, then, with the expectation taken over the randomness in [WeightedRPCholesky](#),*

$$\mathbb{E} [\text{tr}((K - FF^{\top})w)] \leq H_{n,r}, \quad (45)$$

where $H_{n,r}$ is defined as

$$H_{n,r} \triangleq \begin{cases} 8 \sum_{\ell=\mathfrak{U}(r)}^n \left(\frac{L_{\mathbf{k}}(R_n)}{\ell}\right)^{\frac{2}{\alpha}} & \text{POLYGROWTH}(\alpha, \beta), \\ 8 \sum_{\ell=\mathfrak{U}(r)}^n \exp\left(1 - \left(\frac{\ell}{L_{\mathbf{k}}(R_n)}\right)^{\frac{1}{\alpha}}\right) & \text{LOGGROWTH}(\alpha, \beta), \end{cases} \quad (46)$$

for $L_{\mathbf{k}}$ defined in (6) and

$$\mathfrak{U}(r) \triangleq \left\lfloor \sqrt{\frac{r + \frac{1}{\log 2}}{\log 2}} - \frac{1}{\log 2} \right\rfloor. \quad (47)$$

Moreover, $H_{n,r}$ satisfies the bounds in Thm. 4.

Proof of Prop. F.1. Recall the notation $L_{\mathbf{k}}(R_n) = \frac{c_d R_n^{\beta}}{\log 2}$ from (6). Define $q \triangleq \mathfrak{U}(r)$ so that q is the biggest integer for which $r \geq 2q + q^2 \log 2$. The lower bound assumption of r is chosen such that $q > L_{\mathbf{k}}(R_n) > 0$. By [Chen et al. \(2022, Theorem 3.1\)](#) with $\epsilon = 1$, we have

$$\mathbb{E} [\text{tr}((K - FF^{\top})w)] \leq 2 \sum_{\ell=q+1}^n \lambda_{\ell}(K^w). \quad (48)$$

Since $q > L_{\mathbf{k}}(R_n)$, we can apply Cor. B.1 to bound $\lambda_{\ell}(K^w)$ for $\ell \geq q + 1$ and obtain (45) since $H_{n,r}$ (46) is constructed to match the bounds when applying Cor. B.1 to (48). It remains to justify the bounds for $H_{n,r}$ in Thm. 4.

If \mathbf{k} is [POLYGROWTH](#)(α, β), by Assum. (α, β) -kernel we have $\alpha < 2$. Hence

$$H_{n,r} = 8 \sum_{\ell=q}^n \left(\frac{L_{\mathbf{k}}(R_n)}{\ell}\right)^{\frac{2}{\alpha}} \leq 8 L_{\mathbf{k}}(R_n)^{\frac{2}{\alpha}} \int_{q-1}^{\infty} \ell^{-\frac{2}{\alpha}} d\ell = 8 L_{\mathbf{k}}(R_n)^{\frac{2}{\alpha}} (q-1)^{1-\frac{2}{\alpha}} = O\left(\sqrt{r} \left(\frac{R_n^{2\beta}}{r}\right)^{\frac{1}{\alpha}}\right),$$

where we used the fact that $\int_{q-1}^{\infty} \ell^{-\frac{2}{\alpha}} d\ell = (q-1)^{1-\frac{2}{\alpha}}$ for $\alpha < 2$, $L_{\mathbf{k}}(R_n) = O(R_n^{\beta})$, and $q = \Theta(\sqrt{r})$.

If \mathbf{k} is [LOGGROWTH](#)(α, β), then

$$H_{n,r} = 8 \sum_{\ell=q}^n \exp\left(1 - \left(\frac{\ell}{L_{\mathbf{k}}(R_n)}\right)^{\frac{1}{\alpha}}\right) = 8e \sum_{\ell=q}^n c^{\ell^{1/\alpha}} \leq 8e \int_{\ell=q-1}^{\infty} c^{\ell^{1/\alpha}},$$

where $c \triangleq \exp(-L_{\mathbf{k}}(R_n)^{-1/\alpha}) \in (0, 1)$. Defining $m \triangleq -\log c > 0$ and $q' = q - 1$, we have

$$\int_{x=q'}^{\infty} c^{x^{1/\alpha}} dx = \int_{x=q'}^{\infty} \exp(-mx^{1/\alpha}) dx = \alpha q' (mq'^{1/\alpha})^{-\alpha} \Gamma(\alpha, mq'^{1/\alpha}) = \alpha m^{-\alpha} \Gamma(\alpha, mq'^{1/\alpha}), \quad (49)$$

where $\Gamma(\alpha, x) \triangleq \int_x^{\infty} t^{\alpha-1} e^{-t} dt$ is the incomplete gamma function. Since $\alpha > 0$, by [Pinelis \(2020, Thm. 1.1\)](#), we have

$$\Gamma(\alpha, mq'^{1/\alpha}) \leq \frac{(mq'^{1/\alpha} + b)^\alpha - (mq'^{1/\alpha})^\alpha}{\alpha b} e^{-mq'^{1/\alpha}},$$

where b is a known constant depending only on α . By the equivalence of norms on \mathbb{R}^2 , there exists $C_\alpha > 0$ such that $(x + y)^\alpha \leq C_\alpha(x^\alpha + y^\alpha)$ for any $x, y > 0$. Hence

$$\Gamma(\alpha, mq'^{1/\alpha}) \leq \frac{(mq'^{1/\alpha} + b)^\alpha}{\alpha b} e^{-mq'^{1/\alpha}} \leq \frac{C_\alpha(m^\alpha q' + b^\alpha)}{\alpha b} e^{-mq'^{1/\alpha}}.$$

Hence from (49) we deduce

$$\sum_{\ell=q'}^{\infty} c^{\ell^{1/\alpha}} \leq C_\alpha(q'b^{-1} + b^{\alpha-1}m^{-\alpha})e^{-mq'^{1/\alpha}}. \quad (50)$$

Since $m = -\log c = L_{\mathbf{k}}(R_n)^{-1/\alpha}$, we can bound the exponent by

$$-mq'^{1/\alpha} = -(L_{\mathbf{k}}(R_n)^{-1}q')^{1/\alpha} = -\left(\frac{q' \log 2}{\mathfrak{C}_d R_n^\beta}\right)^{1/\alpha} \leq -\left(\frac{0.83\sqrt{r}-2.39}{\mathfrak{C}_d R_n^\beta}\right)^{1/\alpha},$$

where we used the fact that $q' \log 2 = (q - 1) \log 2 \geq \left(\sqrt{\frac{r + \frac{1}{\log 2}}{\log 2}} - \frac{1}{\log 2} - 2\right) \log 2 \geq 0.83\sqrt{r} - 2.39$. On the other hand, since $q' = q - 1 \geq L(R_n) = m^{-\alpha}$, we can absorb the $b^{\alpha-1}m^{-1}$ term in (50) into q and finally obtain the bounds for $H_{n,r}$ in [Thm. 4](#). \square

The last piece of our analysis involves bounding the optimization error $(w^\top K' w - w'^\top K' w')$ in (44).

Lemma F.3 (AMD guarantee for debiasing). *Let $K \in \mathbb{R}^{n \times n}$ be an SPSD matrix. Let $f(w) \triangleq w^\top K w$. Then the final iterate x_T of Nesterov's 1-memory method ([Wang et al., 2023, Algorithm 14](#)) after T steps with objective function $f(w)$, norm $\|\cdot\| = \|\cdot\|_1$, distance-generating function $\phi(x) = \sum_{i=1}^n x_i \log x_i$, and initial point $w_0 = (\frac{1}{n}, \dots, \frac{1}{n}) \in \Delta_{n-1}$ satisfies*

$$f(w_T) - f(w_{\text{OPT}}) \leq \frac{16 \log n \max_{i \in [n]} K_{ii}}{T^2},$$

where $w_{\text{OPT}} \in \operatorname{argmin}_{x \in \mathbb{R}^n} f(x)$.

Proof of Lem. F.3. We apply [Wang et al. \(2023, Theorem 14\)](#). Hence it remains to determine the smoothness constant $L > 0$ such that, for all $x, y \in \Delta_{n-1}$,

$$\|\nabla f(x) - \nabla f(y)\|_\infty \leq L \|x - y\|_1,$$

and an upper bound for the Bregman divergence $D_{w_0}^\phi(w_{\text{OPT}}) = \sum_{i=1}^n w_{\text{OPT}i} \log \frac{w_{\text{OPT}i}}{(w_0)_i} = \sum_{i=1}^n w_{\text{OPT}i} \log n w_{\text{OPT}i}$. To determine L , note $\nabla f(w) = 2Kw$, so we have, for any $x, y \in \Delta_{n-1}$,

$$\begin{aligned} \|\nabla f(x) - \nabla f(y)\|_\infty &= 2 \|K(x - y)\|_\infty = 2 \max_{i \in [n]} |K_{i,:}(x - y)| \\ &\leq 2 \max_{i \in [n]} \|K_{i,:}\|_\infty \|x - y\|_1 = 2 (\max_{i \in [n]} K_{ii}) \|x - y\|_1 = 2 (\max_{i \in [n]} K_{ii}) \|x - y\|_1, \end{aligned}$$

where we used the fact that the largest entry in an SPSD matrix appears on its diagonal. Thus we can take the smoothness constant to be

$$L = 2 \max_{i \in [n]} K_{ii}.$$

To bound $D_{w_0}^\phi(w_{\text{OPT}})$, note that by Jensen's inequality,

$$D_{w_0}^\phi(w) = \sum_{i=1}^n w_i \log n w_i \leq \log \left(\sum_{i=1}^n n w_i^2 \right) = \log n + \log \|w\|_2^2 \leq \log n,$$

where we used the fact that $\|w\|_2^2 \leq \|w\|_1 = 1$ for $w \in \Delta_{n-1}$. \square

With these tools in hand, we turn to the proof of Thm. 4. For the runtime of LD, it follows from the fact that [WeightedRPCholesky](#) takes $O((d_{\mathbf{k}_P} + r)nr)$ time and one step of [AMD](#) takes $O(nr)$ time.

The error analysis is different for the first adaptive iteration and the ensuing adaptive iterations. Roughly speaking, we will show that the first adaptive iteration brings the MMD gap to the desired level, while the ensuing iterations do not introduce an excessive amount of error.

Step 1. Bound $\Delta \text{MMD}_{\mathbf{k}_P}(w^{(1)})$

Let $K \triangleq \mathbf{k}_P(\mathcal{S}_n, \mathcal{S}_n)$ and F denote the low-rank approximation factor generated by [WeightedRPCholesky](#). Denote $\widehat{K} \triangleq FF^\top$. Then $K' = \widehat{K} + \text{diag}(K - \widehat{K})$. First, note that since $w^{(0)} = (\frac{1}{n}, \dots, \frac{1}{n})$, [Resample](#) returns $\tilde{w} = w^{(0)}$ with probability one. By Lem. F.2, we have, using $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$ repeatedly and Lem. F.1 that $\text{tr}((K - \widehat{K})^w) \leq \lambda_1(K - \widehat{K})$ and $\max_{i \in [n]} (K - \widehat{K})_{ii} \leq \lambda_1(K - \widehat{K})$,

$$\begin{aligned} \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(1)}}, \mathbb{P}) &\leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{2\lambda_1(K - \widehat{K})} + \sqrt{w^{(1)\top} K' w^{(1)} - w'^\top K' w'} \\ &\leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{2\lambda_1(K - \widehat{K})} + \sqrt{\frac{16 \log n \|\mathbf{k}_P\|_n}{T^2}}, \end{aligned}$$

where we applied Lem. F.1 and Lem. F.3 in the last inequality. Fix $\delta \in (0, 1)$. By Markov's inequality, we have

$$\Pr\left(\sqrt{\lambda_1(K - \widehat{K})} > \sqrt{\frac{\mathbb{E}[\lambda_1(K - \widehat{K})]}{\delta}}\right) \leq \delta.$$

This means that with probability at least $1 - \delta$, we have

$$\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(1)}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{\frac{2\mathbb{E}[\lambda_1(K - \widehat{K})]}{\delta}} + \sqrt{\frac{16 \log n \|\mathbf{k}_P\|_n}{T^2}}.$$

Note that the lower bound condition on r in Assum. [\(\alpha, \beta\)](#)-params implies the lower bound condition in Prop. F.1. Hence, by Prop. F.1 with $w = (\frac{1}{n}, \dots, \frac{1}{n})$ and using the identity $\lambda_1(K - \widehat{K}) \leq \text{tr}(K - \widehat{K})$ while noting that a factor of n appears, we have

$$\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(1)}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{\frac{2nH_{n,r}}{\delta}} + \sqrt{\frac{16\|\mathbf{k}_P\|_n \log n}{T^2}}.$$

Step 2. Bound the error of the remaining iterations

Fix $\delta > 0$. The previous step shows that, with probability at least $1 - \frac{\delta}{2}$,

$$\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(1)}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w_{\text{OPT}}}, \mathbb{P}) + \sqrt{\frac{4nH_{n,r}}{\delta}} + \sqrt{\frac{16\|\mathbf{k}_P\|_n \log n}{T^2}}.$$

Fix $q > 1$, and let \tilde{w} be the resampled weight defined in the q -th iteration in Alg. 2. Without loss of generality, we assume $\tilde{w}_i > 0$ for all $i > 0$, since if $w_i = 0$ then index i is irrelevant for the rest of the algorithm. Thus, thanks to [Resample](#), we have $\tilde{w}_i \geq 1/n$ for all $i \in [n]$. Let a/b denote the entry-wise division between two vectors. As in the previous step of the proof, we let $K \triangleq \mathbf{k}_P(\mathcal{S}_n, \mathcal{S}_n)$, F be the low-rank factor output by [WeightedRPCholesky](#)($\mathbf{k}_P, \mathcal{S}_n, \tilde{w}, r$), and $\widehat{K} = FF^\top$. For any $w \in \Delta_{n-1}$, recall the notation $K^w \triangleq \text{diag}(\sqrt{w})K \text{diag}(\sqrt{w})$. Then we have

$$\begin{aligned} w^\top K w &= (w/\sqrt{\tilde{w}})^\top \text{diag}(K^{\tilde{w}})(w/\sqrt{\tilde{w}}) \\ &= (w/\sqrt{\tilde{w}})^\top (\text{diag}(\sqrt{\tilde{w}})\widehat{K} \text{diag}(\sqrt{\tilde{w}})) + \text{diag}(\sqrt{\tilde{w}})(K - \widehat{K}) \text{diag}(\sqrt{\tilde{w}}))(w/\sqrt{\tilde{w}}) \\ &= w^\top \widehat{K} w + (w/\sqrt{\tilde{w}})^\top (\text{diag}(\sqrt{\tilde{w}})(K - \widehat{K}) \text{diag}(\sqrt{\tilde{w}}))(w/\sqrt{\tilde{w}}) \\ &\leq w^\top \widehat{K} w + \max_{i \in [n]} (1/\tilde{w}_i) \text{tr}(\text{diag}(\sqrt{\tilde{w}})(K - \widehat{K}) \text{diag}(\sqrt{\tilde{w}})) \\ &\leq w^\top \widehat{K} w + n \text{tr}((K - \widehat{K})^{\tilde{w}}). \end{aligned} \tag{51}$$

Note that

$$K' = \widehat{K} + \text{diag}(K - \widehat{K}) = K + (\widehat{K} - K) + \text{diag}(K - \widehat{K}).$$

Since $K' \succeq \widehat{K}$, we have

$$w^{(q)\top} \widehat{K} w^{(q)} \leq w^{(q)\top} K' w^{(q)} \leq \tilde{w}^\top K' \tilde{w}, \quad (52)$$

where the last inequality follows from the if conditioning at the end of Alg. 2. In addition,

$$\begin{aligned} \tilde{w}^\top K' \tilde{w} &= \tilde{w}^\top (K + (\widehat{K} - K) + \text{diag}(K - \widehat{K})) \tilde{w} \\ &\leq \tilde{w}^\top K \tilde{w} + \tilde{w}^\top \text{diag}(K - \widehat{K}) \tilde{w} \\ &= \tilde{w}^\top K \tilde{w} + \sqrt{\tilde{w}}^\top \text{diag}((K - \widehat{K})^{\tilde{w}}) \sqrt{\tilde{w}} \\ &\leq w^\top K \tilde{w} + \text{tr}((K - \widehat{K})^{\tilde{w}}), \end{aligned}$$

where we used the fact that $K \succeq \widehat{K}$ and $\|\sqrt{\tilde{w}}\|_2 = 1$. Plugging the previous inequality into (52) and then into (51) with $w = w^{(q)}$, we get

$$w^{(q)\top} K w^{(q)} \leq \tilde{w}^\top K \tilde{w} + (n+1) \text{tr}((K - \widehat{K})^{\tilde{w}}). \quad (53)$$

Taking square-root on both sides using $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$ and the triangle inequality, we get

$$\begin{aligned} \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q)}}, \mathbb{P}) &\leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{\tilde{w}}, \mathbb{P}) + \sqrt{(n+1) \text{tr}((K - \widehat{K})^{\tilde{w}})} \\ &\leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{P}) + \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{S}_n^{\tilde{w}}) + \sqrt{(n+1) \text{tr}((K - \widehat{K})^{\tilde{w}})}. \end{aligned}$$

By Markov's inequality, we have

$$\begin{aligned} \Pr\left(\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{S}_n^{\tilde{w}}) > \sqrt{\frac{4Q\mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{S}_n^{\tilde{w}})]}{\delta}}\right) &\leq \frac{\delta}{4Q} \\ \Pr\left(\sqrt{\text{tr}((K - \widehat{K})^{\tilde{w}})} > \sqrt{\frac{4Q\mathbb{E}[\text{tr}((K - \widehat{K})^{\tilde{w}})]}{\delta}}\right) &\leq \frac{\delta}{4Q}. \end{aligned}$$

By Prop. E.1(c), we have

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{S}_n^{\tilde{w}})] = \mathbb{E}\left[\mathbb{E}[\text{MMD}_{\mathbf{k}_P}^2(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{S}_n^{\tilde{w}}) \mid w^{(q-1)}]\right] \leq \frac{\|\mathbf{k}_P\|_n}{n}.$$

Thus by the union bound, with probability at least $1 - \frac{\delta}{2Q}$, using Prop. F.1 (recall low-rank approximation \widehat{K} is obtained using \tilde{w}), we have

$$\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q)}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{(q-1)}}, \mathbb{P}) + \sqrt{\frac{4Q\|\mathbf{k}_P\|_n}{n\delta}} + \sqrt{\frac{4Q(n+1)H_{n,r}}{\delta}}. \quad (54)$$

Finally, applying union bound and summing up the bounds for $q = 1, \dots, Q$, we get, with probability at least $1 - \delta$,

$$\Delta \text{MMD}_{\mathbf{k}_P}(w^{(q)}) \leq \sqrt{\frac{2nH_{n,r}}{\delta}} + \sqrt{\frac{16\|\mathbf{k}_P\|_n \log n}{T^2}} + (Q-1) \left(\sqrt{\frac{4Q\|\mathbf{k}_P\|_n}{n\delta}} + \sqrt{\frac{4Q(n+1)H_{n,r}}{\delta}} \right).$$

This matches the stated asymptotic bound in Thm. 4. \square

F.2. Thinning with KT-Compress++

For compression with target distribution \mathbb{P} , we modify the original KT-Compress++ algorithm of (Shetty et al., 2022, Ex. 6): in HALVE and THIN of Compress++, we use KT-SPLIT with kernel \mathbf{k}_P without KT-SWAP (so our version of Compress++ outputs 2^q coresets, each of size \sqrt{n}), followed by KT-Swap to obtain a size \sqrt{n} coreset. We call the resulting thinning algorithm **KT-Compress++**. We show in Lem. F.4 and Cor. F.1 that **KT-Compress++** satisfies an MMD guarantee similar to that of quadratic-time kernel thinning.

Algorithm F.3 *KT-Compress++* (modified [Shetty et al. \(2022, Alg. 2\)](#) to minimize MMD to \mathbb{P})

Input: kernel $\mathbf{k}_{\mathbb{P}}$ with zero-mean under \mathbb{P} , input points $\mathcal{S}_n = (x_i)_{i \in [n]}$, multiplicity n' with $n' \in 4^{\mathbb{N}}$, weight $w \in \Delta_{n-1} \cap (\frac{\mathbb{N}_0}{n'})^n$, thinning parameter \mathfrak{g} , failure probability δ
 $\mathcal{S} \leftarrow$ index sequence where $k \in [n]$ appears $n'w_k$ times
 $(\mathbf{I}^{(\ell)})_{\ell \in [2^{\mathfrak{g}}]} \leftarrow \text{Compress++}(\mathfrak{g}, \mathcal{S}_n[\mathcal{S}]) \triangleright$ [Shetty et al. \(2022, Ex. 6\)](#) with *KT* substituted with *KT-SPLIT* in *HALVE* and *THIN*.
 $\mathbf{I}^{(\ell)} \leftarrow \mathcal{S}[\mathbf{I}^{(\ell)}]$ for each $\ell \in [2^{\mathfrak{g}}]$
 $\mathbf{I} \leftarrow \text{KT-Swap}(\mathbf{k}_{\mathbb{P}}, \mathcal{S}_n, (\mathbf{I}^{(\ell)})_{\ell \in [2^{\mathfrak{g}}]})$
 $w_{\text{C++}} \leftarrow$ simplex weights corresponding to $\mathbf{I} \triangleright w_i = \frac{\text{number of occurrences of } i \text{ in } \mathbf{I}}{|\mathbf{I}|}$
Return: $w_{\text{C++}} \in \Delta_{n-1} \cap (\frac{\mathbb{N}_0}{\sqrt{n}})^n \triangleright$ Hence $\|w_{\text{C++}}\|_0 \leq \sqrt{n}$

Lemma F.4 (Sub-gaussian guarantee for *Compress++*). *Let \mathcal{S}_n be a sequence of n points with $n \in 4^{\mathbb{N}}$. For any $\delta \in (0, 1)$ and integer $\mathfrak{g} \geq \lceil \log_2 \log(n+1) + 3.1 \rceil$, consider the *Compress++* algorithm ([Shetty et al., 2022, Algorithm 2](#)) with thinning parameter \mathfrak{g} , halving algorithm $\text{HALVE}^{(k)} \triangleq \text{symmetrized}^7(\text{KT-SPLIT}(\mathbf{k}, \cdot, 1, \frac{n_k^2}{4n2^{\mathfrak{g}}(\mathfrak{g}+(\beta_n+1)2^{\mathfrak{g}})}\delta))$ for an input of $n_k \triangleq 2^{\mathfrak{g}+1+k}\sqrt{n}$ points and $\beta_n \triangleq \log_2(\frac{n}{n_0})$, and with thinning algorithm $\text{THIN} \triangleq \text{KT-SPLIT}(\mathbf{k}, \cdot, \mathfrak{g}, \frac{\mathfrak{g}}{\mathfrak{g}+(\beta_n+1)2^{\mathfrak{g}}}\delta)$. Then this instantiation of *Compress++* compresses \mathcal{S}_n to $2^{\mathfrak{g}}$ coresets $(\mathcal{S}_{\text{out}}^{(i)})_{i \in [2^{\mathfrak{g}]}}$ of \sqrt{n} points each. Denote the signed measure $\phi^{(i)} \triangleq \frac{1}{n} \sum_{x \in \mathcal{S}_n} \delta_x - \frac{1}{\sqrt{n}} \sum_{x \in \mathcal{S}_{\text{out}}^{(i)}} \delta_x$. Then for each $i \in [2^{\mathfrak{g}}]$, on an event $\mathcal{E}_{\text{equi}}^{(i)}$ with $\Pr(\mathcal{E}_{\text{equi}}^{(i)}) \geq 1 - \frac{\delta}{2}$, $\phi^{(i)} = \tilde{\phi}^{(i)}$ for a random signed measure $\tilde{\phi}^{(i)}$ such that, for any $\delta' \in (0, 1)$,*

$$\Pr\left(\left\|\tilde{\phi}^{(i)}\mathbf{k}\right\|_{\mathcal{H}_{\mathbf{k}}} \geq a'_n \left(1 + \sqrt{\log\left(\frac{1}{\delta'}\right)}\right)\right) \leq \delta',$$

where

$$a'_n = \frac{4}{\sqrt{n}} \left(2 + \sqrt{\frac{8}{3}\|\mathbf{k}\|_n \log\left(\frac{6\sqrt{n}(\mathfrak{g} + (\frac{\log_2 n}{2} - \mathfrak{g})2^{\mathfrak{g}})}{\delta}\right) \log(4\mathcal{N}_{\mathbf{k}}(\mathcal{B}_2(R_n), n^{-1/2}))}\right).$$

Proof of Lem. F.4. This proof is similar to the one for [Shetty et al. \(2022, Ex. 6\)](#) but with explicit constant tracking and is self-contained, invoking only [Shetty et al. \(2022, Thm. 4\)](#) which gives MMD guarantees for *Compress++* given the sub-Gaussian parameters of *HALVE* and *THIN*.

Recall that n_k is the number of input points for the halving subroutine at recursion level k in *Compress++*, and β_n is the total number of recursion levels. Let \mathcal{S}_{C} denote the output of *COMPRESS* ([Shetty et al., 2022, Alg. 1](#)) of size $2^{\mathfrak{g}}\sqrt{n}$. Fix $\delta, \delta' \in (0, 1)$. Suppose we use $\text{HALVE}^{(k)} \triangleq \text{symmetrized}(\text{KT-SPLIT}(\mathbf{k}, \cdot, 1, \gamma_k\delta))$ for an input of n_k points for γ_k to be determined. Suppose we use $\text{THIN} \triangleq \text{KT-SPLIT}(\mathbf{k}, \cdot, \mathfrak{g}, \gamma'\delta)$ for γ' to be determined; this is the kernel thinning stage that thins $2^{\mathfrak{g}}\sqrt{n}$ points to $2^{\mathfrak{g}}$ coresets, each with \sqrt{n} points. Since the analysis is the same for all coresets, we will fix an arbitrary coreset without superscript in the notation.

By [Lem. D.1](#), with notation $t \triangleq \log \frac{1}{\delta'}$, there exist events $\mathcal{E}_{k,j}$, \mathcal{E}_{T} , and random signed measures $\phi_{k,j}$, $\tilde{\phi}_{k,j}$, ϕ_{T} , $\tilde{\phi}_{\text{T}}$ for $0 \leq k \leq \beta_n$ and $j \in [4^k]$ such that

- (a) $\Pr(\mathcal{E}_{k,j}^c) \leq \frac{\gamma_k\delta}{2}$ and $\Pr(\mathcal{E}_{\text{T}}^c) \leq \frac{\gamma'\delta}{2}$,
- (b) $\mathbb{1}_{\mathcal{E}_{k,j}}\phi_{k,j} = \mathbb{1}_{\mathcal{E}_{k,j}}\tilde{\phi}_{k,j}$ and $\mathbb{1}_{\mathcal{E}_{\text{T}}}\phi_{\text{T}} = \mathbb{1}_{\mathcal{E}_{\text{T}}}\tilde{\phi}_{\text{T}}$,
- (c) We have

$$\Pr\left(\left\|\tilde{\phi}_{k,j}\mathbf{k}\right\|_{\mathcal{H}_{\mathbf{k}}} \geq a_{n_k} + v_{n_k}\sqrt{t} \left|\left\{\tilde{\phi}_{k',j'}\right\}_{k'>k,j' \geq 1}, \left\{\tilde{\phi}_{k',j'}\right\}_{k',j'<j}\right)\right) \leq e^{-t}$$

$$\Pr\left(\left\|\tilde{\phi}_{\text{T}}\mathbf{k}\right\|_{\mathcal{H}_{\mathbf{k}}} \geq a'_n + v'_n\sqrt{t} \left|S_{\text{C}}\right.\right) \leq e^{-t},$$

⁷Any halving algorithm can be converted into an unbiased one by symmetrization, i.e., returning either the output half or its complement with equal probability ([Shetty et al., 2022, Remark 3](#)).

where, by Lem. D.1, and by increasing the sub-Gaussian constants if necessary, we have

$$a_{n_k} \triangleq v_{n_k} \triangleq a_{n_k, n_k/2} = \frac{2}{n_k} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{3n_k}{\gamma_k \delta}\right) \log\left(4\mathcal{N}_{\mathbf{k}}\left(\mathcal{B}_2(R_n), \frac{2}{n_k}\right)\right)} \right),$$

$$a'_n \triangleq v'_n \triangleq a_{2^g \sqrt{n}, \sqrt{n}} = \frac{1}{\sqrt{n}} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6g\sqrt{n}}{\gamma' \delta}\right) \log\left(4\mathcal{N}_{\mathbf{k}}\left(\mathcal{B}_2(R_n), n^{-1/2}\right)\right)} \right), \quad \text{and}$$

$$(d) \mathbb{E} \left[\tilde{\phi}_{k,j} \mathbf{k} \left| \{\tilde{\phi}_{k',j'}\}_{k'>k, j' \geq 1}, \{\tilde{\phi}_{k',j'}\}_{k', j' < j} \right. \right] = 0.$$

Hence on the event $\mathcal{E} = \bigcap_{k,j} \mathcal{E}_{k,j} \cap \mathcal{E}_T$, these properties hold simultaneously. We will choose $\{\gamma_k\}_k$ and γ' such that $\Pr(\mathcal{E}^c) \leq \frac{\delta}{2}$. By the union bound,

$$\Pr(\mathcal{E}^c) \leq \Pr(\mathcal{E}_T^c) + \sum_{k=0}^{\beta_n} \sum_{j=1}^{4^k} \Pr(\mathcal{E}_{k,j}^c) \leq \frac{\gamma' \delta}{2} + \sum_{k=0}^{\beta_n} 4^k \frac{\gamma_k \delta}{2}. \quad (55)$$

On the event \mathcal{E} , we apply Shetty et al. (2022, Thm. 4, Rmk. 7) to get a sub-Gaussian guarantee for $\text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}})$. We want to choose γ_k, γ' such that the rescaled quantities $\tilde{\zeta}_H \triangleq \frac{n_0}{2} a_{n_0}$ and $\tilde{\zeta}_T \triangleq \sqrt{n} a'_n$ satisfy $\tilde{\zeta}_H = \tilde{\zeta}_T$ (Shetty et al., 2022, Eq. (13)), which implies that we need

$$\frac{3n_0}{\gamma_0 \delta} = \frac{6g\sqrt{n}}{\gamma' \delta} \iff \frac{\gamma_0}{\gamma'} = \frac{2^g}{g}. \quad (56)$$

Hence if we take $\gamma' = \frac{g}{g + (\beta_n + 1)2^g}$ and $\gamma_k = \frac{n_k^2}{4n2^{2g}(g + (\beta_n + 1)2^g)}$, then (56) holds and the upper bound in (55) becomes $\frac{\delta}{2}$. Note that $n_k a_{n_k}$ is non-decreasing in n_k , so by Shetty et al. (2022, Theorem 4, Remark 7), Compress++(δ, \mathbf{g}) outputs a signed measure ϕ that, on the event \mathcal{E} with $\Pr(\mathcal{E}^c) \leq \frac{\delta}{2}$, equals another signed measure $\tilde{\phi}$ that satisfies, for any $\delta' \in (0, 1)$,

$$\Pr\left(\left\|\tilde{\phi}\mathbf{k}\right\|_{\mathcal{H}_{\mathbf{k}}} \geq \hat{a}_n + \hat{v}_n \sqrt{\log\left(\frac{1}{\delta'}\right)}\right) \leq \delta',$$

where \hat{a}_n, \hat{v}_n satisfy $\max(\hat{a}_n, \hat{v}_n) \leq 4a'_n$ whenever $g \geq \lceil \log_2 \log(n+1) + 3.1 \rceil$. \square

Corollary F.1 (MMD guarantee for Compress++). *Let \mathcal{S}_∞ be an infinite sequence of points in \mathbb{R}^d and \mathbf{k} a kernel. For any $\delta \in (0, 1)$ and $n \in \mathbb{N}$ such that $n \in 4^{\mathbb{N}}$, consider the Compress++ with the same parameters as in Lem. F.4 with $g \geq \lceil \log_2 \log(n+1) + 3.1 \rceil$. Then for any $i \in \lceil \sqrt{n} \rceil$, with probability at least $1 - \delta$,*

$$\text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}}^{(i)}) \leq \frac{4}{\sqrt{n}} \left(2 + \sqrt{\frac{8}{3} \|\mathbf{k}\|_n \log\left(\frac{6\sqrt{n}(g + (\frac{\log_2 n}{2} - g)2^g)}{\delta}\right) \log\left(4\mathcal{N}_{\mathbf{k}}\left(\mathcal{B}_2(R_n), n^{-1/2}\right)\right)} \right) \left(1 + \sqrt{\log\left(\frac{2}{\delta}\right)} \right).$$

Proof. After applying Lem. F.4 with $\delta' = \frac{\delta}{2}$ and following the same argument as in the proof of Cor. D.1, we have, with probability at least $1 - \delta$,

$$\text{MMD}_{\mathbf{k}}(\mathbb{S}_n, \mathbb{S}_{\text{out}}^{(i)}) \leq a'_n \left(1 + \sqrt{\log\left(\frac{2}{\delta}\right)} \right).$$

Plugging in the expression of a'_n from Lem. F.4 gives the claimed bound. \square

F.3. Proof of Thm. 5: MMD guarantee for LSKT

First of all, the claimed runtime follows from the runtime of LD (Thm. 4), the $O(d_{\mathbf{k}_P} 4^g n \log n)$ runtime of Compress++, and the $O(d_{\mathbf{k}_P} n^{1.5})$ runtime of KT-Swap.

Without loss of generality assume $n \in 4^{\mathbb{N}}$. Fix $\delta \in (0, 1)$. Let w^\diamond denote the output of LD, and w^{sr} denote the output of Resample, both regarded as random variables. By Thm. 4, we have, with probability at least $1 - \frac{\delta}{3}$,

$$\text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^\diamond}, \mathbb{P}) = \text{MMD}_{\mathbf{k}_P}(\mathbb{S}_n^{w^{\text{OPT}}}, \mathbb{P}) + O\left(\sqrt{\frac{nH_{n,r}}{\delta}}\right) + O\left(\sqrt{\frac{\|\mathbf{k}_P\|_n \max(\log n, 1/\delta)}{n}}\right). \quad (57)$$

By Prop. E.1(c) with $\mathbf{k} = \mathbf{k}_{\mathbb{P}}$, we have the upper bound

$$\mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^{w^{\diamond}})] = \mathbb{E} [\mathbb{E} [\text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^{w^{\diamond}}) | w^{\diamond}]] \leq \frac{\|\mathbf{k}_{\mathbb{P}}\|_n}{n}.$$

Thus, by Markov's inequality,

$$\Pr \left(\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^{w^{\diamond}}) \geq \sqrt{\frac{3\|\mathbf{k}_{\mathbb{P}}\|_n}{n\delta}} \right) \leq \frac{\delta}{3}.$$

Hence, with probability at least $1 - \frac{\delta}{3}$, we have

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^{w^{\diamond}}) \leq \sqrt{\frac{3\|\mathbf{k}_{\mathbb{P}}\|_n}{n\delta}}. \quad (58)$$

Let $\mathbb{S}_{\text{out}}^{(i)}$ denote the i -th coresot output by THIN in **KT-Compress++** (Alg. F.3). By Cor. F.1 with $\mathbf{k} = \mathbf{k}_{\mathbb{P}}$, we have, with probability at least $1 - \frac{\delta}{3}$,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_{\text{out}}^{(i)}) = O \left(\sqrt{\frac{\|\mathbf{k}_{\mathbb{P}}\|_n \log n \log(\epsilon \mathcal{N}_{\mathbf{k}_{\mathbb{P}}}(\mathcal{B}_2(R_n), n^{-1/2}))}{n}} \log \frac{\epsilon}{\delta} \right).$$

Since **KT-Swap** can never increase $\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\cdot, \mathbb{P})$, we have, by the triangle inequality,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{\text{wLSKT}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_{\text{out}}^{(1)}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_{\text{out}}^{(1)}, \mathbb{S}_n^{w^{\text{sr}}}) + \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\text{sr}}}, \mathbb{S}_n^{w^{\diamond}}) + \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\diamond}}, \mathbb{P}). \quad (59)$$

By the union bound, with probability at least $1 - \delta$, the bounds (57), (58), (59) hold, so that the claim is shown by adding together the right-hand sides of these bounds and applying Assum. (α, β) -kernel. \square

G. Simplex-Weighted Debiased Compression

In this section, we provide deferred analyses for **RT** and **SR/LSR**, as well as the algorithmic details of **Recombination** (Alg. G.1) and **KT-Swap-LS** (Alg. G.2).

G.1. MMD guarantee for RT

We start by stating the MMD guarantee of **RT**, a result that might be of independent interest.

Proposition G.1 (RT guarantee). *Under Assums. 1 and (α, β) -kernel, given $w \in \Delta_{n-1}$ and that $m \geq \left(\frac{\mathfrak{c}_d R_n^2 + 1}{\sqrt{\log 2}} + \sqrt{\log 2}\right)^2 - \frac{1}{\log 2} + 1$, **RecombinationThinning** (Alg. 4) outputs $w_{\text{RT}} \in \Delta_{n-1}$ with $\|w_{\text{RT}}\|_0 \leq m$ in $O((d_{\mathbf{k}_{\mathbb{P}}} + m)nm + m^3 \log n)$ time such that with probability at least $1 - \delta$,*

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w_{\text{RT}}}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^w, \mathbb{P}) + \sqrt{\frac{2\|\mathbf{k}_{\mathbb{P}}\|_n}{n\delta}} + \sqrt{\frac{2nH_{n,m-1}}{\delta}}, \quad (60)$$

where $H_{n,r}$ is defined in (46).

Proof of Prop. G.1. The runtime follows from the $O((d_{\mathbf{k}_{\mathbb{P}}} + m)nm)$ runtime of **WeightedRPCholesky**, the $O(d_{\mathbf{k}_{\mathbb{P}}}nm)$ runtime of **KT-Swap-LS**, and the $O(m^3 \log n)$ runtime of **Recombination** (Tchernychova, 2016) which dominates the $O(m^3)$ weight optimization step.

Recall $w' \in \Delta_{n-1}$ from **RT**. The formation of F in Alg. 4 is identical to the formation of F (with $r = m - 1$) in Alg. 2 for $q > 1$. Thus by (51) with $w = w'$, $K = \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)$,

$$w'^{\top} K w' \leq w'^{\top} F F^{\top} w' + n \text{tr}((K - F F^{\top})^{\tilde{w}}),$$

where $K = \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)$. By construction of w' using **Recombination**, we have $F^{\top} \tilde{w} = F^{\top} w'$. Since $K \succeq F F^{\top}$, we have

$$w'^{\top} K w' \leq \tilde{w}^{\top} F F^{\top} \tilde{w} + n \text{tr}((K - F F^{\top})^{\tilde{w}}) \leq \tilde{w}^{\top} K \tilde{w} + n \text{tr}((K - F F^{\top})^{\tilde{w}}).$$

We recognize the right-hand side is precisely the right-hand side of (53) aside from having a multiplier of n instead of $n + 1$ in front of the trace and that F is rank $m - 1$. Now applying (54) with $Q = \frac{1}{2}$, $w^{(q)} = w'$, $w^{(q-1)} = w$, $r = m - 1$, and noticing that **KT-Swap-LS** and the quadratic-programming solve at the end cannot decrease the objective, we obtain (60) with probability at least $1 - \delta$. Note that the lower bound of m in Assum. (α, β) -params makes $r = m - 1$ satisfy the lower bound for r in Prop. F.1. \square

Algorithm G.1 Recombination (rephrasing of [Tchernychova \(2016, Alg. 1\)](#) that takes $O(m^3 \log n)$ time)

Input: matrix $A \in \mathbb{R}^{m \times n}$ with $m < n$ and one row of A all positive, a nonnegative vector $x_0 \in \mathbb{R}_{\geq 0}^n$.

function FindBFS(A, x_0)

▷ The requirement of A and x_0 are the same as the input. This subroutine takes $O(n^3)$ time.

$x \leftarrow x_0$

$U, S, V^\top \leftarrow \text{SVD}(A)$ ▷ any $O(n^3)$ -time SVD algorithm that gives $USV^\top = A$

$V \leftarrow (V^\top)_{m+1:n}$ ▷ $V \in \mathbb{R}^{(n-m) \times n}$ so that the null space of A is spanned by the rows of V

for $i = 1$ **to** $n - m$ **do**

$v \leftarrow V_i$

$k \leftarrow \operatorname{argmin}_{j \in [n]: v_j > 0} \frac{x_j}{v_j}$ ▷ This must succeed because $Av = 0$ and A has an all-positive row, so one of the coordinates of v must be positive.

$x \leftarrow x - \frac{x_k v}{v_k}$ ▷ This zeros out the k -th coordinate of x while still ensuring x is nonnegative.

for $j = i + 1$ **to** $n - m$ **do**

$V_j \leftarrow V_j - \frac{V_{j,k} v}{v_k}$ ▷ $\{V_j\}_{j=i+1}^{n-m}$ remain independent and have 0 on the k -th coordinate.

end for

end for

return: $x \in \mathbb{R}_{\geq 0}^n$ such that $Ax = Ax_0$ and $\|x\|_0 \leq m$.

end function

$x \leftarrow x_0$

while $\|x\|_0 > 2m$ **do**

Divide $\{i \in [n] : x_i > 0\}$ into $2m$ index blocks I_1, \dots, I_{2m} , each of size at most $\left\lfloor \frac{\|x\|_0}{2m} \right\rfloor$.

$A_i \leftarrow A_{:,I_i}, x_{I_i} \in \mathbb{R}^m, \forall i \in [2m]$

Form \hat{A} to be the $m \times 2m$ matrix with columns A_i ▷ Hence, one row of A contains all positive entries.

$\hat{x} \leftarrow \text{FindBFS}(\hat{A}, \mathbf{1}_{2m})$ ▷ $\|\hat{x}\|_0 \leq n$ and $\hat{A}\hat{x} = \sum A_i \hat{x}_i = \sum A_i = \sum A_{:,I_i} x_{I_i} = Ax$.

for $i = 1$ **to** $2m$ **do**

$x_{I_i} \leftarrow \hat{x}_i \cdot x_{I_i}$ **if** $\hat{x}_i > 0$ **else** 0

end for

▷ After the update, the support of x shrinks by 2 while it maintains that $Ax = Ax_0$.

end while

if $\|x\|_0 \geq m + 1$ **then**

$I \leftarrow \{i \in [n] : x_i > 0\}$

$x_I = \text{FindBFS}(A_{:,I}, x_I)$

end if

Return: $x \in \mathbb{R}_{\geq 0}^n$ such that $Ax = Ax_0$ and $\|x\|_0 \leq m$.

G.2. Proof of Thm. 6: MMD guarantee for SR/LSR

The claimed runtime follows from the runtime of [SteinThinning](#) (Alg. D.1) or [LD](#) (Thm. 4) plus the runtime of [RT](#) (Prop. G.1).

Note the lower bound for m in Assum. (α, β) -params implies the lower bound condition in Prop. G.1. For the case of [SR](#), we proceed as in the proof of Thm. 3 and use Prop. G.1. For the case of [LSR](#), we proceed as in the proof of Thm. 5 and use Thm. 4 and Prop. G.1. \square

H. Constant-Preserving Debiased Compression

In this section, we provide deferred analyses for [CT](#) and [SC/LSC](#).

H.1. MMD guarantee for CT

Proposition H.1 (CT guarantee). *Under Assums. 1 and (α, β) -kernel, given $w \in \Delta_{n-1}$ and $m \geq \left(\frac{c_d R_n^2 + 1}{\sqrt{\log 2}} + \frac{2}{\sqrt{\log 2}}\right)^2 - \frac{1}{\log 2}$, CT outputs $w_{\text{CT}} \in \mathbb{R}^n$ with $\mathbf{1}_n^\top w_{\text{CT}} = 1$ and $\|w_{\text{CT}}\|_0 \leq m$ in $O((d_{\text{kp}} + m)nm + m^3)$ time such that, for any*

Algorithm G.2 KT-Swap with Linear Search (KT-Swap-LS)

Input: kernel $\mathbf{k}_{\mathbb{P}}$ with zero-mean under \mathbb{P} , input points $\mathcal{S}_n = (x_i)_{i \in [n]}$, weights $w \in \Delta_{n-1}$, $\text{fmt} \in \{\text{SPLX}, \text{CP}\}$
 $\mathbf{S} \leftarrow \{i \in [n] : w_i \neq 0\}$
 \triangleright Maintain two sufficient statistics: $g = Kw$ and $D = w^\top Kw$.

function Add(g, D, i, t)
 $g \leftarrow g + t\mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, x_i)$
 $D \leftarrow D + 2tg_i + t^2\mathbf{k}_{\mathbb{P}}(x_i, x_i)$
return: (g, D)

end function

function Scale(g, D, α)
 $g \leftarrow \alpha g$
 $D \leftarrow \alpha^2 D$
return: (g, D)

end function

$\text{Kdiag} \leftarrow \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)$
 $g \leftarrow \mathbf{0} \in \mathbb{R}^n$
 $D \leftarrow 0$

for i **in** \mathbf{S} **do**
 $(g, D) \leftarrow \text{Add}(g, D, i, w_i)$

end for

for i **in** \mathbf{S} **do**
if $w_i = 1$ **then continue;** \triangleright We cannot swap i out if $\sum_{j \neq i} w_j = 0!$
 \triangleright First zero out w_i .
 $(g, D) \leftarrow \text{Add}(g, D, i, -w_i)$
 $(g, D) \leftarrow \text{Scale}(g, D, \frac{1}{1-w_i})$
 $w_i = 0$
 \triangleright Next perform line search to add back a point.
 $\alpha = (D - g) ./ (D - 2g + \text{Kdiag}); \triangleright \alpha_i = \text{argmin}_t \text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{te_i + (1-t)w}, \mathbb{P}) = \text{argmin}_t (1-t)^2 D + 2t(1-t)g + t^2 \text{Kdiag}$
if $\text{fmt} = \text{SPLX}$ **then**
 $\alpha = \text{clip}(\alpha, 0, 1); \triangleright$ Clipping α to $[0, 1]$. This corresponds to $\text{argmin}_{t \in [0, 1]} \text{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{te_i + (1-t)w}, \mathbb{P})$.
end if
 $D' \leftarrow (1 - \alpha)^2 D + 2\alpha(1 - \alpha)g + \alpha^2 \text{Kdiag} \triangleright$ multiplications are element-wise
 $k \leftarrow \text{argmin}_i D'_i$
 $(g, D) \leftarrow \text{Scale}(g, D, 1 - \alpha_k)$
 $(g, D) \leftarrow \text{Add}(g, D, k, \alpha_k)$

end for

Return: $w \in \Delta_{n-1}$

$\delta \in (0, 1)$, with probability $1 - \delta$,

$$\text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w_{\text{CT}}}, \mathbb{P}) \leq 2 \text{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^w, \mathbb{P}) + \sqrt{\frac{4H_{n,m'}}{\delta}},$$

where $H_{n,m}$ is defined in (46) and $m' \triangleq m + \log 2 - 2\sqrt{m \log 2 + 1}$.

Proof of Prop. H.1. The runtime follows from the $O((d_{\mathbf{k}_{\mathbb{P}}} + m)nm)$ runtime of **WeightedRPCholesky**, the $O(nm)$ runtime of **KT-Swap-LS**, and the $O(m^3)$ runtime of matrix inversion in solving the two minimization problems using (64).

To improve the clarity of notation, we will use w^\diamond to denote the input weight w to **CT**. For index sequences $\mathbf{I}, \mathbf{J} \subset [n]$ and a kernel \mathbf{k} , we use $\mathbf{k}(\mathbf{I}, \mathbf{J})$ to indicate the matrix $\mathbf{k}(\mathcal{S}_n[\mathbf{I}], \mathcal{S}_n[\mathbf{J}]) = [\mathbf{k}(x_i, x_j)]_{i \in \mathbf{I}, j \in \mathbf{J}}$, and similarly for a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we use $f(\mathbf{I})$ to denote the vector $(f(x_i))_{i \in \mathbf{I}}$.

Recall the regularized kernel is $\mathbf{k}_c \triangleq \mathbf{k}_{\mathbb{P}} + c$. Suppose for now that $c > 0$ is an arbitrary constant. Let \mathbf{I} denote the indices

output by [WeightedRPCholesky](#) in CT. Let

$$w^c \triangleq \operatorname{argmin}_{w: \operatorname{supp}(w) \subset \mathbf{I}} \operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^w, \mathbb{S}_n^{w^\diamond}).$$

Note that w^c is not a probability vector and may not sum to 1.

Step 1. Bound $\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond})$ in terms of [WeightedRPCholesky](#) approximation error

We start by using an argument similar to that of [Epperly and Moreno \(2024, Prop. 3\)](#) to exploit the optimality condition of w^c . Since

$$\operatorname{argmin}_{w: \operatorname{supp}(w) \subset \mathbf{I}} \operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^w, \mathbb{S}_n^{w^\diamond}) = \operatorname{argmin}_{w: \operatorname{supp}(w) \subset \mathbf{I}} w_{\mathbf{I}}^\top \mathbf{k}_c(\mathbf{I}, \mathbf{I}) w_{\mathbf{I}} - 2w^{\diamond\top} \mathbf{k}_c(\mathcal{S}_n, \mathbf{I}) w_{\mathbf{I}},$$

by optimality, w^c satisfies,

$$\mathbf{k}_c(\mathbf{I}, \mathbf{I}) w_{\mathbf{I}}^c = \mathbb{S}_n^{w^\diamond} \mathbf{k}_c(\mathbf{I}).$$

We comment that the index sequence \mathbf{I} returned by [WeightedRPCholesky](#) makes $\mathbf{k}_c(\mathbf{I}, \mathbf{I})$ invertible with probability 1: by the Guttman rank additivity formula of Schur complement ([Zhang, 2006, Eq. \(6.0.4\)](#)), each iteration of [WeightedRPCholesky](#) chooses a pivot with a non-zero diagonal and thus increases the rank of the low-rank approximation matrix, which is spanned by the columns of pivots, by 1. Hence

$$\begin{aligned} \mathbb{S}_n^{w^c} \mathbf{k}_c(\cdot) &= \mathbf{k}_c(\cdot, \mathcal{S}_n) w^c = \mathbf{k}_c(\cdot, \mathbf{I}) w_{\mathbf{I}}^c = \mathbf{k}_c(\cdot, \mathbf{I}) \mathbf{k}_c(\mathbf{I}, \mathbf{I})^{-1} \mathbf{k}_c(\mathbf{I}, \mathbf{I}) w_{\mathbf{I}}^c \\ &= \mathbf{k}_c(\cdot, \mathbf{I}) \mathbf{k}_c(\mathbf{I}, \mathbf{I})^{-1} \mathbb{S}_n^{w^\diamond} \mathbf{k}_c(\mathbf{I}) = \mathbb{S}_n^{w^\diamond} \mathbf{k}_{c\mathbf{I}}(\cdot), \end{aligned}$$

where $\mathbf{k}_{c\mathbf{I}}(x, y) \triangleq \mathbf{k}_c(x, \mathbf{I}) \mathbf{k}_c(\mathbf{I}, \mathbf{I})^{-1} \mathbf{k}_c(\mathbf{I}, y)$. Then

$$\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond}) = \|\mathbb{S}_n^{w^\diamond} \mathbf{k}_c - \mathbb{S}_n^{w^c} \mathbf{k}_c\|_{\mathbf{k}_c}^2 = \|\mathbb{S}_n^{w^\diamond} \mathbf{k}_c - \mathbb{S}_n^{w^\diamond} \mathbf{k}_{c\mathbf{I}}\|_{\mathbf{k}_c}^2 = w^{\diamond\top} (\mathbf{k}_c - \mathbf{k}_{c\mathbf{I}})(\mathcal{S}_n, \mathcal{S}_n) w^\diamond.$$

Recall the index set \mathbf{I} consists of the m pivots selected by [WeightedRPCholesky](#) on the input matrix

$$K_c^\diamond \triangleq \mathbf{k}_c(\mathcal{S}_n, \mathcal{S}_n) w^\diamond.$$

Define

$$\widehat{K}_c^\diamond \triangleq \mathbf{k}_{c\mathbf{I}}(\mathcal{S}_n, \mathcal{S}_n) w^\diamond.$$

Thus, by [Lem. F.1](#),

$$\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond}) = w^{\diamond\top} (\mathbf{k}_c - \mathbf{k}_{c\mathbf{I}})(\mathcal{S}_n, \mathcal{S}_n) w^\diamond = \sqrt{w^{\diamond\top}} (K_c^\diamond - \widehat{K}_c^\diamond) \sqrt{w^\diamond} \leq \lambda_1(K_c^\diamond - \widehat{K}_c^\diamond) \leq \operatorname{tr}(K_c^\diamond - \widehat{K}_c^\diamond).$$

Step 2. Bound $\operatorname{tr}(K_c^\diamond - \widehat{K}_c^\diamond)$ using the trace bound of the unregularized kernel

Let $\llbracket A \rrbracket_r$ denote the best rank- r approximation of an SPSD matrix $A \in \mathbb{R}^{n \times n}$ in the sense that

$$\llbracket A \rrbracket_r \triangleq \operatorname{argmin}_{\substack{X \in \mathbb{R}^{n \times n} \\ X = X^\top \\ A \succeq X \succeq 0 \\ \operatorname{rank}(X) \leq r}} \operatorname{tr}(A - X). \quad (61)$$

By the Eckart-Young-Mirsky theorem applied to symmetric matrices ([Dax, 2014, Theorem 19](#)), the solution to (61) is given by r -truncated eigenvalue decomposition of A , so that

$$\operatorname{tr}(A - \llbracket A \rrbracket_r) = \sum_{\ell=r+1}^n \lambda_\ell(A).$$

Let $q \triangleq \mathfrak{U}(m)$ where \mathfrak{U} is defined in (47), so that by [Chen et al. \(2022, Thm. 3.1\)](#) with $\epsilon = 1$, we have

$$\mathbb{E} \left[\operatorname{tr}(K_c^\diamond - \widehat{K}_c^\diamond) \right] \leq 2 \operatorname{tr}(K_c^\diamond - \llbracket K_c^\diamond \rrbracket_q).$$

We know one specific rank- q approximation of K_c^\diamond :

$$\tilde{K}_c^\diamond \triangleq \llbracket K^\diamond \rrbracket_{q-1} + \text{diag}(\sqrt{w^\diamond}) c \mathbf{1}_n \mathbf{1}_n^\top \text{diag}(\sqrt{w^\diamond}),$$

which satisfies

$$K_c^\diamond - \tilde{K}_c^\diamond = K^\diamond + \text{diag}(\sqrt{w^\diamond}) c \mathbf{1}_n \mathbf{1}_n^\top \text{diag}(\sqrt{w^\diamond}) - \tilde{K}_c^\diamond = K^\diamond - \llbracket K^\diamond \rrbracket_{q-1}.$$

Thus by the variational definition in (61), we have

$$\text{tr}(K_c^\diamond - \llbracket K_c^\diamond \rrbracket_q) \leq \text{tr}(K_c^\diamond - \tilde{K}_c^\diamond) = \text{tr}(K^\diamond - \llbracket K^\diamond \rrbracket_{q-1}) = \sum_{\ell=q}^n \lambda_\ell(K^\diamond).$$

Note the last bound does not depend on c . The tail sum of eigenvalues in the last expression is the same (up to a constant multiplier) as the one in (48) except for an off-by-1 difference in the summation index. A simple calculation shows that for $m' \triangleq m + \log 2 - 2\sqrt{m \log 2} + 1$, we have $\mathfrak{U}(m') = \mathfrak{U}(m) - 1$. Another simple calculation shows that $m \geq (\frac{\mathfrak{C}_d R_n^2 + 1}{\sqrt{\log 2}} + \frac{2}{\sqrt{\log 2}})^2 - \frac{1}{\log 2}$ implies that m' satisfies the lower bound requirement of r in Prop. F.1. Thus, arguing as in the proof that follows (48), we get

$$\mathbb{E} \left[\text{tr}(K_c^\diamond - \hat{K}_c^\diamond) \right] \leq H_{n,m'}.$$

Thus so far we have shown

$$\mathbb{E}[\text{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond})] \leq \mathbb{E} \left[\text{tr}(K_c^\diamond - \hat{K}_c^\diamond) \right] \leq H_{n,m'}.$$

By Markov's inequality, with probability at least $1 - \delta$, we have

$$\text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond}) \leq \sqrt{\frac{H_{n,m'}}{\delta}}.$$

Recall that $\text{MMD}_{\mathbf{k}}(\mu, \nu) = \|(\mu - \nu)\mathbf{k}\|_{\mathbf{k}}$ for signed measures μ, ν . By the triangle inequality, we have

$$\begin{aligned} \text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{P}) &\leq \text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond}) + \text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^\diamond}, \mathbb{P}) \\ &= \text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{S}_n^{w^\diamond}) + \text{MMD}_{\mathbf{k}_\mathbb{P}}(\mathbb{S}_n^{w^\diamond}, \mathbb{P}), \end{aligned}$$

where we used that fact that $\sum_{i \in [n]} w_i^\diamond = 1$ to get the identity $\text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^\diamond}, \mathbb{P}) = \text{MMD}_{\mathbf{k}_\mathbb{P}}(\mathbb{S}_n^{w^\diamond}, \mathbb{P})$. Hence, with probability at least $1 - \delta$,

$$\text{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{P}) \leq \text{MMD}_{\mathbf{k}_\mathbb{P}}(\mathbb{S}_n^{w^\diamond}, \mathbb{P}) + \sqrt{\frac{H_{n,m'}}{\delta}}. \quad (62)$$

Step 3. Incorporating sum-to-one constraint

We now turn w^c into a constant-preserving weight while not inflating the MMD by much. Define

$$w^1 \triangleq \underset{w: \text{supp}(w) \subset \mathbb{I}, \sum_{i \in [n]} w_i = 1}{\text{argmin}} \text{MMD}_{\mathbf{k}_\mathbb{P}}^2(\mathbb{S}_n^w, \mathbb{P}). \quad (63)$$

Note w^1 is the weight right before **KT-Swap-LS** step in **CT**. Let $K_{\mathbb{I}} = \mathbf{k}_\mathbb{P}(\mathbb{I}, \mathbb{I})$. Let $\mathbf{1}_{\mathbb{I}}$ denote the $|\mathbb{I}|$ -dimensional all-one vector. The Karush-Kuhn-Tucker condition (Ghojogh et al., 2021, Sec. 4.7) applied to (63) implies that, the solution w^1 is a stationary point of the Lagrangian function

$$L(w_{\mathbb{I}}, \lambda) \triangleq w_{\mathbb{I}}^\top K_{\mathbb{I}} w_{\mathbb{I}} + \lambda (\mathbf{1}_{\mathbb{I}}^\top w_{\mathbb{I}} - 1).$$

Then $\nabla_{w_{\mathbb{I}}} L(w_{\mathbb{I}}^1, \lambda) = 0$ implies $2K_{\mathbb{I}} w_{\mathbb{I}}^1 - \lambda \mathbf{1}_{\mathbb{I}} = 0$, so $w_{\mathbb{I}}^1 = \frac{\lambda K_{\mathbb{I}}^{-1} \mathbf{1}_{\mathbb{I}}}{2}$. The Lagrangian multiplier λ is determined by the constraint $\mathbf{1}_{\mathbb{I}}^\top w_{\mathbb{I}}^1 = 1$, so we find

$$w_{\mathbb{I}}^1 = \frac{K_{\mathbb{I}}^{-1} \mathbf{1}_{\mathbb{I}}}{\mathbf{1}_{\mathbb{I}}^\top K_{\mathbb{I}}^{-1} \mathbf{1}_{\mathbb{I}}}. \quad (64)$$

Define

$$w^{c,\mathbb{P}} \triangleq \operatorname{argmin}_{w:\operatorname{supp}(w)\subset\mathbb{I}} \operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^w, \mathbb{P}).$$

Since $w^{c,\mathbb{P}}$ is optimized to minimize $\operatorname{MMD}_{\mathbf{k}_c}$ to \mathbb{P} on the same support as w^c , we have

$$\operatorname{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}) \leq \operatorname{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{P}).$$

The optimality condition for $w^{c,\mathbb{P}}$ is

$$(K_{\mathbb{I}} + c\mathbf{1}_{\mathbb{I}}\mathbf{1}_{\mathbb{I}}^{\top})w - c\mathbf{1}_{\mathbb{I}} = 0,$$

and hence by the Sherman–Morrison formula,

$$w_{\mathbb{I}}^{c,\mathbb{P}} = (K_{\mathbb{I}} + c\mathbf{1}_{\mathbb{I}}\mathbf{1}_{\mathbb{I}}^{\top})^{-1}c\mathbf{1}_{\mathbb{I}} = \left(K_{\mathbb{I}}^{-1} - \frac{cK_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}}{1+c\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}}\right)c\mathbf{1}_{\mathbb{I}} = \frac{K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}}{1/c+\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}}.$$

Let $\rho_c \triangleq \frac{\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}}{1/c+\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}}$, so that $w_{\mathbb{I}}^{c,\mathbb{P}} = \rho_c w_{\mathbb{I}}^1$. In particular, w^1 and $w^{c,\mathbb{P}}$ are scalar multiples of one another. To relate $\operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^1}, \mathbb{P})$ and $\operatorname{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P})$, note that

$$\begin{aligned} \operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^1}, \mathbb{P}) &= w_{\mathbb{I}}^{\top} K_{\mathbb{I}} w_{\mathbb{I}}^1 = \frac{w_{\mathbb{I}}^{c,\mathbb{P}\top} K_{\mathbb{I}} w_{\mathbb{I}}^{c,\mathbb{P}}}{\rho_c^2} = \frac{w_{\mathbb{I}}^{c,\mathbb{P}\top} (K_{\mathbb{I}} + c\mathbf{1}_{\mathbb{I}}\mathbf{1}_{\mathbb{I}}^{\top}) w_{\mathbb{I}}^{c,\mathbb{P}} - c(\mathbf{1}_{\mathbb{I}}^{\top} w_{\mathbb{I}}^c)^2}{\rho_c^2} \\ &= \frac{\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}) + 2c\mathbf{1}_{\mathbb{I}}^{\top} w_{\mathbb{I}}^c - c - c(\mathbf{1}_{\mathbb{I}}^{\top} w_{\mathbb{I}}^c)^2}{\rho_c^2} = \frac{\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}) - c(\rho_c - 1)^2}{\rho_c^2}. \end{aligned}$$

So far the argument does not depend on any particular choice of $c > 0$. Let us now discuss how to choose c . Note that

$$\mathbf{1}_{\mathbb{I}}^{\top} K_{\mathbb{I}}^{-1} \mathbf{1}_{\mathbb{I}} = m \frac{\mathbf{1}_{\mathbb{I}}}{\sqrt{m}}^{\top} K_{\mathbb{I}}^{-1} \frac{\mathbf{1}_{\mathbb{I}}}{\sqrt{m}} \geq m \lambda_m(K_{\mathbb{I}}^{-1}) \geq \frac{m}{\lambda_1(K_{\mathbb{I}})} \geq \frac{m}{\operatorname{tr}(K_{\mathbb{I}})} \geq \frac{m}{\sum_{i \in [m]} \operatorname{diag}(K)_{\downarrow}^i},$$

where $\operatorname{diag}(K)_{\downarrow}$ denote the diagonal entries of $K = \mathbf{k}_{\mathbb{P}}(\mathcal{S}_n, \mathcal{S}_n)$ sorted in descending order. Thus

$$\rho_c = \frac{1}{\frac{1}{c\mathbf{1}_{\mathbb{I}}^{\top}K_{\mathbb{I}}^{-1}\mathbf{1}_{\mathbb{I}}} + 1} \geq \frac{1}{\frac{\sum_{i \in [m]} \operatorname{diag}(K)_{\downarrow}^i}{mc} + 1}.$$

Hence we can choose c to make sure ρ_c is bounded from below by a positive value. Recall in [CT](#), we take

$$c = \frac{\sum_{i \in [m]} \operatorname{diag}(K)_{\downarrow}^i}{m},$$

so that $\rho_c \geq \frac{1}{2}$ and

$$\operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}^2(\mathbb{S}_n^{w^1}, \mathbb{P}) = \frac{\operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}) - c(\rho_c - 1)^2}{\rho_c^2} \leq 4 \operatorname{MMD}_{\mathbf{k}_c}^2(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}).$$

Hence by [\(62\)](#) and the fact that [KT-Swap-LS](#) and the final reweighting in [CT](#) only improves MMD, we have, with probability at least $1 - \delta$,

$$\operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^{\text{CT}}}, \mathbb{P}) \leq \operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{w^1}, \mathbb{P}) \leq 2 \operatorname{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^{c,\mathbb{P}}}, \mathbb{P}) \leq 2 \operatorname{MMD}_{\mathbf{k}_c}(\mathbb{S}_n^{w^c}, \mathbb{P}) \leq 2 \operatorname{MMD}_{\mathbf{k}_{\mathbb{P}}}(\mathbb{S}_n^{\circ}, \mathbb{P}) + 2\sqrt{\frac{H_{n,m'}}{\delta}},$$

where we use [\(62\)](#) in the last inequality. \square

H.2. Proof of [Thm. 7](#): MMD guarantee for [SC](#) / [LSC](#)

The claimed runtime follows from the runtime of [SteinThinning](#) ([Alg. D.1](#)) or [LD](#) ([Thm. 4](#)) plus the runtime of [CT](#) ([Prop. H.1](#)).

Note the lower bound for m in [Assum. \$\(\alpha, \beta\)\$ -params](#) implies the lower bound condition in [Prop. H.1](#). For the case of [SC](#), we proceed as in the proof of [Thm. 3](#) and use [Prop. H.1](#). For the case of [LSC](#), we proceed as in the proof of [Thm. 5](#) by invoking [Thm. 4](#) and [Prop. H.1](#). \square

I. Implementation and Experimental Details

In this section, we collect experimental details that were deferred from Sec. 5.

I.1. $O(d)$ -time Stein kernel evaluation

In this section, we show that for $\mathcal{S}_n = (x_i)_{i \in [n]}$, each Stein kernel evaluation $\mathbf{k}_p(x_i, x_j)$ for a radially analytic base kernel (Def. B.3) can be done in $O(d)$ time after computing certain sufficient statistics in $O(nd^2 + d^3)$ time. Let $M \in \mathbb{R}^{d \times d}$ be a positive definite preconditioning matrix for \mathbf{k}_p . Let L be the Cholesky decomposition of M which can be done in $O(d^3)$ time so that $M = LL^\top$. From the expression (15), we can achieve $O(d)$ time evaluation if we can compute $\|x - y\|_M^2$ and $M\nabla \log p(x)$ in $O(d)$ time. For $M\nabla \log p(x)$, we can simply precompute $M\nabla \log p(x_i)$ for all $i \in [n]$. For $\|x - y\|_M^2$, we have

$$\|x - y\|_M^2 = (x - y)^\top M^{-1}(x - y) = (x - y)^\top (LL^\top)^{-1}(x - y) = \|L^{-1}(x - y)\|_2^2.$$

Hence it suffices to precompute $L^{-1}x_i$ for all $i \in [n]$, and we can precompute the inverse L^{-1} in $O(d^3)$ time.

I.2. Default parameters for algorithms

For LD, we always use $Q = 3$. To ensure that the guarantees of Lem. F.3 and Thm. 4 hold while achieving fast convergence in practice, we take the step size of AMD to be $1/(8\|\mathbf{k}_\mathbb{P}\|_n)$ in the first adaptive round and $1/(8\sum_{i \in [n]} w_i^{(q-1)} \mathbf{k}_\mathbb{P}(x_i, x_i))$ in subsequent adaptive rounds. We use $T = 7\sqrt{n_0}$ for AMD in all experiments.

We implemented our modified versions of KernelThinning and KT-Compress++ in JAX (Bradbury et al., 2018) so that certain subroutines can achieve a speedup using just-in-time compilation and the parallel computation power of GPUs. For Compress++, we use $g = 4$ in all experiments as in Shetty et al. (2022). For both KernelThinning and KT-Compress++, we use choose $\delta = 1/2$ as in the goodpoints library.

Each experiment was run with a single NVIDIA RTX 6000 GPU and an AMD EPYC 7513 32-Core CPU.

I.3. Correcting for burn-in details

We use the four MCMC chains provided by Riabiz et al. (2022) that include both the sample points and their scores. The reference chain used to compute the energy distance is the same one used in Riabiz et al. (2022) for the energy distance and was kindly provided by the authors.

In Tab. I.1, we collect the runtime for the burn-in correction experiments.

Fig. I.1, Fig. I.2, Fig. I.3, display the results of the burn-in correction experiment of Sec. 5 repeated with three other MCMC algorithms: MALA without preconditioning, random walk (RW), and adaptive random walk (ADA-RW). The results of P-MALA from Sec. 5 are also included for completeness. For all four chains, our methods reliably achieve better quality coresets when compared with the baseline methods.

n_0	ST	LD (0.5)	LD (0.4)	KT	KT-Compress++	RT (0.5)	RT (0.4)	CT (0.5)	CT (0.4)
2^{14}	2.50	13.22	12.88	7.31	3.49	0.79	0.60	2.06	1.96
2^{16}	8.48	16.15	15.82	20.77	5.90	2.59	1.68	3.66	3.04
2^{18}	111.06	32.14	20.60	193.03	11.73	11.16	2.63	6.48	3.67
2^{20}	-	314.67	131.31	-	35.99	113.71	11.06	51.14	8.42

Table I.1: **Breakdown of runtime (in seconds)** for the burn-in correction experiment ($d = 4$) of Sec. 5. n_0 is the input size after standard thinning from the length $n = 2 \times 10^6$ chain (Rem. 2). Each runtime is the median of 3 runs. KT and KT-Compress++ output $m = \sqrt{n_0}$ equal-weighted points. RT and CT respectively output $m = n_0^\tau$ points with simplex or constant-preserving weights for τ shown in parentheses. In addition, LD, RT, and CT use the rank n_0^τ . ST and KT took longer than 30 minutes for $n_0 = 2^{20}$ and hence their numbers are not reported.

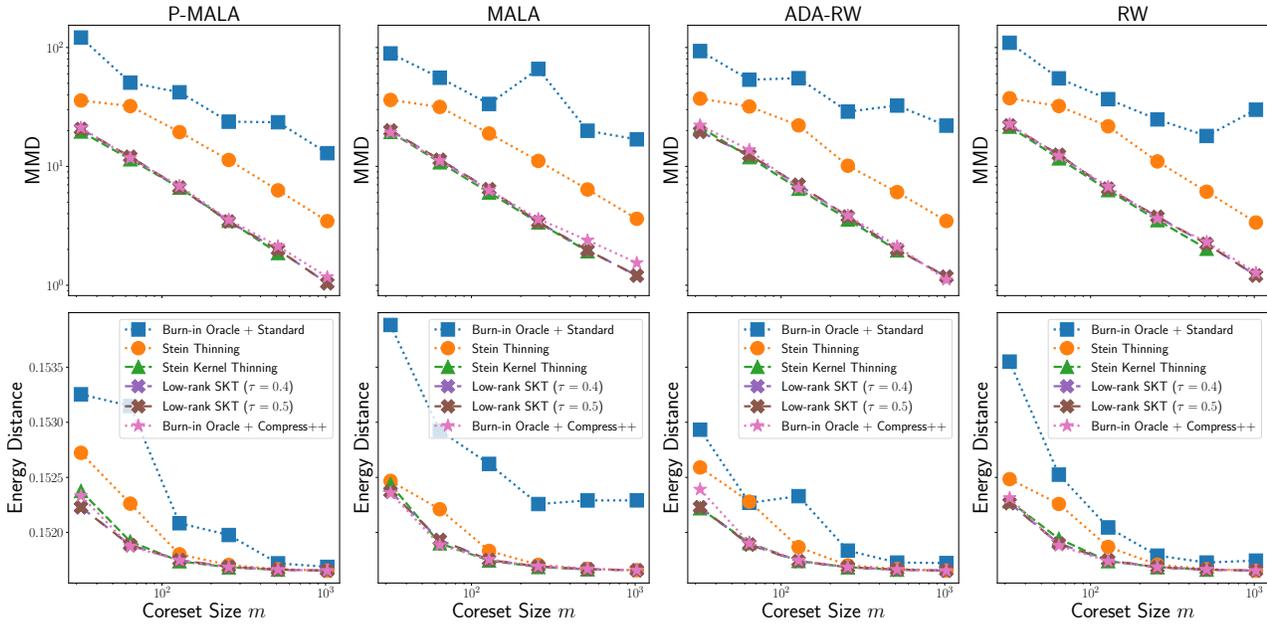


Figure I.1: **Correcting for burn-in with equal-weighted compression.** For each of four MCMC algorithms and using only one chain, our methods consistently outperform the Stein and standard thinning baselines and match the 6-chain oracle.

I.4. Correcting for approximate MCMC details

Surrogate ground truth Following Liu and Lee (2017), we took the first 10,000 data points and generated 2^{20} surrogate ground truth sample points using NUTS (Hoffman and Gelman, 2014) for the evaluation. To generate the surrogate ground truth using NUTS, we used `numpyro` (Phan et al., 2019). It took 12 hours to generate the surrogate ground truth points using the GPU implementation, and we estimate it would have taken 200 hours using the CPU implementation.

SGFS For SGFS, we used batch size 32 and the step size schedule $\eta/(1+t)^{0.55}$ where t is the step count and η is the initial step size. We chose η from $\{10.0, 5.0, 1.0, 0.5, 0.1, 0.05, 0.01\}$, found $\eta = 1.0$ gave the best standard thinning MMD to get a coreset size of $m = 2^{10}$, and hence we fixed $\eta = 1.0$ in all experiments. We used the version of SGFS (Ahn et al., 2012, SGFS-f) that involves inversion of $d \times d$ matrices — we found the faster version (SGFS-d) that inverts only the diagonal resulted in significantly worse mixing. We implemented SGFS in `numpy` and ran it on the CPU.

Runtime The SGFS chain of length 2^{24} took approximately 2 hours to generate using the CPU. Remarkably, all of our low-rank methods finish within 10 minutes for $n_0 = 2^{20}$, which is orders of magnitude faster than the time taken to generate the NUTS surrogate ground truth.

Additional results In Fig. I.4, we plot the posterior mean mean-squared error (MSE) for each compression method in the approximate MCMC experiment of Sec. 5.

I.5. Correcting for tempering details

In the data release of Riabiz et al. (2020), we noticed there were 349 sample points for which the provided scores were NaNs, so we removed those points at the recommendation of the authors.

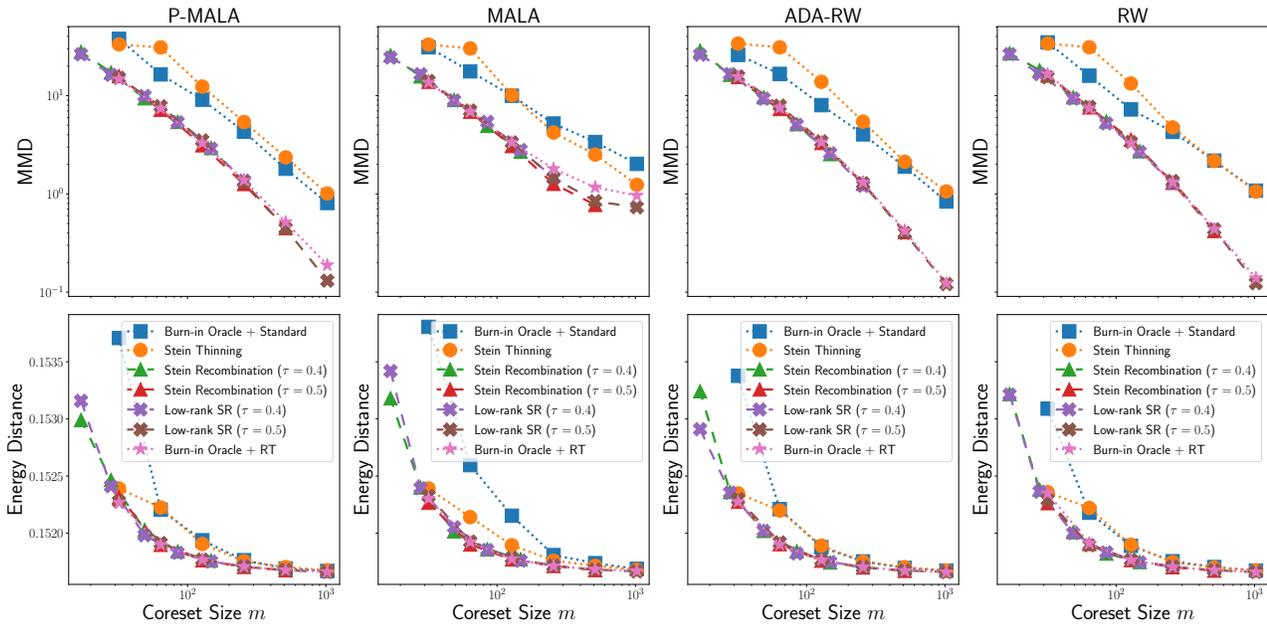


Figure I.2: **Correcting for burn-in with simplex-weighted compression.** For each of four MCMC algorithms and using only one chain, our methods consistently outperform the Stein and standard thinning baselines and match the 6-chain oracle.

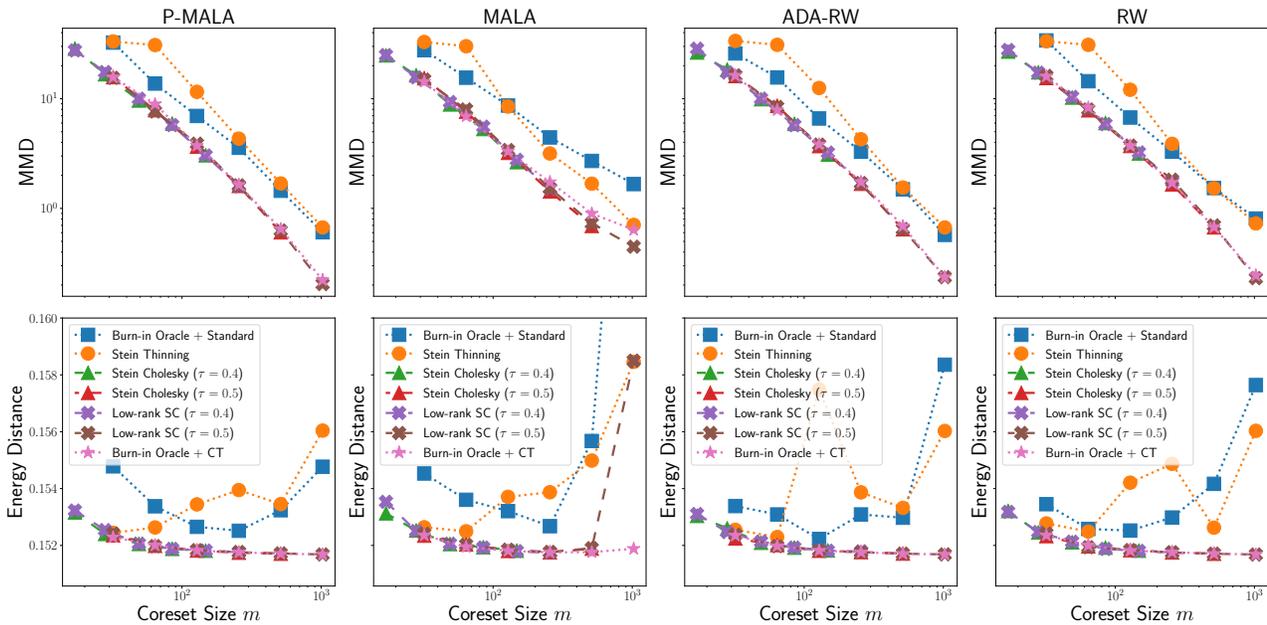


Figure I.3: **Correcting for burn-in with constant-preserving compression.** For each of four MCMC algorithms and using only one chain, our methods consistently outperform the Stein and standard thinning baselines and match the 6-chain oracle.

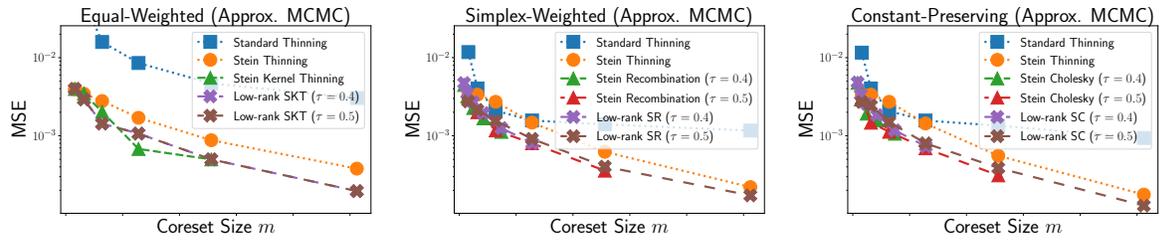


Figure I.4: **Posterior mean mean-squared error (MSE)** for the approximate MCMC compression experiment of Sec. 5. MSE is computed as $\|\hat{\mathbb{E}}_{\mathbb{P}} Z - \sum_{i \in [n_0]} w_i x_i\|_M^2 / d$ where $\hat{\mathbb{E}}_{\mathbb{P}} Z$ is the mean of the surrogate ground truth NUTS sample.