
Low-Rank Thinning

Annabelle Michael Carrell¹ Albert Gong² Abhishek Shetty³ Raaz Dwivedi² Lester Mackey⁴

Abstract

The goal in thinning is to summarize a dataset using a small set of representative points. Remarkably, sub-Gaussian thinning algorithms like Kernel Halving and Compress can match the quality of uniform subsampling while substantially reducing the number of summary points. However, existing guarantees cover only a restricted range of distributions and kernel-based quality measures and suffer from pessimistic dimension dependence. To address these deficiencies, we introduce a new low-rank analysis of sub-Gaussian thinning that applies to any distribution and any kernel, guaranteeing high-quality compression whenever the kernel or data matrix is approximately low-rank. To demonstrate the broad applicability of the techniques, we design practical sub-Gaussian thinning approaches that improve upon the best known guarantees for approximating attention in transformers, accelerating stochastic gradient training through reordering, and distinguishing distributions in near-linear time.

1. Introduction

This work is about thinning, finding a small set of representative points to accurately summarize a larger dataset. Here, we use the term “dataset” liberally to refer to any collection of points, be they experimental observations (as in Sec. 6), stochastic gradients (as in Sec. 5), or key-value pairs (as in Sec. 4). State-of-the-art thinning techniques provably improve upon uniform subsampling but only for restricted classes of kernel-based quality measures and with pessimistic dependence on the data dimension (see, e.g.,

¹University of Cambridge ²Cornell Tech ³MIT ⁴Microsoft Research New England. Correspondence to: Annabelle Carrell <ac2411@cam.ac.uk>, Albert Gong <agong@cs.cornell.edu>, Abhishek Shetty <ashetty1995@gmail.com>, Raaz Dwivedi <dwivedi@cornell.edu>, Lester Mackey <lmackey@microsoft.com>.

Harvey & Samadi, 2014; Phillips & Tai, 2020; Alweiss et al., 2021; Dwivedi & Mackey, 2024; 2022; Shetty et al., 2022; Li et al., 2024). We introduce a new analysis for sub-Gaussian thinning algorithms that applies to any kernel and shows that one can efficiently identify a better-than-uniform set of representative points whenever the kernel or data matrix is nearly low-rank. This opens the door to a variety of impactful applications including approximate dot-product attention in transformers, accelerating model training through gradient reordering, and distinguishing distributions with deep kernels in near-linear time.

In Sec. 2, we introduce our formal definition of thinning, two kernel-based measures of thinning quality, and a suite of candidate thinning algorithms. Sec. 3 presents our main theorem relating thinning quality to low-rank properties of a dataset and its induced kernel matrix. In Sec. 4, we translate the problem of approximating attention into a thinning problem and develop a practical solution, Thinformer, with state-of-the-art quality guarantees. In Sec. 5, we develop a thinned stochastic gradient reordering rule that provably accelerates training and bridges the theory-practice gaps left open by prior work. Finally, in Sec. 6 we derive new and improved power guarantees for practical thinned hypothesis tests that distinguish distributions in near-linear time. Each section also includes its own discussion of related work.

Notation. For each $n \in \mathbb{N}$ and $a, b \in \mathbb{R}$, we define $[n] \triangleq \{1, \dots, n\}$, $a \wedge b \triangleq \min(a, b)$, and $a \vee b \triangleq \max(a, b)$. We let $\|\mathbf{A}\|_{\text{op}}$, $\|\mathbf{A}\|_{\text{max}}$, and $\|\mathbf{A}\|_{2,\infty}$ respectively represent the maximum singular value, absolute entry, and row Euclidean norm of a matrix \mathbf{A} and let $\lambda_r(\mathbf{K})$ denote the r -th largest eigenvalue of a suitable matrix \mathbf{K} . We also define the Euclidean norm balls $\mathbb{B}^m \triangleq \{\mathbf{u} \in \mathbb{R}^m : \|\mathbf{u}\|_2 \leq 1\}$ and $\mathbb{B}^m(R) \triangleq R\mathbb{B}^m$ for each $m \in \mathbb{N}$ and $R > 0$. For an event \mathcal{E} and an integrable random variable X , we define $\mathbb{E}_{\mathcal{E}}[X] \triangleq \mathbb{E}[X \cdot \mathbf{1}[\mathcal{E}]]$. We write $a_n \leq \tilde{O}(b_n)$ to mean $a_n \leq b_n \text{ polylog}(n)$.

2. Sub-Gaussian Thinning

We begin with a formal definition of our problem setting. Consider a fixed collection of n_{in} input points \mathcal{X}_{in} belonging to a potentially larger universe of datapoints

Table 1: **Examples of $(\mathbf{K}, \nu, \delta)$ -sub-Gaussian thinning algorithms.** For input size n_{in} , output size $n_{\text{out}} \geq \sqrt{n_{\text{in}}}$, and $\|\mathbf{K}\|_{\text{max}} = 1$ we report each sub-Gaussian parameter ν and runtime up to constants independent of $(n_{\text{in}}, n_{\text{out}}, \delta, \mathbf{K})$.

Algorithm	SUBSAMPLING Prop. B.1	KH(δ) Prop. B.2	KH-COMPRESS(δ) Prop. B.5	GS-THIN Prop. B.6	GS-COMPRESS Prop. B.10
Sub-Gaussian parameter ν	$\frac{1}{\sqrt{n_{\text{out}}}}$	$\frac{\sqrt{\log(n_{\text{out}}/\delta)}}{n_{\text{out}}}$	$\frac{\sqrt{\log(n_{\text{out}}) \log(n_{\text{out}}/\delta)}}{n_{\text{out}}}$	$\frac{1}{n_{\text{out}}}$	$\frac{\sqrt{\log(n_{\text{out}})}}{n_{\text{out}}}$
Runtime	n_{out}	n_{in}^2	n_{out}^2	n_{in}^3	n_{out}^3

$\mathcal{X} \triangleq \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. The aim of a thinning algorithm is to select n_{out} points from \mathcal{X}_{in} that together accurately summarize \mathcal{X}_{in} . This is formalized by the following definition.

Definition 1 (Thinning algorithms). A thinning algorithm ALG takes as input \mathcal{X}_{in} and returns a possibly random subset \mathcal{X}_{out} of size n_{out} . We denote the input and output empirical distributions by $\mathbb{P}_{\text{in}} \triangleq \frac{1}{n_{\text{in}}} \sum_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \delta_{\mathbf{x}}$ and $\mathbb{P}_{\text{out}} \triangleq \frac{1}{n_{\text{out}}} \sum_{\mathbf{x} \in \mathcal{X}_{\text{out}}} \delta_{\mathbf{x}}$ and define the induced probability vectors $\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}} \in \Delta_{n-1}$ over the indices $[n]$ by

$$\mathbf{p}_{\text{in},i} = \frac{1[\mathbf{x}_i \in \mathcal{X}_{\text{in}}]}{n_{\text{in}}} \text{ and } \mathbf{p}_{\text{out},i} = \frac{1[\mathbf{x}_i \in \mathcal{X}_{\text{out}}]}{n_{\text{out}}} \text{ for all } i \in [n].$$

When $\mathcal{X} \subset \mathbb{R}^d$, we use $\mathbf{X} \triangleq [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ to denote the input point matrix so that

$$\mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{\text{in}}}[\mathbf{x}] = \mathbf{X}^\top \mathbf{p}_{\text{in}} \text{ and } \mathbb{E}_{\mathbf{x} \sim \mathbb{P}_{\text{out}}}[\mathbf{x}] = \mathbf{X}^\top \mathbf{p}_{\text{out}}.$$

We will make use of two common measures of summarization quality.

Definition 2 (Kernel MMD and max seminorm). Given two distributions $\mu, \tilde{\mu}$ and a reproducing kernel \mathbf{k} (Steinwart & Christmann, 2008, Def. 4.18), the associated kernel maximum mean discrepancy (MMD) is the worst-case difference in means for functions in the unit ball $\mathbb{B}_{\mathbf{k}}$ of the associated reproducing kernel Hilbert space:

$$\text{MMD}_{\mathbf{k}}(\mu, \tilde{\mu}) \triangleq \sup_{f \in \mathbb{B}_{\mathbf{k}}} |\mathbb{E}_{\mathbf{x} \sim \mu} f(\mathbf{x}) - \mathbb{E}_{\mathbf{x} \sim \tilde{\mu}} f(\mathbf{x})|.$$

When $\mu = \mathbb{P}_{\text{in}}$ and $\tilde{\mu} = \mathbb{P}_{\text{out}}$ as in Def. 1 and $\mathbf{K} \triangleq (\mathbf{k}(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n \in \mathbb{R}^{n \times n}$ denotes the induced kernel matrix, then the MMD can be expressed as a Mahalanobis distance between \mathbf{p}_{in} and \mathbf{p}_{out} :

$$\begin{aligned} \text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{P}_{\text{out}}) &= \sqrt{(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})^\top \mathbf{K} (\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})} \\ &\triangleq \text{MMD}_{\mathbf{K}}(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}). \end{aligned}$$

For any indices $\mathcal{I} \subseteq [n]$, we further define the kernel max seminorm (KMS)

$$\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \triangleq \max_{i \in \mathcal{I}} |e_i^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})|. \quad (1)$$

Notably, when the input points lie in \mathbb{R}^d and $\mathbf{k}(\mathbf{x}_i, \mathbf{x}_j)$ is the linear kernel $\langle \mathbf{x}_i, \mathbf{x}_j \rangle$ (so that $\mathbf{K} = \mathbf{X}\mathbf{X}^\top$), MMD measures the Euclidean discrepancy in datapoint means between the input and output distributions:

$$\text{MMD}_{\mathbf{K}}(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) = \|\mathbf{X}^\top \mathbf{p}_{\text{in}} - \mathbf{X}^\top \mathbf{p}_{\text{out}}\|_2.$$

A common strategy for bounding the error of a thinning algorithm is to establish its sub-Gaussianity.

Definition 3 (Sub-Gaussian thinning algorithm). We write $\text{ALG} \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$ and say ALG is $(\mathbf{K}, \nu, \delta)$ -sub-Gaussian, if ALG is a thinning algorithm, \mathbf{K} is a symmetric positive semidefinite (SPSD) matrix, $\nu > 0$, $\delta \in [0, 1)$, and there exists an event \mathcal{E} with probability at least $1 - \delta/2$ such that, the input and output probability vectors satisfy

$$\mathbb{E}_{\mathcal{E}}[\exp(\langle \mathbf{u}, \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}) \rangle)] \leq \exp\left(\frac{\nu^2}{2} \mathbf{u}^\top \mathbf{K} \mathbf{u}\right), \forall \mathbf{u} \in \mathbb{R}^n.$$

Here, the sub-Gaussian parameter ν controls the summarization quality of the thinning algorithm, and we see from Tab. 1 that a variety of practical thinning algorithms are $(\mathbf{K}, \nu, \delta)$ -sub-Gaussian for varying levels of ν .

2.1. Examples of sub-Gaussian thinning algorithms

Perhaps the simplest sub-Gaussian thinning algorithm is *uniform subsampling*: by Prop. B.1, selecting n_{out} points from \mathcal{X}_{in} uniformly at random (without replacement) is $(\mathbf{K}, \nu, 0)$ -sub-Gaussian with $\nu = \sqrt{\|\mathbf{K}\|_{\text{max}}}/\sqrt{n_{\text{out}}}$. Unfortunately, uniform subsampling suffers from relatively poor summarization quality. As we prove in App. B.1.1, its root-mean-squared MMD and KMS are both $\Omega(1/\sqrt{n_{\text{out}}})$ meaning that $n_{\text{out}} = 10000$ points are needed to achieve 1% relative error.

Proposition 1 (Quality of uniform subsampling). For any $\mathcal{I} \subseteq [n]$, a uniformly subsampled thinning satisfies

$$\begin{aligned} \mathbb{E}[\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}})] &= \frac{1}{n_{\text{out}}} \frac{n_{\text{in}} - n_{\text{out}}}{n_{\text{in}} - 1} C_{\mathbf{K}} \text{ and} \\ \mathbb{E}[\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}}^2] &\geq \frac{1}{n_{\text{out}}} \frac{n_{\text{in}} - n_{\text{out}}}{n_{\text{in}} - 1} \max_{i \in \mathcal{I}} C_{\mathbf{K} e_i e_i^\top \mathbf{K}} \end{aligned}$$

for any SPSP \mathbf{K} with $C_{\mathbf{K}} \triangleq \sum_{i=1}^n \mathbf{p}_{\text{in},i} \mathbf{K}_{ii} - \mathbf{p}_{\text{in}}^\top \mathbf{K} \mathbf{p}_{\text{in}}$.

Fortunately, uniform subsampling is not the only sub-Gaussian thinning algorithm available. For example, the Kernel Halving (KH(δ)) algorithm of Dwivedi & Mackey (2024) provides a substantially smaller sub-Gaussian parameter, $\nu = O(\sqrt{\log(n_{\text{out}}/\delta)/n_{\text{out}}})$, in n_{in}^2 time by selecting one out of every two points and biasing selection toward the point that yields smaller approximation error. Similarly, the KH-COMPRESS(δ) algorithm of Shetty et al. (2022, Ex. 3) delivers $\nu = O(\sqrt{\log(n_{\text{out}})\log(n_{\text{out}}/\delta)/n_{\text{out}}})$ in only n_{out}^2 time by halving and concatenating coresets of geometrically increasing size until the target output size is met. We derive simplified versions of these algorithms with identical sub-Gaussian constants in Apps. B.2 and B.5 and a linear-kernel variant (LKH(δ)) with $n_{\text{in}}d$ runtime in App. B.3. To round out our set of examples, we show in App. B.6.1 that two new thinning algorithms based on the Gram-Schmidt walk of Bansal et al. (2018) yield even smaller ν at the cost of increased runtime. We call these algorithms Gram-Schmidt Thinning (GS-THIN) and GS-COMPRESS.

3. Low-rank Sub-Gaussian Thinning

One might hope that the improved sub-Gaussian constants of Tab. 1 would also translate into improved quality metrics. Our main result, proved in App. C, shows that this is indeed the case whenever the inputs are approximately low-rank.

Theorem 1 (Low-rank sub-Gaussian thinning). *Fix any $\delta' \in (0, 1)$, $r \leq n$, and $\mathcal{I} \subseteq [n]$. If $\text{ALG} \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$, then the following bounds hold individually with probability at least $1 - \delta/2 - \delta'$:*

$$\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \leq \nu^2 [e^2 r + e \log(\frac{1}{\delta'})] + \lambda_{r+1} \left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}} \right) \quad \text{and} \quad (2)$$

$$\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \leq \nu D_{\mathcal{I}} \sqrt{2 \log(\frac{2|\mathcal{I}|}{\delta'})}. \quad (3)$$

Here, λ_j denotes the j -th largest eigenvalue of \mathbf{K} , $\lambda_{n+1} \triangleq 0$, and $D_{\mathcal{I}} \triangleq \max_{i \in \mathcal{I}} \sqrt{\mathbf{K}_{ii}}$.

Suppose that, in addition, $\mathcal{X} \subset \mathbb{R}^d$ and $|\mathbf{K}_{il} - \mathbf{K}_{jl}| \leq L_{\mathbf{K}} \|\mathbf{x}_i - \mathbf{x}_j\|_2$ for some $L_{\mathbf{K}} > 0$ and all $i, j \in \mathcal{I}$ and $l \in \text{supp}(\mathbf{p}_{\text{in}})$. Then, with probability at least $1 - \delta/2 - \delta'$,

$$\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \leq \nu D_{\mathcal{I}} \sqrt{2 \log(4/\delta')} (1 + \frac{32}{\sqrt{3}}) + \nu D_{\mathcal{I}} 32 \sqrt{\frac{2}{3} \text{rank}(\mathbf{X}_{\mathcal{I}}) \log(\frac{3e^2 R_{\mathcal{I}} L_{\mathbf{K}}}{D_{\mathcal{I}}^2 \wedge (R_{\mathcal{I}} L_{\mathbf{K}})})} \quad (4)$$

for $R_{\mathcal{I}} \triangleq \max_{i \in \mathcal{I}} \|\mathbf{x}_i\|_2$ and $\mathbf{X}_{\mathcal{I}} \triangleq [\mathbf{x}_i]_{i \in \mathcal{I}}^{\top}$.

Let us unpack the three components of this result. First, Thm. 1 provides a high-probability $O(\nu \sqrt{\log(|\mathcal{I}|)})$ bound (3) on the KMS for any kernel and any sub-Gaussian thinning algorithm on any space. In particular, the non-uniform

algorithms of Tab. 1 all enjoy $O(\log(n_{\text{out}}) \sqrt{\log(|\mathcal{I}|)/n_{\text{out}}})$ KMS, a significant improvement over the $\Omega(1/\sqrt{n_{\text{out}}})$ KMS of uniform subsampling. The bound (3) follows from the sub-Gaussianity of the thinning algorithm (Def. 3) and the union bound over $\pm e_i$ for each $i \in \mathcal{I}$.

For datapoints in \mathbb{R}^d , Thm. 1 also provides a refined $O(\nu \sqrt{\text{rank}(\mathbf{X}_{\mathcal{I}}) \log(R_{\mathcal{I}} L_{\mathbf{K}})})$ bound (4) on KMS. For bounded data, this trades an explicit dependence on the number of query points $|\mathcal{I}|$ for a rank factor that is never larger (and sometimes significantly smaller) than d . This refinement follows from a more elaborate chaining argument that frames $(e_i^{\top} K(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))_{i \in \mathcal{I}}$ as a sub-Gaussian process (Lem. C.5) and uses the Lipschitzness of K to control its entropy integral. We will make use of these results when approximating dot-product attention in Sec. 4.

Notably, Thm. 3.1 of Phillips & Tai (2020) implies that any thinning algorithm must incur at least $\Omega(\sqrt{d}/n_{\text{out}})$ KMS error for some dataset in \mathbb{R}^d and many common kernels. Meanwhile, our Tab. 1 and Thm. 1 imply that GS-THIN has $\nu = O(1/n_{\text{out}})$ and hence KMS $O(\sqrt{d}/n_{\text{out}})$. Taken together, these results imply that no algorithm can have sub-Gaussian constant $\nu = o(1/n_{\text{out}})$ and that GS-THIN enjoys minimax rate-optimal KMS and a minimax rate-optimal sub-Gaussian constant.

Finally, Thm. 1 provides an $O(\nu \sqrt{r} + \sqrt{\lambda_{r+1}/n_{\text{out}}})$ high-probability bound on kernel MMD, where the approximate rank parameter r can be freely optimized. We establish this result by projecting $K^{1/2}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})$ onto the first r eigenvectors of K , bounding the residual error in terms of λ_{r+1} , and bounding the projection magnitude with high probability using the sub-Gaussianity of ALG and a union bound over a finite cover of a Euclidean ball in R^r .

When $\mathbf{K} = (\mathbf{k}(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n$ is generated by a finite-rank kernel \mathbf{k} , like a linear kernel $\langle \mathbf{x}_i, \mathbf{x}_j \rangle$, a polynomial kernel $(1 + \langle \mathbf{x}_i, \mathbf{x}_j \rangle)^p$, or a random Fourier feature kernel (Rahimi & Recht, 2007), this guarantee becomes $O(\nu)$ and improves upon uniform subsampling whenever $\nu = o(1/\sqrt{n_{\text{out}}})$. In this case, the non-uniform algorithms of Tab. 1 all enjoy $O(\log(n_{\text{out}})/n_{\text{out}})$ MMD, a significant improvement over the $\Omega(1/\sqrt{n_{\text{out}}})$ MMD of uniform subsampling. We will revisit this finite-rank setting when studying stochastic gradient acceleration strategies in Sec. 5.

More generally, Thm. 1 guarantees improved MMD even for full-rank \mathbf{K} , provided that the eigenvalues of \mathbf{K} decay sufficiently rapidly. For example, optimizing over the approximate rank parameter r yields an $O(\nu \log^{p/2}(n_{\text{out}}))$ bound under exponential eigenvalue decay $\lambda_{r+1} = O(ne^{-cr^{1/p}})$ and an $O(\nu^{\frac{p}{p+1}} (\frac{n}{n_{\text{out}}})^{\frac{1}{2(p+1)}})$ bound under polynomial eigenvalue decay $\lambda_{r+1} = O(n/r^p)$. Fortunately, some of the most commonly-used kernels generate kernel matrices with rapid eigenvalue decay.

For example, the popular Gaussian kernel on \mathbb{R}^d ,

$$\text{GAUSS}(\eta) : \mathbf{k}(x, y) = \exp(-\eta\|x - y\|_2^2) \text{ for } \eta > 0, \quad (5)$$

generates $\mathbf{K} = (\mathbf{k}(x_i, x_j))_{i,j=1}^n$ satisfying

$$\lambda_{r+1} \leq ne^{-\frac{d}{2e}r^{1/d} \log\left(\frac{dr^{1/d}}{4e^2\eta R^2}\right)} \text{ for } (2e)^d \leq r < n \quad (6)$$

whenever $\mathcal{X} \subset \mathbb{B}^d(R)$ (Altschuler et al., 2019, Thm. 3). Combined with Thm. 1, this fact immediately yields an MMD guarantee for each algorithm in Tab. 1. We present a representative guarantee for KH(δ).

Corollary 1 (Gaussian MMD of KH). *If $\mathcal{X}_{\text{in}} \subset \mathbb{B}^d(R)$ for $R > 0$, then KH(δ) with $\mathbf{k} = \text{GAUSS}(\eta)$, and $n = n_{\text{in}}$ delivers*

$$\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \leq O\left(\frac{\log(n_{\text{out}}/\delta)}{n_{\text{out}}^2} \left(\left(\frac{\log(n_{\text{out}}) \vee (R^2\eta)}{d}\right)^d + \log\left(\frac{1}{\delta'}\right)\right)\right)$$

with probability at least $1 - \delta/2 - \delta'$.

The proof in App. D provides a fully explicit and easily computed bound on the Gaussian MMD. Under the same assumptions, the distinct analysis of Dwivedi & Mackey (2022, Thm. 2, Prop. 3) provides a squared MMD bound of size $\Theta\left(\frac{\log(n_{\text{out}}/\delta)}{n_{\text{out}}^2} \left(\frac{\log^{d+1}(n_{\text{out}})R^d\eta^{d/2}}{(\log \log(n_{\text{out}}))^d} + \log\left(\frac{1}{\delta'}\right)\right)\right)$. Notably, Cor. 1 improves upon this best known KH(δ) guarantee whenever the datapoint radius $R = O(\log n_{\text{out}})$, a property that holds almost surely for any bounded, sub-Gaussian, or subexponential data sequence (see Dwivedi & Mackey, 2024, Prop. 2).

Altschuler et al. (2019, Thm. 4) additionally showed that Gaussian kernel matrix eigenvalues satisfy

$$\lambda_{r+1} \leq ne^{-cr^{2/(5d^*)}} \text{ for } 1 \leq r < n \quad (7)$$

for a constant c independent of \mathcal{X} when \mathcal{X} belongs to a smooth compact manifold of dimension $d^* < d$. In this case, our low-rank analysis yields adaptive MMD guarantees that scale with the potentially much smaller intrinsic dimension d^* . We use Thm. 1 to prove the first such intrinsic-dimension guarantee for KH(δ) in App. E.

Corollary 2 (Intrinsic Gaussian MMD of KH). *If \mathcal{X}_{in} lies on a smooth manifold $\Omega \subset \mathbb{B}^d$ of dimension $d^* < d$ (Assump. E.1), then KH(δ) with $\mathbf{k} = \text{GAUSS}(\eta)$, and $n = n_{\text{in}}$ delivers*

$$\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \leq O\left(\frac{\log(n_{\text{out}}/\delta)}{n_{\text{out}}^2} \left(\left(\frac{\log(n_{\text{out}})}{c}\right)^{\frac{5d^*}{2}} + \log\left(\frac{1}{\delta'}\right)\right)\right)$$

with probability at least $1 - \frac{\delta}{2} - \delta'$ for c independent of \mathcal{X}_{in} .

In Sec. 6, we will use Cors. 1 and 2 to establish new guarantees for distinguishing distributions in near-linear time.

With our core theory in hand, we now turn our attention to a series of impactful applications.

4. Approximating Attention

We will first use our analysis to accelerate attention approximation in transformers. Dot-product attention lies at the heart of the transformer neural network architecture that has revolutionized natural language processing, computer vision, and speech recognition over the last decade (Vaswani et al., 2017; Dosovitskiy et al., 2021; Dong et al., 2018). Given a collection of query, key, and value vectors $(\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i)_{i=1}^n$ each in \mathbb{R}^d , dot-product attention computes the softmax matrix

$$\mathbf{T} \triangleq \text{ATTENTION}((\mathbf{q}_i)_{i=1}^n, (\mathbf{k}_j, \mathbf{v}_j)_{j=1}^n) \triangleq \mathbf{D}^{-1} \mathbf{A} \mathbf{V} \quad (8)$$

for $\mathbf{A}_{ij} \triangleq \exp(\frac{\langle \mathbf{q}_i, \mathbf{k}_j \rangle}{\sqrt{d}})$, $\mathbf{D} = \text{diag}(\mathbf{A} \mathbf{1}_n)$, and $\mathbf{V}_{ij} \triangleq \mathbf{v}_{ij}$.

While attention has enjoyed unprecedented success in capturing long-range dependencies amongst datapoints, its computation is expensive, requiring $\Theta(dn^2)$ time to construct and multiply the matrix \mathbf{A} . This quadratic-time bottleneck has inspired a plethora of practical approximate attention mechanisms (e.g., Kitaev et al., 2020; Choromanski et al., 2021; Chen et al., 2021), but, to our knowledge, only two guarantee accurate reconstruction of the softmax matrix \mathbf{T} (Zandieh et al., 2023; Han et al., 2024).¹ In this section, we design a new fast attention approximation based on sub-Gaussian thinning and derive guarantees that improve upon the prior art.

4.1. Thinning attention in theory

Algorithm 1: Thinformer

Input: Queries, keys, and values $(\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i)_{i=1}^n$ in \mathbb{R}^d , n_{out}

// Define key-value attention kernel

$\mathbf{k}_{\text{att}}((\tilde{\mathbf{k}}, \tilde{\mathbf{v}}), (\tilde{\mathbf{k}}', \tilde{\mathbf{v}}')) \triangleq \exp(\langle \tilde{\mathbf{k}}, \tilde{\mathbf{k}}' \rangle) \langle \tilde{\mathbf{v}}, \tilde{\mathbf{v}}' \rangle$

// Thin augmented key-value pairs using \mathbf{k}_{att}

$v_{\text{max}} \leftarrow \max_{i \in [n]} \|\mathbf{v}_i\|_{\infty}$; $(\tilde{\mathbf{k}}_i, \tilde{\mathbf{v}}_i)_{i=1}^n \leftarrow (\mathbf{k}_i/d^{\frac{1}{4}}, (\mathbf{v}_i, v_{\text{max}}))_{i=1}^n$

$\mathcal{X}_{\text{out}} \leftarrow \text{KH-COMPRESS}(0.5)(\mathcal{X}_{\text{in}} = (\tilde{\mathbf{k}}_i, \tilde{\mathbf{v}}_i)_{i=1}^n, \mathbf{k}_{\text{att}}, n_{\text{out}})$

// Return exact attention on selected key-value subset

return $\hat{\mathbf{T}} \triangleq \text{ATTENTION}((\mathbf{q}_i)_{i=1}^n, \{(\mathbf{k}, \mathbf{v}) : (\tilde{\mathbf{k}}, \tilde{\mathbf{v}}) \in \mathcal{X}_{\text{out}}\})$

Alg. 1 summarizes our new *Thinformer* module. At its heart is a new key-value attention kernel \mathbf{k}_{att} that mimics the special structure of the softmax matrix \mathbf{T} . Alg. 1 uses the attention kernel and a high-quality thinning algorithm, KH-COMPRESS(0.5), to subselect key-value pairs and then computes exact attention (8) for the key-value subset. In total, this requires only $O(dn_{\text{out}}^2)$ time to run KH-COMPRESS(0.5) and $O(dn_{\text{out}})$ time to compute

¹A third remarkable work (Alman & Song, 2024) establishes upper and lower bounds for attention approximation but without a practical implementation.

Table 2: **Practical approximations with guarantees.** For each approximation $\hat{\mathbf{T}} \in \mathbb{R}^{n \times d}$ to the softmax matrix \mathbf{T} (8), we report, up to a constant factor, the best worst-case error guarantee for $\|\hat{\mathbf{T}} - \mathbf{T}\|_{\max}$ given $O(dn^{1+a})$ running time and γ -bounded (9) queries and keys. Here, the ratio $\|\mathbf{V}\|_{\text{op}}/\|\mathbf{V}\|_{2,\infty}$ lies in $[1, \sqrt{n}]$ and $\tau = 0.173 + o(1)$.

Approximation	Guarantee
Thinformer	$\frac{n^{2\gamma} \sqrt{d \log(n\ \mathbf{V}\ _{\max}) \log n}}{n^a} \cdot \ \mathbf{V}\ _{2,\infty}$
KDEformer	$\frac{n^{2\gamma + \frac{\tau}{2}(1 + \frac{\gamma}{2})}}{n^{a/2}} \cdot \ \mathbf{V}\ _{\text{op}}$
HyperAttention	$\frac{n^{\frac{17\gamma}{3}(\log n)^{\frac{1}{6}}}}{n^{a/6}} \cdot \ \mathbf{V}\ _{\text{op}}$

ATTENTION with n queries and n_{out} key-value pairs. In contrast, computing the exact softmax matrix \mathbf{T} with standard matrix multiplication requires $\Theta(dn^2)$ time. Our next result, proved in App. F, shows that Alg. 1 also admits a strong quality guarantee for approximating \mathbf{T} .

Theorem 2 (Quality of Thinformer). *With probability at least $\frac{1}{2}$, Thinformer (Alg. 1) yields*

$$\|\hat{\mathbf{T}} - \mathbf{T}\|_{\max} \leq \frac{c \exp(\frac{2R^2}{\sqrt{d}}) \|\mathbf{V}\|_{2,\infty} \sqrt{\log_2(n_{\text{out}}) \log(12n_{\text{out}} \log_2 \frac{n_{\text{in}}}{n_{\text{out}}})}}{n_{\text{out}}}$$

$$\text{for } c \triangleq \frac{128}{\sqrt{3}} \sqrt{(d+1) \log(3e^2(\frac{R^2}{\sqrt{d}} + 2) \|\mathbf{V}\|_{\max})} + \sqrt{\log(8)(4 + \frac{128}{\sqrt{3}})} \text{ and } R \triangleq \max_{i \in [n]} \max(\|\mathbf{k}_i\|_2, \|\mathbf{q}_i\|_2).$$

To put this result into context, let us compare with the existing guarantees for practical attention approximation, summarized in Tab. 2. Under the γ -boundedness assumption,

$$\max_{i \in [n]} \max(\|\mathbf{k}_i\|_2^2, \|\mathbf{q}_i\|_2^2) \leq \gamma \sqrt{d} \log n, \quad (9)$$

the KDEformer approximation (Zandieh et al., 2023, Cor. 3.6) with $\tau = 0.173 + o(1)$, the HyperAttention approximation (Han et al., 2024, Thm. 1) with no masking, and the Thinformer approximation (Thm. 2) guarantee the $\|\hat{\mathbf{T}} - \mathbf{T}\|_{\max}$ bounds of Tab. 2 with $O(dn^{1+a})$ runtime and probability at least $\frac{1}{2}$. The Thinformer guarantee exhibits four improvements over its predecessors. First, it establishes a significantly faster error decay rate (n^{-a} versus $n^{-a/2}$ or $n^{-a/6}$) for a given subquadratic runtime n^{1+a} . Second, it reduces the dependence on the error inflation factor γ . Third, like the HyperAttention guarantee, it eliminates all dependence on the KDEformer penalty parameter τ . Finally, it reduces dependence on the value matrix by a factor of $\frac{\|\mathbf{V}\|_{\text{op}}}{\|\mathbf{V}\|_{2,\infty}} \in [1, \sqrt{n}]$.

Put otherwise, with bounded $\|\mathbf{V}\|_{2,\infty}$, $\hat{\mathbf{T}}_{\text{thin}}$ can provide consistent (i.e., $\|\hat{\mathbf{T}}_{\text{thin}} - \mathbf{T}\|_{\max} \rightarrow 0$ as $n \rightarrow \infty$) subquadratic estimation whenever γ is bounded away from $1/2$

and guarantee, for example, $O(\frac{1}{\sqrt{n}})$ error in $\tilde{O}(dn^{\frac{3}{2}+2\gamma})$ time. In contrast, the $\hat{\mathbf{T}}_{\text{kde}}$ and $\hat{\mathbf{T}}_{\text{hyp}}$ bounds require quadratic runtime to guarantee $O(\frac{1}{\sqrt{n}})$ error in the best case ($\|\mathbf{V}\|_{\text{op}} = O(1)$) and cannot guarantee consistent subquadratic estimation in the worst case ($\|\mathbf{V}\|_{\text{op}} = \Omega(\sqrt{n})$).

4.2. Thinning attention in practice

To gauge the practical effectiveness of Alg. 1, we recreate the benchmark Tokens-To-Token Vision Transformer (T2T-ViT) and BigGAN image generation experiments of Zandieh et al. (2023). In the T2T-ViT experiment, attention approximations are scored on their ImageNet classification accuracy and computational expense when used as drop-in replacements for the two most expensive attention layers in a pretrained T2T-ViT neural network (Yuan et al., 2021). In the BigGAN experiment, approximations are scored on their computational expense and two popular measures of image generation quality, the Frechet Inception Distance (FID, Heusel et al., 2017) and Inception Score (IS, Salimans et al., 2016). Using the exact implementations and settings provided by Zandieh et al. (2023), we benchmark our PyTorch implementation of Thinformer against exact attention and four leading attention approximations: Performer (Choromanski et al., 2021), Reformer (Kitaev et al., 2020), ScatterBrain (Chen et al., 2021), and KDEformer.

In Tab. 3, we find that Thinformer ($g = 2$) provides the highest Top-1 accuracy on the ImageNet 2012 validation set (Russakovsky et al., 2015), while running faster than all of the alternatives. In Tab. 4, Thinformer ($g = 2$) yields better FID and IS than all of the alternatives while running significantly faster than exact, KDEformer, Reformer, and ScatterBrain. Performer runs faster still but at the expense of substantially worse FID and IS. The final attention call of Thinformer can also be combined with optimized attention implementations like FlashAttention (Dao et al., 2022; Dao, 2024) to further reduce the time and memory footprint. We provide PyTorch code replicating this experiment at <https://github.com/microsoft/thinformer> and supplementary experiment details in App. L.1.

5. Faster SGD Training

We now turn to a second application, accelerating training through gradient reordering. To train a machine learning model parameterized by $\mathbf{w} \in \mathbb{R}^d$, a standard approach is to minimize the empirical risk $f(\mathbf{w}) \triangleq \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{w})$ using stochastic gradient descent (SGD) updates,

$$\mathbf{w}^{k+\frac{i}{n}} = \mathbf{w}^{k+\frac{i-1}{n}} - \alpha \nabla f_{\pi_k(i)}(\mathbf{w}^{k+\frac{i-1}{n}}), \quad (10)$$

for each epoch $k \in [K]$ and datapoint $i \in [n]$. Here, $\alpha > 0$ is a step size, each f_i is a datapoint-specific loss function,

Table 3: **Quality of T2T-ViT attention approximations on ImageNet.** We report mean Top-1 accuracy ± 1 standard deviation across five random seeds and mean forward pass runtime ± 1 standard deviation across 50 batches of 64 images.

Attention Algorithm	Top-1 Accuracy (%)	Layer 1 Runtime (ms)	Layer 2 Runtime (ms)
Exact	82.55 ± 0.00	18.48 ± 0.12	1.40 ± 0.01
Performer	80.56 ± 0.30	2.54 ± 0.01	0.60 ± 0.01
Reformer	81.47 ± 0.06	7.84 ± 0.03	1.53 ± 0.01
KDEformer	82.00 ± 0.07	5.39 ± 0.03	2.28 ± 0.03
Scatterbrain	82.05 ± 0.08	6.86 ± 0.02	1.55 ± 0.03
Thinformer (Ours)	82.18 ± 0.05	2.06 ± 0.01	0.54 ± 0.00

 Table 4: **Quality of BigGAN attention approximations for image generation.** We report Frechet Inception Distance (FID) with the ImageNet validation set, Inception Scores (IS), and mean forward pass runtime ± 1 standard deviation across 10 batches of 32 images. A lower FID or higher IS indicates better image generation quality.

Attention Algorithm	FID (\downarrow)	IS (\uparrow)	Runtime (ms)
Exact	32.18	58.37 ± 4.21	5.83 ± 0.09
Performer	33.58	38.07 ± 3.43	2.28 ± 0.00
Reformer	72.23	19.14 ± 2.09	12.10 ± 0.01
KDEformer	30.71	57.09 ± 4.05	6.40 ± 0.34
ScatterBrain	38.78	36.93 ± 2.97	3.13 ± 0.02
Thinformer (Ours)	30.54	57.12 ± 3.96	2.66 ± 0.01

and π_k is a permutation of $[n]$ representing the order in which datapoints are processed in the k -th epoch.

Algorithm 2: Thinned Reordering

Input: Stochastic gradients $(\mathbf{x}_i^k \triangleq \nabla f_{\pi_k(i)}(\mathbf{w}^{k+\frac{i-1}{n}}))_{i=1}^n$, prior ordering π_k , thinning algorithm ALG

// Select half of points using linear kernel
 $\mathcal{X}_{\text{out}}^k \leftarrow \text{ALG}(\mathcal{X}_{\text{in}} = (\mathbf{x}_i^k)_{i=1}^n, n_{\text{out}} = \frac{n}{2}, \mathbf{k}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle)$
 $\Pi \leftarrow []; \quad \Pi' \leftarrow []$ // Initialize empty start and end lists

for $i = 1, \dots, n$ **do**
 $\Pi.\text{append}(\pi_k(i))$ if $\mathbf{x}_i^k \in \mathcal{X}_{\text{out}}^k$ else $\Pi'.\text{prepend}(\pi_k(i))$
end

return $\pi_{k+1} = \text{concatenate}(\Pi, \Pi')$

Typically, one selects the orderings π_k uniformly at random, but recent work has demonstrated faster convergence using non-uniform, adaptively selected orderings. Specifically, Lu et al. (2022); Cooper et al. (2023) show that any sufficiently accurate thinning algorithm can be efficiently transformed into a reordering rule that improves the convergence rate of SGD by a substantial $\tilde{O}(n^{-1})$ factor. Their approach, distilled in Alg. 2, uses an elegant construction of Harvey & Samadi (2014, Thm. 10) to translate a high-quality thinning of stochastic gradients into a higher-quality reordering. However, these prior studies leave two

problems unaddressed.

First, while the established convergence rates of Lu et al. (2022) nearly match the minimax lower bounds for permuted SGD algorithms (Cha et al., 2023, Thm. 4.5), a multiplicative gap of size $\Theta(d)$ remains in the worst case. This led Cha et al. (2023) to declare, “It is an open problem whether there exists a permutation-based SGD algorithm that gives a dimension-free upper bound while maintaining the same dependency on other factors.”

Second, Lu et al. (2022) carry out their analysis using the self-balancing walk (SBW) thinning algorithm of Alweiss et al. (2021) but find its overhead to be too high in practice. Hence, in all experiments they instead employ a greedy thinning algorithm that often works well in practice but is not covered by their analysis.

5.1. Bridging the dimension gap

To address the first problem, we derive a new guarantee for SGD with LKH reordering that replaces the typical $\Theta(d)$ penalty with a soft notion of rank.

Definition 4 (ϵ -rank). *The ϵ -rank, $\text{rank}_\epsilon(\mathbf{X})$, of a matrix \mathbf{X} is the number of singular values greater than ϵ .*

Theorem 3 (LKH-SGD convergence). *Suppose that, for*




Figure 1: **Train and test convergence trajectories for mortgage classification with reordered SGD variants.** We display mean values ± 1 standard deviation across 5 random seeds. See Sec. 5.2 for more details.

all $i \in [n]$ and $\mathbf{w}, \mathbf{v} \in \mathbb{R}^d$, the losses f and f_i satisfy

$$\begin{aligned} \|\nabla f_i(\mathbf{w}) - \nabla f(\mathbf{w})\|_2^2 &\leq \sigma^2 \text{ (bounded noise),} \\ \|\nabla f_i(\mathbf{w}) - \nabla f_i(\mathbf{v})\|_2 &\leq L\|\mathbf{w} - \mathbf{v}\|_2 \text{ (smoothness), and} \\ f(\mathbf{w}) - f^* &\leq \frac{1}{2\mu}\|\nabla f(\mathbf{w})\|_2^2 \text{ (PL) for } f^* \triangleq \inf_{\mathbf{v} \in \mathbb{R}^d} f(\mathbf{v}). \end{aligned}$$

Then, with probability at least $\frac{1}{2}$, SGD (10) with $\text{LKH}(\frac{1}{2K})$ reordering (Alg. 2) and step size α given in App. G satisfies

$$\begin{aligned} f(\mathbf{w}_K) - f^* &\leq \tilde{O}\left(\frac{r}{n^2 K^2}\right) \text{ for} \\ r &\triangleq \max_{k \in [K]} \text{rank}_{\epsilon_k}([\mathbf{x}_1^k, \dots, \mathbf{x}_n^k]), \quad \bar{\mathbf{x}}^k \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^k, \\ \text{and } \epsilon_k &\triangleq \max_{i \in [n]} \frac{\sqrt{9e \log(4Kn \log(en/2)) \log(4Kn)} \|\mathbf{x}_i^k - \bar{\mathbf{x}}^k\|_2}{\sqrt{n}}. \end{aligned}$$

The proof of Thm. 3 in App. G simply uses Thm. 1 to bound the thinning quality of $\text{LKH}(\frac{1}{2K})$ and then adapts the prior SGD analysis of Cooper et al. (2023). Notably, the standard practice of *random reshuffling*, i.e., SGD with uniform reordering, can only guarantee a significantly slower $\Omega(\frac{1}{nK^2})$ rate under these assumptions (Rajput et al., 2020, Thm. 2), while Lu et al. (2022, Thm. 4) implies a similar but dimension-dependent $\tilde{O}(\frac{d}{n^2 K^2})$ rate for SBW reordering. Thm. 3 matches the minimax lower bound of Cha et al. (2023, Thm. 4.5) up to the ϵ -rank parameter and shows that dimension dependence can be avoided when the gradient update matrices $[\mathbf{x}_1^k, \dots, \mathbf{x}_n^k]$ are low-rank, or, more generally, $\epsilon = O(\frac{\log(Kn)}{\sqrt{n}})$ -approximable by low-rank matrices.

5.2. Bridging the theory-practice gap

Two criticisms levied by Lu et al. (2022) against the SBW algorithm were the need to estimate the maximum Euclidean norm of any possible gradient vector in advance and the need to tune its free hyperparameter. $\text{LKH}(\frac{1}{2K})$ has neither of these drawbacks as it automatically adapts

to the scale of each input and has no hyperparameters to tune. Moreover, with a linear kernel, $\text{LKH}(\frac{1}{2K})$ can be run online in $O(nd)$ time. Hence, $\text{LKH}(\frac{1}{2K})$ is a promising substitute for the greedy thinning of Lu et al. (2022); Cooper et al. (2023). Indeed, when we recreate the Home Mortgage Disclosure Act logistic regression experiment of Cooper et al. (2023) with a single worker (Fig. 1), we find that LKH-SGD strongly outperforms the standard practice of random reshuffling (RR) and the theoretically justified but overly conservative CD-GraB: SBW variant. In addition, LKH-SGD matches the state-of-the-art test accuracy of CD-GraB: Greedy and lags only slightly in terms of training convergence.

The accelerated convergence rate of LKH-SGD over the standard slow SGD rate of RR provides a direct verification of the Thm. 3 guarantee, and we further verify in Fig. L.1 that the singular values of the gradient update matrices drop off steeply, resulting in relatively small ϵ_k -ranks (see Fig. L.1). See <https://github.com/microsoft/khsgd> for PyTorch code replicating this experiment and App. L.2 for supplementary experiment details.

6. Cheap Two-Sample Testing

Our final application is two-sample testing, determining whether two datasets are drawn from the same underlying distribution. We observe independent samples $\mathcal{X} \triangleq (\mathbf{x}_i)_{i=1}^m$ and $\mathcal{Y} \triangleq (\mathbf{y}_j)_{j=1}^n$ from the unknown distributions \mathbb{P} and \mathbb{Q} respectively, and we seek to accept or reject the null hypothesis that $\mathbb{P} = \mathbb{Q}$. Standard kernel MMD tests tackle this task by computing the empirical MMD

$$\text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{Q}_{\text{in}}) \text{ for } \mathbb{P}_{\text{in}}, \mathbb{Q}_{\text{in}} \triangleq \frac{1}{m} \sum_{\mathbf{x} \in \mathcal{X}} \delta_{\mathbf{x}}, \frac{1}{n} \sum_{\mathbf{y} \in \mathcal{Y}} \delta_{\mathbf{y}}$$

for an appropriate kernel \mathbf{k} and rejecting the null hypothesis whenever $\text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{Q}_{\text{in}})$ is sufficiently large (Gretton et al., 2012). Such tests are prized both for their broad

applicability and for their high discriminating power, that is, their probability of rejecting the null when $\mathbb{P} \neq \mathbb{Q}$. A standard way to summarize the power properties of a test is through its *detectable separation rate*.

Definition 5 (Detectable separation rate). *We say a two-sample test has detectable separation rate $\epsilon_{\mathbf{k},m,n}$ if, for any detection probability $1 - \beta \in (0, 1)$, there exists a constant $c_{\mathbf{k},\beta} > 0$ such that the test has power at least $1 - \beta$ of rejecting the null whenever $\text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{Q}) \geq c_{\mathbf{k},\beta} \cdot \epsilon_{\mathbf{k},m,n}$.*

Standard MMD tests can detect distributional differences on the order of $\epsilon_{\mathbf{k},m,n} = \frac{1}{\sqrt{\min(m,n)}}$ (Gretton et al., 2012, Cor. 9), and this detectable separation rate is known to be the best possible for MMD tests (Domingo-Enrich et al., 2023, Prop. 2) and minimax optimal for translation invariant kernels (Kim & Schrab, 2023, Thm. 8). However, standard MMD tests also suffer from the $\Theta((m+n)^2)$ time burden of computing the empirical MMD. Recently, Domingo-Enrich et al. (2023) showed that one can improve scalability while preserving power by compressing \mathbb{P}_{in} and \mathbb{Q}_{in} using a high-quality thinning algorithm. However, their analysis applies only to a restricted class of distributions and kernels and exhibits a pessimistic dimension dependence on \mathbb{R}^d . Here, we offer a new analysis of their Compress Then Test approach that applies to any bounded kernel on any domain and, as an application, develop the first non-asymptotic power guarantees for testing with learned deep neural network kernels.

6.1. Low-rank analysis of Compress Then Test

Algorithm 3: Compress Then Test (CTT)

Input: Samples $(\mathcal{X}, \mathcal{Y})$, # coresets s , compression level \mathbf{g} , kernel \mathbf{k} , failure probability δ , # replicates \mathcal{B} , level α

Partition \mathcal{X} into $s_m = \frac{sm}{m+n}$ equal-sized bins $(\mathcal{X}^{(i)})_{i=1}^{s_m}$

Partition \mathcal{Y} into $s_n = \frac{sn}{m+n}$ equal-sized bins $(\mathcal{Y}^{(i)})_{i=1}^{s_n}$

// Identify coreset of size $n_{\text{out}} = 2^{\mathbf{g}} \sqrt{\frac{m+n}{s}}$ for each bin

for $i = 1, \dots, s_m$ **do** $\mathbb{P}_{\text{out}}^{(i)} \leftarrow \text{KT-COMPRESS}(\delta)(\mathcal{X}^{(i)}, \mathbf{g}, \mathbf{k})$

for $i = 1, \dots, s_n$ **do** $\mathbb{Q}_{\text{out}}^{(i)} \leftarrow \text{KT-COMPRESS}(\delta)(\mathcal{Y}^{(i)}, \mathbf{g}, \mathbf{k})$

// Compute CORESETMMD test statistic

$M_{\mathcal{B}+1} \leftarrow \text{MMD}_{\mathbf{k}}(\frac{1}{s_m} \sum_{i=1}^{s_m} \mathbb{P}_{\text{out}}^{(i)}, \frac{1}{s_n} \sum_{i=1}^{s_n} \mathbb{Q}_{\text{out}}^{(i)})$ (11)

// Simulate null by randomly permuting the s coresets \mathcal{B} times

for $b = 1, \dots, \mathcal{B}$ **do**

$(\mathbb{P}_{\text{out},b}^{(i)})_{i=1}^{s_m}, (\mathbb{Q}_{\text{out},b}^{(i)})_{i=1}^{s_n} \leftarrow \text{PERMUTE}((\mathbb{P}_{\text{out}}^{(i)})_{i=1}^{s_m}, (\mathbb{Q}_{\text{out}}^{(i)})_{i=1}^{s_n})$

$M_b \leftarrow \text{MMD}_{\mathbf{k}}(\frac{1}{s_m} \sum_{i=1}^{s_m} \mathbb{P}_{\text{out},b}^{(i)}, \frac{1}{s_n} \sum_{i=1}^{s_n} \mathbb{Q}_{\text{out},b}^{(i)})$

end

// Threshold test statistic

$R \leftarrow$ position of $M_{\mathcal{B}+1}$ in an increasing ordering of $(M_b)_{b=1}^{\mathcal{B}+1}$ with ties broken uniformly at random

return Reject with prob. $\min(1, \max(0, R - (1 - \alpha)(\mathcal{B} + 1)))$

Alg. 3 details the Compress Then Test (CTT) approach of Domingo-Enrich et al. (2023, Alg. 1). Given a coreset count $s \geq 2$, a compression level $\mathbf{g} \geq 0$, and a nominal level $\alpha \in (0, 1)$, CTT divides \mathcal{X} and \mathcal{Y} into data-point bins of size $n_{\text{in}} \triangleq \frac{m+n}{s}$, thins each bin down to size $n_{\text{out}} \triangleq 2^{\mathbf{g}} \sqrt{n_{\text{in}}}$ using KT-COMPRESS(δ) (a refinement of KH-COMPRESS(δ) detailed in App. H), and uses the thinned coresets to cheaply approximate $\text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{Q}_{\text{in}})$ and permuted versions thereof. Domingo-Enrich et al. (2023, (8)) showed that the total runtime of CTT is dominated by

$$O(4^{\mathbf{g}}(m+n)(s + \log_4(\frac{m+n}{s} - \mathbf{g}))),$$

kernel evaluations, yielding a near-linear $O((m+n) \log^c(m+n))$ time algorithm whenever $s = O(\log_4(m+n))$ and $\mathbf{g} \leq c \log_4 \log(m+n)$. Moreover, Prop. 1 of Domingo-Enrich et al. (2023) ensures that CTT has probability at most α of falsely rejecting the null hypothesis.

Our next, complementary result shows that CTT also matches the detectable separation rate of standard MMD tests up to an inflation factor $\mathbf{R}_{\mathbf{k}}/2^{\mathbf{g}}$ depending on the compression level \mathbf{g} .

Theorem 4 (Low-rank analysis of CTT power). *Suppose the parameters of CTT (Alg. 3) satisfy $m \leq n$,*

$$s_m \geq \frac{32}{9} \log(\frac{2e}{\gamma}), \quad \text{and} \quad \delta = \min(\frac{\tilde{\beta}}{6}, (\frac{\tilde{\beta}}{2})^{1/\lfloor \alpha(\mathcal{B}+1) \rfloor} \frac{\alpha}{30e s})$$

for $\tilde{\beta} \triangleq \frac{\beta}{1+\beta/2}$ and $\gamma \triangleq \frac{\alpha}{4e} (\frac{\tilde{\beta}}{4})^{1/\lfloor \alpha(\mathcal{B}+1) \rfloor}$. Then CTT has detectable separation rate (Def. 5)

$$\epsilon_{\mathbf{k},m,n} = (1 + \mathbf{R}_{\mathbf{k}}/2^{\mathbf{g}})/\sqrt{m},$$

where $\mathbf{R}_{\mathbf{k}}^2$ denotes the $(1 - \frac{\tilde{\beta}}{20s_n})$ -th quantile of

$$\hat{\mathbf{R}}_{\mathbf{k}}^2 \triangleq \log(\frac{m+n}{s}) \log(\frac{n}{\beta}). \quad (12)$$

$$\min_{r \leq 2n_{\text{out}}} \{ \|\mathbf{k}\|_{\infty} r \log(\frac{n}{\beta}) + (\lambda_{r+1}(\mathbf{K}) + \lambda_{r+1}(\mathbf{K}')) n_{\text{out}} \}.$$

for $\mathbf{K} \triangleq (\mathbf{k}(x_i, x_j))_{i,j=1}^m$, $\mathbf{K}' \triangleq (\mathbf{k}(y_i, y_j))_{i,j=1}^n$, and $\|\mathbf{k}\|_{\infty} \triangleq \sup_{x,y \in \text{supp}(\mathbb{P}+\mathbb{Q})} |\mathbf{k}(x, y)|$.

The proof in App. I combines the low-rank sub-Gaussian error bounds of Thm. 1 with the generic compressed power analysis of Domingo-Enrich et al. (2023, App. B.1) to yield power guarantees for bounded kernels on any domain. Notably, when $\text{rank}(\mathbf{K})$ and $\text{rank}(\mathbf{K}')$ are bounded or, more generally, $\text{polylog}(n)$ one can choose the compression level $\mathbf{g} = \Theta(\log_4 \log(m+n))$ to exactly match the optimal quadratic-time detectable separation rates with a near-linear time CTT test. Moreover, the inflation factors remain well-controlled whenever the induced kernel matrices exhibit rapid eigenvalue decay.

As a concrete example, consider the learned deep neural network kernel of Liu et al. (2020),

$$\mathbf{k}_{\text{deep}}(\mathbf{x}, \mathbf{y}) \triangleq [(1 - \epsilon)\kappa(\phi(\mathbf{x}), \phi(\mathbf{y})) + \epsilon]q(\mathbf{x}, \mathbf{y}), \quad (13)$$

where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^{d_{\text{embd}}}$ is a pretrained neural network, q and κ are GAUSS(η) kernels (5) on \mathbb{R}^d and $\mathbb{R}^{d_{\text{embd}}}$ respectively, and $\epsilon \in (0, 1)$. This deep kernel generates full-rank kernel matrices (Liu et al., 2020, Prop. 5) but induces exponential eigenvalue decay due to its decomposition as a mixture of Gaussian kernels. Hence, as we show in App. J, CTT with \mathbf{k}_{deep} , $\mathbf{g} = \Theta(\log_4 \log(m + n))$, and sub-Gaussian inputs matches the detection quality of a quadratic-time MMD test in near-linear time.

Corollary 3 (Power of deep kernel CTT). *Instantiate the assumptions of Thm. 4 with $\mathbf{k} = \mathbf{k}_{\text{deep}}$ (13). If the inputs $(\phi(\mathbf{x}_1), \mathbf{x}_1, \phi(\mathbf{y}_1), \mathbf{y}_1)$ are sub-Gaussian, that is,*

$$\mathbb{E}[e^{c\|(\phi(\mathbf{x}_1), \mathbf{x}_1, \phi(\mathbf{y}_1), \mathbf{y}_1)\|_2^2}] < \infty \quad (14)$$

for some $c > 0$, then CTT satisfies the conclusions of Thm. 4 with $d' \triangleq d_{\text{embd}} + d$ and

$$\mathbf{R}_{\mathbf{k}_{\text{deep}}} = O(\log^{\frac{d'}{2} + \frac{3}{2}}(\frac{n}{\beta})).$$

Moreover, when the input and neural features lie on smooth compact manifolds (as, e.g., in Zhu et al., 2018), the error inflation of CTT adapts to the smaller intrinsic manifold dimension, enabling an improved trade-off between runtime and detection power. See App. K for our proof.

Corollary 4 (Power of deep manifold kernel CTT). *Under the assumptions of Cor. 3, if $\mathbf{x}_1, \mathbf{y}_1, (\mathbf{x}_1, \phi(\mathbf{x}_1))$, and $(\mathbf{y}_1, \phi(\mathbf{y}_1))$ belong to smooth compact manifolds (Assump. E.1) with dimension $d^* < d'$ then CTT satisfies the conclusions of Thm. 4 with*

$$\mathbf{R}_{\mathbf{k}_{\text{deep}}} = O(\log^{\frac{5d^*}{4} + \frac{3}{2}}(\frac{n}{\beta})).$$

Cors. 3 and 4 follow from explicitly bounding the eigenvalues of the generated deep kernel matrices as in (6) and (7). By Kim & Schrab (2023, Thm. 8), the separation rate of Thm. 4 is minimax optimal up to the inflation factor $\mathbf{R}_{\mathbf{k}}/2^g$ and, hence, those of Cors. 3 and 4 are minimax optimal up to log factors.

One could alternatively bound the compression error of KT-COMPRESS(δ) using the covering number approach of Dwivedi & Mackey (2022, Thm. 2, Prop. 3). In the setting of Cor. 3, the argument of App. J combined with this distinct analysis would yield an alternative error inflation factor $\tilde{\mathbf{R}}_{\mathbf{k}_{\text{deep}}}/2^g$ with worse dimension dependence,

$$\tilde{\mathbf{R}}_{\mathbf{k}_{\text{deep}}} = \Theta(\log^{\frac{3d'}{4} + 2}(\frac{n}{\beta})),$$

and without known adaptivity to an intrinsic manifold dimension.




Figure 2: **Time-power trade-off curves for detecting Higgs bosons with deep kernel MMD tests.** We plot mean values ± 1 standard error across 1000 independent trials with level $\alpha = 0.05$ and $B = 100$ permutations.

6.2. Powerful deep kernel testing in near-linear time

To evaluate the practical utility of deep kernel CTT, we follow the Higgs mixture experiment of Domingo-Enrich et al. (2023, Sec. 5) and use the deep kernel training procedure of Liu et al. (2020, Tab. 1). Here, the aim is to distinguish a Higgs boson signal process \mathbb{P} from a background process \mathbb{Q} given $m = n = 16384$ observations, $d = 2$ particle-detector features, and a five-layer fully-connected neural network ϕ with softplus activations and embedding dimension $d_{\text{embd}} = 20$.

Fig. 2 compares the time-power trade-off curves induced by three fast kernel testing approaches to this problem: SUBSAMPLING, a standard wild-bootstrap MMD test (Chwialkowski et al., 2014) that simply evaluates empirical $\text{MMD}_{\mathbf{k}_{\text{deep}}}$ using $n_{\text{out}} = m_{\text{out}}$ uniformly subsampled points; W-BLOCK, a wild-bootstrap test that averages $\frac{n}{B}$ subsampled squared $\text{MMD}_{\mathbf{k}_{\text{deep}}}$ estimates based on $n_{\text{out}} = m_{\text{out}} = B$ points (Zaremba et al., 2013); and CTT with $s = 32$ bins and varying g . We find that the CTT curve dominates that of the alternative methods and matches the power of an exact MMD test (SUBSAMPLING with $n_{\text{out}} = n$) in a fraction of the time. The improvements of CTT over the standard power-runtime trade-off of SUBSAMPLING provides a direct verification of the Thm. 4 guarantee, and we additionally verify in Fig. L.2 that the empirical inflation factor $\hat{\mathbf{R}}_{\mathbf{k}_{\text{deep}}} = O(\log^5(n))$ in this setting due to approximate low-rankness. See <https://github.com/microsoft/deepctt> for PyTorch code replicating this experiment and App. L.3 for supplementary experiment details.

Impact Statement

This work introduced a new analysis of thinning algorithms that adapts to low-rank structures. We exploited this adaptivity to design fast algorithms with strong quality guarantees for three key applications in machine learning: dot-product attention in Transformers, stochastic gradient training in optimization, and deep kernel testing for distinguishing distributions. More broadly, our techniques provide a general framework for reducing computational resource use in machine learning. Such tools have the potential to reduce energy costs and environmental harms from model training, inference, and evaluation and to improve accessibility in resource-constrained settings, all while provably maintaining high quality.

Acknowledgments We thank Insu Han, A. Feder Cooper, and Wentao Guo for their assistance with their code bases and datasets.

References

- Alman, J. and Song, Z. Fast attention requires bounded entries. *Advances in Neural Information Processing Systems*, 36, 2024. (Cited on page 4.)
- Altschuler, J., Bach, F., Rudi, A., and Niles-Weed, J. Massively scalable sinkhorn distances via the nystrom method. *Advances in Neural Information Processing Systems*, 32, 2019. (Cited on pages 4, 29, and 30.)
- Alweiss, R., Liu, Y. P., and Sawhney, M. Discrepancy minimization via a self-balancing walk. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 14–20, 2021. (Cited on pages 1 and 6.)
- Bansal, N., Dadush, D., Garg, S., and Lovett, S. The gram-schmidt walk: a cure for the banaszczyk blues. In *Proceedings of the 50th annual acm sigact symposium on theory of computing*, pp. 587–597, 2018. (Cited on pages 3, 19, and 22.)
- Cha, J., Lee, J., and Yun, C. Tighter lower bounds for shuffling sgd: Random permutations and beyond. In *International Conference on Machine Learning*, pp. 3855–3912. PMLR, 2023. (Cited on pages 6 and 7.)
- Chen, B., Dao, T., Winsor, E., Song, Z., Rudra, A., and Ré, C. Scatterbrain: unifying sparse and low-rank attention approximation. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NeurIPS ’21, Red Hook, NY, USA, 2021. Curran Associates Inc. ISBN 97817113845393. (Cited on pages 4 and 5.)
- Choromanski, K. M., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J. Q., Mohiuddin, A., Kaiser, L., Belanger, D. B., Colwell, L. J., and Weller, A. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ua6zuk0WRH>. (Cited on pages 4 and 5.)
- Chwialkowski, K. P., Sejdinovic, D., and Gretton, A. A wild bootstrap for degenerate kernel tests. *Advances in Neural Information Processing Systems*, 27, 2014. (Cited on page 9.)
- Cooper, A. F., Guo, W., Pham, K., Yuan, T., Ruan, C. F., Lu, Y., and De Sa, C. Coordinating distributed example orders for provably accelerated training. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. (Cited on pages 6, 7, 32, 33, and 37.)
- Dao, T. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=mZn2Xyh9Ec>. (Cited on page 5.)
- Dao, T., Fu, D., Ermon, S., Rudra, A., and Ré, C. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022. (Cited on page 5.)
- Domingo-Enrich, C., Dwivedi, R., and Mackey, L. Compress then test: Powerful kernel testing in near-linear time. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research. PMLR, 25–27 Apr 2023. (Cited on pages 8, 9, and 34.)
- Dong, L., Xu, S., and Xu, B. Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5884–5888, 2018. doi: 10.1109/ICASSP.2018.8462506. (Cited on page 4.)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>. (Cited on page 4.)
- Dwivedi, R. and Mackey, L. Generalized kernel thinning. In *International Conference on Learning Representations*, 2022. (Cited on pages 1, 4, and 9.)
- Dwivedi, R. and Mackey, L. Kernel thinning. *Journal of Machine Learning Research*, 25(152):1–77, 2024. (Cited on pages 1, 3, 4, 15, 16, 17, 18, 33, 34, and 36.)
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. (Cited on pages 7 and 8.)
- Han, I., Jayaram, R., Karbasi, A., Mirrokni, V., Woodruff, D., and Zandieh, A. Hyperattention: Long-context attention in near-linear time. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Eh0Od2BJIM>. (Cited on pages 4 and 5.)
- Harshaw, C., Sävje, F., Spielman, D. A., and Zhang, P. Balancing covariates in randomized experiments with the gram-schmidt walk design. *Journal of the American Statistical Association*, pp. 1–13, 2024. (Cited on page 22.)
- Harvey, N. and Samadi, S. Near-Optimal Herding. In *Proceedings of The 27th Conference on Learning Theory*, volume 35, pp. 1165–1182, 2014. (Cited on pages 1, 6, and 32.)

- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. (Cited on page 5.)
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding*, pp. 409–426, 1994. (Cited on page 15.)
- Horn, R. A. and Johnson, C. R. *Matrix Analysis*. Cambridge University Press, 1985. (Cited on page 36.)
- Kim, I. and Schrab, A. Differentially private permutation tests: Applications to kernel methods. *arXiv preprint arXiv:2310.19043*, 2023. (Cited on pages 8 and 9.)
- Kitaev, N., Kaiser, L., and Levskaya, A. Reformer: The efficient transformer. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgNkKhtvB>. (Cited on pages 4 and 5.)
- Li, L., Dwivedi, R., and Mackey, L. Debiased distribution compression. In *Proceedings of the 41st International Conference on Machine Learning*, volume 203 of *Proceedings of Machine Learning Research*. PMLR, 21–27 Jul 2024. (Cited on page 1.)
- Liu, F., Xu, W., Lu, J., Zhang, G., Gretton, A., and Sutherland, D. J. Learning deep kernels for non-parametric two-sample tests. In *International conference on machine learning*, pp. 6316–6326. PMLR, 2020. (Cited on pages 9 and 37.)
- Lu, Y., Guo, W., and De Sa, C. M. Grab: Finding provably better data permutations than random reshuffling. *Advances in Neural Information Processing Systems*, 35:8969–8981, 2022. (Cited on pages 6 and 7.)
- Markov, A. On certain applications of algebraic continued fractions. *Unpublished Ph. D. thesis, St Petersburg*, 1884. (Cited on pages 26 and 27.)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019. (Cited on page 36.)
- Phillips, J. M. and Tai, W. M. Near-optimal coresets of kernel density estimates. *Discrete & Computational Geometry*, 63 (4):867–887, 2020. (Cited on pages 1 and 3.)
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems*, 20, 2007. (Cited on page 3.)
- Rajput, S., Gupta, A., and Papailiopoulos, D. Closing the convergence gap of sgd without replacement. In *International Conference on Machine Learning*, pp. 7964–7973. PMLR, 2020. (Cited on page 7.)
- Rudin, W. *Functional Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1991. ISBN 9780070542365. URL https://books.google.com/books?id=Sh_vAAAAMAAJ. (Cited on page 14.)
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y. (Cited on page 5.)
- Saadetoglu, M. and Dinsev, S. M. Inverses and determinants of $n \times n$ block matrices. *Mathematics*, 11(17), 2023. ISSN 2227-7390. doi: 10.3390/math11173784. URL <https://www.mdpi.com/2227-7390/11/17/3784>. (Cited on page 23.)
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016. (Cited on page 5.)
- Sherman, J. and Morrison, W. J. Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix. *The Annals of Mathematical Statistics*, 21(1):124 – 127, 1950. doi: 10.1214/aoms/1177729893. URL <https://doi.org/10.1214/aoms/1177729893>. (Cited on page 23.)
- Shetty, A., Dwivedi, R., and Mackey, L. Distribution compression in near-linear time. In *International Conference on Learning Representations*, 2022. (Cited on pages 1, 3, 18, 19, 24, 34, and 35.)
- Steinwart, I. and Christmann, A. Support vector machines. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1, 2008. URL <https://api.semanticscholar.org/CorpusID:661123>. (Cited on page 2.)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964. (Cited on page 4.)
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019. (Cited on pages 25, 28, and 29.)
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.-H., Tay, F. E., Feng, J., and Yan, S. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 558–567, 2021. (Cited on page 5.)
- Zandieh, A., Han, I., Daliri, M., and Karbasi, A. Kdeformer: Accelerating transformers via kernel density estimation. In *International Conference on Machine Learning*, pp. 40605–40623. PMLR, 2023. (Cited on pages 4, 5, and 36.)
- Zaremba, W., Gretton, A., and Blaschko, M. B-test: A non-parametric, low variance kernel two-sample test. *Advances in Neural Information Processing Systems*, 26, 2013. (Cited on page 9.)
- Zhu, W., Qiu, Q., Huang, J., Calderbank, R., Sapiro, G., and Daubechies, I. Ldmnet: Low dimensional manifold regularized neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2743–2751, 2018. (Cited on page 9.)

Appendix Contents

A	Appendix Notation and Definitions	13
B	Proof of Tab. 1: Sub-Gaussian Thinning Examples	14
B.1	SUBSAMPLING	14
B.2	KH(δ)	15
B.3	LKH(δ)	16
B.4	RKH(δ)	17
B.5	KH-COMPRESS(δ)	18
B.6	GS-THIN	19
B.7	GS-COMPRESS	24
C	Proof of Thm. 1: Low-rank sub-Gaussian thinning	24
C.1	Proof of MMD bound (2)	25
C.2	Proof of kernel max seminorm bound (3)	26
C.3	Proof of Lipschitz kernel max seminorm bound (4)	27
C.4	Proof of Lem. C.4: Bounded-Hölder sub-Gaussian process	29
D	Proof of Cor. 1: Gaussian MMD of KH	29
E	Proof of Cor. 2: Intrinsic Gaussian MMD of KH	30
F	Proof of Thm. 2: Quality of Thinformer	30
F.1	Proof of Lem. F.1: Decomposing attention approximation error	31
F.2	Proof of Lem. F.2: KMS bound on attention approximation error	31
F.3	Proof of Lem. F.3: Thinned attention problem parameters	31
G	Proof of Thm. 3: LKH-SGD convergence	32
H	KT-COMPRESS(δ)	33
I	Proof of Thm. 4: Low-rank analysis of CTT power	34
I.1	Proof of Thm. I.1: Low-rank analysis of CTT power, detailed	34
I.2	Proof of Lem. I.1: (K, ν, δ) -sub-Gaussian thinning algorithms are k-sub-Gaussian	35
J	Proof of Cor. 3: Power of deep kernel CTT	35
K	Proof of Cor. 4: Power of deep manifold kernel CTT	36
L	Supplementary Experiment Details	36

L.1	Approximating attention experiments	36
L.2	Faster SGD training experiments	37
L.3	Cheap two-sample testing experiment	37

A. Appendix Notation and Definitions

We often use the shorthand $(a)_+ \triangleq \max(a, 0)$ as well as the shorthand $\mathbf{k}(\mathcal{X}, \mathcal{X})$ to represent the matrix $(\mathbf{k}(\mathbf{x}_i, \mathbf{x}_j))_{i,j=1}^n$. In addition, for each kernel \mathbf{k} , we let $\mathcal{H}_{\mathbf{k}}$ and $\|\cdot\|_{\mathbf{k}}$ represent the associated reproducing kernel Hilbert space (RKHS) and RKHS norm, so that $\mathbb{B}_{\mathbf{k}} = \{f \in \mathcal{H}_{\mathbf{k}} : \|f\|_{\mathbf{k}} \leq 1\}$ and define

$$(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \triangleq \frac{1}{n_{\text{in}}} \sum_{x \in \mathcal{X}_{\text{in}}} \mathbf{k}(x, \cdot) - \frac{1}{n_{\text{out}}} \sum_{x \in \mathcal{X}_{\text{out}}} \mathbf{k}(x, \cdot).$$

We also relate our definition of a sub-Gaussian thinning algorithm (Def. 3) to several useful notions of sub-Gaussianity.

Definition A.1 (Sub-Gaussian vector). We say that a random vector $\mathbf{w} \in \mathbb{R}^n$ is (\mathbf{K}, ν) -sub-Gaussian on an event \mathcal{E} if \mathbf{K} is SPSPD and $\nu > 0$ satisfies

$$\mathbb{E}_{\mathcal{E}}[\exp(\mathbf{u}^\top \mathbf{K} \mathbf{w})] \leq \exp\left(\frac{\nu^2}{2} \cdot \mathbf{u}^\top \mathbf{K} \mathbf{u}\right) \quad \text{for all } \mathbf{u} \in \mathbb{R}^n. \quad (15)$$

If, in addition, the event has probability 1, we say that \mathbf{w} is (\mathbf{K}, ν) -sub-Gaussian.

Notably, a thinning algorithm is $(\mathbf{K}, \nu, \delta)$ -sub-Gaussian if and only if its associated vector $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on an event \mathcal{E} of probability at least $1 - \delta/2$.

Definition A.2 (Sub-Gaussian function). For a kernel \mathbf{k} , we say that a random function $\phi \in \mathcal{H}_{\mathbf{k}}$ is (\mathbf{k}, ν) -sub-Gaussian on an event \mathcal{E} if $\nu > 0$ satisfies

$$\mathbb{E}_{\mathcal{E}}[\exp(\langle f, \phi \rangle_{\mathbf{k}})] \leq \exp\left(\frac{\nu^2}{2} \cdot \|f\|_{\mathbf{k}}^2\right) \quad \text{for all } f \in \mathcal{H}_{\mathbf{k}}. \quad (16)$$

If, in addition, the event has probability 1, we say that ϕ is (\mathbf{k}, ν) -sub-Gaussian.

Our next two lemmas show that for finitely-supported signed measures like $\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}$, this notion of functional sub-Gaussianity is equivalent to the prior notion of vector sub-Gaussianity, allowing us to use the two notions interchangeably. Hereafter, we say that \mathbf{k} generates a SPSPD matrix \mathbf{K} if $\mathbf{k}(\mathcal{X}, \mathcal{X}) = \mathbf{K}$.

Lemma A.1 (Functional sub-Gaussianity implies vector sub-Gaussianity). In the notation of Def. 3, if $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on an event \mathcal{E} and \mathbf{k} generates \mathbf{K} , then the vector $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on \mathcal{E} .

Proof. Suppose $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on an event \mathcal{E} , fix a vector $\mathbf{u} \in \mathbb{R}^n$, and define the function

$$f_{\mathbf{u}} \triangleq \sum_{i=1}^n u_i \mathbf{k}(\cdot, x_i) \in \mathcal{H}_{\mathbf{k}}.$$

By the reproducing property,

$$\mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}) = \langle f_{\mathbf{u}}, (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \rangle_{\mathbf{k}} \quad \text{and} \quad \|f_{\mathbf{u}}\|_{\mathbf{k}}^2 = \mathbf{u}^\top \mathbf{K} \mathbf{u}. \quad (17)$$

Invoking the representations (17) and the functional sub-Gaussianity condition (16) we therefore obtain

$$\mathbb{E}_{\mathcal{E}}[\exp(\mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))] = \mathbb{E}_{\mathcal{E}}[\exp(\langle f_{\mathbf{u}}, (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \rangle_{\mathbf{k}})] \leq \exp(\|f_{\mathbf{u}}\|_{\mathbf{k}}^2 \cdot \frac{\nu^2}{2}) = \exp(\mathbf{u}^\top \mathbf{K} \mathbf{u} \cdot \frac{\nu^2}{2}),$$

so that $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on the event \mathcal{E} as claimed. \square

Lemma A.2 (Vector sub-Gaussianity implies functional sub-Gaussianity). In the notation of Def. 3, if $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on an event \mathcal{E} and \mathbf{k} generates \mathbf{K} , then $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on \mathcal{E} .

Proof. Suppose $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on an event \mathcal{E} , fix a function $f \in \mathcal{H}_{\mathbf{k}}$, and consider the set

$$\mathcal{L} \triangleq \left\{ f_{\mathbf{u}} \triangleq \sum_{i=1}^n u_i \mathbf{k}(\cdot, x_i) : \mathbf{u} \in \mathbb{R}^n \right\}.$$

Since \mathcal{L} is a closed linear subspace of $\mathcal{H}_{\mathbf{k}}$, we can decompose f as $f = f_{\mathbf{u}} + f_{\perp}$, where $\mathbf{u} \in \mathbb{R}^n$ and f_{\perp} is orthogonal to \mathcal{L} (Rudin, 1991, Theorem 12.4), so that

$$\|f\|_{\mathbf{k}}^2 = \|f_{\mathbf{u}}\|_{\mathbf{k}}^2 + \|f_{\perp}\|_{\mathbf{k}}^2 \quad \text{and} \quad \|f_{\mathbf{u}}\|_{\mathbf{k}}^2 = \mathbf{u}^{\top} \mathbf{K} \mathbf{u}. \quad (18)$$

Invoking the orthogonality of f_{\perp} and $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \in \mathcal{L}$, the reproducing property representations (17), and the vector sub-Gaussianity condition (15), we find that

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\exp(\langle f, (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \rangle_{\mathbf{k}})] &= \mathbb{E}_{\mathcal{E}}[\exp(\langle f_{\mathbf{u}} + f_{\perp}, (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} \rangle_{\mathbf{k}})] = \mathbb{E}_{\mathcal{E}}[\exp(\mathbf{u}^{\top} \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))] \\ &\leq \exp(\mathbf{u}^{\top} \mathbf{K} \mathbf{u} \cdot \frac{\nu^2}{2}) \stackrel{(18)}{\leq} \exp(\|f\|_{\mathbf{k}}^2 \cdot \frac{\nu^2}{2}), \end{aligned}$$

so that $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on the event \mathcal{E} as claimed. \square

We end our discussion about the versions of sub-Gaussianity considered above by presenting the standard fact about the additivity of sub-Gaussianity parameters under summation of independent sub-Gaussian random vectors, adapted to our setting.

Lemma A.3 (Vector sub-Gaussian additivity). *Suppose that, for each $j \in [m]$, $\Delta_j \in \mathbb{R}^n$ is (\mathbf{K}, ν_j) on an event \mathcal{E}_j given $\Delta_{1:(j-1)} \triangleq (\Delta_1, \dots, \Delta_{j-1})$ and $\mathcal{E}_{\leq j-1} \triangleq \bigcap_{i=1}^{j-1} \mathcal{E}_i$. Then $\sum_{j=1}^m \Delta_j$ is $(\mathbf{K}, (\sum_{j=1}^m \nu_j^2)^{1/2})$ -sub-Gaussian on $\mathcal{E}_{\leq m}$.*

Proof. Let $\mathcal{E}_{\leq s} = \bigcap_{j=1}^s \mathcal{E}_j$ for each $s \in [m]$. We prove the result for $\mathcal{Z}_s = \sum_{j=1}^s \Delta_j$ by induction on $s \in [m]$. The result holds for the base case of $s = 1$ by assumption. For the inductive case, suppose the result holds for $s \in [m-1]$. Fixing $\mathbf{u} \in \mathbb{R}^n$, we may apply the tower property, our conditional sub-Gaussianity assumption, and our inductive hypothesis in turn to conclude

$$\begin{aligned} \mathbb{E}[\exp(\langle \mathbf{u}, \mathbf{K} \sum_{j=1}^{s+1} \Delta_j \rangle) \mathbf{1}[\mathcal{E}_{\leq s+1}]] &= \mathbb{E}[\exp(\langle \mathbf{u}, \mathbf{K} \sum_{j=1}^s \Delta_j \rangle) \mathbf{1}[\mathcal{E}_{\leq s}] \mathbb{E}[\exp(\langle \mathbf{u}, \Delta_{s+1} \rangle) \mathbf{1}[\mathcal{E}_{s+1}] \mid \Delta_{1:s}, \mathcal{E}_{\leq s}]] \\ &\leq \mathbb{E}[\exp(\langle \mathbf{u}, \mathbf{K} \sum_{j=1}^s \Delta_j \rangle) \mathbf{1}[\mathcal{E}_{\leq s}]] \exp\left(\frac{\nu_{s+1}^2}{2} \cdot \mathbf{u}^{\top} \mathbf{K} \mathbf{u}\right) \leq \exp\left(\frac{\sum_{j=1}^{s+1} \nu_j^2}{2} \cdot \mathbf{u}^{\top} \mathbf{K} \mathbf{u}\right). \end{aligned}$$

Hence, \mathcal{Z}_{s+1} is $(\mathbf{K}, (\sum_{j=1}^{s+1} \nu_j^2)^{1/2})$ -sub-Gaussian on $\mathcal{E}_{\leq s+1}$, and the proof is complete. \square

B. Proof of Tab. 1: Sub-Gaussian Thinning Examples

This section provides supplementary details for each of the sub-Gaussian thinning algorithms of Tab. 1.

B.1. SUBSAMPLING

B.1.1. PROOF OF PROP. 1: QUALITY OF UNIFORM SUBSAMPLING

We begin by computing the first and second moments of \mathbf{p}_{out} : $\mathbb{E}[\mathbf{p}_{\text{out}}] = \mathbf{p}_{\text{in}}$ and

$$\mathbb{E}[\mathbf{p}_{\text{out}} \mathbf{p}_{\text{out}}^{\top}] = \frac{1}{n_{\text{out}}} \text{diag}(\mathbf{p}_{\text{in}}) + \frac{n_{\text{in}}(n_{\text{out}}-1)}{n_{\text{out}}(n_{\text{in}}-1)} (\mathbf{p}_{\text{in}} \mathbf{p}_{\text{in}}^{\top} - \frac{1}{n_{\text{in}}} \text{diag}(\mathbf{p}_{\text{in}})) = \frac{1}{n_{\text{out}}} \left(\frac{n_{\text{in}}-n_{\text{out}}}{n_{\text{in}}-1} \right) \text{diag}(\mathbf{p}_{\text{in}}) + \frac{n_{\text{in}}(n_{\text{out}}-1)}{n_{\text{out}}(n_{\text{in}}-1)} \mathbf{p}_{\text{in}} \mathbf{p}_{\text{in}}^{\top}.$$

Hence,

$$\begin{aligned} \mathbb{E}[\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}})] &= \mathbf{p}_{\text{in}}^{\top} \mathbf{K} \mathbf{p}_{\text{in}} - 2 \mathbf{p}_{\text{in}}^{\top} \mathbf{K} \mathbb{E}[\mathbf{p}_{\text{out}}] + \mathbb{E}[\mathbf{p}_{\text{out}}^{\top} \mathbf{K} \mathbf{p}_{\text{out}}] = \text{tr}(\mathbf{K} \mathbb{E}[\mathbf{p}_{\text{out}} \mathbf{p}_{\text{out}}^{\top}]) - \mathbf{p}_{\text{in}}^{\top} \mathbf{K} \mathbf{p}_{\text{in}} \\ &= \frac{1}{n_{\text{out}}} \left(\frac{n_{\text{in}}-n_{\text{out}}}{n_{\text{in}}-1} \right) (\text{tr}(\mathbf{K} \text{diag}(\mathbf{p}_{\text{in}})) - \mathbf{p}_{\text{in}}^{\top} \mathbf{K} \mathbf{p}_{\text{in}}) = \frac{1}{n_{\text{out}}} \left(\frac{n_{\text{in}}-n_{\text{out}}}{n_{\text{in}}-1} \right) C_{\mathbf{K}}. \end{aligned} \quad (19)$$

To derive the second advertised result, we note that

$$\mathbb{E}[\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}}^2] \geq \max_{i \in \mathcal{I}} \mathbb{E}[(\mathbf{e}_i^{\top} \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))^2] = \max_{i \in \mathcal{I}} \mathbb{E}[\text{MMD}_{\mathbf{K} \mathbf{e}_i \mathbf{e}_i^{\top} \mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}})]$$

and invoke the initial result (19) to conclude.

B.1.2. SUB-GAUSSIANITY OF SUBSAMPLING

Proposition B.1 (Sub-Gaussianity of uniform subsampling). *For any SPSPD $\mathbf{K} \in \mathbb{R}^{n \times n}$, uniform subsampling (without replacement) is a $(\mathbf{K}, \nu, 0)$ -sub-Gaussian thinning algorithm with*

$$\nu \triangleq \frac{\sqrt{\|\mathbf{K}\|_{\max}}}{\sqrt{n_{\text{out}}}}.$$

Proof. Fix any vector $\mathbf{u} \in \mathbb{R}^n$, and let $J_1, \dots, J_{n_{\text{out}}}$ be the random indices in $[n]$ selected by uniform subsampling. Since $\mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}) = \frac{1}{n_{\text{out}}} \sum_{i=1}^{n_{\text{out}}} \mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{e}_{J_i})$ is an average of mean-centered scalars drawn without replacement and satisfying

$$|\mathbf{u}^\top \mathbf{K} \mathbf{e}_{J_i}| \leq \sqrt{\mathbf{u}^\top \mathbf{K} \mathbf{u}} \sqrt{\mathbf{e}_{J_i}^\top \mathbf{K} \mathbf{e}_{J_i}} \leq \sqrt{\|\mathbf{K}\|_{\max}} \sqrt{\mathbf{u}^\top \mathbf{K} \mathbf{u}} \quad \text{with probability 1}$$

by Cauchy-Schwarz, Thm. 4 and equations (1.8) and (4.16) of [Hoeffding \(1994\)](#) imply that

$$\mathbb{E}[\exp(\mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))] \leq \exp\left(\frac{\|\mathbf{K}\|_{\max}}{2n_{\text{out}}} \mathbf{u}^\top \mathbf{K} \mathbf{u}\right).$$

□

 B.2. KH(δ)

Algorithm B.1: KH(δ): Kernel Halving with simplified swapping thresholds and failure probability $\delta/2$

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$ with even n_{in} , kernel \mathbf{k}

```

 $\mathcal{S}^{(1)}, \mathcal{S}^{(2)} \leftarrow \{\}; \quad \tilde{\psi}_0 \leftarrow \mathbf{0} \in \mathcal{H}_{\mathbf{k}} \quad // \text{Initialize empty coresets: } \mathcal{S}^{(1)}, \mathcal{S}^{(2)} \text{ have size } i \text{ after round } i$ 
 $\mathbf{b}_{\max, i} \leftarrow 0 \quad // \text{Max function norm so far}$ 
for  $i = 1, 2, \dots, n_{\text{in}}/2$  do
     $// \text{Construct kernel difference function using next two points}$ 
     $(\mathbf{x}, \mathbf{x}') \leftarrow (\mathbf{x}_{2i-1}, \mathbf{x}_{2i}); \quad f_i \leftarrow \mathbf{k}(\mathbf{x}_{2i-1}, \cdot) - \mathbf{k}(\mathbf{x}_{2i}, \cdot); \quad \eta_i \leftarrow -1$ 
     $// \text{Compute swapping threshold } \alpha_i$ 
     $\mathbf{b}_i^2 = \|f_i\|_{\mathbf{k}}^2 = \mathbf{k}(\mathbf{x}, \mathbf{x}) + \mathbf{k}(\mathbf{x}', \mathbf{x}') - 2\mathbf{k}(\mathbf{x}, \mathbf{x}'); \quad \mathbf{b}_{\max, i} = \max(\mathbf{b}_i, \mathbf{b}_{\max, i-1})$ 
     $\alpha_i \leftarrow \mathbf{b}_i \mathbf{b}_{\max, i} (\frac{1}{2} + \log(2n_{\text{in}}/\delta))$ 
     $// \text{Compute RKHS inner product } \langle \tilde{\psi}_{i-1}, f_i \rangle_{\mathbf{k}}, \text{ which has a simple form}$ 
     $\alpha_i \leftarrow \sum_{j=1}^{2i-2} (\mathbf{k}(\mathbf{x}_j, \mathbf{x}) - \mathbf{k}(\mathbf{x}_j, \mathbf{x}')) - 2 \sum_{\mathbf{z} \in \mathcal{S}^{(1)}} (\mathbf{k}(\mathbf{z}, \mathbf{x}) - \mathbf{k}(\mathbf{z}, \mathbf{x}'))$ 
     $// \text{Assign one point to each coreset after probabilistic swapping}$ 
     $(\mathbf{x}, \mathbf{x}') \leftarrow (\mathbf{x}', \mathbf{x}) \text{ and } \eta_i \leftarrow 1 \quad \text{with probability } \min(1, \frac{1}{2}(1 - \frac{\alpha_i}{\alpha_i})_+)$ 
     $\mathcal{S}^{(1)}.append(\mathbf{x}); \quad \mathcal{S}^{(2)}.append(\mathbf{x}'); \quad \tilde{\psi}_i \leftarrow \tilde{\psi}_{i-1} + \eta_i f_i \quad // \psi_i = \sum_{\mathbf{x}' \in \mathcal{S}^{(2)}} \mathbf{k}(\mathbf{x}', \cdot) - \sum_{\mathbf{x} \in \mathcal{S}^{(1)}} \mathbf{k}(\mathbf{x}, \cdot)$ 
end
return  $\mathcal{X}_{\text{out}} \triangleq \mathcal{S}^{(1)}$ , coreset of size  $n_{\text{out}} = n_{\text{in}}/2$ 

```

In this section, we analyze KH(δ) (Alg. B.1), a variant of the Kernel Halving algorithm ([Dwivedi & Mackey, 2024](#), Alg. 2) with simplified swapping thresholds. Prop. B.2, proved in App. B.2.1, establishes the sub-Gaussianity of KH(δ) and its intermediate iterates.

Proposition B.2 (Sub-Gaussianity of KH(δ)). *Suppose $n_{\text{in}} \geq 2$. In the notation of Alg. B.1, on a common event \mathcal{E} of probability at least $1 - \delta/2$, for all $i \in [n_{\text{in}}/2]$, $\frac{1}{2i} \tilde{\psi}_i$ is (\mathbf{k}, ν_i) -sub-Gaussian with*

$$\begin{aligned} \nu_i &= \mathbf{b}_{\max, i} \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{2i} = \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{2i} \max_{j \in [i]} \text{MMD}_{\mathbf{k}}(\delta_{\mathbf{x}_{2j-1}}, \delta_{\mathbf{x}_{2j}}) \leq \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{2i} \max_{j \in [i]} \text{MMD}_{\mathbf{k}}(\delta_{\mathbf{x}_{2j-1}}, \delta_{\mathbf{x}_{2j}}) \\ &\leq \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{2i} 2 \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}(\mathbf{x}, \mathbf{x})}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\delta_{\mathbf{x}}, \mathbb{P}_{\text{in}})). \end{aligned}$$

Prop. B.2 and the triangle inequality imply that $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} = \frac{1}{n_{\text{in}}} \psi_{n_{\text{in}}/2}$ is (\mathbf{k}, ν) -sub-Gaussian on \mathcal{E} with

$$\nu = \mathbf{b}_{\max, n_{\text{in}}/2} \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{n_{\text{in}}} \leq \frac{\sqrt{\log(2n_{\text{in}}/\delta)}}{n_{\text{in}}} 2 \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}(\mathbf{x}, \mathbf{x})}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\delta_{\mathbf{x}}, \mathbb{P}_{\text{in}})).$$

By Lem. A.1, we thus have that the $\text{KH}(\delta)$ output $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on \mathcal{E} for \mathbf{K} generated by \mathbf{k} and that $\text{KH}(\delta) \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$.

B.2.1. PROOF OF PROP. B.2: SUB-GAUSSIANITY OF $\text{KH}(\delta)$

We begin by studying the sub-Gaussian properties of a related algorithm, the self-balancing Hilbert walk (SBHW) of Dwivedi & Mackey (2024, Alg. 3). By Dwivedi & Mackey (2024, Thm. 3(i)), when the SBHW is run on the RKHS $\mathcal{H}_{\mathbf{k}}$ with the same f_i and \mathbf{a}_i sequences employed in $\text{KH}(\delta)$, the output ψ_i of each round is (\mathbf{k}, σ_i) -sub-Gaussian for

$$\sigma_0^2 \triangleq 0 \quad \text{and} \quad \sigma_i^2 \triangleq \sigma_{i-1}^2 + \|f_i\|_{\mathbf{k}}^2 \left(1 + \frac{\sigma_{i-1}^2}{\mathbf{a}_i^2} (\|f_i\|_{\mathbf{k}}^2 - 2\mathbf{a}_i)\right)_+ \quad \forall i \geq 1. \quad (20)$$

The following lemma bounds the growth of the sub-Gaussian constants σ_i in terms of the swapping thresholds \mathbf{a}_i .

Lemma B.1 (Growth of SBHW sub-Gaussian constants). *For each i , the SBHW sub-Gaussian constants (20) satisfy*

$$\sigma_i^2 \leq c_i \quad \text{for} \quad c_i \triangleq \max_{j \in [i]} \max(\mathbf{b}_j^2, r_j) \quad \text{and} \quad r_i \triangleq \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \mathbf{b}_{\max, i}^2} \leq \frac{\mathbf{a}_i^2}{2\mathbf{a}_i - \mathbf{b}_{\max, i}^2}.$$

Proof. We will prove the result by induction on i .

Base case. $\sigma_1^2 = \mathbf{b}_1^2 \leq c_1$ as desired.

Inductive case. Suppose $\sigma_{i-1}^2 \leq c_{i-1}$. Then $\sigma_i^2 = g(\sigma_{i-1}^2)$ for $g(x) = x + \mathbf{b}_i^2(1 - x/r_i)_+$. Note that the slope of g is $1 - \mathbf{b}_i^2/r_i$ for $x < r_i$ and 1 for $x > r_i$. If $1 - \mathbf{b}_i^2/r_i \geq 0$, then g is increasing and its maximum value over $[0, c_i]$ is at c_i . If, on the other hand, $1 - \mathbf{b}_i^2/r_i < 0$, then g first decreases and then increases so its maximum value over $[0, c_i]$ is either at 0 or at c . Since $c_i \geq \max(r_i, c_{i-1})$, $\sigma_i^2 \leq \max(g(0), g(c_i)) = \max(\mathbf{b}_i^2, c_i) = c_i$. The proof is complete. \square

Invoking Lem. B.1, the assumption $n_{\text{in}} \geq 2$, and the fact that $\delta \mapsto \frac{\frac{1}{2} + \log(2/\delta)}{\log(2/\delta)}$ is increasing on $(0, 1]$, we find that

$$\sigma_i^2 \leq \mathbf{b}_{\max, i}^2 \log(2n_{\text{in}}/\delta) \frac{(\frac{1}{2} + \log(2n_{\text{in}}/\delta))^2}{2(\log(2n_{\text{in}}/\delta))^2} \leq \mathbf{b}_{\max, i}^2 \log(2n_{\text{in}}/\delta) \frac{(\frac{1}{2} + \log(4))^2}{2(\log(4))^2} \leq \mathbf{b}_{\max, i}^2 \log(2n_{\text{in}}/\delta). \quad (21)$$

The first inequality in (21) and the definition (20) further imply that

$$\mathbf{a}_i = \mathbf{b}_i \mathbf{b}_{\max, i} (\frac{1}{2} + \log(2n_{\text{in}}/\delta)) \geq \sigma_i \mathbf{b}_i \sqrt{2 \log(2n_{\text{in}}/\delta)} \geq \sigma_{i-1} \mathbf{b}_i \sqrt{2 \log(2n_{\text{in}}/\delta)}.$$

Hence, by Dwivedi & Mackey (2024, Thm. 3(iii)), for each $i \in [n_{\text{in}}/2]$, the vector $\tilde{\psi}_i$ of $\text{KH}(\delta)$ coincides with the vector ψ_i of SBHW on a common event \mathcal{E} of probability at least $1 - \delta/2$. Therefore, each $\frac{1}{2i} \tilde{\psi}_i$ is $(\mathbf{k}, \frac{1}{2i} \sigma_i)$ -sub-Gaussian on \mathcal{E} , implying the result.

B.3. LKH(δ)

In this section, we analyze LKH(δ) (Alg. B.2), the Kernel Halving algorithm of (Dwivedi & Mackey, 2024, Alg. 2) with a linear kernel, $\mathbf{k}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$, on \mathbb{R}^d and failure probability $\delta/2$. Notably, Alg. B.2 can be carried out in only $O(nd)$ time thanks to the linear kernel structure. Prop. B.3, proved in App. B.3.1, establishes the sub-Gaussianity of LKH(δ) and its intermediate iterates.

Proposition B.3 (Sub-Gaussianity of LKH(δ)). *Suppose $n_{\text{in}} \geq 2$. In the notation of Alg. B.2, on a common event \mathcal{E} of probability at least $1 - \delta/2$, for all $i \in [n_{\text{in}}/2]$, $\frac{1}{2i} \tilde{\psi}_i$ is (\mathbf{k}, ν_i) -sub-Gaussian with $\mathbf{k}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$ and*

$$\begin{aligned} \nu_i &= \frac{\sqrt{\log(2n_{\text{in}}(\log(n_{\text{in}}/2)+1)/\delta)}}{2i} \max_{j \in [i]} \|\mathbf{x}_{2j-1} - \mathbf{x}_{2j}\|_2 \\ &\leq \frac{\sqrt{\log(2n_{\text{in}}(\log(n_{\text{in}}/2)+1)/\delta)}}{2i} 2 \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\|\mathbf{x}\|_2}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \|\mathbf{x} - \bar{\mathbf{x}}\|_2) \quad \text{for} \quad \bar{\mathbf{x}} = \frac{1}{n_{\text{in}}} \sum_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \delta_{\mathbf{x}}. \end{aligned}$$

Prop. B.3 and the triangle inequality imply that $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} = \frac{1}{n_{\text{in}}} \psi_{n_{\text{in}}/2}$ is (\mathbf{k}, ν) -sub-Gaussian on \mathcal{E} with

$$\begin{aligned} \nu &= \frac{\sqrt{\log(2n_{\text{in}}(\log(n_{\text{in}}/2)+1)/\delta)}}{n_{\text{in}}} \max_{j \in [n_{\text{in}}/2]} \|\mathbf{x}_{2j-1} - \mathbf{x}_{2j}\|_2 \\ &\leq \frac{\sqrt{\log(2n_{\text{in}}(\log(n_{\text{in}}/2)+1)/\delta)}}{n_{\text{in}}} 2 \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\|\mathbf{x}\|_2}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \|\mathbf{x} - \bar{\mathbf{x}}\|_2) \quad \text{for} \quad \bar{\mathbf{x}} = \frac{1}{n_{\text{in}}} \sum_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \delta_{\mathbf{x}}. \end{aligned}$$

Algorithm B.2: LKH(δ): Kernel Halving with linear kernel and failure probability $\delta/2$

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$ with even n_{in} and $\mathbf{x}_i \in \mathbb{R}^d$

$\mathcal{S}^{(1)}, \mathcal{S}^{(2)} \leftarrow \{\}; \quad \psi_0 \leftarrow \mathbf{0} \in \mathbb{R}^d$ // Initialize empty coresets: $\mathcal{S}^{(1)}, \mathcal{S}^{(2)}$ have size i after round i
 $\sigma_0 \leftarrow 0$ // Keep track of sub-Gaussian constant

for $i = 1, 2, \dots, n_{\text{in}}/2$ **do**
 // Consider two points
 $(\mathbf{x}, \mathbf{x}') \leftarrow (\mathbf{x}_{2i-1}, \mathbf{x}_{2i}); \quad \eta_i \leftarrow -1$
 // Compute swapping threshold \mathbf{a}_i
 $\mathbf{b}_i^2 = \langle \mathbf{x} - \mathbf{x}', \mathbf{x} - \mathbf{x}' \rangle; \quad \delta_i = \frac{\delta}{2i(\log(n_{\text{in}}/2)+1)}$
 $(\mathbf{a}_i, \sigma_i) \leftarrow \text{get_swap_params}(\sigma_{i-1}, \mathbf{b}_i, \delta_i)$
 // Compute inner product
 $\alpha_i \leftarrow \langle \psi_{i-1}, \mathbf{x} - \mathbf{x}' \rangle$
 // Assign one point to each coreset after probabilistic swapping
 $(\mathbf{x}, \mathbf{x}') \leftarrow (\mathbf{x}', \mathbf{x})$ and $\eta_i \leftarrow 1$ with probability $\min(1, \frac{1}{2}(1 - \frac{\alpha_i}{\mathbf{a}_i})_+)$
 $\mathcal{S}^{(1)}.append(\mathbf{x}); \quad \mathcal{S}^{(2)}.append(\mathbf{x}'); \quad \tilde{\psi}_i \leftarrow \tilde{\psi}_{i-1} + \eta_i f_i$
end

return $\mathcal{X}_{\text{out}} \triangleq \mathcal{S}^{(1)}$, coreset of size $n_{\text{out}} = n_{\text{in}}/2$

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

By Lem. A.1, we thus have that the LKH(δ) output $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on \mathcal{E} for \mathbf{K} generated by \mathbf{k} and that $\text{LKH}(\delta) \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$.

B.3.1. PROOF OF PROP. B.3: SUB-GAUSSIANITY OF LKH(δ)

We begin by studying the sub-Gaussian properties of a related algorithm, the self-balancing Hilbert walk (SBHW) of Dwivedi & Mackey (2024, Alg. 3). By Dwivedi & Mackey (2024, Thm. 3(i)), when the SBHW is run on the RKHS $\mathcal{H}_{\mathbf{k}}$ with the same f_i and α_i sequences employed in LKH(δ), the output ψ_i of each round is (\mathbf{k}, σ_i) -sub-Gaussian. Moreover, since

$$\alpha_i \geq \sigma_{i-1} \mathbf{b}_i \sqrt{2\log(2/\delta_i)} \quad \text{for each } i \in [n_{\text{in}}/2],$$

Dwivedi & Mackey (2024, Thm. 3(iii)) implies that, for each $i \in [n_{\text{in}}/2]$, the vector $\tilde{\psi}_i$ of LKH(δ) coincides with the vector ψ_i of SBHW on a common event \mathcal{E} of probability at least $1 - \delta/2$. Therefore, each $\frac{1}{2i}\tilde{\psi}_i$ is $(\mathbf{k}, \frac{1}{2i}\sigma_i)$ -sub-Gaussian on \mathcal{E} . Finally, Dwivedi & Mackey (2024, (46)) shows that $\sigma_i \leq \nu_i$ for each $i \in [n_{\text{in}}/2]$, yielding the result.

B.4. RKH(δ)

Algorithm B.3: RKH(δ): Repeated KH(δ)

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$, kernel \mathbf{k} , output size $n_{\text{out}} \in n_{\text{in}}/2^{\mathbb{N}}$

// Repeatedly divide coreset size in half
 $m \leftarrow \log_2(n_{\text{in}}/n_{\text{out}})$
for $\ell = 1, 2, \dots, m$ **do** $\mathcal{X}_{\text{in}} \leftarrow \text{KH}(\delta/m)(\mathcal{X}_{\text{in}}, \mathbf{k})$;
return $\mathcal{X}_{\text{out}} \triangleq \mathcal{X}_{\text{in}}$, coreset of size $n_{\text{out}} = n_{\text{in}}/2^m$

In this section, we analyze repeated KH(δ) (RKH(δ), Alg. B.3), a variant of the KT-SPLIT algorithm (Dwivedi & Mackey, 2024, Alg. 1a) with simplified swapping thresholds. Our next result, proved in App. B.4.1, establishes the sub-Gaussianity of RKH(δ).

Algorithm B.4: KH-COMPRESS(δ): Compress with KH halving and failure probability δ

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$, kernel \mathbf{k} , $n_{\text{out}} \in \sqrt{n_{\text{in}}} \cdot 2^{\mathbb{N}}$

$g \leftarrow \log_2(n_{\text{out}}/\sqrt{n_{\text{in}}})$ // identify compression level

function compress(\mathcal{S}):

if $|\mathcal{S}| = 4^g$ **then return** \mathcal{S}

 Partition \mathcal{S} into four arbitrary subsequences $\{\mathcal{S}_i\}_{i=1}^4$ each of size $|\mathcal{S}|/4$

for $i = 1, 2, 3, 4$ **do**

$\tilde{\mathcal{S}}_i \leftarrow \text{compress}(\mathcal{S}_i)$ // return coresets of size $2^g \cdot \sqrt{\frac{|\mathcal{S}|}{4}}$

end

$\tilde{\mathcal{S}} \leftarrow \text{CONCATENATE}(\tilde{\mathcal{S}}_1, \tilde{\mathcal{S}}_2, \tilde{\mathcal{S}}_3, \tilde{\mathcal{S}}_4); \quad \ell \leftarrow 2 \cdot 2^g \cdot \sqrt{|\mathcal{S}|}$ // coreset of size ℓ

return KH($\frac{\ell^2}{n_{\text{in}}^{4g+1}(\log_4 n_{\text{in}} - g)}\delta$)($\tilde{\mathcal{S}}, \mathbf{k}$) // coreset of size $2^g \sqrt{|\mathcal{S}|}$

return compress(\mathcal{X}_{in}) // coreset of size $n_{\text{out}} = 2^g \sqrt{n_{\text{in}}}$

Proposition B.4 (Sub-Gaussianity of RKH(δ)). *If $n_{\text{out}} \in n_{\text{in}}/2^{\mathbb{N}}$ then RKH(δ) (Alg. B.3) is (\mathbf{k}, ν) -sub-Gaussian with*

$$\nu = \frac{2}{n_{\text{out}}\sqrt{3}} \sqrt{\log\left(\frac{6n_{\text{out}}\log_2(n_{\text{in}}/n_{\text{out}})}{\delta}\right)} \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}(\mathbf{x}, \mathbf{x})}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\boldsymbol{\delta}_{\mathbf{x}}, \mathbb{P}_{\text{in}}))$$

on an event \mathcal{E} of probability at least $1 - \delta/2$.

By Lem. A.1, we thus have that the RKH(δ) output $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on \mathcal{E} for \mathbf{K} generated by \mathbf{k} and that $\text{RKH}(\delta) \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$. Finally, $\nu = O(\frac{\sqrt{\log(n_{\text{out}}/\delta)}}{n_{\text{out}}})$ when $n_{\text{out}} \geq \sqrt{n_{\text{in}}}$.

B.4.1. PROOF OF PROP. B.4: SUB-GAUSSIANITY OF RKH(δ)

Let $c = 2 \min(\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}(\mathbf{x}, \mathbf{x})}, \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\boldsymbol{\delta}_{\mathbf{x}}, \mathbb{P}_{\text{in}}))$, and, for each $\ell \in [m]$, let $\tilde{\psi}^{(\ell)}$ represent the vector $\tilde{\psi}_{n_{\text{in}}/2^\ell}$ produced at the end of the ℓ -th call to KH(δ). By the proof of Prop. B.2 and the union bound, on an event \mathcal{E} of probability at least $1 - \delta/2$, $(\tilde{\psi}^{(\ell)})_{\ell \in [m]} = (\psi^{(\ell)})_{\ell \in [m]}$, where each $\frac{2^{\ell-1}}{n_{\text{in}}} \psi^{(\ell)}$ is $(\mathbf{k}, \nu^{(\ell)})$ -sub-Gaussian given $(\psi^{(j)})_{j \in [\ell-1]}$ for

$$\nu^{(\ell)} = c \frac{\sqrt{\log(2n_{\text{in}}m/(2^{\ell-1}\delta))}}{n_{\text{in}}/2^{\ell-1}}.$$

Hence, on \mathcal{E} , the weighted sum

$$(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k} = \sum_{\ell \in [m]} \frac{2^{\ell-1}}{n_{\text{in}}} \tilde{\psi}^{(\ell)} = \sum_{\ell \in [m]} \frac{2^{\ell-1}}{n_{\text{in}}} \psi^{(\ell)}$$

is $(\mathbf{k}, \sqrt{\sum_{\ell \in [m]} (\nu^{(\ell)})^2})$ -sub-Gaussian by Dwivedi & Mackey (2024, Lem. 14). Finally, by Dwivedi & Mackey (2024, Eq. (63)), $\sqrt{\sum_{\ell \in [m]} (\nu^{(\ell)})^2} \leq \nu$.

B.5. KH-COMPRESS(δ)

In this section, we analyze KH-COMPRESS(δ) (Alg. B.4), a variant of the KT-SPLIT-COMPRESS algorithm (Shetty et al., 2022, Ex. 3) with simplified swapping thresholds.

Proposition B.5 (Sub-Gaussianity of KH-COMPRESS(δ)). *If $n_{\text{out}} \in \sqrt{n_{\text{in}}} 2^{\mathbb{N}}$ then KH-COMPRESS(δ) (Alg. B.4) is (\mathbf{k}, ν) -sub-Gaussian with*

$$\nu = \frac{1}{n_{\text{out}}} \sqrt{\log_2(n_{\text{out}}) \log\left(\frac{4n_{\text{out}}\log_2(n_{\text{in}}/n_{\text{out}})}{\delta}\right)} \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}(\mathbf{x}, \mathbf{x})}$$

on an event \mathcal{E} of probability at least $1 - \delta/2$.

Proof. Since the original Kernel Halving algorithm of Dwivedi & Mackey (2024, Alg. 2) is equal to the KT-SPLIT algorithm of Dwivedi & Mackey (2024, Alg. 1a) with $m = 1$ halving round, KH-COMPRESS(δ) is simply the KT-SPLIT-COMPRESS algorithm of (Shetty et al., 2022, Ex. 3) with KH(δ) of Alg. B.1 substituted for KT-SPLIT($\delta, m = 1$). The

result now follows immediately from the $\text{KH}(\delta)$ sub-Gaussian constant of Prop. B.2 and the argument of Shetty et al. (2022, Rem. 2, Ex. 3). \square

Now fix any SPSP \mathbf{K} and any kernel \mathbf{k} that generates \mathbf{K} . By Lem. A.1, we have that $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian on \mathcal{E} and hence that $\text{KH-COMPRESS}(\delta) \in \mathcal{G}_{\nu, \delta}(\mathbf{K})$. In addition, $\nu = O(\frac{\sqrt{\log(n_{\text{out}}) \log(n_{\text{out}}/\delta)}}{n_{\text{out}}})$ when $n_{\text{out}} \geq \sqrt{n_{\text{in}}}$. Furthermore, Shetty et al. (2022, Rem. 1) implies that $\text{KH-COMPRESS}(\delta)$ has a runtime less than $4^{g+1} n_{\text{in}} (\log_4(n_{\text{in}}) - g) = 4n_{\text{out}}^2 \log_2(n_{\text{in}}/n_{\text{out}}) = O(n_{\text{out}}^2)$ when $n_{\text{out}} \geq \sqrt{n_{\text{in}}}$.

B.6. GS-THIN

The section introduces and analyzes the Gram-Schmidt Thinning algorithm (GS-THIN, Alg. B.5). GS-THIN repeatedly divides an input sequence in half using, GS-HALVE (Alg. B.6), a symmetrized and kernelized version of the Gram-Schmidt (GS) Walk of Bansal et al. (2018). We will present two different implementations of GS-HALVE: a quartic-time implementation (Alg. B.6) based on the GS Walk description of Bansal et al. (2018) and a cubic-time implementation based on local updates to the matrix inverse (Alg. B.7). While both the algorithms lead to the same output given the same source of randomness, we present the original implementation² for conceptual clarity and the optimized implementation for improved runtime. Throughout, for a matrix \mathbf{Q} and vector \mathbf{u} , we use the notation $\mathbf{Q}_{\mathcal{I} \times \mathcal{J}}$ and $\mathbf{u}_{\mathcal{I}}$ to represent the submatrix $(\mathbf{Q}_{ij})_{i \in \mathcal{I}, j \in \mathcal{J}}$ and subvector $(\mathbf{u}_i)_{i \in \mathcal{I}}$.

Algorithm B.5: GS-THIN: Gram-Schmidt Thinning

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$, kernel \mathbf{k} , output size $n_{\text{out}} \in n_{\text{in}}/2^{\mathbb{N}}$, $\text{HALVE} \in \{\text{GS-HALVE}, \text{GS-HALVE-CUBIC}\}$
 // Repeatedly divide coreset size in half
 $m \leftarrow \log_2(n_{\text{in}}/n_{\text{out}})$
for $\ell = 1, 2, \dots, m$ **do** $\mathcal{X}_{\text{in}} \leftarrow \text{HALVE}(\mathcal{X}_{\text{in}}, \mathbf{k})$;
return $\mathcal{X}_{\text{out}} \triangleq \mathcal{X}_{\text{in}}$, coreset of size $n_{\text{out}} = n_{\text{in}}/2^m$

Our first result, proved in App. B.6.1, shows that GS-THIN is a sub-Gaussian thinning algorithm.

Proposition B.6 (GS-THIN sub-Gaussianity). *For \mathbf{K} generated by \mathbf{k} , GS-THIN (Alg. B.5) is a $(\mathbf{K}, \nu, 0)$ -sub-Gaussian thinning algorithm with parameter*

$$\nu \triangleq \frac{2}{\sqrt{3}} \frac{\sqrt{\|\mathbf{K}\|_{\max}}}{n_{\text{out}}}. \quad (22)$$

Our second result, proved in App. B.6.2, shows that GS-THIN with the GS-HALVE implementation has $O(n_{\text{in}}^4)$ runtime.

Proposition B.7 (Runtime of GS-THIN with GS-HALVE). *The runtime of GS-THIN with implementation GS-HALVE (Alg. B.6) is $O(n_{\text{in}}^4)$.*

Our third result, proved in App. B.6.3, establishes the equivalence between GS-HALVE and GS-HALVE-CUBIC. More precisely, we show that the sequence of partial assignment vectors generated by `kernel_gs_walk`(\cdot) of Alg. B.6 and `kernel_gs_walk_cubic`(\cdot) of Alg. B.7 are identical given identical inputs, an invertible induced kernel matrix, and an identical source of randomness.

Proposition B.8 (Agreement of GS-HALVE and GS-HALVE-CUBIC). *Let $\mathbf{z}_1, \mathbf{z}_2, \dots$ be the fractional assignment sequence generated by `kernel_gs_walk`((\mathbf{x}_i) _{$i=1$} ^{n_{in}}) in Alg. B.6 and $\mathbf{z}'_1, \mathbf{z}'_2, \dots$ be the fractional assignment sequence generated by `kernel_gs_walk_cubic`((\mathbf{x}_i) _{$i=1$} ^{n_{in}}) in Alg. B.7 with an identical source of randomness. If the pairwise difference matrix*

$$\mathbf{Q} \triangleq (\mathbf{k}(x_{2i-1}, x_{2j-1}) + \mathbf{k}(x_{2i}, x_{2j}) - \mathbf{k}(x_{2i-1}, x_{2j}) - \mathbf{k}(x_{2i}, x_{2j-1}))_{i, j \in [n_{\text{in}}/2]}$$

is positive definite, then $\mathbf{z}_t = \mathbf{z}'_t$ for all t .

Our fourth result, proved in App. B.6.4, shows that GS-THIN with the GS-HALVE-CUBIC implementation has $O(n_{\text{in}}^3)$ runtime.

² Towards making this equivalence clear, Alg. B.6 has been expressed with the same variables that Alg. B.7 uses. Alg. B.6 can be slightly simplified if it were to be considered independently.

Algorithm B.6: GS-HALVE: Gram-Schmidt Halving

Input: point sequence $\mathcal{X}_{\text{in}} = (x_i)_{i=1}^{n_{\text{in}}}$ with even n_{in} , kernel \mathbf{k}

$\mathcal{X}_{\text{out}} \leftarrow \{\}$ // Initialize empty coreset

// Select one point to keep from each consecutive pair using kernelized GS Walk

$\mathbf{z} \leftarrow \text{kernel_gs_walk}(\mathcal{X}_{\text{in}})$

for $i = 1, \dots, n_{\text{in}}/2$ **do**

if $z_i = 1$ **then**

$\mathcal{X}_{\text{out}}.\text{append}(x_{2i-1})$

else

$\mathcal{X}_{\text{out}}.\text{append}(x_{2i})$

end

end

return \mathcal{X}_{out} , coreset of size $n_{\text{in}}/2$

function $\text{kernel_gs_walk}((\mathbf{x}_i)_{i=1}^{n_{\text{in}}})$:

$t \leftarrow 1$; $\mathbf{z}_t \leftarrow (0, 0, \dots, 0) \in \mathbb{R}^{n_{\text{in}}/2}$ // Initialize fractional assignment vector

$\mathcal{A} \leftarrow [n_{\text{in}}/2]$ // Initialize set of active coordinates

$p \sim \mathcal{A}$ // Select a pivot uniformly at random

while $\mathbf{z}_t \notin \{\pm 1\}^{n_{\text{in}}/2}$ **do**

$\mathcal{A}' \leftarrow \mathcal{A} \setminus \{\min(\{i \in [n_{\text{in}}/2] : |z_{ti}| = 1\} \setminus ([n_{\text{in}}/2] \setminus \mathcal{A}))\}$

 // Update set of active coordinates by removing smallest index set to ± 1

if $p \notin \mathcal{A}'$ **then**

$p' \sim \text{Unif}(\mathcal{A}')$ // Select a new pivot from \mathcal{A}' uniformly at random

else

$p' \leftarrow p$

end

 // Compute step direction in which to update fractional assignment vector

$\mathbf{u}_t \leftarrow \text{argmin}_{\mathbf{u} \in \mathbb{R}^{n_{\text{in}}/2}} \mathbf{u}^\top \mathbf{Q} \mathbf{u}$ subject to $\mathbf{u}_{p'} = 1$ and $\mathbf{u}_i = 0$ for all $i \notin \mathcal{A}'$,

 where $\mathbf{Q} \in \mathbb{R}^{(n_{\text{in}}/2) \times (n_{\text{in}}/2)}$ has entries $\mathbf{Q}_{ij} \triangleq \mathbf{k}(x_{2i-1}, x_{2j-1}) + \mathbf{k}(x_{2i}, x_{2j}) - \mathbf{k}(x_{2i-1}, x_{2j}) - \mathbf{k}(x_{2i}, x_{2j-1})$

$\delta^+ \leftarrow |\max \Delta|$ and $\delta^- \leftarrow |\min \Delta|$, where $\Delta = \left\{ \delta \in \mathbb{R} : \mathbf{z}_t + \delta \mathbf{u}_t \in [-1, +1]^{n_{\text{in}}/2} \right\}$ // Select candidate step sizes

$\delta_t \leftarrow \delta^+$ with probability $\delta^- / (\delta^+ + \delta^-)$; otherwise $\delta_t \leftarrow -\delta^-$ // Choose step size and sign at random

$\mathbf{z}_{t+1} \leftarrow \mathbf{z}_t + \delta_t \mathbf{u}_t$ // Update fractional assignments

$t \leftarrow t + 1$; $\mathcal{A} \leftarrow \mathcal{A}'$; $p \leftarrow p'$

end

return \mathbf{z}_t , sign vector in $\{\pm 1\}^{n_{\text{in}}/2}$

Algorithm B.7: GS-HALVE-CUBIC: Gram-Schmidt Halving with cubic runtime

Input: point sequence $\mathcal{X}_{\text{in}} = (x_i)_{i=1}^{n_{\text{in}}}$ with even n_{in} , kernel \mathbf{k} with positive definite $\mathbf{k}(\mathcal{X}_{\text{in}}, \mathcal{X}_{\text{in}})$

$\mathcal{X}_{\text{out}} \leftarrow \{\}$ // Initialize empty coreset

// Select one point to keep from each consecutive pair using kernelized GS Walk

$\mathbf{z} \leftarrow \text{kernel_gs_walk_cubic}(\mathcal{X}_{\text{in}})$

for $i = 1, \dots, n_{\text{in}}/2$ **do**

if $\mathbf{z}_i = 1$ **then**

$\mathcal{X}_{\text{out}}.\text{append}(x_{2i-1})$

else

$\mathcal{X}_{\text{out}}.\text{append}(x_{2i})$

end

end

return \mathcal{X}_{out} , coreset of size $n_{\text{in}}/2$

function $\text{kernel_gs_walk_cubic}((x_i)_{i=1}^{n_{\text{in}}})$:

$t \leftarrow 1$; $\mathbf{z}_t \leftarrow (0, 0, \dots, 0) \in \mathbb{R}^{n_{\text{in}}/2}$ // Initialize fractional assignment vector

$\mathcal{A} \leftarrow [n_{\text{in}}/2]$ // Initialize set of active coordinates

$p \sim \mathcal{A}$ // Select pivot uniformly at random

$\mathbf{Q} \leftarrow (\mathbf{k}(x_{2i-1}, x_{2j-1}) + \mathbf{k}(x_{2i}, x_{2j}) - \mathbf{k}(x_{2i-1}, x_{2j}) - \mathbf{k}(x_{2i}, x_{2j-1}))_{i,j=1}^{n_{\text{in}}/2}$ // Form paired difference kernel matrix

$\mathbf{C} \leftarrow (\mathbf{Q}_{\mathcal{A} \setminus \{p\} \times \mathcal{A} \setminus \{p\}})^{-1}$

while $\mathbf{z}_t \notin \{\pm 1\}^{n_{\text{in}}/2}$ **do**

$\mathcal{A}' \leftarrow \mathcal{A} \setminus \{\min(\{i \in [n_{\text{in}}/2] : |\mathbf{z}_{ti}| = 1\} \setminus ([n_{\text{in}}/2] \setminus \mathcal{A}))\}$

 // Update set of active coordinates by removing smallest index set to ± 1

if $p \notin \mathcal{A}'$ **then**

$p' \sim \text{Unif}(\mathcal{A}')$ // Select a new pivot from \mathcal{A}' uniformly at random

else

$p' \leftarrow p$

end

$\mathcal{A}_1 \leftarrow \mathcal{A} \setminus \{p\}$

$\mathcal{A}_2 \leftarrow \mathcal{A}' \setminus \{p'\}$

$i \leftarrow \mathcal{A}_1 \setminus \mathcal{A}_2$ // Choose i as the (unique) index that was removed from the active coordinates

 // Compute $(\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2})^{-1}$ using block matrix inversion and the Sherman-Morrison formula

$\mathbf{D} \leftarrow \mathbf{C}_{\mathcal{A}_2 \times \mathcal{A}_2}$

$\mathbf{C} \leftarrow \mathbf{D} - \frac{\mathbf{D}\mathbf{Q}_{\mathcal{A}_2 \times \{i\}}\mathbf{Q}_{\{i\} \times \mathcal{A}_2}\mathbf{D}}{\mathbf{Q}_{ii} + \mathbf{Q}_{\{i\} \times \mathcal{A}_2}\mathbf{D}\mathbf{Q}_{\mathcal{A}_2 \times \{i\}}}$

 // Compute step direction in which to update fractional assignment vector

 Compute \mathbf{u}_t as $(\mathbf{u}_t)_{\mathcal{A}_2} = -\mathbf{C}\mathbf{Q}_{\mathcal{A}_2 \times \{p'\}}$, $\mathbf{u}_{tp'} = 1$, and $\mathbf{u}_{ti} = 0$ for $i \notin \mathcal{A}'$

$\delta^+ \leftarrow |\max \Delta|$ and $\delta^- \leftarrow |\min \Delta|$, where $\Delta = \{\delta \in \mathbb{R} : \mathbf{z}_t + \delta \mathbf{u}_t \in [-1, +1]^{n_{\text{in}}/2}\}$ // Select candidate step sizes

$\delta_t \leftarrow \delta^+$ with probability $\delta^-/(\delta^+ + \delta^-)$; otherwise $\delta_t \leftarrow -\delta^-$ // Choose step size and sign at random

$\mathbf{z}_{t+1} \leftarrow \mathbf{z}_t + \delta_t \mathbf{u}_t$ // Update fractional assignments

$t \leftarrow t + 1$; $\mathcal{A} \leftarrow \mathcal{A}'$; $p \leftarrow p'$

end

return \mathbf{z}_t , sign vector in $\{\pm 1\}^{n_{\text{in}}/2}$

Proposition B.9 (Runtime of GS-THIN with GS-HALVE-CUBIC). *The runtime of GS-THIN with implementation GS-HALVE-CUBIC (Alg. B.7) is $O(n_{\text{in}}^3)$.*

B.6.1. PROOF OF PROP. B.6: GS-THIN SUB-GAUSSIANITY

Our first lemma bounds the sub-Gaussian constant of GS-HALVE (Alg. B.6).

Lemma B.2 (GS-HALVE sub-Gaussianity). *In the notation of Def. 1, consider the input and output vectors $\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}} \in \mathbb{R}^n$ of GS-HALVE (Alg. B.6) for $\mathcal{X} \supseteq \mathcal{X}_{\text{in}}$ with $|\mathcal{X}| = n \geq n_{\text{in}}$. If $\mathbf{K} = \mathbf{k}(\mathcal{X}, \mathcal{X})$, then $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian with*

$$\nu \triangleq \frac{2\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{in}}} = \frac{\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{out}}}.$$

Proof. Since \mathbf{K} is SPSD, there exists a matrix $\Phi \in \mathbb{R}^{n \times d}$ such that $\mathbf{K} = \Phi\Phi^\top$. Let $\mathbf{B} \in \mathbb{R}^{d \times (n_{\text{in}}/2)}$ be the matrix with entries

$$\mathbf{B}_{j,i} \triangleq \Phi_{2i-1,j} - \Phi_{2i,j} \quad \text{for } i \in [n_{\text{in}}/2] \quad \text{and } j \in [d].$$

Note that, for each $i \in [n_{\text{in}}/2]$,

$$\sum_{j \in [d]} \mathbf{B}_{j,i}^2 = \mathbf{K}_{2i-1,2i-1} + \mathbf{K}_{2i,2i} - \mathbf{K}_{2i-1,2i} - \mathbf{K}_{2i,2i-1} \leq 4\|\mathbf{K}\|_{\max}.$$

Hence, by Harshaw et al. (2024, Thm. 6.6), $\frac{1}{n_{\text{in}}}\mathbf{B}\mathbf{z}$ is (\mathbf{I}, ν) -sub-Gaussian where \mathbf{I} is the identity matrix in $\mathbb{R}^{d \times d}$.

Now fix any $\mathbf{u} \in \mathbb{R}^d$. Since $\frac{1}{n_{\text{in}}}\mathbf{B}\mathbf{z} = -\Phi^\top(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})$ by construction,

$$\mathbb{E}[\exp(\mathbf{u}^\top \mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}))] \leq \mathbb{E}\left[\exp\left(-\langle \Phi^\top \mathbf{u}, \frac{1}{n_{\text{in}}}\mathbf{B}\mathbf{z} \rangle\right)\right] \leq \exp\left(\frac{\nu^2}{2} \cdot \|\Phi^\top \mathbf{u}\|_2^2\right) = \exp\left(\frac{\nu^2}{2} \cdot \mathbf{u}^\top \mathbf{K} \mathbf{u}\right).$$

□

Now, for $\ell \in [m]$, let $\mathbf{p}_\ell \in \mathbb{R}^n$ denote the output probability vector produced by the ℓ -th call to GS-HALVE. Defining $\mathbf{p}_0 \triangleq \mathbf{p}_{\text{in}}$ and $\mathbf{p}_{\text{out}} \triangleq \mathbf{p}_m$, we have

$$\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}} = \sum_{i=1}^m \Delta_i, \quad \text{for } \Delta_i \triangleq \mathbf{p}_{i-1} - \mathbf{p}_i \quad \text{for } i \in [m].$$

By Lem. B.2, each $\mathbf{p}_{i-1} - \mathbf{p}_i$ is $(\mathbf{K}, \frac{2\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{in}}/2^{i-1}})$ -sub-Gaussian conditional on $(\Delta_1, \dots, \Delta_{i-1})$. Applying Lem. A.3 to the sequence $(\Delta_j)_{j=1}^m$, we find that $\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}}$ is (\mathbf{K}, ν) -sub-Gaussian with parameter

$$\nu = \left(\sum_{j=1}^m \frac{4\|\mathbf{K}\|_{\max}}{(n_{\text{in}}/2^{j-1})^2}\right)^{1/2} = \frac{2\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{in}}} \left(\sum_{j=1}^m 4^j\right)^{1/2} \leq \frac{\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{in}}} \sqrt{\frac{4}{3}4^m}.$$

Simplifying the above using the fact that $n_{\text{out}} = n_{\text{in}}/2^m$ yields our desired result (22).

B.6.2. PROOF OF PROP. B.7: RUNTIME OF GS-THIN WITH GS-HALVE

We essentially reproduce the argument from Bansal et al. (2018) for the runtime of the GS-HALVE algorithm in our kernelized context.

The main computational cost of GS-HALVE is the execution of the `kernel_gs_walk`(\cdot) subroutine in Alg. B.6. The number of iterations in while loop for \mathbf{z}_t is at most $n_{\text{in}}/2$. This is due to the fact that in each iteration, at least one new variable is set to $\{\pm 1\}$. Further, in each iteration, the main computational cost is the computation of

$$\mathbf{u}_t \leftarrow \arg\min_{\mathbf{u} \in \mathbb{R}^{n_{\text{in}}/2}} \mathbf{u}^\top \mathbf{Q} \mathbf{u}$$

under the constraints that $\mathbf{u}_p = 1$ and $\mathbf{u}_i = 0$ for all $i \notin \mathcal{A}$. Since this can be implemented in $O(n_{\text{in}}^3)$ time using standard convex optimization techniques, GS-HALVE has total runtime

$$r_H(\ell) \leq C\ell^4$$

for an input sequence of size ℓ and a constant C independent of ℓ . Now, note that GS-THIN calls GS-HALVE iteratively on inputs of size $n_{\text{in}}2^{-i}$ for $i = 0, 1, \dots, m-1$ where $m = \log_2(n_{\text{in}}/n_{\text{out}})$. Thus, GS-THIN has runtime

$$\sum_{i=0}^{m-1} r_H(n_{\text{in}}/2^i) \leq \sum_{i=0}^{m-1} C(n_{\text{in}}/2^i)^4 = O(n_{\text{in}}^4).$$

B.6.3. PROOF OF PROP. B.8: AGREEMENT OF GS-HALVE AND GS-HALVE-CUBIC

We want to reason that any round of partial coloring leads to the same output across the two algorithms. Fix any fractional assignment update round. Recall that $\mathcal{A}_1 = \mathcal{A} \setminus \{p\}$ and $\mathcal{A}_2 = \mathcal{A}' \setminus \{p'\}$. These represent the active set coordinates without the pivot before and after the update respectively.

The main difference between Algs. B.6 and B.7 is in the computation of the step direction \mathbf{u}_t , which is the solution of the program

$$\mathbf{u}_t \leftarrow \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^n} \mathbf{u}^\top \mathbf{Q} \mathbf{u} \quad \text{subject to} \quad \mathbf{u}_{p'} = 1 \quad \text{and} \quad \mathbf{u}_i = 0 \quad \text{for all} \quad i \notin \mathcal{A}'.$$

\mathbf{u}_t has a closed form with entries

$$(\mathbf{u}_t)_{\mathcal{A}_2} = -(\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2})^{-1} \cdot \mathbf{Q}_{\mathcal{A}_2 \times \{p'\}}.$$

Note that the invertibility of $\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2}$ follows from the positive-definiteness of \mathbf{Q} , as, for any $\mathbf{w} \in \mathbb{R}^{|\mathcal{A}_2|}$,

$$\mathbf{w}^\top \mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2} \mathbf{w} = \tilde{\mathbf{w}}^\top \mathbf{Q} \tilde{\mathbf{w}} > 0$$

for a second vector $\tilde{\mathbf{w}}$ with $\tilde{\mathbf{w}}_{\mathcal{A}_2} = \mathbf{w}$ and all other entries equal to zero. Therefore, to compute \mathbf{u}_t , it suffices to keep track of the inverse of $\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2}$ as \mathcal{A}' across iterations.

Let i be the unique element in $\mathcal{A}_1 \setminus \mathcal{A}_2$. Writing $\mathbf{Q}_{\mathcal{A}_1 \times \mathcal{A}_1}$ in block form, we have

$$\mathbf{Q}_{\mathcal{A}_1 \times \mathcal{A}_1} = \begin{bmatrix} \mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2} & \mathbf{Q}_{\mathcal{A}_2 \times \{i\}} \\ \mathbf{Q}_{\{i\} \times \mathcal{A}_2} & \mathbf{Q}_{ii} \end{bmatrix}.$$

By block matrix inversion (see, e.g., [Saadetoglu & Dinsev, 2023](#), Thm. 2), the leading size $|\mathcal{A}_2| \times |\mathcal{A}_2|$ principal submatrix of $(\mathbf{Q}_{\mathcal{A}_1 \times \mathcal{A}_1})^{-1}$ equals

$$\mathbf{D} \triangleq \left(\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2} - \frac{\mathbf{Q}_{\mathcal{A}_2 \times \{i\}} \mathbf{Q}_{\{i\} \times \mathcal{A}_2}}{\mathbf{Q}_{ii}} \right)^{-1}.$$

Thus, by the Sherman-Morrison formula ([Sherman & Morrison, 1950](#)),

$$(\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2})^{-1} = \left(\mathbf{D}^{-1} + \frac{\mathbf{Q}_{\mathcal{A}_2 \times \{i\}} \mathbf{Q}_{\{i\} \times \mathcal{A}_2}}{\mathbf{Q}_{ii}} \right)^{-1} = \mathbf{D} - \frac{\mathbf{D} \mathbf{Q}_{\mathcal{A}_2 \times \{i\}} \mathbf{Q}_{\{i\} \times \mathcal{A}_2} \mathbf{D}}{\mathbf{Q}_{ii} + \mathbf{Q}_{\{i\} \times \mathcal{A}_2} \mathbf{D} \mathbf{Q}_{\mathcal{A}_2 \times \{i\}}}. \quad (23)$$

Hence, if we already have access to a matrix $\mathbf{C} = (\mathbf{Q}_{\mathcal{A}_1 \times \mathcal{A}_1})^{-1}$, we can compute \mathbf{D} by dropping the row and column of \mathbf{C} corresponding to i and then compute $(\mathbf{Q}_{\mathcal{A}_2 \times \mathcal{A}_2})^{-1}$ using (23). Since in Alg. B.7 we begin by explicitly computing the inverse of $\mathbf{Q}_{\mathcal{A}' \times \mathcal{A}'}$, the update step in Alg. B.7 maintains the required inverse and thus its partial assignment updates match those of Alg. B.6.

B.6.4. PROOF OF PROP. B.9: RUNTIME OF GS-THIN WITH GS-HALVE-CUBIC

We begin by establishing the runtime of `kernel_gs_walk_cubic(·)`.

Lemma B.3 (Running time of `kernel_gs_walk_cubic(·)`). *The routine `kernel_gs_walk_cubic(·)` runs in $O(\ell^3)$ time given a point sequence of size ℓ .*

Proof. First, the initialization of \mathbf{C} costs $O(\ell^3)$ time using standard matrix inversion algorithms. Second, the number of iterations in the while loop is at most $\ell/2$ since, in each iteration, at least one new variable is assigned a permanent sign in $\{\pm 1\}$. In each while loop iteration, the main computational costs are the update of \mathbf{C} and the computation of the step direction \mathbf{u}_t , both of which cost $O(\ell^2)$ time using standard matrix-vector multiplication. Hence, together, all while loop iterations cost $O(\ell^3)$ time. \square

Given the above lemma, we have that GS-HALVE-CUBIC, on input of size ℓ , has a running time

$$r_H(\ell) \leq C\ell^3$$

for some C independent of ℓ . When used in GS-THIN this yields the runtime

$$\sum_{i=0}^{m-1} r_H(n_{\text{in}}/2^i) = \sum_{i=0}^{m-1} C(n_{\text{in}}/2^i)^3 = O(n_{\text{in}}^3).$$

Algorithm B.8: GS-COMPRESS: Compress with GS-HALVE-CUBIC halving

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$, kernel \mathbf{k} , $n_{\text{out}} \in \sqrt{n_{\text{in}}} \cdot 2^{\mathbb{N}}$

$g \leftarrow \log_2(n_{\text{out}}/\sqrt{n_{\text{in}}})$ // identify compression level

function `compress` (\mathcal{S}):

if $|\mathcal{S}| = 4^g$ **then return** \mathcal{S}

 Partition \mathcal{S} into four arbitrary subsequences $\{\mathcal{S}_i\}_{i=1}^4$ each of size $|\mathcal{S}|/4$

for $i = 1, 2, 3, 4$ **do**

$\tilde{\mathcal{S}}_i \leftarrow \text{compress}(\mathcal{S}_i)$ // return coresets of size $2^g \cdot \sqrt{\frac{|\mathcal{S}|}{4}}$

end

$\tilde{\mathcal{S}} \leftarrow \text{CONCATENATE}(\tilde{\mathcal{S}}_1, \tilde{\mathcal{S}}_2, \tilde{\mathcal{S}}_3, \tilde{\mathcal{S}}_4)$; $\ell \leftarrow 2 \cdot 2^g \cdot \sqrt{|\mathcal{S}|}$ // coreset of size ℓ

return GS-HALVE-CUBIC($\tilde{\mathcal{S}}, \mathbf{k}$) // coreset of size $2^g \sqrt{|\mathcal{S}|}$

return `compress` (\mathcal{X}_{in}) // coreset of size $n_{\text{out}} = 2^g \sqrt{n_{\text{in}}}$

B.7. GS-COMPRESS

This section introduces and analyzes the new GS-COMPRESS algorithm (Alg. B.8) which combines the COMPRESS meta-algorithm of Shetty et al. (2022) with the GS-HALVE-CUBIC halving algorithm (Alg. B.7). The following result bounds the sub-Gaussian constant and runtime of GS-COMPRESS.

Proposition B.10 (GS-COMPRESS sub-Gaussianity and runtime). *If \mathbf{K} is generated by \mathbf{k} , then GS-COMPRESS is $(\mathbf{K}, \nu, 0)$ -sub-Gaussian with*

$$\nu \triangleq \frac{1}{n_{\text{out}}} \sqrt{\log_2(n_{\text{out}}) \|\mathbf{K}\|_{\max}}.$$

Moreover, GS-COMPRESS has an $O(n_{\text{out}}^3)$ runtime.

Proof. By Lem. B.2 and Prop. B.8, GS-HALVE-CUBIC is $(\mathbf{K}, \nu_H(\ell))$ -sub-Gaussian for an input point sequence of size ℓ and $\nu_H(\ell) = 2\sqrt{\|\mathbf{K}\|_{\max}/\ell}$. Hence, by Lem. A.2, GS-HALVE-CUBIC is also $\nu_H(\ell)$ f -sub-Gaussian in the sense of Shetty et al. (2022, Def. 2) for each $f \in \mathcal{H}_{\mathbf{k}}$. By Shetty et al. (2022, Rmk. 2), GS-COMPRESS is therefore f -sub-Gaussian with parameter

$$\nu \leq \sqrt{\log_2(n_{\text{in}}/n_{\text{out}})} \nu_H(2n_{\text{out}}) \leq \sqrt{\log_2(n_{\text{out}})} \frac{\|\mathbf{K}\|_{\max}^{1/2}}{n_{\text{out}}}$$

for each $f \in \mathcal{H}_{\mathbf{k}}$. Hence, Lem. A.1 implies that GS-COMPRESS is a $(\mathbf{K}, \nu, 0)$ -sub-Gaussian thinning algorithm.

Furthermore, Shetty et al. (2022, Thm. 1) implies that GS-COMPRESS has a runtime of

$$\sum_{i=0}^{\log_2(n_{\text{in}}/(2n_{\text{out}}))} 4^i \cdot r_H(2n_{\text{out}}2^{-i}).$$

where the GS-HALVE-CUBIC runtime $r_H(\ell) \leq C\ell^3$ for C independent of the input size ℓ by Lem. B.3. Therefore, the GS-COMPRESS runtime is bounded by

$$\sum_{i=0}^{\log_2(n_{\text{in}}/(2n_{\text{out}}))} 4^i \cdot (2n_{\text{out}})^3 2^{-3i} = O(n_{\text{out}}^3).$$

□

Remark 1 (COMPRESS with GS-HALVE). If the GS-HALVE implementation were used in place of GS-HALVE-CUBIC, parallel reasoning would yield an $O(n_{\text{out}}^4)$ runtime for GS-COMPRESS.

C. Proof of Thm. 1: Low-rank sub-Gaussian thinning

We establish the MMD bound (2) in App. C.1, the first kernel max seminorm bound (3) in App. C.2, and the Lipschitz kernel max seminorm bound (4) in App. C.3. Throughout, we use the notation $\mathbb{P}_{\mathcal{E}}(\mathcal{E}') \triangleq \mathbb{P}(\mathcal{E}, \mathcal{E}')$ for events $(\mathcal{E}, \mathcal{E}')$.

C.1. Proof of MMD bound (2)

Without loss of generality, we suppose that $r \leq \text{rank}(\mathbf{K})$. Let $\mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$ be an eigendecomposition of \mathbf{K} with orthonormal $\mathbf{V} \in \mathbb{R}^{n \times n}$ and diagonal $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$. Let \mathbf{V}_r represent the first r columns of \mathbf{V} , and let \mathbf{V}_{-r} represent the last $n - r$ columns of \mathbf{V} . Introduce the shorthand

$$\mathbf{w} \triangleq \mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}} \in \mathbb{R}^n \quad \text{and} \quad \mathbf{\Phi} \triangleq \mathbf{V}\mathbf{\Lambda}^{1/2}\mathbf{V}^\top \in \mathbb{R}^{n \times n}. \quad (24)$$

We can directly verify that

$$\mathbf{V}\mathbf{V}^\top = \mathbf{V}^\top\mathbf{V} = \mathbf{I}, \quad \mathbf{V}\mathbf{V}^\top = \mathbf{V}_r\mathbf{V}_r^\top + \mathbf{V}_{-r}\mathbf{V}_{-r}^\top, \quad \text{and} \quad \mathbf{K} = \mathbf{\Phi}\mathbf{\Phi}^\top. \quad (25)$$

Using the above equalities, we decompose the squared MMD into two components,

$$\begin{aligned} \text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) &= \mathbf{w}^\top \mathbf{K} \mathbf{w} = \mathbf{w}^\top \mathbf{\Phi} \mathbf{\Phi}^\top \mathbf{w} = \mathbf{w}^\top \mathbf{\Phi} \mathbf{V} \mathbf{V}^\top \mathbf{\Phi}^\top \mathbf{w} = \mathbf{w}^\top \mathbf{\Phi} \mathbf{V}_r \mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w} + \mathbf{w}^\top \mathbf{\Phi} \mathbf{V}_{-r} \mathbf{V}_{-r}^\top \mathbf{\Phi}^\top \mathbf{w} \\ &= \|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2 + \|\mathbf{V}_{-r}^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2. \end{aligned} \quad (26)$$

In Apps. C.1.1 and C.1.2 respectively, we will establish the bounds

$$\mathbb{P}(\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2 \leq e\nu^2(er + \log(1/\delta')) \geq 1 - \delta/2 - \delta' \quad \text{and} \quad (27)$$

$$\mathbb{P}(\|\mathbf{V}_{-r}^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2 \leq \lambda_{r+1}(\frac{1}{n_{\text{out}}} - \frac{1}{n})) = 1, \quad (28)$$

which when combined with (26) yield the advertised claim (2) on the squared MMD.

C.1.1. PROOF OF (27): BOUNDING $\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2$

Our first lemma bounds the Euclidean norm of a vector in terms of a finite number of inner products.

Lemma C.1 (Euclidean norm cover). *For any $\mathbf{v} \in \mathbb{R}^r$ and $\varepsilon \in (0, 1)$,*

$$\|\mathbf{v}\|_2 \leq \frac{1}{1-\varepsilon} \max_{\mathbf{u} \in \mathcal{C}_{\varepsilon,r}} \langle \mathbf{u}, \mathbf{v} \rangle \quad (29)$$

for a set $\mathcal{C}_{\varepsilon,r}$ contained in the ball \mathbb{B}^r with $|\mathcal{C}_{\varepsilon,r}| \leq (1 + 2/\varepsilon)^r$.

Proof. Fix any $\varepsilon \in (0, 1)$, and let $\mathcal{C}_{\varepsilon,r}$ be a set of minimum cardinality satisfying

$$\mathcal{C}_{\varepsilon,r} \subset \mathbb{B}^r \quad \text{and} \quad \sup_{\mathbf{u} \in \mathbb{B}^r} \min_{\mathbf{u}' \in \mathcal{C}_{\varepsilon,r}} \|\mathbf{u} - \mathbf{u}'\|_2 \leq \varepsilon.$$

By Wainwright (2019, Lem. 5.2), $|\mathcal{C}_{\varepsilon,r}| \leq (1 + 2/\varepsilon)^r$. Now we invoke the variational representation of $\|\cdot\|_2$ and the Cauchy-Schwarz inequality to conclude that

$$\begin{aligned} \|\mathbf{v}\|_2 &= \sup_{\mathbf{u} \in \mathbb{B}^r} \langle \mathbf{u}, \mathbf{v} \rangle = \sup_{\mathbf{u} \in \mathbb{B}^r} \min_{\mathbf{u}' \in \mathcal{C}_{\varepsilon,r}} [\langle \mathbf{u} - \mathbf{u}', \mathbf{v} \rangle + \langle \mathbf{u}', \mathbf{v} \rangle] \\ &\leq \sup_{\mathbf{u} \in \mathbb{B}^r} \min_{\mathbf{u}' \in \mathcal{C}_{\varepsilon,r}} \|\mathbf{u} - \mathbf{u}'\|_2 \|\mathbf{v}\|_2 + \max_{\mathbf{u}' \in \mathcal{C}_{\varepsilon,r}} \langle \mathbf{u}', \mathbf{v} \rangle \\ &\leq \varepsilon \|\mathbf{v}\|_2 + \max_{\mathbf{u}' \in \mathcal{C}_{\varepsilon,r}} \langle \mathbf{u}', \mathbf{v} \rangle. \end{aligned}$$

Rearranging terms yields the claimed bound (29). □

Our next lemma uses this covering estimate to bound the exponential moments of $\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2$.

Lemma C.2 (Norm sub-Gaussianity). *For any $\varepsilon > 0$ and any $t > 0$,*

$$\mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2)] \leq (1 + \frac{2}{\varepsilon})^r \exp(\frac{\nu^2 t^2}{2(1-\varepsilon)^2}).$$

Proof. Fix any $t > 0$. Since $x \mapsto \exp(tx)$ is increasing, Lem. C.1 implies that

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2)] &\leq \mathbb{E}_{\mathcal{E}}[\exp(t \cdot \frac{1}{1-\varepsilon} \max_{\mathbf{u} \in \mathcal{C}_{\varepsilon,r}} \langle \mathbf{u}, \mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w} \rangle)] \\ &= \mathbb{E}_{\mathcal{E}}[\max_{\mathbf{u} \in \mathcal{C}_{\varepsilon,r}} \exp(\frac{t}{1-\varepsilon} \langle \mathbf{V}_r \mathbf{u}, \mathbf{\Phi}^\top \mathbf{w} \rangle)] \\ &\leq \sum_{\mathbf{u} \in \mathcal{C}_{\varepsilon,r}} \mathbb{E}_{\mathcal{E}}[\exp(\frac{t}{1-\varepsilon} \langle \mathbf{V}_r \mathbf{u}, \mathbf{\Phi}^\top \mathbf{w} \rangle)] \end{aligned}$$

for a subset $\mathcal{C}_{\varepsilon,r}$ with $|\mathcal{C}_{\varepsilon,r}| \leq (1 + \frac{2}{\varepsilon})^r$ and $\|\mathbf{u}\|_2 \leq 1$ for each $\mathbf{u} \in \mathcal{C}_{\varepsilon,r}$.

Now fix any $\mathbf{u} \in \mathcal{C}_{\varepsilon,r}$ and let $\mathbf{\Lambda}_r = \text{diag}(\lambda_1, \dots, \lambda_r)$. Using (24) and (25), we have

$$\begin{aligned} \mathbf{V}_r &= \mathbf{\Phi}^\top \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \quad \text{and therefore} \\ \langle \mathbf{V}_r \mathbf{u}, \mathbf{\Phi}^\top \mathbf{w} \rangle &= \langle \mathbf{\Phi}^\top \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u}, \mathbf{\Phi}^\top \mathbf{w} \rangle = \langle \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u}, \mathbf{K} \mathbf{w} \rangle. \end{aligned}$$

In addition, we have

$$(\mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u})^\top \mathbf{K} (\mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u}) = \mathbf{u}^\top \mathbf{\Lambda}_r^{-1/2} \mathbf{V}_r^\top \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u} = \mathbf{u}^\top \mathbf{u}.$$

Next, we can invoke our sub-Gaussianity assumption (Def. 3) to conclude that

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\exp(\frac{t}{1-\varepsilon} \langle \mathbf{V}_r \mathbf{u}, \mathbf{\Phi}^\top \mathbf{w} \rangle)] &= \mathbb{E}_{\mathcal{E}}[\exp(\frac{t}{1-\varepsilon} \langle \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u}, \mathbf{K} \mathbf{w} \rangle)] \leq \exp(\frac{\nu^2 t^2}{2(1-\varepsilon)^2} \langle \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u}, \mathbf{K} \mathbf{V}_r \mathbf{\Lambda}_r^{-1/2} \mathbf{u} \rangle) \\ &\leq \exp(\frac{\nu^2 t^2}{2(1-\varepsilon)^2} \|\mathbf{u}\|_2^2). \end{aligned}$$

Since $\|\mathbf{u}\|_2 \leq 1$ and $|\mathcal{C}_{\varepsilon,r}| \leq (1 + \frac{2}{\varepsilon})^r$, the advertised result now follows. \square

By Markov's inequality (Markov, 1884) and Lem. C.2, for any $\alpha > 0$,

$$\begin{aligned} \mathbb{P}(\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2 > \alpha) &= \mathbb{P}_{\mathcal{E}}(\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2 > \alpha) + \mathbb{P}(\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2 > \alpha, \mathcal{E}^c) \\ &\leq \mathbb{P}_{\mathcal{E}}(\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2 > \alpha) + \mathbb{P}(\mathcal{E}^c) \\ &\leq \inf_{t>0} \mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2)] / \exp(t\alpha) + \delta/2 \\ &\leq (1 + \frac{2}{\varepsilon})^r \inf_{t>0} \exp(\frac{\nu^2 t^2}{2(1-\varepsilon)^2} - t\alpha) + \delta/2 \\ &= (1 + \frac{2}{\varepsilon})^r \exp(\frac{-(1-\varepsilon)^2 \alpha^2}{2\nu^2}) + \delta/2. \end{aligned}$$

Next, we have

$$(1 + \frac{2}{\varepsilon})^r \exp(\frac{-(1-\varepsilon)^2 \alpha^2}{2\nu^2}) \leq \delta' \quad \text{if} \quad \alpha \geq \frac{\nu\sqrt{2}}{1-\varepsilon} \sqrt{\log(\frac{1}{\delta'}) + r \log(1 + \frac{2}{\varepsilon})}$$

Since this bound holds for any ε , choosing $\varepsilon = 1 - \sqrt{2/e}$, we find that

$$\|\mathbf{V}_r^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2 \leq e\nu^2 \left[r \log(1 + 2/(1 - \sqrt{2/e})) + \log(1/\delta') \right] \leq e\nu^2 [er + \log(1/\delta')]$$

with probability at least $1 - \delta/2 - \delta'$ as claimed.

C.1.2. PROOF OF (28): BOUNDING $\|\mathbf{V}_{-r}^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2$

Since

$$\|\mathbf{w}\|_2^2 = \mathbf{p}_{\text{in}}^\top \mathbf{p}_{\text{in}} + \mathbf{p}_{\text{out}}^\top \mathbf{p}_{\text{out}} - 2\mathbf{p}_{\text{in}}^\top \mathbf{p}_{\text{out}} = \frac{n_{\text{in}}}{n_{\text{in}}^2} + \frac{n_{\text{out}}}{n_{\text{out}}^2} - \frac{2n_{\text{out}}}{n_{\text{in}}n_{\text{out}}} = \frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}, \quad (30)$$

we have, for $\mathbf{\Lambda}_{-r} \triangleq \text{diag}(\lambda_{r+1}, \dots, \lambda_n)$ and λ_{\max} the maximum eigenvalue of a SPSD matrix,

$$\|\mathbf{V}_{-r}^\top \mathbf{\Phi}^\top \mathbf{w}\|_2^2 = \mathbf{w}^\top \mathbf{V}_{-r} \mathbf{\Lambda}_{-r} \mathbf{V}_{-r}^\top \mathbf{w} \leq \lambda_{\max}(\mathbf{V}_{-r} \mathbf{\Lambda}_{-r} \mathbf{V}_{-r}^\top) \|\mathbf{w}\|_2^2 \stackrel{(30)}{=} \lambda_{r+1} \left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}} \right).$$

C.2. Proof of kernel max seminorm bound (3)

We begin by establishing a general bound on the maximum discrepancy between input and output expectations over a collection of test functions admitting a finite cover.

Lemma C.3 (Discrepancy cover bound). *Fix any kernel \mathbf{k} , subset $\mathcal{F} \subset \mathcal{H}_{\mathbf{k}}$, and scalars $\varepsilon \geq 0$ and $\delta' \in (0, 1)$. Define*

$$a \triangleq \sup_{f \in \mathcal{F}} \|f\|_{\mathbf{k}} \quad \text{and} \quad \mathbb{B}_{\mathcal{F}} \triangleq \{f \in \mathcal{H}_{\mathbf{k}} : \|f\|_{\mathbf{k}} \leq a\},$$

and let $\mathcal{C}_{\epsilon, \mathcal{F}}$ be a set of minimum cardinality satisfying

$$\mathcal{C}_{\epsilon, \mathcal{F}} \subset \mathbb{B}_{\mathcal{F}} \quad \text{and} \quad \sup_{f \in \mathcal{F}} \min_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} |f(\mathbf{x}) - f'(\mathbf{x})| \leq \epsilon. \quad (31)$$

If $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on an event \mathcal{E} (Def. A.2), then, on \mathcal{E} ,

$$\|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{F}} \triangleq \sup_{f \in \mathcal{F}} (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})f \leq 2\epsilon + \nu a \sqrt{2 \log(|\mathcal{C}_{\epsilon, \mathcal{F}}|/\delta')} \quad \text{with probability at least } 1 - \delta'.$$

Proof. The triangle inequality and the covering property (31) together imply that, with probability 1,

$$\begin{aligned} (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})f &\leq \min_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})f' + |(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})(f - f')| \\ &\leq \|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{C}_{\epsilon, \mathcal{F}}} + \min_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} |\mathbb{P}_{\text{in}}(f - f')| + |\mathbb{P}_{\text{out}}(f - f')| \\ &\leq \|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{C}_{\epsilon, \mathcal{F}}} + 2 \min_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} |f(\mathbf{x}) - f'(\mathbf{x})| \\ &\leq \|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{C}_{\epsilon, \mathcal{F}}} + 2\epsilon \end{aligned} \quad (32)$$

for each $f \in \mathcal{F}$. Since $s \mapsto e^{ts}$ is increasing, the bound (32), the assumed sub-Gaussianity (Def. A.2), and the fact that $\mathcal{C}_{\epsilon, \mathcal{F}}$ belongs to $\mathbb{B}_{\mathcal{F}}$ imply that

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{F}})] &\leq e^{2t\epsilon} \mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{C}_{\epsilon, \mathcal{F}}})] \\ &\leq \sum_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} e^{2t\epsilon} \mathbb{E}_{\mathcal{E}}[\exp(t(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})f')] \\ &\leq \sum_{f' \in \mathcal{C}_{\epsilon, \mathcal{F}}} \exp\left(\frac{t^2 \nu^2 \|f'\|_{\mathbf{k}}^2}{2} + 2t\epsilon\right) \leq |\mathcal{C}_{\epsilon, \mathcal{F}}| \exp\left(\frac{t^2 \nu^2 a^2}{2} + 2t\epsilon\right). \end{aligned}$$

Now, by Markov's inequality (Markov, 1884), for any $\alpha > 0$,

$$\begin{aligned} \mathbb{P}_{\mathcal{E}}(\sup_{f \in \mathcal{F}} (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})f > \alpha + 2\epsilon) &\leq \inf_{t>0} \mathbb{E}_{\mathcal{E}}[\exp(t\|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{F}})] / \exp(t(\alpha + 2\epsilon)) \\ &\leq |\mathcal{C}_{\epsilon, \mathcal{F}}| \inf_{t>0} \exp\left(\frac{t^2 \nu^2 a^2}{2} - t\alpha\right) = |\mathcal{C}_{\epsilon, \mathcal{F}}| \exp\left(\frac{-\alpha^2}{2\nu^2 a^2}\right). \end{aligned}$$

Finally, choosing $\alpha = \nu a \sqrt{2 \log(|\mathcal{C}_{\epsilon, \mathcal{F}}|/\delta')}$ yields the desired claim. \square

Now fix any $\epsilon \geq 0$, $\delta' \in (0, 1)$, and kernel \mathbf{k} that generates \mathbf{K} , and consider the subset $\mathcal{F} = \{\pm \mathbf{k}(x_i, \cdot) : i \in \mathcal{I}\}$. Since $\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} = \|\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}}\|_{\mathcal{F}}$ and $\sup_{f \in \mathcal{F}} \|f\|_{\mathbf{k}} = D_{\mathcal{I}}$, Lem. C.3 implies that, on the event \mathcal{E} ,

$$\|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \leq 2\epsilon + \nu D_{\mathcal{I}} \sqrt{2 \log(|\mathcal{C}_{\epsilon, \mathcal{F}}|/\delta')} \quad \text{with probability at least } 1 - \delta'.$$

Since $\mathbb{P}(\mathcal{E}^c) \leq \delta/2$ and $|\mathcal{F}| \leq 2|\mathcal{I}|$, we use the estimate $|\mathcal{C}_{0, \mathcal{F}}| \leq 2|\mathcal{I}|$ with $\epsilon = 0$ to obtain the advertised bound (3).

C.3. Proof of Lipschitz kernel max seminorm bound (4)

Introduce the query point set $\mathcal{Z} \triangleq \{\mathbf{x}_i : i \in \mathcal{I}\}$, fix any $\delta' \in (0, 1)$ and $\mathbf{z}_0 \in \mathcal{Z}$, and define the symmetrized seminorm

$$\|(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}\|_{\mathcal{Z}, \mathcal{Z}} \triangleq \sup_{\mathbf{z}, \mathbf{z}' \in \mathcal{Z}} |(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z}) - (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z}')|.$$

By the triangle inequality and the derivation of App. C.2, we have, on the event \mathcal{E} ,

$$\begin{aligned} \|\mathbf{K}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} &\leq \|(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}\|_{\mathcal{Z}, \mathcal{Z}} + |(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z}_0)| \\ &\leq \|(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}\|_{\mathcal{Z}, \mathcal{Z}} + \nu \sqrt{\mathbf{k}(\mathbf{z}_0, \mathbf{z}_0)} \sqrt{2 \log(4/\delta')} \quad \text{with probability at least } 1 - \delta'/2. \end{aligned} \quad (33)$$

Since $\mathbb{P}(\mathcal{E}^c) \leq \delta/2$, it only remains to upper bound $\|(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}\|_{\mathcal{Z}, \mathcal{Z}}$ on \mathcal{E} with probability at least $1 - \delta'/2$.

To this end, we first establish that $((\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z}))_{\mathbf{z} \in \mathcal{Z}}$ is a sub-Gaussian process on \mathcal{E} with respect to a particular bounded-Hölder metric ρ .

Definition C.1 (Sub-Gaussian process on an event). *We say an indexed collection of random variables $(X_{\theta})_{\theta \in \Theta}$ is a sub-Gaussian process with respect to ρ on an event \mathcal{E} if ρ is a metric on Θ and*

$$\mathbb{E}_{\mathcal{E}} \left[\exp \left(\frac{(X_{\theta} - X_{\theta'})^2}{\rho(\theta, \theta')^2} \right) \right] \leq 2 \quad \text{for all } \theta, \theta' \in \Theta.$$

Lemma C.4 (Bounded-Hölder sub-Gaussian process). *Consider a kernel \mathbf{k} on $\mathcal{X} = \mathbb{R}^d$ satisfying $|\mathbf{k}(\mathbf{z}, \mathbf{x}) - \mathbf{k}(\mathbf{z}', \mathbf{x})| \leq L_{\mathbf{k}} \|\mathbf{z} - \mathbf{z}'\|_2$ for all $\mathbf{z}, \mathbf{z}' \in \mathcal{Z} \subset \mathcal{X}$ and $\mathbf{x} \in \mathcal{X}_{\text{in}}$. If $(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}$ is (\mathbf{k}, ν) -sub-Gaussian on an event \mathcal{E} (Def. A.2), then $((\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z}))_{\mathbf{z} \in \mathcal{Z}}$ is a sub-Gaussian process on \mathcal{E} with respect to the metric*

$$\rho(\mathbf{z}, \mathbf{z}') \triangleq \nu \sqrt{8/3} \min(2 \sup_{\mathbf{z} \in \mathcal{Z}} \sqrt{\mathbf{k}(\mathbf{z}, \mathbf{z})}, \sqrt{2L_{\mathbf{k}} \|\mathbf{z} - \mathbf{z}'\|_2}). \quad (34)$$

The proof of Lem. C.4 can be found in App. C.4. Our next lemma, a slight modification of Wainwright (2019, Thm. 5.36), bounds the suprema of symmetrized sub-Gaussian processes on an event in terms of covering numbers.

Lemma C.5 (Sub-Gaussian process tails). *Suppose $(X_{\theta})_{\theta \in \Theta}$ is a sub-Gaussian process with respect to ρ on an event \mathcal{E} , and define the diameter $\text{diam}(\Theta, \rho) \triangleq \sup_{\theta, \theta' \in \Theta} \rho(\theta, \theta')$, the covering number*

$$\mathcal{N}(u; \Theta, \rho) \triangleq \min\{|\mathcal{C}_u| : \mathcal{C}_u \subseteq \Theta, \max_{\theta \in \Theta} \min_{\theta' \in \mathcal{C}_u} \rho(\theta, \theta') \leq u\} \quad \text{for all } u > 0,$$

and the entropy integral $\mathcal{J}(\Theta, \rho) \triangleq \int_0^{\text{diam}(\Theta, \rho)} \sqrt{\log(1 + \mathcal{N}(u; \Theta, \rho))} du$. Then,

$$\mathbb{P}_{\mathcal{E}}(\sup_{\theta, \theta' \in \Theta} |X_{\theta} - X_{\theta'}| \geq 8(\mathcal{J}(\Theta, \rho) + t)) \leq 2 \exp(-t^2 / \text{diam}(\Theta, \rho)^2) \quad \text{for all } t > 0.$$

Proof. Since $\sqrt{\log(1 + xy)} \leq \sqrt{\log((1+x)(1+y))} \leq \sqrt{\log(1+x)} + \sqrt{\log(1+y)}$ for all $x, y > 0$, the proof is identical to that of Wainwright (2019, Thm. 5.36) with $c_1 = 8$ and $(\mathbb{E}_{\mathcal{E}}, \mathbb{P}_{\mathcal{E}})$ substituted for (\mathbb{E}, \mathbb{P}) . \square

Our final lemma bounds the diameter, covering numbers, and entropy integral of \mathcal{Z} using the metric ρ .

Lemma C.6 (Covering properties of bounded-Hölder metric). *Consider the bounded-Hölder metric ρ (34) for a kernel \mathbf{k} on $\mathcal{X} = \mathbb{R}^d$ and a finite set $\mathcal{Z} \subset \mathcal{X}$. If \mathbf{Z} is a matrix with one row corresponding to each element of \mathcal{Z} , $r = \text{rank}(\mathbf{Z})$, and $R = \max_{\mathbf{z} \in \mathcal{Z}} \|\mathbf{z}\|_2$, then, in the notation of Lem. C.5,*

$$\mathcal{N}(u; \mathcal{Z}, \rho) \leq (1 + c^2/u^2)^r \quad \text{for } c \triangleq \nu \sqrt{\frac{32}{3} R L_{\mathbf{k}}} \quad \text{and all } u > 0, \quad (35)$$

$$\text{diam}(\mathcal{Z}, \rho) \leq D \triangleq \min(c, \nu \sqrt{\frac{32}{3}} \max_{\mathbf{z} \in \mathcal{Z}} \sqrt{\mathbf{k}(\mathbf{z}, \mathbf{z})}), \quad \text{and} \quad (36)$$

$$\mathcal{J}(\mathcal{Z}, \rho) \leq D \sqrt{2r \log(\sqrt{3}ec/D)}.$$

Proof. The diameter bound (36) follows directly from the definition of ρ (34) and the fact $\max_{\mathbf{z}, \mathbf{z}' \in \mathcal{Z}} \|\mathbf{z} - \mathbf{z}'\|_2 \leq 2R$.

To establish the covering number bound (35), we let $\mathbf{U}\Sigma\mathbf{V}^{\top}$ be a compact singular value decomposition of \mathbf{Z} so that

$$\mathbf{V} \in \mathbb{R}^{d \times r}, \quad \mathbf{Z} = \mathbf{Z}\mathbf{V}\mathbf{V}^{\top}, \quad \text{and} \quad \max_{\mathbf{z} \in \mathcal{Z}} \|\mathbf{V}^{\top} \mathbf{z}\|_2 = \max_{\mathbf{z} \in \mathcal{Z}} \|\mathbf{z}\|_2 = R.$$

Fix any $\epsilon > 0$, and let \mathcal{C} and \mathcal{C}_{ext} be a sets of minimum cardinality satisfying

$$\begin{aligned} \mathcal{C} &\subset \mathbb{B}^r(R), & \max_{\mathbf{v} \in \mathbb{B}^r(R)} \min_{\mathbf{v}' \in \mathcal{C}} \|\mathbf{v}' - \mathbf{v}\|_2 &\leq \epsilon^2/2, \\ \mathcal{C}_{\text{ext}} &\subset \mathbb{B}^d(R), & \text{and } \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{z}' \in \mathcal{C}_{\text{ext}}} \|\mathbf{z}' - \mathbf{z}\|_2 &\leq \epsilon^2/2. \end{aligned} \quad (37)$$

Since $\mathbf{V}^{\top} \mathbf{z} \in \mathbb{B}^r(R)$ for each $\mathbf{z} \in \mathcal{Z}$ and $\mathbf{V}\mathbf{v}' \in \mathbb{B}^d$ for each $\mathbf{v}' \in \mathbb{B}^r$, we have

$$\begin{aligned} \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{v}' \in \mathcal{C}} \|\mathbf{V}\mathbf{v}' - \mathbf{z}\|_2 &= \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{v}' \in \mathcal{C}} \|\mathbf{V}(\mathbf{v}' - \mathbf{V}^{\top} \mathbf{z})\|_2 \\ &= \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{v}' \in \mathcal{C}} \|\mathbf{v}' - \mathbf{V}^{\top} \mathbf{z}\|_2 \leq \epsilon^2/2, \end{aligned}$$

so that $\mathbf{V}\mathcal{C}$ satisfies the criteria of (37). Since $|\mathbf{V}\mathcal{C}| \leq |\mathcal{C}| \leq (1 + 4R/\epsilon^2)^r$ by Wainwright (2019, Lem. 5.2), we must also have $|\mathcal{C}_{\text{ext}}| \leq (1 + 4R/\epsilon^2)^r$.

Now, since \mathcal{C}_{ext} has minimum cardinality amongst sets satisfying (37), for each $\mathbf{z}' \in \mathcal{C}_{\text{ext}}$, there is some $\mathbf{z} \in \mathcal{Z}$ satisfying $\|\mathbf{z}' - \mathbf{z}\|_2 \leq \epsilon^2/2$ (or else \mathbf{z}' would be superfluous). Hence, there exists a set $\mathcal{C}_{\text{int}} \subseteq \mathcal{Z}$ satisfying

$$|\mathcal{C}_{\text{int}}| \leq |\mathcal{C}_{\text{ext}}| \leq (1 + 4R/\epsilon^2)^r \quad \text{and} \quad \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{z}' \in \mathcal{C}_{\text{int}}} \|\mathbf{z}' - \mathbf{z}\|_2 \leq \epsilon^2.$$

Moreover, by our metric definition (34),

$$\max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{z}' \in \mathcal{C}_{\text{int}}} \rho(\mathbf{z}, \mathbf{z}') \leq \frac{c}{2\sqrt{R}} \max_{\mathbf{z} \in \mathcal{Z}} \min_{\mathbf{z}' \in \mathcal{C}_{\text{int}}} \sqrt{\|\mathbf{z} - \mathbf{z}'\|_2} \leq \frac{c\epsilon}{2\sqrt{R}}.$$

Hence, for $u = \frac{c\epsilon}{2\sqrt{R}}$, $\mathcal{N}(u; \mathcal{Z}, \rho) \leq |\mathcal{C}_{\text{int}}| \leq (1 + c^2/u^2)^r$. Since $\epsilon > 0$ was arbitrary, we have established (35).

Finally, we bound the entropy integral using the inequality $1 \leq c^2/u^2$ for $u \in [0, D]$, the concavity of the square-root function, and Jensen's inequality:

$$\begin{aligned} \mathcal{J}(\mathcal{Z}, \rho) &\leq \int_0^D \sqrt{\log(1 + (1 + c^2/u^2)^r)} du \leq \int_0^D \sqrt{\log((3c^2/u^2)^r)} du = \int_0^D \sqrt{2r \log(\sqrt{3}c/u)} du \\ &\leq D \sqrt{\frac{1}{D} \int_0^D 2r \log(\sqrt{3}c/u) du} = D \sqrt{2r \log(\sqrt{3}ec/D)}. \end{aligned}$$

□

Together, Lems. C.4, C.5, and C.6 imply that, in the notation of Lem. C.6,

$$\|(\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}\|_{\mathcal{Z}, \mathcal{Z}} \leq 8D \sqrt{2r \log(\sqrt{3}ec/D)} + 8D \sqrt{\log(4/\delta')}$$

on \mathcal{E} with probability at least $1 - \delta'/2$. Combining this bound with the inequality (33) yields the result.

C.4. Proof of Lem. C.4: Bounded-Hölder sub-Gaussian process

Define $X_{\mathbf{z}} = (\mathbb{P}_{\text{in}} - \mathbb{P}_{\text{out}})\mathbf{k}(\mathbf{z})$ for each $\mathbf{z} \in \mathcal{Z}$, and fix any $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}$. Our sub-Gaussianity assumption implies

$$\mathbb{E}_{\mathcal{E}}[\exp(\lambda(X_{\mathbf{z}} - X_{\mathbf{z}'}))] \leq \exp\left(\frac{\nu^2 \lambda^2}{2} \|\mathbf{k}(\mathbf{z}, \cdot) - \mathbf{k}(\mathbf{z}', \cdot)\|_{\mathbf{k}}^2\right) \quad \text{for all } \lambda \in \mathbb{R}.$$

Moreover, by our Lipschitz assumption,

$$\|\mathbf{k}(\mathbf{z}, \cdot) - \mathbf{k}(\mathbf{z}', \cdot)\|_{\mathbf{k}}^2 = \mathbf{k}(\mathbf{z}, \mathbf{z}) - \mathbf{k}(\mathbf{z}, \mathbf{z}') + \mathbf{k}(\mathbf{z}', \mathbf{z}') - \mathbf{k}(\mathbf{z}', \mathbf{z}) \leq \min(4 \max_{\mathbf{z} \in \mathcal{Z}} \mathbf{k}(\mathbf{z}, \mathbf{z}), 2L_{\mathbf{k}} \|\mathbf{z} - \mathbf{z}'\|_2).$$

Finally, Lem. C.7 shows that $\mathbb{E}_{\mathcal{E}}[\exp(\frac{(X_{\mathbf{z}} - X_{\mathbf{z}'})^2}{\rho(\mathbf{z}, \mathbf{z}')^2})] \leq 2$ so that $(X_{\mathbf{z}})_{\mathbf{z} \in \mathcal{Z}}$ is a sub-Gaussian process on \mathcal{E} with respect to ρ .

Lemma C.7 (Squared exponential moment bound). *If $\mathbb{E}_{\mathcal{E}}[\exp(\lambda X)] \leq \exp(\frac{\nu^2 \lambda^2}{2})$ for all $\lambda \in \mathbb{R}$, then $\mathbb{E}_{\mathcal{E}}[\exp(\frac{3X^2}{8\nu^2})] \leq 2$.*

Proof. The proof is identical to that in Wainwright (2019, Sec. 2.4) with $\mathbb{E}_{\mathcal{E}}$ substituted for \mathbb{E} . □

D. Proof of Cor. 1: Gaussian MMD of KH

Cor. 1 follows immediately from the following explicit, non-asymptotic bound.

Corollary D.1 (Detailed Gaussian MMD of KH). *If $\mathcal{X}_{\text{in}} \subset \mathbb{B}^d(R)$ for $R > 0$, then $\text{KH}(\delta)$ with $\mathbf{k} = \text{GAUSS}(\eta)$, $n = n_{\text{in}} \geq (2e)^d$, and $b \triangleq \frac{1}{2}$ delivers*

$$\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \leq \frac{1}{n_{\text{out}}^2} \log\left(\frac{4n_{\text{out}}}{\delta}\right) \left[e^2 \max\left\{ \left[\frac{2e}{d} \log(n_{\text{in}} n_{\text{out}} b) \right]^d, \left(\frac{R^2 \eta e^3 4}{d} \right)^d \right\} + e \log\left(\frac{1}{\delta'}\right) \right] + \frac{1}{n_{\text{out}}} \left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}} \right)$$

with probability at least $1 - \delta/2 - \delta'$.

Proof. Consider the approximate rank parameter

$$r^* \triangleq \max\left\{ \left[\frac{2e}{d} \log(n_{\text{in}} n_{\text{out}} b) \right]^d, (R^2 \eta e^3 4/d)^d \right\}.$$

The assumption $n_{\text{in}} \geq (2e)^d$ and the fact that $b \geq 1/(2^d n_{\text{out}})$ ensure that $\log(n_{\text{in}} n_{\text{out}} b) \geq d + \log(n_{\text{out}} b/2^d) \geq d$ and therefore that $r^* \geq (2e)^d$. Hence, by Altschuler et al. (2019, Thm. 3), the $(r^* + 1)$ -th eigenvalue of \mathbf{K} satisfies

$$\begin{aligned} \lambda_{r^*+1} &\leq n_{\text{in}} \exp\left\{ -\frac{d}{2e} \max\left\{ \frac{2e}{d} \log(n_{\text{in}} n_{\text{out}} b), (R^2 \eta e^3 4/d) \right\} \log\left(\frac{d \max\left\{ \frac{2e}{d} \log(n_{\text{in}} n_{\text{out}} b), (R^2 \eta e^3 4/d) \right\}}{4e^2 \eta R^2} \right) \right\} \\ &\leq n_{\text{in}} \exp\{-\log(n_{\text{in}} n_{\text{out}} b) \log(e)\} \leq n_{\text{in}} \left(\frac{1}{n_{\text{in}} n_{\text{out}} b} \right) = \frac{1}{n_{\text{out}} b}. \end{aligned}$$

Since $\|\mathbf{K}\|_{\text{max}} = 1$ and $\text{KH}(\delta) \in \mathcal{G}_{\nu}(\mathbf{K})$ with ν defined in Prop. B.2, the result now follows from Thm. 1. □

E. Proof of Cor. 2: Intrinsic Gaussian MMD of KH

Assumption E.1 (d^* -manifold with Q -smooth atlas (Altschuler et al., 2019, Assum. 1)). Let $\Omega \subset \mathbb{R}^d$ be a smooth compact manifold without boundary of dimension $d^* < d$. Let $(\Psi_j, U_j)_{j \in [T]}$ for $T \in \mathbb{N}$ be an atlas for Ω , where $(U_j)_j$ are open sets covering Ω and $\Psi_j : U_j \mapsto \mathbb{B}^{d^*}(r_j)$ are smooth maps with smooth inverses, mapping U_j bijectively to $\mathbb{B}^{d^*}(r_j)$. Assume that there exists $Q > 0$ such that $\sup_{u \in \mathbb{B}^{d^*}(r_j)} \|D^\alpha \Psi_j^{-1}(u)\| \leq Q^{|\alpha|}$ for all $\alpha \in \mathbb{N}^{d^*}$ and $j \in [T]$, where $|\alpha| \triangleq \sum_{j=1}^{d^*} \alpha_j$ and $D^\alpha = \frac{\partial^{|\alpha|}}{\partial u_1^{\alpha_1} \dots \partial u_{d^*}^{\alpha_{d^*}}}$ for $\alpha \in \mathbb{N}^{d^*}$.

Cor. 2 follows immediately from the following more detailed result.

Corollary E.1 (Detailed Intrinsic Gaussian MMD of KH). Suppose \mathcal{X}_{in} lies on a manifold $\Omega \subset \mathbb{B}^d$ satisfying Assump. E.1. Then $\text{KH}(\delta)$ with $\mathbf{k} = \text{GAUSS}(\eta)$ and $n = n_{\text{in}}$ delivers

$$\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \leq \frac{1}{n_{\text{out}}^2} \log\left(\frac{4n_{\text{out}}}{\delta}\right) \left(\frac{e^2}{c^{5d^*/2}} \log^{\frac{5d^*}{2}}(n_{\text{in}}n_{\text{out}}) + e \log\left(\frac{1}{\delta'}\right)\right) + \frac{1}{n_{\text{out}}} \left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}\right)$$

with probability at least $1 - \frac{\delta}{2} - \delta'$ for c independent of \mathcal{X}_{in} .

Proof. Altschuler et al. (2019, Thm. 4) showed that the $(r+1)$ -th eigenvalue of \mathbf{K} satisfies (7) for a constant c independent of $\mathcal{X} = \mathcal{X}_{\text{in}}$. Since $\|\mathbf{K}\|_{\max} = 1$ and $\text{KH}(\delta) \in \mathcal{G}_\nu(\mathbf{K})$ with ν defined in Prop. B.2, the result now follows from Thm. 1 with $r = (\log(n_{\text{in}}n_{\text{out}})/c)^{5d^*/2}$. \square

F. Proof of Thm. 2: Quality of Thinformer

Throughout we will make use of the convenient representation

$$\hat{\mathbf{T}} = \hat{\mathbf{D}}^{-1} \hat{\mathbf{A}} \mathbf{V} \text{ for } \mathcal{I}_{\text{out}} \triangleq \{i \in [n] : (\tilde{\mathbf{k}}_i, \tilde{\mathbf{v}}_i) \in \mathcal{X}_{\text{out}}\}, \quad \hat{\mathbf{A}} \triangleq \frac{n}{n_{\text{out}}} (\exp(\frac{\langle \mathbf{q}_i, \mathbf{k}_j \rangle}{\sqrt{d}}) \mathbf{1}[j \in \mathcal{I}_{\text{out}}])_{i,j=1}^n, \text{ and } \hat{\mathbf{D}} \triangleq \hat{\mathbf{A}} \mathbf{1}_n. \quad (38)$$

Our proof makes use of three lemmas. The first, proved in App. F.1, bounds the approximation error for the attention matrix \mathbf{T} in terms of the approximation error for \mathbf{AV} and $\mathbf{A1}_n$.

Lemma F.1 (Decomposing attention approximation error). In the notation of Alg. 1 and (38),

$$\|\hat{\mathbf{D}}^{-1} \hat{\mathbf{A}} \mathbf{V} - \mathbf{D}^{-1} \mathbf{A} \mathbf{V}\|_{\max} \leq \min\left(\|(\frac{1}{n} \mathbf{D})^{-1}\|_{\max}, \|(\frac{1}{n} \hat{\mathbf{D}})^{-1}\|_{\max}\right) \left(\frac{1}{n} \|\hat{\mathbf{A}} \mathbf{V} - \mathbf{A} \mathbf{V}\|_{\max} + \frac{1}{n} \|\mathbf{A1}_n - \hat{\mathbf{A}} \mathbf{1}_n\|_{\infty} \|\mathbf{V}\|_{\max}\right).$$

The second, proved in App. F.2, bounds the approximation error for \mathbf{AV} and $\mathbf{A1}_n$ in terms of the KMS (1) for a specific choice of attention kernel matrix.

Lemma F.2 (KMS bound on attention approximation error). Instantiate the notation of Alg. 1 and (38) and define the query set

$$\mathcal{X}' \triangleq \{\mathbf{x}_{i+nj} \triangleq (\tilde{\mathbf{q}}_i, \mathbf{e}_j^{d+1}) : i \in [n], j \in [d+1]\} \text{ where } \tilde{\mathbf{q}}_i \triangleq \mathbf{q}_i/d^{\frac{1}{4}}$$

and \mathbf{e}_j^{d+1} is the j -th standard basis vector in \mathbb{R}^{d+1} . If $\mathbf{K}_{\text{att}} \triangleq \mathbf{k}_{\text{att}}(\mathcal{X}, \mathcal{X})$ for $\mathcal{X} \triangleq \mathcal{X}' \cup \mathcal{X}_{\text{in}}$, then

$$\max\left(\frac{1}{n} \|\hat{\mathbf{A}} - \mathbf{A}\|_{\max}, \frac{1}{n} \|(\hat{\mathbf{A}} - \mathbf{A}) \mathbf{1}_n\|_{\infty} \|\mathbf{V}\|_{\max}\right) = \|\mathbf{K}_{\text{att}}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \text{ for } \mathcal{I} \triangleq [n(d+1)].$$

Our third lemma, proved in App. F.3, bounds the size of key parameters of the thinned attention problem.

Lemma F.3 (Thinned attention problem parameters). Instantiate the notation of Lem. F.2, and define $R \triangleq \max_{i \in [n]} \max(\|\mathbf{q}_i\|_2, \|\mathbf{k}_i\|_2)$. Then, for all $i, j \in \mathcal{I}$ and $l \in \text{supp}(\mathbf{p}_{\text{in}})$,

$$\begin{aligned} \|(\frac{1}{n} \mathbf{D})^{-1}\|_{\max} &\leq \exp\left(\frac{R^2}{\sqrt{d}}\right), \quad \max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}_{\text{att}}(\mathbf{x}, \mathbf{x})} \leq \exp\left(\frac{R^2}{2\sqrt{d}}\right) \sqrt{\|\mathbf{V}\|_{2,\infty}^2 + \|\mathbf{V}\|_{\max}^2}, \\ R_{\mathcal{I}} &\triangleq \max_{i \in \mathcal{I}} \|\mathbf{x}_i\|_2 \leq \sqrt{\frac{R^2}{\sqrt{d}} + 1}, \quad D_{\mathcal{I}} \triangleq \max_{i \in \mathcal{I}} \sqrt{\mathbf{K}_{\text{att},ii}} \leq \exp\left(\frac{R^2}{2\sqrt{d}}\right), \\ \text{rank}(\mathbf{X}_{\mathcal{I}}) &\leq d+1 \text{ for } \mathbf{X}_{\mathcal{I}} \triangleq [\mathbf{x}_i]_{i \in \mathcal{I}}^{\top}, \text{ and} \\ |\mathbf{K}_{\text{att},il} - \mathbf{K}_{\text{att},jl}| &\leq L_{\mathbf{K}_{\text{att}}} \|\mathbf{x}_i - \mathbf{x}_j\|_2 \text{ for } L_{\mathbf{K}_{\text{att}}} \triangleq \exp\left(\frac{R^2}{\sqrt{d}}\right) \sqrt{\frac{R^2}{\sqrt{d}} + 2} \|\mathbf{V}\|_{\max}. \end{aligned}$$

Now instantiate the notation of Lem. F.2, and define the coefficient

$$c \triangleq 2\sqrt{2} \left(32\sqrt{\frac{2}{3}} (d+1) \log(3e^2(\frac{R^2}{\sqrt{d}} + 2)\|\mathbf{V}\|_{\max}) + \sqrt{2\log(8)}(1 + \frac{32}{\sqrt{3}}) \right).$$

Together, Lem. F.3, the KMS quality bound of Thm. 1, and the KH-COMPRESS(0.5) sub-Gaussian constant ν of Prop. B.5 imply that, with probability at least $\frac{1}{2}$,

$$\|\mathbf{K}_{\text{att}}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} \leq \frac{c}{2\sqrt{2}} \exp(\frac{R^2}{\sqrt{d}}) \sqrt{\|\mathbf{V}\|_{2,\infty}^2 + \|\mathbf{V}\|_{\max}^2} \frac{\sqrt{\log_2(n_{\text{out}}) \log(8n_{\text{out}} \log_2 \frac{n_{\text{in}}}{n_{\text{out}}})}}{n_{\text{out}}}.$$

Hence, by Lems. F.1 and F.2, with probability at least $\frac{1}{2}$,

$$\begin{aligned} \|\hat{\mathbf{D}}^{-1}\hat{\mathbf{A}}\mathbf{V} - \mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max} &\leq \frac{c}{\sqrt{2}} \exp(\frac{2R^2}{\sqrt{d}}) \sqrt{\|\mathbf{V}\|_{2,\infty}^2 + \|\mathbf{V}\|_{\max}^2} \frac{\sqrt{\log_2(n_{\text{out}}) \log(8n_{\text{out}} \log_2 \frac{n_{\text{in}}}{n_{\text{out}}})}}{n_{\text{out}}} \\ &\leq c \exp(\frac{2R^2}{\sqrt{d}}) \|\mathbf{V}\|_{2,\infty} \frac{\sqrt{\log_2(n_{\text{out}}) \log(8n_{\text{out}} \log_2 \frac{n_{\text{in}}}{n_{\text{out}}})}}{n_{\text{out}}}. \end{aligned}$$

F.1. Proof of Lem. F.1: Decomposing attention approximation error

By the triangle inequality, we have

$$\|\hat{\mathbf{D}}^{-1}\hat{\mathbf{A}}\mathbf{V} - \mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max} \leq \|\hat{\mathbf{D}}^{-1}\hat{\mathbf{A}}\mathbf{V} - \hat{\mathbf{D}}^{-1}\mathbf{A}\mathbf{V}\|_{\max} + \|\hat{\mathbf{D}}^{-1}\mathbf{A}\mathbf{V} - \mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max}.$$

We bound the first term on the right-hand side using the submultiplicativity of the max norm under diagonal rescaling:

$$\|\hat{\mathbf{D}}^{-1}\hat{\mathbf{A}}\mathbf{V} - \hat{\mathbf{D}}^{-1}\mathbf{A}\mathbf{V}\|_{\max} \leq \|\hat{\mathbf{D}}^{-1}\|_{\max} \|\hat{\mathbf{A}}\mathbf{V} - \mathbf{A}\mathbf{V}\|_{\max} = \|(\frac{1}{n}\hat{\mathbf{D}})^{-1}\|_{\max} \frac{1}{n} \|\hat{\mathbf{A}}\mathbf{V} - \mathbf{A}\mathbf{V}\|_{\max}.$$

To bound the second term we use the same submultiplicativity property and the fact that each entry of $\mathbf{D}^{-1}\mathbf{A}\mathbf{V}$ is the average of values in \mathbf{V} :

$$\begin{aligned} \|\hat{\mathbf{D}}^{-1}\mathbf{A}\mathbf{V} - \mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max} &= \|\hat{\mathbf{D}}^{-1}(\mathbf{D} - \hat{\mathbf{D}})\mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max} \leq \|\hat{\mathbf{D}}^{-1}\|_{\max} \|\mathbf{D} - \hat{\mathbf{D}}\|_{\max} \|\mathbf{D}^{-1}\mathbf{A}\mathbf{V}\|_{\max} \\ &= \|(\frac{1}{n}\hat{\mathbf{D}})^{-1}\|_{\max} \frac{1}{n} \|\mathbf{A}\mathbf{1}_n - \hat{\mathbf{A}}\mathbf{1}_n\|_{\infty} \|\mathbf{V}\|_{\max}. \end{aligned}$$

An identical argument reversing the roles of (\mathbf{D}, \mathbf{A}) and $(\hat{\mathbf{D}}, \hat{\mathbf{A}})$ yields the second bound.

F.2. Proof of Lem. F.2: KMS bound on attention approximation error

Define the augmented value matrix $\tilde{\mathbf{V}} = [\mathbf{V}, \|\mathbf{V}\|_{\max}\mathbf{1}_n] \in \mathbb{R}^{d+1}$. By the definition of \mathbf{K}_{att} and $\hat{\mathbf{A}}$,

$$\|\mathbf{K}_{\text{att}}(\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})\|_{\mathcal{I}} = \max_{i \in [n], j \in [d+1]} |\sum_{\ell \in [n]} \mathbf{A}_{i\ell} \tilde{\mathbf{V}}_{\ell j} (\mathbf{p}_{\text{in}} - \mathbf{p}_{\text{out}})_{\ell}| = \frac{1}{n} \|(\mathbf{A} - \hat{\mathbf{A}}) \tilde{\mathbf{V}} e_j^d\|_{\infty} = \frac{1}{n} \|(\mathbf{A} - \hat{\mathbf{A}}) \tilde{\mathbf{V}}\|_{\max}.$$

F.3. Proof of Lem. F.3: Thinned attention problem parameters

First, by the Cauchy-Schwarz inequality and the nonnegativity of $\mathbf{D} = \mathbf{A}\mathbf{1}_n$ we have

$$\|(\frac{1}{n}\mathbf{D})^{-1}\|_{\max} = \frac{1}{\min_{i \in [n]} \frac{1}{n} \sum_{j \in [n]} \mathbf{A}_{ij}} \leq \frac{1}{\min_{i \in [n], j \in [n]} \exp(\frac{\langle \mathbf{q}_i, \mathbf{k}_j \rangle}{\sqrt{d}})} \leq \frac{1}{\min_{i \in [n], j \in [n]} \exp(\frac{-\|\mathbf{q}_i\|_2 \|\mathbf{k}_j\|_2}{\sqrt{d}})} \leq \exp(\frac{R^2}{\sqrt{d}}).$$

Second, the $\max_{\mathbf{x} \in \mathcal{X}_{\text{in}}} \sqrt{\mathbf{k}_{\text{att}}(\mathbf{x}, \mathbf{x})}$ inequality follows as

$$\mathbf{k}_{\text{att}}((\tilde{\mathbf{k}}_i, \tilde{\mathbf{v}}_i), (\tilde{\mathbf{k}}_i, \tilde{\mathbf{v}}_i)) = \exp(\frac{\|\mathbf{k}_i\|_2^2}{\sqrt{d}}) (\|\mathbf{v}_i\|_2^2 + \|\mathbf{V}\|_{\max}^2) \leq \exp(\frac{R^2}{\sqrt{d}}) (\|\mathbf{V}\|_{2,\infty}^2 + \|\mathbf{V}\|_{\max}^2).$$

Third, the $R_{\mathcal{I}}$ inequality follows as

$$\|(\tilde{\mathbf{q}}_i, e_j^{d+1})\|_2 = \sqrt{\|\tilde{\mathbf{q}}_i\|_2^2 + 1} \leq \sqrt{\frac{R^2}{\sqrt{d}} + 1} \quad \text{for all } i \in [n], j \in [d+1].$$

Fourth, the $D_{\mathcal{T}}$ inequality follows as

$$\max_{i \in \mathcal{I}} \mathbf{K}_{\text{att}, ii} = \max_{i \in [n]} \exp\left(\frac{\|\mathbf{q}_i\|_2^2}{\sqrt{d}}\right) \leq \exp\left(\frac{R^2}{\sqrt{d}}\right).$$

Fifth, the rank inequality follows as $\mathbf{x}_i \in \mathbb{R}^{d+1}$ for $i \in \mathcal{I}$. Finally, the Lipschitz inequality follows as, for any $i, k, l \in [n]$ and $j, m \in [d+1]$,

$$\begin{aligned} & |\exp\left(\frac{\langle \mathbf{q}_i, \mathbf{k}_l \rangle}{\sqrt{d}}\right) \langle \mathbf{e}_j^{d+1}, \tilde{\mathbf{v}}_l \rangle - \exp\left(\frac{\langle \mathbf{q}_k, \mathbf{k}_l \rangle}{\sqrt{d}}\right) \langle \mathbf{e}_m^{d+1}, \tilde{\mathbf{v}}_l \rangle| \\ & \leq \exp\left(\frac{\langle \mathbf{q}_i, \mathbf{k}_l \rangle}{\sqrt{d}}\right) |\tilde{\mathbf{v}}_{lj} - \tilde{\mathbf{v}}_{lm}| + |\exp\left(\frac{\langle \mathbf{q}_i, \mathbf{k}_l \rangle}{\sqrt{d}}\right) - \exp\left(\frac{\langle \mathbf{q}_k, \mathbf{k}_l \rangle}{\sqrt{d}}\right)| |\tilde{\mathbf{v}}_{lm}| \\ & \leq \exp\left(\frac{\|\mathbf{q}_i\|_2 \|\mathbf{k}_l\|_2}{\sqrt{d}}\right) \|\mathbf{e}_j^{d+1} - \mathbf{e}_m^{d+1}\|_2 \frac{|\tilde{\mathbf{v}}_{lj} - \tilde{\mathbf{v}}_{lm}|}{\sqrt{2}} + \exp\left(\frac{\max(\|\mathbf{q}_i\|_2, \|\mathbf{q}_k\|_2) \|\mathbf{k}_l\|_2}{\sqrt{d}}\right) \left| \frac{\langle \mathbf{q}_i - \mathbf{q}_k, \mathbf{k}_l \rangle}{\sqrt{d}} \right| |\tilde{\mathbf{v}}_{lm}| \\ & \leq \exp\left(\frac{R^2}{\sqrt{d}}\right) \|\mathbf{e}_j^{d+1} - \mathbf{e}_m^{d+1}\|_2 \frac{|\tilde{\mathbf{v}}_{lj} - \tilde{\mathbf{v}}_{lm}|}{\sqrt{2}} + \exp\left(\frac{R^2}{\sqrt{d}}\right) \frac{\|\mathbf{q}_i - \mathbf{q}_k\|_2 R}{\sqrt{d}} |\tilde{\mathbf{v}}_{lm}| \\ & \leq \exp\left(\frac{R^2}{\sqrt{d}}\right) \|\mathbf{e}_j^{d+1} - \mathbf{e}_m^{d+1}\|_2 \sqrt{2} \|\mathbf{V}\|_{\max} + \exp\left(\frac{R^2}{\sqrt{d}}\right) \frac{\|\mathbf{q}_i - \mathbf{q}_k\|_2 R}{\sqrt{d}} \|\mathbf{V}\|_{\max} \\ & \leq \exp\left(\frac{R^2}{\sqrt{d}}\right) \sqrt{\frac{R^2}{\sqrt{d}} + 2} \|\mathbf{V}\|_{\max} \|(\tilde{\mathbf{q}}_i, \mathbf{e}_j^{d+1}) - (\tilde{\mathbf{q}}_k, \mathbf{e}_m^{d+1})\|_2 \end{aligned}$$

by the triangle inequality, multiple applications of Cauchy-Schwarz, and the mean-value theorem applied to $x \mapsto e^x$.

G. Proof of Thm. 3: LKH-SGD convergence

Our proof makes use of three intermediate results. The first, inspired by [Harvey & Samadi \(2014, Thm. 10\)](#) and [Cooper et al. \(2023, Lem. 1\)](#), relates the quality of the ordering produced by Alg. 2 to the quality of the thinning.

Lemma G.1 (Quality of thinned reordering). *The output of thinned reordering (Alg. 2) satisfies*

$$\max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{x}_{\pi_{k+1}^{-1}(\pi_k^{-1}(i))}^k \right\|_2 \leq \frac{1}{2} \max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{x}_i^k \right\|_2 + \frac{1}{2} \max_{j \in [n]} \left\| \sum_{i=1}^j \eta_i^k \mathbf{x}_i^k \right\|_2 + \left\| \sum_{i=1}^n \mathbf{x}_i^k \right\|_2$$

where π_k^{-1} is the inverse permutation of π_k and $\eta_i^k \triangleq 2(\mathbf{1}[\mathbf{x}_i^k \in \mathcal{X}_{\text{out}}^k] - 1)$.

Proof. Fix any $j^* \in \arg\max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{x}_{\pi_{k+1}^{-1}(\pi_k^{-1}(i))}^k \right\|_2$. If $j^* \leq n/2$, then

$$2 \left\| \sum_{i=1}^{j^*} \mathbf{x}_{\pi_{k+1}^{-1}(\pi_k^{-1}(i))}^k \right\|_2 \leq 2 \max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{1}[\eta_i^k = 1] \mathbf{x}_i^k \right\|_2 \leq \max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{x}_i^k \right\|_2 + \max_{j \in [n]} \left\| \sum_{i=1}^j \eta_i^k \mathbf{x}_i^k \right\|_2$$

by the triangle inequality. Similarly, if $j^* > n/2$, then,

$$\begin{aligned} 2 \left(\left\| \sum_{i=1}^{j^*} \mathbf{x}_{\pi_{k+1}^{-1}(\pi_k^{-1}(i))}^k \right\|_2 - \left\| \sum_{i=1}^n \mathbf{x}_i^k \right\|_2 \right) & \leq 2 \left\| \sum_{i > j^*} \mathbf{x}_{\pi_{k+1}^{-1}(\pi_k^{-1}(i))}^k \right\|_2 \leq 2 \max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{1}[\eta_i^k = -1] \mathbf{x}_i^k \right\|_2 \\ & \leq \max_{j \in [n]} \left\| \sum_{i=1}^j \mathbf{x}_i^k \right\|_2 + \max_{j \in [n]} \left\| - \sum_{i=1}^j \eta_i^k \mathbf{x}_i^k \right\|_2. \end{aligned}$$

□

The second, a mild adaptation of [Cooper et al. \(2023, Thms. 2 and 3\)](#), bounds the convergence rate of SGD with thinned reordering in terms of the thinning quality.

Theorem G.1 (Convergence of SGD with thinned reordering). *Suppose that, for all $i \in [n]$ and $\mathbf{w}, \mathbf{v} \in \mathbb{R}^d$,*

$$\|\nabla f_i(\mathbf{w}) - \nabla f(\mathbf{w})\|_2^2 \leq \sigma^2 \quad \text{and} \quad \|\nabla f_i(\mathbf{w}) - \nabla f_i(\mathbf{v})\|_2 \leq L \|\mathbf{w} - \mathbf{v}\|_2$$

and that SGD (10) with thinned reordering (Alg. 2) satisfies the prefix discrepancy bound

$$\max_{j \in [n]} \left\| \sum_{i=1}^j \eta_i^k \mathbf{x}_i^k \right\|_2 \leq 2\tilde{A} \max_{i \in [n]} \|\mathbf{x}_i^k - \bar{\mathbf{x}}^k\|_2 \quad \text{for} \quad \eta_i^k \triangleq 2(\mathbf{1}[\mathbf{x}_i^k \in \mathcal{X}_{\text{out}}^k] - 1), \quad \bar{\mathbf{x}}^k \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^k, \quad (39)$$

and each epoch $k \in [K]$. Then the step size setting

$$\alpha = \min \left\{ \frac{1}{16L(2n+\tilde{A})}, \left(\frac{4F_1}{42L^2\sigma^2\tilde{A}^2nK+18L^2n^3\sigma^2} \right)^{1/3} \right\} \quad \text{with} \quad F_1 \triangleq f(\mathbf{w}_1) - f^* \quad \text{and} \quad f^* \triangleq \inf_{\mathbf{v} \in \mathbb{R}^d} f(\mathbf{v})$$

yields the convergence bound

$$\frac{1}{K} \sum_{k=1}^K \|\nabla f(\mathbf{w}_k)\|^2 \leq \frac{9(F_1 L \sigma \tilde{A})^{2/3}}{(nK)^{2/3}} + \frac{(72F_1 L \sigma)^{2/3} + 64F_1 L(2 + \tilde{A}/(n))}{K}.$$

If, in addition, f satisfies the μ -Polyak-Łojasiewicz (PL) condition,

$$\mu(f(\mathbf{w}) - f^*) \leq \frac{1}{2} \|\nabla f(\mathbf{w})\|_2^2 \quad \text{for all } \mathbf{w} \in \mathbb{R}^d,$$

and the number of epochs satisfies

$$K \geq 10 + \frac{1}{\mu} 32L(2 + \tilde{A}/n)\tilde{W} \quad \text{for } \tilde{W} \triangleq W_0(K^2 n^2 C_3) \quad \text{and} \quad C_3 \triangleq \frac{(F_1 + \sigma^2/L)\mu^2}{224L^2\sigma^2\tilde{A}^2},$$

where W_0 denotes the Lambert W function, then the step size setting $\alpha = \frac{2\tilde{W}}{Kn\mu}$ yields the convergence bound

$$f(\mathbf{w}_K) - f^* \leq \frac{1}{(nK)^2} \left(\frac{(F_1 + L^2\sigma^2)\tilde{W}}{C_3} + \frac{112L^2\sigma^2\tilde{A}^2\tilde{W}^2}{\mu^3} \right).$$

Proof. The proof is identical to that of Cooper et al. (2023, Thms. 2 and 3) with $m = 1$ worker once each instance of $\|\cdot\|_\infty$ is replaced with $\|\cdot\|_2$, each instance of $L_{2,\infty}$ is replaced with L , each instance of T is replaced with K , and Lem. G.1 is substituted for Cooper et al. (2023, Lem. 1). \square

The final result uses Thm. 1 to bound the prefix discrepancy of LKH(δ).

Lemma G.2 (LKH(δ) prefix discrepancy). *Fix any epoch $k \in [K]$. With probability at least $1 - \frac{\delta}{2} - \delta'$, thinned reordering (Alg. 2) with LKH(δ) satisfies the prefix discrepancy bound (39) with*

$$\tilde{A} = \sqrt{\log\left(\frac{2n(\log(n/2)+1)}{\delta}\right) [e^2 \text{rank}_{\epsilon_k}(\mathbf{X}^k) + 10e \log(\frac{n}{\delta'})]}$$

for $\mathbf{X}^k \triangleq [\mathbf{x}_1^k, \dots, \mathbf{x}_n^k]^\top$, $\epsilon_k \triangleq \max_{i \in [n]} \frac{\sqrt{9e \log(2n \log(en/2)/\delta) \log(n/\delta') \|\mathbf{x}_i^k - \bar{\mathbf{x}}^k\|_2}}{\sqrt{n}}$, and $\bar{\mathbf{x}}^k \triangleq \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^k$.

Proof. Define $\mathcal{X}^k = \{\mathbf{x}_1^k, \dots, \mathbf{x}_n^k\}$, $c = 2 \max_{i \in [n]} \|\mathbf{x}_i^k - \bar{\mathbf{x}}^k\|_2$, and $r = \text{rank}_{\epsilon_k}(\mathbf{X}^k)$. For any $j \in [n]$, we can write

$$\|\sum_{i=1}^j \eta_i^k \mathbf{x}_i^k\|_2 = \|\sum_{i=1}^j \mathbf{x}_i^k - \sum_{i=1}^j \mathbf{1}[\mathbf{x}_i^k \in \mathcal{X}_{\text{out},j}^k] \mathbf{x}_i^k\|_2 = 2j \|(\mathbf{X}^k)^\top (\mathbf{p}_{\text{in}}^j - \mathbf{p}_{\text{out}}^j)\|_2 = 2j \text{MMD}_{\mathbf{X}^k(\mathbf{X}^k)^\top}(\mathbf{p}_{\text{in}}^j, \mathbf{p}_{\text{out}}^j)$$

where \mathbf{p}_{in}^j and $\mathbf{p}_{\text{out}}^j$ are the empirical distributions over $\mathcal{X}_{\text{in},j}^k = (\mathbf{x}_i^k)_{i=1}^j$ and $\mathcal{X}_{\text{out},j}^k = \{\mathbf{x}_i^k \in \mathcal{X}_{\text{out}}^k : i \in [j]\}$.

Since LKH(δ) is an online algorithm that assigns signs $(\eta_i^k, \eta_{i+1}^k = 1 - \eta_i^k)$ to the points $(\mathbf{x}_i^k, \mathbf{x}_{i+1}^k)$ sequentially, we can view $\mathcal{X}_{\text{out},j}^k$ as the output of LKH(δ) applied to $\mathcal{X}_{\text{in},j}^k$ with $n_{\text{out}} = \frac{j}{2}$ and the linear kernel $\mathbf{k}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for each $j \in [n]$. Therefore, we may invoke the established LKH(δ) sub-Gaussian constants ν_j of Prop. B.3, Thm. 1, the union bound, and the definition of ϵ -rank (Def. 4) to deduce that

$$\max_{j \in [n]} \|\sum_{i=1}^j \eta_i^k \mathbf{x}_i^k\|_2^2 \leq \max_{j \in [n]} 4j^2 \nu_j^2 [e^2 r + e \log(\frac{n}{\delta'})] + \sigma_{r+1}(\mathbf{X}^k)^2 \frac{4j^2}{j} \leq c^2 \tilde{A}^2$$

with probability at least $1 - \frac{\delta}{2} - \delta'$. \square

Thm. 3 now follows directly from Thm. G.1 and Lem. G.2 applied to LKH($\frac{1}{2K}$) with $\delta' = \frac{1}{4K}$ and a union bound over epochs.

H. KT-COMPRESS(δ)

We describe the thinning algorithm KT-COMPRESS(δ) used in Alg. 3. We use KH(δ) for every halving round except for the last round, which thins a point sequence of size $2n_{\text{out}}$ to n_{out} . For this final halving round we use KH-REFINE(δ) (Alg. H.1) derived from the KT-SWAP algorithm of Dwivedi & Mackey (2024, Alg. 1a). The refinement stage of Alg. H.1 greedily improves the MMD of the initial KH(δ) output. Hence, $\text{MMD}_{\mathbf{k}}(\mathcal{X}_{\text{in}}, \mathcal{X}_{\text{out}}) \leq \text{MMD}_{\mathbf{k}}(\mathcal{X}_{\text{in}}, \mathcal{S}^{(1)})$ with probability 1.

Algorithm H.1: KH-REFINE(δ): KH(δ) with greedy refinement (Dwivedi & Mackey, 2024, Alg. 1a)

Input: point sequence $\mathcal{X}_{\text{in}} = (\mathbf{x}_i)_{i=1}^{n_{\text{in}}}$, kernel \mathbf{k} , input size $n_{\text{in}} \in 2\mathbb{N}$

$\mathcal{S} \leftarrow \text{KH}(\delta)(\mathcal{X}_{\text{in}}, \mathbf{k}); \quad n_{\text{out}} \triangleq n_{\text{in}}/2$

// Swap out each point in \mathcal{X}_{out} for the best alternative in \mathcal{X}_{in}

$\mathcal{X}_{\text{out}} \leftarrow \mathcal{S}.\text{copy}()$

for $\mathbf{x} \in \mathcal{S}$ **do**

$\mathcal{X}_{\text{out}} \leftarrow \mathcal{X}_{\text{out}} \setminus \{\mathbf{x}\} \cup \{\text{argmin}_{\mathbf{x}' \in \mathcal{X}_{\text{in}}} \text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{P}_{\text{out}} + \frac{1}{n_{\text{out}}}(\delta_{\mathbf{x}'} - \delta_{\mathbf{x}}))\}$

end

return \mathcal{X}_{out} , refined coreset of size $n_{\text{in}}/2$

I. Proof of Thm. 4: Low-rank analysis of CTT power

Thm. 4 follows from the following more detailed statement, proved in App. I.1 as

$$\mathbf{R}_{\mathbf{K}}^2(n_{\text{in}}, \frac{\tilde{\beta}}{20s_n}, \mathfrak{g}) + \mathbf{R}_{\mathbf{K}'}^2(n_{\text{in}}, \frac{\tilde{\beta}}{20s_n}, \mathfrak{g}) = O(\hat{\mathbf{R}}_{\mathbf{k}}^2).$$

Theorem I.1 (Low-rank analysis of CTT power, detailed). *Under the assumptions of Thm. 4 with $n_{\text{in}} \triangleq \frac{m+n}{s}$, CTT (Alg. 3) rejects with probability at least $1 - \beta$ whenever $c' \text{MMD}_{\mathbf{k}}(\mathbb{P}, \mathbb{Q})/\sqrt{\log(1/\gamma)}$ exceeds*

$$2c_{\tilde{\beta}/(20s)} \frac{\|\mathbf{k}\|_{\infty}^{\frac{1}{2}}}{\sqrt{m}} + \frac{\mathbf{R}_{\mathbf{k}}(\mathbb{P}, n_{\text{in}}, \frac{\tilde{\beta}}{20s_m}, \mathfrak{g}) + \mathbf{R}_{\mathbf{k}}(\mathbb{Q}, n_{\text{in}}, \frac{\tilde{\beta}}{20s_n}, \mathfrak{g})}{2^{\mathfrak{g}} \sqrt{m}}.$$

Here, $c' > 0$ is a universal constant, $c_{\delta} \triangleq 2 + \sqrt{2 \log(\frac{2}{\delta})}$, and $\mathbf{R}_{\mathbf{k}}(\mathbb{P}, n_{\text{in}}, \delta, \mathfrak{g})$ and $\mathbf{R}_{\mathbf{k}}(\mathbb{Q}, n_{\text{in}}, \delta, \mathfrak{g})$ respectively denote the $(1 - \frac{\delta}{2})$ -th quantiles of $\mathbf{R}_{\mathbf{K}}(n_{\text{in}}, \delta, \mathfrak{g})$ and $\mathbf{R}_{\mathbf{K}'}(n_{\text{in}}, \delta, \mathfrak{g})$, where

$$\begin{aligned} \mathbf{R}_{\tilde{\mathbf{K}}}^2(n_{\text{in}}, \delta, \mathfrak{g}) &\triangleq 256(\log_4 n_{\text{in}} - \mathfrak{g} - 1)(\sqrt{\log(n_{\text{in}} + 1)} + \sqrt{\log(2/\delta)})^2 \\ &\cdot \left(\frac{2\sqrt{\|\tilde{\mathbf{K}}\|_{\max}}}{\sqrt{3}} \left[\sqrt{e \log\left(\frac{6 \cdot 2^{\mathfrak{g}} \sqrt{n_{\text{in}}}(\log_4 n_{\text{in}} - \mathfrak{g})}{\delta}\right)} + \sqrt{\log\left(\frac{3n_{\text{in}}(\log_4 n_{\text{in}} - \mathfrak{g} - 1)}{\delta}\right)} \right] \right. \\ &\quad \left. + \min_{r \leq 2^{\mathfrak{g}+1} \sqrt{n_{\text{in}}}} \left\{ \frac{2\sqrt{\|\tilde{\mathbf{K}}\|_{\max}}}{\sqrt{3}} \sqrt{e^2 r \log\left(\frac{6 \cdot 2^{\mathfrak{g}} \sqrt{n_{\text{in}}}(\log_4 n_{\text{in}} - \mathfrak{g})}{\delta}\right)} + \sqrt{\lambda_{r+1}(\tilde{\mathbf{K}}) \cdot 2^{\mathfrak{g}-1} \sqrt{n_{\text{in}}}} \right\} \right)^2. \end{aligned} \quad (40)$$

I.1. Proof of Thm. I.1: Low-rank analysis of CTT power, detailed

Recall the following definition from Shetty et al. (2022, Def. 3).

Definition I.1 (k-sub-Gaussian thinning algorithm). *We say a thinning algorithm ALG (satisfying Def. 1) is k-sub-Gaussian on an event \mathcal{E} with shift a and parameter v if*

$$\mathbb{P}_{\mathcal{E}}(\text{MMD}_{\mathbf{k}}(\mathbb{P}_{\text{in}}, \mathbb{P}_{\text{out}}) \geq a + v\sqrt{t} \mid \mathcal{X}_{\text{in}}) \leq e^{-t} \quad \text{for all } t \geq 0.$$

Fix $\tilde{\mathbf{K}} \in \{\mathbf{K}, \mathbf{K}'\}$. To conclude our power result, it suffices, by Domingo-Enrich et al. (2023, Rmk. 2, App. B.1) and the failure probability setting of Domingo-Enrich et al. (2023, Lem. 11), to establish that

$$\begin{aligned} \mathbf{R}_{\tilde{\mathbf{K}}}^2(n_{\text{in}}, \delta, \mathfrak{g}) &= 256(\log_4 n_{\text{in}} - \mathfrak{g} - 1)(C_{\tilde{\mathbf{K}}}(\delta, 2^{\mathfrak{g}+1} \sqrt{n_{\text{in}}}) + \mathfrak{M}_{\tilde{\mathbf{K}}}(\delta, 2^{\mathfrak{g}+1} \sqrt{n_{\text{in}}}) \sqrt{\log\left(\frac{3n_{\text{in}}(\log_4 n_{\text{in}} - \mathfrak{g} - 1)}{\delta}\right)})^2 \\ &\cdot (\sqrt{\log(n_{\text{in}} + 1)} + \sqrt{\log(2/\delta)})^2, \end{aligned} \quad (41)$$

for any scalars $C_{\tilde{\mathbf{K}}}(\delta, 2^{\mathfrak{g}+1} \sqrt{n_{\text{in}}})$ and $\mathfrak{M}_{\tilde{\mathbf{K}}}(\delta, 2^{\mathfrak{g}+1} \sqrt{n_{\text{in}}})$ satisfying the property that, on an event of probability at least $1 - \delta/2$, every call to HALVE \triangleq KT-SPLIT($\frac{\ell^2}{n_{\text{in}} 4^{\mathfrak{g}+1} (\log_4 n_{\text{in}} - \mathfrak{g})} \delta$) with input size ℓ and output size $\ell/2$ is k-sub-Gaussian (Def. I.1) with shift $a_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ and parameter $v_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ satisfying

$$a_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}} = \frac{C_{\tilde{\mathbf{K}}}(\delta, \ell)}{\ell/2} \quad \text{and} \quad v_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}} = \frac{\mathfrak{M}_{\tilde{\mathbf{K}}}(\delta, \ell)}{\ell/2} \sqrt{\log\left(\frac{12n_{\text{in}} 4^{\mathfrak{g}} (\log_4 n_{\text{in}} - \mathfrak{g})}{\ell \delta}\right)}.$$

Substituting $\mathfrak{M}_{\tilde{\mathbf{K}}}(\delta, 2^{g+1}\sqrt{n_{\text{in}}}) = (2^g\sqrt{n_{\text{in}}})v_{2^{g+1}\sqrt{n_{\text{in}}}, n_{\text{in}}, \tilde{\mathbf{K}}} \left[\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{2^{g+1}\sqrt{n_{\text{in}}}\delta}\right) \right]^{-\frac{1}{2}}$ and $C_{\tilde{\mathbf{K}}}(\delta, 2^{g+1}\sqrt{n_{\text{in}}}) = (2^g\sqrt{n_{\text{in}}})a_{2^{g+1}\sqrt{n_{\text{in}}}, n_{\text{in}}, \tilde{\mathbf{K}}}$ into (41), we obtain the sufficient condition

$$\begin{aligned} \mathbf{R}_{\tilde{\mathbf{K}}}^2(n_{\text{in}}, \delta, g) &= 256(\log_4 n_{\text{in}} - g - 1) \cdot (2^g\sqrt{n_{\text{in}}})^2 \cdot (\sqrt{\log(n_{\text{in}} + 1)} + \sqrt{\log(2/\delta)})^2 \\ &\quad \cdot \left(a_{2^{g+1}\sqrt{n_{\text{in}}}, n_{\text{in}}, \tilde{\mathbf{K}}} + v_{2^{g+1}\sqrt{n_{\text{in}}}, n_{\text{in}}, \tilde{\mathbf{K}}} \left[\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{2^{g+1}\sqrt{n_{\text{in}}}\delta}\right) \right]^{-\frac{1}{2}} \sqrt{\log\left(\frac{3n_{\text{in}}(\log_4 n_{\text{in}} - g - 1)}{\delta}\right)} \right)^2. \end{aligned} \quad (42)$$

We now identify suitable $a_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ and $v_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ with the aid of the following lemma, proved in App. I.2.

Lemma I.1 ((\mathbf{K}, ν, δ)-sub-Gaussian thinning algorithms are k-sub-Gaussian). Suppose ALG is a (\mathbf{K}, ν, δ)-sub-Gaussian thinning algorithm, satisfying Def. 3 with an event \mathcal{E} of probability at least $1 - \delta/2$. Then ALG is k-sub-Gaussian (Def. I.1) on \mathcal{E} with shift $a_{n_{\text{out}}, n_{\text{in}}, \mathbf{K}}$ and parameter $v_{n_{\text{out}}, n_{\text{in}}, \mathbf{K}}$ defined as

$$a_{n_{\text{out}}, n_{\text{in}}, \mathbf{K}} \triangleq \nu\sqrt{e} + \min_{r \leq n_{\text{in}}} \left\{ \nu\sqrt{e^2 r} + \sqrt{\lambda_{r+1}(\mathbf{K})\left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}\right)} \right\} \quad \text{and} \quad v_{n_{\text{out}}, n_{\text{in}}, \mathbf{K}} \triangleq \nu\sqrt{e}.$$

By Prop. B.4 and Lem. A.1, $\text{KH}\left(\frac{\ell^2}{n_{\text{in}}4^{g+1}(\log_4 n_{\text{in}} - g)}\delta\right)$ with input size ℓ and output size $\ell/2$ is a ($\mathbf{K}, \nu, \frac{\ell^2}{n_{\text{in}}4^{g+1}(\log_4 n_{\text{in}} - g)}\delta$)-sub-Gaussian thinning algorithm with

$$\nu \leq \frac{2}{(\ell/2)\sqrt{3}} \sqrt{\log\left(\frac{6(\ell/2)\log_2(\ell/(\ell/2))}{\delta} \cdot \frac{n_{\text{in}}4^{g+1}(\log_4 n_{\text{in}} - g)}{\ell^2}\right)} \|\mathbf{K}\|_{\max} = \frac{2}{(\ell/2)\sqrt{3}} \sqrt{\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{\ell\delta}\right)} \|\mathbf{K}\|_{\max}.$$

By Lem. I.1, on an event of probability at least $1 - \frac{\ell^2}{2n_{\text{in}}4^{g+1}(\log_4 n_{\text{in}} - g)}\delta$, $\text{KH}\left(\frac{\ell^2}{n_{\text{in}}4^{g+1}(\log_4 n_{\text{in}} - g)}\delta\right)$ with input size ℓ and output size $\ell/2$ is a k-sub-Gaussian thinning algorithm with shift $a_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ and parameter $v_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}}$ defined as

$$\begin{aligned} a_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}} &= \frac{2}{(\ell/2)\sqrt{3}} \sqrt{\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{\ell\delta}\right)} \|\tilde{\mathbf{K}}\|_{\max} \sqrt{e \log 2} \\ &\quad + \min_{r \leq \ell} \left\{ \frac{2}{(\ell/2)\sqrt{3}} \sqrt{\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{\ell\delta}\right)} \|\tilde{\mathbf{K}}\|_{\max} \sqrt{e^2 r} + \sqrt{\lambda_{r+1}(\tilde{\mathbf{K}})\left(\frac{1}{\ell/2} - \frac{1}{\ell}\right)} \right\} \quad \text{and} \end{aligned} \quad (43)$$

$$v_{\ell, n_{\text{in}}, \tilde{\mathbf{K}}} = \frac{2}{(\ell/2)\sqrt{3}} \sqrt{\log\left(\frac{12n_{\text{in}}4^g(\log_4 n_{\text{in}} - g)}{\ell\delta}\right)} \|\tilde{\mathbf{K}}\|_{\max} \sqrt{e}. \quad (44)$$

Moreover, by the union bound, as detailed in Shetty et al. (2022, App. F.1), every HALVE call made by KT-COMPRESS is simultaneously k-sub-Gaussian with these input-size-dependent parameters on a common event of probability at least $1 - \frac{\delta}{2}$. Substituting (43) and (44) with $\ell = 2^{g+1}\sqrt{n_{\text{in}}}$ into (42), we obtain our error inflation factor expression (40), completing the proof.

I.2. Proof of Lem. I.1: (\mathbf{K}, ν, δ)-sub-Gaussian thinning algorithms are k-sub-Gaussian

Fix any $t \geq 0$, and let $\delta' = e^{-t}$. By our sub-Gaussian assumption, Thm. 1 implies that, as advertised,

$$\begin{aligned} e^{-t} &\geq \mathbb{P}_{\mathcal{E}}\left(\text{MMD}_{\mathbf{K}}^2(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \geq \min_{r \leq n_{\text{in}}} \nu^2[e^2 r + et] + \lambda_{r+1}(\mathbf{K})\left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}\right)\right) \\ &= \mathbb{P}_{\mathcal{E}}\left(\text{MMD}_{\mathbf{K}}(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \geq \min_{r \leq n_{\text{in}}} \sqrt{\nu^2[e^2 r + et] + \lambda_{r+1}(\mathbf{K})\left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}\right)}\right) \\ &\geq \mathbb{P}_{\mathcal{E}}\left(\text{MMD}_{\mathbf{K}}(\mathbf{p}_{\text{in}}, \mathbf{p}_{\text{out}}) \geq \nu\sqrt{e}\sqrt{t} + \min_{r \leq n_{\text{in}}} \nu\sqrt{e^2 r} + \sqrt{\lambda_{r+1}(\mathbf{K})\left(\frac{1}{n_{\text{out}}} - \frac{1}{n_{\text{in}}}\right)}\right). \end{aligned}$$

J. Proof of Cor. 3: Power of deep kernel CTT

Define the radius

$$R' \triangleq \max_{\mathbf{y} \in \mathcal{Y} \cup \mathcal{X}} \|(\phi(\mathbf{y}), \mathbf{y})\|_2,$$

the augmented vectors $\mathcal{Y}' \triangleq \{(\phi(\mathbf{y}), \mathbf{y})\}_{\mathbf{y} \in \mathcal{Y}}$, and the augmented kernel

$$q'((\phi(\mathbf{x}), \mathbf{x}), (\phi(\mathbf{y}), \mathbf{y})) \triangleq \kappa(\phi(\mathbf{x}), \phi(\mathbf{y}))q(\mathbf{x}, \mathbf{y}) = \exp(-\eta\|\phi(\mathbf{x}), \mathbf{x}) - (\phi(\mathbf{y}), \mathbf{y})\|_2^2).$$

Since the deep kernel (13) takes the form

$$\mathbf{k}_{\text{deep}}(\mathbf{x}, \mathbf{y}) = (1 - \epsilon)q'((\phi(\mathbf{x}), \mathbf{x}), (\phi(\mathbf{y}), \mathbf{y})) + \epsilon q(\mathbf{x}, \mathbf{y})$$

we also have

$$\mathbf{K}_{\text{deep}} \triangleq \mathbf{k}_{\text{deep}}(\mathcal{Y}, \mathcal{Y}) = (1 - \epsilon)\mathbf{Q}' + \epsilon\mathbf{Q} \quad \text{for} \quad \mathbf{Q}' \triangleq q'(\mathcal{Y}', \mathcal{Y}') \quad \text{and} \quad \mathbf{Q} \triangleq q(\mathcal{Y}, \mathcal{Y}).$$

Hence, by Weyl's inequality (Horn & Johnson, 1985, Thm. 4.3.1) and the Gaussian kernel matrix eigenvalue bound (6),

$$\lambda_{2r+1}(\mathbf{K}_{\text{deep}}) \leq (1 - \epsilon)\lambda_{r+1}(\mathbf{Q}') + \epsilon\lambda_{r+1}(\mathbf{Q}) \leq ne^{-\frac{d'}{2e}r^{1/d'} \log\left(\frac{d'x^{1/d'}}{4e^2\eta R'^2}\right)} \quad \text{for} \quad (2e)^{d'} \leq r < n.$$

Parallel reasoning and the assumption $m \leq n$ yield the same bound for $\lambda_{2r+1}(\mathbf{k}_{\text{deep}}(\mathcal{X}, \mathcal{X}))$ and $(2e)^{d'} \leq r < m$. Now consider the approximate rank parameter

$$r^* \triangleq \max\left\{\left\lceil \frac{2e}{d'} \log(nn_{\text{out}}b) \right\rceil^{d'}, (R'^2\eta e^3 4/d')^{d'}\right\} \quad \text{for} \quad b \triangleq \frac{1}{2}. \quad (45)$$

Then, for $n \geq (2e)^{d'}$, we have, exactly as in App. D,

$$\lambda_{2r^*+1}(\mathbf{K}_{\text{deep}}) + \lambda_{2r^*+1}(\mathbf{k}_{\text{deep}}(\mathcal{X}, \mathcal{X})) \leq \frac{2}{n_{\text{out}}b}$$

and therefore

$$\hat{\mathbf{R}}_{\mathbf{K}} = O\left(\sqrt{\log\left(\frac{n}{s}\right)} \log\left(\frac{n}{\beta}\right) \max\left\{\left\lceil \frac{2e}{d'} \log(nn_{\text{out}}b) \right\rceil^{d'/2}, (R'^2\eta e^3 4/d')^{d'/2}\right\}\right).$$

Our final step is to bound the quantile of the sole remaining data-dependent term, R' . Since the inputs are c -sub-Gaussian (14), Lem. 1 of Dwivedi & Mackey (2024) with $\psi^{-1}(r) = \frac{\sqrt{\log r}}{\sqrt{c}}$ implies that the $1 - \frac{\tilde{\beta}}{20s_n}$ quantile of R' is $O\left(\sqrt{\log\left(\frac{n}{\beta}\right)}\right)$, yielding the result.

K. Proof of Cor. 4: Power of deep manifold kernel CTT

Our reasoning is identical to that in App. J with the manifold Gaussian kernel matrix eigenvalue bound (7) now substituted for the Euclidean ball bound (6) and the approximate rank setting $r^* = (\log(nn_{\text{out}})/c)^{5d^*/2}$ substituted for (45).

L. Supplementary Experiment Details

L.1. Approximating attention experiments

The experiment of Tab. 3 was carried out using Python 3.12.9, PyTorch 2.8.0.dev20250407+cu128 (Paszke et al., 2019), and an Ubuntu 22.04.5 LTS server with an AMD EPYC 7V13 64-Core Processor, 220 GB RAM, and a single NVIDIA A100 GPU (80 GB memory, CUDA 12.8, driver version 570.124.04). For reference, attention layer 1 has $(n, d) = (3136, 64)$ and attention layer 2 has $(n, d) = (784, 64)$. For each of the first 50 ImageNet 2012 validation set batches of size 64, we measured the time required to complete a forward pass through each the approximate softmax matrix (8) layer using CUDA events following 10 warm-up batches to initialize the GPU. Tab. L.1 provides the hyperparameter settings for each attention approximation in Tab. 3.

The experiment of Tab. 4 was carried out using Python 3.12.9, PyTorch 2.6.0, and an Ubuntu 22.04.5 LTS server with an Intel(R) Xeon(R) Gold 5218 CPU Processor, 1000 GB RAM, and a single NVIDIA A6000 GPU (48 GB memory, CUDA 12.1, driver version 530.30.02). For reference, the BigGAN model contains a single attention layer with 4096 queries in \mathbb{R}^{64} , 1024 keys in \mathbb{R}^{64} , and 1024 values in \mathbb{R}^{256} . We measured the time required to complete a forward pass through the approximate softmax matrix (8) layer using CUDA events following 10 warm-up batches to initialize the GPU. Tab. L.2 provides the hyperparameter settings for each attention approximation in Tab. 4.

The settings and implementations for all methods other than Thinformer were provided by Zandieh et al. (2023), and our experiment code builds on their open-source repository <https://github.com/majid-daliri/kdeforner>.

Table L.1: Configurations for the attention approximation methods of Tab. 3.

Attention Algorithm	Layer 1 Configuration	Layer 2 Configuration
Performer	num_features=49	num_features=12
Reformer	bucket_size=49 n_hashes=2	bucket_size=12 n_hashes=2
ScatterBrain	local_context=49 num_features=48	local_context=12 num_features=6
KDEformer	sample_size=64 bucket_size=32	sample_size=56 bucket_size=32
Thinformer (Ours)	g=2	g=4

Table L.2: Configurations for the attention approximation methods of Tab. 4.

Attention Algorithm	Layer Configuration
Performer	num_features=128
Reformer	bucket_size=64 n_hashes=8
ScatterBrain	local_context=32 num_features=128
KDEformer	sample_size=128 bucket_size=64
Thinformer (Ours)	g=2

L.2. Faster SGD training experiments

The experiment of Sec. 5.2 was carried out using Python 3.10, PyTorch 2.0.1, a Rocky Linux 8.9 server with 64 CPU cores (Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz), and a NVIDIA A100 GPU (40 GB memory, CUDA 12.4, driver version 550.54.15).

Technically, the CD-GraB: SBW algorithm requires an a priori upper bound on the maximum Euclidean norm b_{\max} of any stochastic gradient that it will encounter. To conduct our experiment, we first estimate b_{\max} by calculating the maximum gradient Euclidean encountered across 10 epochs of running SGD with LKH($\frac{1}{2K}$) reordering. One would typically not choose to carry out such a two-step procedure in practice, but the experiment serves to demonstrate that the CD-GraB: SBW leads to overly conservative performance even if reasonable upper bound is known in advance.

The settings and implementation for both random reshuffling (RR) and CD-GraB: Greedy were those used in the original logistic regression on mortgage application experiment of Cooper et al. (2023). Our experiment code builds on the open-source CD-GraB repository <https://github.com/GarlGuo/CD-GraB>. As in Cooper et al. (2023), optimization was carried out with a learning rate of $\alpha = 0.01$, datapoints were loaded in batches of size 16, and stochastic gradients were reordered for each datapoint individually.

L.3. Cheap two-sample testing experiment

The experiment of Sec. 6.2 was carried out using Python 3.12.9, PyTorch 2.6.0+cu124, and an Ubuntu 20.04.6 LTS server with an AMD EPYC 9554 64-Core Processor, 100 GB RAM, and a single NVIDIA H100 GPU (96 GB memory, CUDA 12.2, driver version 535.154.05). Each test is run with replication count $B = 100$, nominal level $\alpha = 0.05$, and failure probability $\delta = 0.5$. The neural network ϕ was trained exactly as in Liu et al. (2020) (with learning rate 5×10^{-5} and batch size equal to the full training sample size), and runtime measurements exclude the time required to train ϕ . Our experiment code builds on the open-source deep kernel testing (<https://github.com/fengliu90/DK-for-TST>) and Compress Then Test (<https://github.com/microsoft/goodpoints>) repositories.




Figure L.1: **Approximate low-rank structure of stochastic gradient matrices.** *Left.* For each epoch k of the LKH($\frac{1}{2K}$) experiment of Sec. 5.2, we record the ϵ_k -rank of the stochastic gradient matrix $\mathbf{X}^k \triangleq [\mathbf{x}_1^k, \dots, \mathbf{x}_n^k]^\top \in \mathbb{R}^{n \times d}$ (see Def. 4 and Thm. 3). Notably, the ϵ_k -ranks are significantly smaller than the ambient dimension $d = 19$. *Right.* We display the singular values of the first stochastic gradient matrix, \mathbf{X}^1 . The singular values drop off steeply, resulting in relatively small ϵ_k -ranks.




Figure L.2: **Slow-growing error inflation.** In the experimental setting of Sec. 6.2, the error inflation factor $\hat{\mathbf{R}}_{\mathbf{k}_{\text{deep}}}$ enjoys $O(\log^5(n))$ growth as the ratio $\hat{\mathbf{R}}_{\mathbf{k}_{\text{deep}}} / \log^5(n)$ decreases with n . Here, $\hat{\mathbf{R}}_{\mathbf{k}_{\text{deep}}}^2$ is defined by (12) with $n = m$, $n_{\text{in}} \triangleq \frac{2n}{s}$ and $n_{\text{out}} \triangleq 2^g \sqrt{n_{\text{in}}}$ for constants $\beta = 0.05$, $s = 32$, and $g = 4$.