

Near-Optimal Coresets of Kernel Density Estimates

Jeff M. Phillips¹

School of Computing, University of Utah
Salt Lake City, USA
jeffp@cs.utah.edu

Wai Ming Tai

School of Computing, University of Utah
Salt Lake City, USA
wmtai@cs.utah.edu

Abstract

We construct near-optimal coresets for kernel density estimate for points in \mathbb{R}^d when the kernel is positive definite. Specifically we show a polynomial time construction for a coreset of size $O(\sqrt{d} \log(1/\varepsilon)/\varepsilon)$, and we show a near-matching lower bound of size $\Omega(\sqrt{d}/\varepsilon)$. The upper bound is a polynomial in $1/\varepsilon$ improvement when $d \in [3, 1/\varepsilon^2)$ (for all kernels except the Gaussian kernel which had a previous upper bound of $O((1/\varepsilon) \log^d(1/\varepsilon))$) and the lower bound is the first known lower bound to depend on d for this problem. Moreover, the upper bound restriction that the kernel is positive definite is significant in that it applies to a wide-variety of kernels, specifically those most important for machine learning. This includes kernels for information distances and the sinc kernel which can be negative.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Coresets, Kernel Density Estimate, Discrepancy

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.66


1 Introduction

Kernel density estimates are pervasive objects in data analysis. They are the classic way to estimate a continuous distribution from a finite sample of points [28, 27]. With some negative weights, they are the prediction function in kernel SVM classifiers [25]. They are the core of many robust topological reconstruction approaches [22, 12, 6]. And they arise in many other applications including mode estimation [1], outlier detection [26], regression [11], and clustering [23].

Generically, consider a dataset $P \subset \mathbb{R}^d$ of size n , and a kernel $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, for instance the Gaussian kernel $K(x, p) = \exp(-\alpha^2 \|x - p\|^2)$ with $1/\alpha$ as a bandwidth parameter. Then a kernel density estimate is defined at any point $x \in \mathbb{R}^d$ as $\text{KDE}_P(x) = \frac{1}{n} \sum_{p \in P} K(x, p)$.

Given that it takes $O(n)$ time to evaluate KDE_P , and that data sets are growing to massive sizes, in order to continue to use these powerful modeling objects, a common approach is to replace P with a much smaller data sets Q so that KDE_Q approximates KDE_P . While statisticians have classically studied various sorts of average deviations (L_2 [28, 27] or L_1 error [9]), for most modern data modeling purposes, a worst-case L_∞ is more relevant (e.g., for preserving classification margins [25], density estimates [30], topology [22], and hypothesis

¹ Thanks to supported by NSF CCF-1350888, IIS-1251019, ACI-1443046, CNS-1514520, and CNS-1564287



testing on distributions [13]). Specifically this error guarantee preserves

$$\| \text{KDE}_P - \text{KDE}_Q \|_\infty = \max_{x \in \mathbb{R}^d} | \text{KDE}_P(x) - \text{KDE}_Q(x) | \leq \varepsilon.$$

We call such a set Q an ε -KDE coreset of P .

Traditionally the approximate set Q has been considered to be constructed as a random sample of P [28, 27, 16], sometimes known as a Nyström approximation [10]. However, in the last decade, a slew of data-aware approaches have been developed that can obtain a set Q with the same L_∞ error guarantee, but with considerably smaller size.

To describe either the random sample results or the data-aware approaches, we first need to be more specific about the properties of the kernel functions. We start with *positive definite* kernels, the central class required for most machine learning approaches to work [15].

Postive definite kernels. Consider a kernel $K : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ defined over some domain \mathcal{D} (often \mathbb{R}^d). It is called a *positive definite kernel* if, for any m points $x_1, x_2, \dots, x_m \in \mathcal{D}$, they are used to define an $m \times m$ Gram matrix G so each i, j entry is defined $G_{i,j} = K(x_i, x_j)$, and the matrix G is positive definite. Recall, a

■ **Table 1** Example Positive Definite Kernels

	$K(x, p) =$	domain
Gaussian	$\exp(-\alpha^2 \ x - p\ ^2)$	\mathbb{R}^d
Laplacian	$\exp(-\alpha \ x - p\)$	\mathbb{R}^d
Exponential	$\exp(-\alpha(1 - \langle x, p \rangle))$	\mathbb{S}^d
JS	$\exp(-\alpha(H(\frac{x+p}{2}) - \frac{H(x)+H(p)}{2}))$	Δ^d
Helinger	$\exp(-\alpha \sum_{i=1}^d (\sqrt{x_i} - \sqrt{p_i})^2)$	Δ^d
Sinc	$\frac{\sin(\alpha \ x-p\)}{\alpha \ x-p\ }$	$\mathbb{R}^{d \leq 3}$

matrix G is positive definite if any vector $z \in \mathbb{R}^n$ that is not all zeros satisfies $z^T G z > 0$. Moreover, a positive definite matrix G can always be decomposed as a product $H^T H$ with real-valued matrix H .

There are many positive definite kernels, and we will next highlight a few. We normalize all kernels so $K(x, x) = 1$ for all $x \in \mathcal{D}$ and therefore $|K(x, y)| \leq 1$ for all $x, y \in \mathcal{D}$. We will use $\alpha > 0$ as a parameter, where $1/\alpha$ represents the bandwidth, or smoothness of the kernel. For $\mathcal{D} = \mathbb{R}^d$ the most common positive definite kernels [29] are the Gaussian (described earlier) and the Laplacian, defined $\exp(-\alpha \|x - y\|)$ for $x, y \in \mathbb{R}^d$. Another common domain is $\Delta^d = \{x \in \mathbb{R}^{d+1} \mid \sum_{i=1}^d x_i = 1, x_i \geq 0\}$, for instance in representing discrete distributions such as normalized counts of words in a text corpus or fractions of tweets per geographic region. Common positive definite kernels for $x, y \in \Delta^d$ include the Hellinger kernel $\exp(-\alpha \sum_{i=1}^d (\sqrt{x_i} - \sqrt{y_i})^2)$ and the Jensen-Shannon (JS) divergence kernel $\exp(-\alpha(H(\frac{x+y}{2}) - \frac{H(x)+H(y)}{2}))$, where $H(x) = \sum_{i=1}^d -x_i \log x_i$ is entropy [14]. In other settings it is more common to normalize data points x to lie on a sphere $\mathbb{S}^d = \{x \in \mathbb{R}^{d+1} \mid \|x\| = 1\}$. Then with $x, y \in \mathbb{S}^d$, the exponential kernel $\exp(-\alpha(1 - \langle x, y \rangle))$ is positive definite [15]. Perhaps surprisingly, positive definite kernels do not need to satisfy $K(x, y) \geq 0$. For $x, y \in \mathbb{R}^d$, the sinc kernel is defined as $\frac{\sin(\alpha \|x-y\|)}{\alpha \|x-y\|}$ and is positive definite for $d = \{1, 2, 3\}$ [24].

Other classes of kernels. There are other ways to characterize kernels, which provide sufficient conditions for various other coreset bounds. For clarity, we describe these for kernels with a domain of \mathbb{R}^d , but they can apply more generally.

We say a kernel K is C_K -Lipschitz if, for any $x, y, z \in \mathbb{R}^d$, $|K(x, z) - K(y, z)| \leq C_K \cdot \|x - y\|$. This ensures that the kernels do not fluctuate too widely, a necessity for robustness, but also prohibits “binary” kernels; for instance the *ball kernel* is defined $K(x, y) = \{1 \text{ if } \|x - y\| \leq r; \text{ and } 0 \text{ otherwise}\}$. Such binary kernels are basically range counting queries (for instance

the ball kernel corresponds with a range defined by a ball), and as we will see, this distinction allows the bounds for ε -KDE coresets to surpass lower bounds for coresets for range counting queries. Aside from the ball kernel, all kernels we discuss in this paper will be C_K -Lipshcitz.

Another way to characterize a kernel is with their shape. We can measure this by considering binary ranges defined by super-level sets of kernels. For instance, given a fixed K and $x \in \mathbb{R}^d$, and a threshold $\tau \in \mathbb{R}$, the *super-level set* is $\{p \in \mathbb{R}^d \mid K(x, p) \geq \tau\}$. For a fixed K , the family of such sets over all choices of x and τ describes a range space with ground set \mathbb{R}^d . For many kernels the VC-dimension of this range space is bounded; in particular, for common kernels, this range is equivalent to those defined by balls in \mathbb{R}^d . Notably, the sinc kernel, which is positive-definite for \mathbb{R}^d with $d \leq 3$ does not correspond to a range space with bounded VC-dimension.

Finally, we mention that kernels being *characteristic* [29] is an important property for many bounds that rely on mappings from data points $x \in \mathbb{R}^d$ to a function space $K(x, \cdot) = \phi_K(x)$. A characteristic kernel requires that the kernel K is positive definite, and the mapping $\phi_K(x)$ is isomorphic and hence its induced distance $D_K(p, x) = \sqrt{\|\phi_K(x)\|_K^2 + \|\phi_K(p)\|_K^2 - 2\langle \phi_K(p), \phi_K(x) \rangle_K}$ is a metric.

Discrepancy-based approaches. Our approach for creating an ε -KDE coreset will follow a technique for creating range counting coresets [8, 19, 5]. It focuses on assigning a coloring $\chi : P \rightarrow \{-1, +1\}$ to P . Then retains either all $P_+ = \{p \in P \mid \chi(p) = +1\}$ or the remainder P_- , and recursively applies this halving until a small enough coreset Q has been retained.

Classically, when the goal is to compute a range counting coreset for a range space (P, \mathcal{R}) , then the specific goal of the coloring is to minimize discrepancy

$$\text{disc}_R(P, \chi) = \left| \sum_{p \in P \cap R} \chi(p) \right|$$

over all choices of ranges $R \in \mathcal{R}$. In the KDE-setting we consider a *kernel range space* (P, \mathcal{K}) where $\mathcal{K} = \{K(x, \cdot) \mid x \in \mathcal{D}\}$ defined by kernel $K : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}$ and a fixed domain \mathcal{D} which is typically assumed, and usually $\mathcal{D} = \mathbb{R}^d$. We instead want to minimize the kernel discrepancy

$$\text{disc}(P, \chi, x) = \left| \sum_{p \in P} \chi(p) K(x, p) \right|.$$

Now in contrast to the case with the binary range space (P, \mathcal{R}) , *each* point $p \in P$ is *partially* inside the “range” where the amount inside is controlled by the kernel K . Understanding the quantity

$$\text{disc}(n, \mathcal{K}) = \max_{P:|P|=n} \min_{\chi} \max_{x \in \mathcal{D}} \text{disc}(P, \chi, x)$$

is key. If for a particular \mathcal{K} we have $\text{disc}(n, \mathcal{K}) = n^\tau$ or $\text{disc}(n, \mathcal{K}) = \log^\eta n$, then applying the recursive halving algorithm obtains an ε -KDE coreset of size $O(1/\varepsilon^{1/(1-\tau)})$ and $O((1/\varepsilon) \log^\eta(1/\varepsilon))$, respectively [20].

Previous results. We now highlight some previous size bounds for ε -KDE coresets as shown in Table 2; see also Phillips and Tai [21] for a more thorough review of the historical progress on this problem. For simplicity, we report the size when only a constant probability of failure bound is required for randomized algorithms.

For kernels with bounded VC-dimension of their super-level sets, $O(d/\varepsilon^2)$ random samples are sufficient to create an ε -KDE coreset [16]. For characteristic kernels [13] a bound of $O(1/\varepsilon^4)$, independent of dimension is possible.

■ **Table 2** Asymptotic ε -KDE coreset sizes in terms of ε and d . SISS = Shift- and rotation invariant, and somewhere-steep (see Section 3).

Paper	Coreset Size	Restrictions and Notes
Joshi <i>et al.</i> [16]	d/ε^2	bounded VC; random sample
Gretton <i>et al.</i> [13]	$1/\varepsilon^4$	characteristic kernels; random sample
Bach <i>et al.</i> [2, 21]	$1/\varepsilon^2$	characteristic kernels; greedy
Phillips [20]	$(1/\varepsilon)^{\frac{2d}{d+2}} \log^{\frac{d}{d+2}}(1/\varepsilon)$	Lipschitz, d is constant; discrepancy-based
Phillips [20]	$\Theta(1/\varepsilon)$	$d = 1$
Phillips and Tai [21]	$(1/\varepsilon) \log^d(1/\varepsilon)$	Gaussian, d is constant; discrepancy-based
Phillips and Tai [21]	$\Omega(1/\varepsilon^2)$	SISS (e.g., Gaussian); $d = \Omega(1/\varepsilon^2)$

Then Bach *et al.* [2] showed that a greedy approach in the function space (based on the Frank-Wolfe technique) can obtain a bound of $O(1/\varepsilon^2)$, independent of dimension. This created a non-uniformly weighted coreset, and Phillips and Tai [21] provided a construction and analysis with uniform weighting.

Joshi *et al.* [16] first showed that sub- $1/\varepsilon^2$ is possible using a discrepancy-based approach. These bounds were superseded by Phillips [20] that presented a size bound of $O(((1/\varepsilon^2) \log(1/\varepsilon))^{d/(d+2)})$ for constant d , or $O(1/\varepsilon)$ for $d = 1$. For $d = 2$, this bound is near-linear in $1/\varepsilon$, specifically $O((1/\varepsilon)\sqrt{\log(1/\varepsilon)})$ and matches our bound. He also showed a lower bound of $\Omega(1/\varepsilon)$. Notably, this upper bound for constant d is strictly smaller than those of the binary range space for balls, based on VC-dimension.

Then recently, Phillips and Tai [21] greatly improved the upper bound to $O((1/\varepsilon) \log^d(1/\varepsilon))$ also based on discrepancy-based approaches, but this only applies to the Gaussian kernel. It is based on a special decomposable structure of the Gaussian, and does not seem to generalize. They also provided a lower bound of $\Omega(1/\varepsilon^2)$, but this construction required $d = \Omega(1/\varepsilon^2)$ dimensions, so does not imply much in the low-dimensional setting.

1.1 Our Results

We first show a new upper bound on the size of an ε -KDE coreset of $O((1/\varepsilon)\sqrt{d \log(1/\varepsilon)})$ in Section 2. The main restriction on the kernel K is that it is positive definite, a weaker bound than the similar characteristic assumption. There are also fairly benign restrictions that K is Lipschitz and its value greater than $1/|P|$ (or $\geq \varepsilon^2$) for a bounded region due to the specifics of some geometric preprocessing. Notably, this upper bound applies to a very wide range of kernels including the sinc kernel, whose super-level sets do not have bounded VC-dimension and is not characteristic, so no non-trivial ε -KDE coreset bound was previously known. Moreover, unlike previous discrepancy-based approaches, we do not need to assume the dimension d is constant.

We then show a nearly-matching lower bound on the size of an ε -KDE coreset of $\Omega(\sqrt{d}/\varepsilon)$, in Section 3. This construction requires a standard restriction that it is shift- and rotation-invariant, and a benign one that it is somewhere-steep (see Section 3), satisfied by all common kernels. This closes the problem for many kernels (e.g., Gaussians, Laplace), except for a $\sqrt{\log(1/\varepsilon)}$ factor when $1 < d < 1/\varepsilon^2$. The gap filled by the new bounds are shown in Table 3.

■ **Table 3** Size bounds for ε -KDE coresets for Gaussian and Laplace kernels; also holds under more general assumption, see text. (★) For $d = 2$, [20] matches upper bound.

d	Upper	Lower	
1	$1/\varepsilon$	$1/\varepsilon$	[20]
$[2, \frac{1}{\varepsilon^2})$	$\sqrt{d}/\varepsilon \cdot \sqrt{\log \frac{1}{\varepsilon}}$	\sqrt{d}/ε	new ★
$\geq \frac{1}{\varepsilon^2}$	$1/\varepsilon^2$	$1/\varepsilon^2$	[2, 21]

Our approach and context. As mentioned above, bounding the size ε -KDE coresets can be reduced to bounding kernel discrepancy. The range space discrepancy problem, for a range space (P, \mathcal{R}) , has been widely studied in multiple areas [17, 7]. For instance, Tuskady's problem restricts \mathcal{R} to represent axis-aligned rectangles in \mathbb{R}^d , has received much recent focus [18]. By reducing the Gaussian kernel discrepancy to this problem, led to the best previous results for these ε -KDE coresets [21]. To achieve their result, Matousek *et al.* [18] use a balancing technique of Banaszczyk [3] on a matrix version of discrepancy, by studying the so-call γ_2 -norm.

Roughly speaking, we are able to show how to directly reduce the kernel discrepancy problem to the γ_2 -norm, and the bound derived from Banaszczyk's Theorem [3] – bypassing the reduction to Tuskady's problem. In particular, the positive definiteness of a kernel, allows us to define a specific gram matrix G which has a real-valued decomposition, which matches the structure studied with the γ_2 norm. Hence, while our positive definite restriction is similar to the characteristic restriction studied for ε -KDE coresets in many other settings [13, 2] it uses a very different aspect of this property: the decomposability, not the embedding.

The lower bound is an extension of a recent lower bound [21] of $\Omega(1/\varepsilon^2)$. The key insight here is that we can reduce the construction which relied on a single $1/\varepsilon^2$ -dimensional simplex to a set of $1/(\sqrt{d}\varepsilon)$ d -dimensional simplices. Through some careful analysis, we show applying a modified variant of the argument from the previous lower bound to a particular simplex yields the desired result. Notably, this is the first lower bound for this problem that depends on the dimension d , and this dependence on d matches our upper bound.

2 Upper bound for KDE coreset

Consider a point set $P \subset \mathbb{R}^d$ as input, but as Section 4 describes, it is possible to apply these arguments to other domains. We assume that P is finite and of size n ; however, as mentioned in the related work, for many settings, we can reduce this to a point set of size independent of n (size $1/\varepsilon^2$ or d/ε^2 , depending on the kernel). Indeed these techniques may start with inputs as continuous distributions as long as we can draw random samples.

To prove our ε -kernel coreset upper bound we introduce two properties that the kernel must have.

- We say a kernel K has c_K -bounded influence if, for any $x \in \mathbb{R}^d$ and $\delta > 0$, $|K(x, y)| < \delta$ for all $y \notin x + [-(1/\delta)^{c_K}, (1/\delta)^{c_K}]^d$ for some constant c_K . By default we set $\delta = 1/n$. If c_K is an absolute constant we simply say K is bounded influence.
- We say a kernel K is C_K -Lipschitz if, for any $x, y, z \in \mathbb{R}^d$, $|K(x, z) - K(y, z)| < C_K \|x - y\|$ for some C_K . If C_K is an absolute constant within the context of the problem, we often just say the kernel is Lipschitz.

Next define a lattice $R = \left\{ \left(\frac{i_1}{\sqrt{dn}}, \frac{i_2}{\sqrt{dn}}, \dots, \frac{i_d}{\sqrt{dn}} \right) \mid i_j \text{ are integers} \right\}$. Also, denote, for each $p \in P$, $S_p = p + R \cap [-n^{c_K}, n^{c_K}]^d$ and $S = \cup_{p \in P} S_p$.

The following lemma explains that we only need to consider the evaluation at a finite set (specifically S) rather than the entire space while preserving the discrepancy asymptotically. The advantage of doing this is we can then use the matrix representation of the formula.

► **Lemma 1.** $\max_{x \in \mathbb{R}^d} \text{disc}(P, \chi, x) \leq \max_{x \in S} \text{disc}(P, \chi, x) + O(1)$

Proof. For any $x \in \mathbb{R}^d$, if $x \notin \cup_{p \in P} (p + [-n^{c_K}, n^{c_K}]^d)$, that is x is not within n^{c_K} in all coordinates of some $p \in P$, then $K(p, x) \leq 1/n$ for all $p \in P$. Hence we have

$$\text{disc}(P, \chi, x) = \left| \sum_{p \in P} \chi(p) K(p, x) \right| \leq O(1).$$

Otherwise, pick $x_0 \in S$ be the closest point to x . We have

$$\begin{aligned}
\text{disc}(P, \chi, x) &= \left| \sum_{p \in P} \chi(p) K(p, x) \right| \\
&= \left| \sum_{p \in P} \chi(p) (K(p, x_0) + K(p, x) - K(p, x_0)) \right| \\
&\leq \left| \sum_{p \in P} \chi(p) K(p, x_0) \right| + \sum_{p \in P} |K(p, x) - K(p, x_0)| \\
&\leq \text{disc}(P, \chi, x_0) + \sum_{p \in P} C_K \cdot \|x - x_0\| \\
&\leq \text{disc}(P, \chi, x_0) + n \cdot C_K \cdot \sqrt{d \left(\frac{1}{\sqrt{dn}} \right)^2} \\
&= \text{disc}(P, \chi, x_0) + O(1). \quad \blacktriangleleft
\end{aligned}$$

Now we discuss the matrix view of discrepancy, known results, and then how to map the discretized kernel discrepancy problem into this setting. Consider any $s \times t$ matrix A , and define

$$\text{disc}(A) = \min_{x \in \{-1, +1\}^t} \|Ax\|_\infty.$$

Following Matousek *et al.* [18] we define $\gamma_2(A) = \min_{BC=A} l_1 \cdot l_2$ where l_1 is largest Euclidean norm of row vectors of B and l_2 is largest Euclidean norm of column vectors of C . There is an equivalent geometric interpretation of γ_2 . Let \mathcal{E}_A be the set of ellipsoids in \mathbb{R}^s that contain all column vectors of A . Then, $\gamma_2(A) = \min_{E \in \mathcal{E}_A} \max_{x \in E} \|x\|_\infty$. It is easy to see that γ_2 is a norm and $\gamma_2(A) \leq \gamma_2(A')$ when the columns of A are subset of the columns of A' . We will apply these properties shortly.

A recent result by Matousek *et al.* [18] shows the following property about connecting discrepancy to γ_2 , which was recently made constructive in polynomial time [4].

► **Lemma 2** (Matousek *et al.* [18]). *For an $s \times t$ matrix A , $\text{disc}(A) \leq O(\sqrt{\log s}) \cdot \gamma_2(A)$.*

Let the size of S be $m = O(n^{O(d)})$, and define an $m \times n$ matrix G so its rows are indexed by $x \in S$ and columns indexed by $p \in P$, and $G_{x,p} = K(p, x)$. By examination, $\text{disc}(G) = \min_\chi \max_{x \in S} \text{disc}(P, \chi, x)$.

► **Lemma 3.** $\gamma_2(G) = 1$.

Proof. Denote G' be a $m \times m$ matrix with both row and column indexed $x, y \in S$ such that $G'_{x,y} = K(x, y)$. Note that columns of G are a subset of columns of G' since $P \subset S$. Since K is positive definite kernel, it means that G' can be expressed as $H^T H$ for some matrix H . Now denote v_x as the x th column of H for all $x \in S$. We have $v_x^T v_x = G'_{x,x} = 1$ which means the norm $\|v_x\| = \sqrt{v_x^T v_x}$ of each column $v_x \in H$ is 1. Hence the same holds for rows in H^T , and this bounds $\gamma_2(G') \leq 1$. Then since $\gamma_2(G) \leq \gamma_2(G')$ we have $\gamma_2(G) \leq 1$.

On the other hand, one of the coordinates in a column of G is 1. By the geometric definition, any ellipsoid containing columns of G has a point inside of it such that one of its coordinate is 1. Hence $\gamma_2(G) \geq 1$. ◀

Combining all above lemmas, for any $P \subset \mathbb{R}^d$ of size n

$$\begin{aligned} \text{disc}(n, \mathcal{K}) &\leq \max_{P:|P|=n} \min_{\chi} \max_{x \in S} \text{disc}(P, \chi, x) + O(1) && \text{Lemma 1} \\ &= \max_{P:|P|=n} \text{disc}(G) + O(1) && \text{Definition of } G \\ &\leq O(\sqrt{d \log n} \cdot \gamma_2(G)) && \text{Lemma 2 [18]} \\ &= O(\sqrt{d \log n}). && \text{Lemma 3} \end{aligned}$$

► **Theorem 4.** *Let $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a bounded influence, Lipschitz, positive definite kernel. For any integer n , $\text{disc}(n, \mathcal{K}_d) = O(\sqrt{d \log n})$.*

► **Corollary 5.** *Let $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a bounded influence, Lipschitz, positive definite kernel. For any set $P \subset \mathbb{R}^d$, there is a subset $Q \subset P$ of size $O(\frac{1}{\epsilon} \sqrt{d \log \frac{1}{\epsilon}})$ such that*

$$\max_{x \in \mathbb{R}^d} |\text{KDE}_P(x) - \text{KDE}_Q(x)| < \epsilon.$$

Proof. In order to apply the standard halving technique [8, 19], we need to make sure the coloring has the property that half of point assigned +1 and the other half of them assigned -1. We adapt a standard idea from combinatorial discrepancy [17].

This can be done by adding an all-one row to the discrepancy matrix G . It guarantees that the difference of number of +1 and -1 is $O(\sqrt{d \log n})$ since γ_2 is a norm and therefore we can apply the triangle inequality. Namely,

$$\gamma_2 \left(\begin{bmatrix} \mathbb{1}_{1 \times n} \\ G \end{bmatrix} \right) \leq \gamma_2 \left(\begin{bmatrix} O_{1 \times n} \\ G \end{bmatrix} \right) + \gamma_2 \left(\begin{bmatrix} \mathbb{1}_{1 \times n} \\ O_{m \times n} \end{bmatrix} \right)$$


where $\mathbb{1}$ is all-one matrix and O is zero matrix. Suppose there are more +1 than -1. Choose $O(\sqrt{d \log n})$ points assigned +1 arbitrarily and flip them to -1 such that it makes the difference zero. Let $P_+ = \{p \in P \mid \chi(p) = +1\}$ and $P_- = \{p \in P \mid \chi(p) = -1\}$. Also, P'_+ and P'_- are defined in the same way after flipping the value. For any $x \in \mathbb{R}^d$,

$$\begin{aligned} &\left| \sum_{p \in P'_+} K(x, p) - \sum_{p \in P'_-} K(x, p) \right| \\ &\leq \left| \sum_{p \in P_+} K(x, p) - \sum_{p \in P_-} K(x, p) \right| + \left| \sum_{p \in P'_+ \setminus P_+} K(x, p) \right| + \left| \sum_{p \in P_- \setminus P'_-} K(x, p) \right| \\ &= O(\sqrt{d \log n}). \end{aligned}$$

Now, we can apply the standard halving technique to achieve

$$\max_{x \in \mathbb{R}^d} |\text{KDE}_P(x) - \text{KDE}_Q(x)| < \epsilon. \quad \blacktriangleleft$$

Implementation. Note that we do not need to decompose the entire matrix G . Instead, we just need a set of vectors $V = \{v_p \mid p \in P\}$ such that the inner product $\langle v_{p_1}, v_{p_2} \rangle = K(p_1, p_2)$ as input to the algorithm in [4]. This set V can be computed in $\text{poly}(n, d) = \text{poly}(n)$ time assuming $d < n$. Using the standard Merge-Reduce framework [19], the coreset with desired size can be constructed in $O(n \text{poly}(1/\epsilon))$ time.



■ **Figure 1** Illustration of the lower bound construction.

3 Lower bound for KDE coreset

In this section, we add two new conditions on our kernel; both of these are common properties of kernels.

- A kernel K is *rotation- and shift-invariant* if there exists a function f such that $K(x, y) = f(\|x - y\|^2)$.
- A rotation- and shift-invariant kernel is *somewhere C_f -steep* if there exist a constant $C_f > 0$, and values $z_f > r_f > 0$ such that $f(z_1) - f(z_2) > C_f \cdot (z_2 - z_1)$ for all $z_1 \in (z_f - r_f, z_f)$ and $z_2 \in (z_f, z_f + r_f)$. When C_f is an absolute constant, we often just say the kernel is *somewhere steep*.

Phillips [20] constructed an example of P of size $1/\varepsilon$ where each point in P is far away from all others. Therefore, if one of them is not picked for a KDE coreset Q , the evaluation of KDE_Q at that point has large error. More recently, Phillips and Tai [21] gave another example of P of size $1/\varepsilon^2$ in an appropriately scaled simplex; that spans $\mathbb{R}^{1/\varepsilon^2}$. They showed that it produces error of $\Omega(1/\sqrt{k})$ at some point, if k is the number of points picked. The following construction combines these two approaches. We divide $n = \frac{\sqrt{d}}{\varepsilon}$ points into n/d groups where each group has d points that form a simplex, and each group is far away from all other groups. It means that there is a group producing $\Omega(1/\sqrt{d})$ error and then, since we have n/d groups, the final error would be $\Omega(\frac{1/\sqrt{d}}{n/d}) = \Omega(\varepsilon)$.

► **Theorem 6.** Consider a rotation- and shift-invariant, somewhere steep, bounded influence kernel K . Assume $d \geq \frac{9z_f^2}{r_f}$, where z_f and r_f are absolute constants that depend on K and are defined as they pertain to the somewhere steep criteria. There is a set of $P \in \mathbb{R}^d$ such that, for any subset Q of size $k \leq \frac{\sqrt{d}}{2\varepsilon}$, there is a point $x \in \mathbb{R}^d$ such that $|\text{KDE}_P(x) - \text{KDE}_Q(x)| > \varepsilon$.

Proof. Let $n = \sqrt{d}/\varepsilon$. We allow weighted coresets of Q ; that is, for each $q \in Q$, there is a real number β_q such that $\text{KDE}_Q(x) = \sum_{q \in Q} \beta_q K(x, q)$.

Let $k \leq n/2$ be the size of the potential coreset we consider. Construct P with size of n in \mathbb{R}^d as follow. Let $\{e_i\}_{i=1}^d$ is the standard basis and L is a very large number. Set

$P_j = \{p_{i,j} = \sqrt{\frac{z_f}{2}}e_i + jLe_1 \mid i = 1, 2, \dots, d\}$ for all $j = 1, 2, \dots, \frac{n}{d}$. Define $P = \cup_{j=1}^{n/d} P_j$. Namely, we divide n points into $\frac{n}{d}$ groups and each group has d points which forms a d -simplex. Also, the groups are sufficiently far away from each other. Suppose $Q = \cup_{j=1}^{n/d} \{p_{i_a,j} \mid a = 1, 2, \dots, k_j\}$ where k_j is the number of points in Q at group j . Denote $Q_j = \{p_{i_a,j} \mid a = 1, 2, \dots, k_j\}$. That is, $Q = \cup_{j=1}^{n/d} Q_j$ and $|Q_j| = k_j \leq d$ with $\sum_{j=1}^{n/d} k_j = |Q| = k$.

Since $\sum_{j=1}^{n/d} |Q_j| = k \leq n/2$, at least one j must satisfy $k_j \leq \frac{d}{2}$. Denote j' to be that j . We can assume $k_{j'} = d/2$, otherwise, pick enough points arbitrarily from $P_j \setminus Q_{j'}$ and place them in $Q_{j'}$ to make $|Q_{j'}| = k_{j'} = d/2$, but set the corresponding weight to be 0. Denote $\bar{p} = \frac{1}{d} \sum_{p \in P_{j'}} p$ the mean of $P_{j'}$; $\bar{q}_+ = \frac{2}{d} \sum_{q \in Q_{j'}} q$ the mean of $Q_{j'}$; and $\bar{q}_- = \frac{2}{d} \sum_{q \in P_{j'} \setminus Q_{j'}} q$ the mean of points in $P_{j'}$ not selected into $Q_{j'}$; see Figure 1. Also, denote $p_+^* = \bar{q}_+ + \sqrt{\frac{z_f}{2}} \frac{\bar{q}_+ - \bar{p}}{\|\bar{q}_+ - \bar{p}\|}$ and $p_-^* = \bar{q}_- + \sqrt{\frac{z_f}{2}} \frac{\bar{q}_- - \bar{p}}{\|\bar{q}_- - \bar{p}\|}$; translates of these points away from the mean \bar{p} by a specific vector. Note that $\|p_+^* - q\|$ is the same for all $q \in Q_{j'}$, denoted by l_1 and $\|p_-^* - q\|$ is same for all $q \in P_{j'} \setminus Q_{j'}$, denoted by l_2 . By symmetry, we also have that $l_1 = \|p_-^* - q\|$ for all $q \in P_{j'} \setminus Q_{j'}$ and $l_2 = \|p_+^* - q\|$ for all $q \in Q_{j'}$.

If $\sum_{q \in Q_{j'}} \beta_q \geq d/n$, we evaluate the error at p_+^* .

$$\begin{aligned} & (\text{KDE}_Q - \text{KDE}_P)(p_+^*) \\ &= \sum_{q \in Q_{j'}} (\beta_q - \frac{1}{n}) f(\|p_+^* - q\|^2) + \sum_{q \in P_{j'} \setminus Q_{j'}} (-\frac{1}{n}) f(\|p_+^* - q\|^2) + s \\ &\geq \frac{d}{2n} (f(l_1^2) - f(l_2^2)) + s \end{aligned}$$

where $|s|$ is arbitrarily small due to the choice of arbitrarily large number L and the fact that K is bounded influence. If $\sum_{q \in Q_{j'}} \beta_q \leq d/n$, we evaluate the error at p_-^* .

$$\begin{aligned} & (\text{KDE}_P - \text{KDE}_Q)(p_-^*) \\ &= \sum_{q \in P_{j'} \setminus Q_{j'}} \frac{1}{n} f(\|p_-^* - q\|^2) + \sum_{q \in Q_{j'}} (\frac{1}{n} - \beta_q) f(\|p_-^* - q\|^2) + s \\ &\geq \frac{d}{2n} (f(l_1^2) - f(l_2^2)) + s \end{aligned}$$

Therefore, in either case, we need to bound $f(l_1^2) - f(l_2^2)$ from below.

By direct computation, we have $l_1^2 = z_f - \frac{z_f}{d}$ and $l_2^2 = z_f + \frac{z_f}{d} + \frac{2z_f}{\sqrt{d}}$. By enforcing that

$$z_f - r_f < z_f - \frac{z_f}{d} = l_1^2 < z_f$$

and

$$z_f < z_f + \frac{z_f}{d} + \frac{2z_f}{\sqrt{d}} = l_2^2 < z_f + \frac{3z_f}{\sqrt{d}} < z_f + r_f,$$

we can invoke the somewhere C_f -steep property that there exists an x in \mathbb{R}^d for which the inequality holds. Therefore,

$$f(l_1^2) - f(l_2^2) > C_f \cdot (l_2^2 - l_1^2) > C_f \cdot z_f \cdot \frac{2}{\sqrt{d}}.$$

Hence, the error is at least

$$\frac{d}{2n} (f(l_1^2) - f(l_2^2)) + s > \frac{d}{2n} \left(C_f \cdot z_f \cdot \frac{2}{\sqrt{d}} \right) + s > \frac{\sqrt{d}}{n} \cdot C_f \cdot z_f + s = \Omega(\sqrt{d}/n) = \Omega(\varepsilon). \blacktriangleleft$$

4 Applications to specific kernels

In this section, we work through the straight-forward application of these bounds to some specific kernels and settings.

Gaussian and Laplace kernels. These kernels are defined over \mathbb{R}^d . They have bounded influence, so $|K(x, p)| \leq \frac{1}{n}$ for all $p \notin [-n^{c_K}, n^{c_K}]^d$ for $c_K = 1$. They are also C_K -Lipschitz with constant $C_K = \alpha$, so $|K(x, z) - K(p, z)| \leq C_K \|x - p\|$ for any $x, p \in \mathbb{R}^d$. These properties imply we can invoke the discrepancy upper bound in Theorem 4.

These kernels are also rotation- and shift-invariant, and somewhere steep with constant $C_f = (\alpha/2) \exp(-\alpha^2)$. Hence we can invoke the lower bound in Theorem 6.

► **Corollary 7.** *For Gaussian or Laplacian kernels, for any set $P \in \mathbb{R}^d$, there is a ε -KDE coreset of size $O((\sqrt{d}/\varepsilon)\sqrt{\log 1/\varepsilon})$, and it cannot have an ε -KDE coreset of size $o(\sqrt{d}/\varepsilon)$.*

The Gaussian kernel has an amazing decomposition property that in \mathbb{R}^d if we fix any d' coordinates in any way, then conditioned on those, the remaining $d - d'$ coordinates still follow a Gaussian distribution. Among other things, this means it is useful to construct kernels for complex scenarios. For instance, consider a large set T of n trajectories, each with k waypoints; e.g., backpacking or road trips or military excursions with k nights, and let the waypoints be the (x, y) -coordinates for the location of each night stay. We can measure the similarity between two trajectories $t = (p_1, p_2, \dots, p_k)$ and $t' = (p'_1, p'_2, \dots, p'_k)$ as the average similarity between the corresponding waypoints, and we can measure the similarity of any two corresponding waypoints p_j and p'_j with a 2-dimensional Gaussian. Then, by the decomposition property, the full similarity between the trajectories is precisely a $(2k)$ -dimensional Gaussian. We can thus define a kernel density estimate over these trajectories KDE_T using this $(2k)$ -dimensional Gaussian kernel. Now, given Corollary 7 we know that to approximate KDE_T with a much smaller data set $S \subset T$ so $\|\text{KDE}_T - \text{KDE}_S\|_\infty \leq \varepsilon$, we can construct S so $|S| = O(\sqrt{k}/\varepsilon \cdot \sqrt{\log 1/\varepsilon})$ but cannot in general achieve $|S| = o(\sqrt{k}/\varepsilon)$.

Jensen-Shannon and Hellinger kernels. In order to apply our technique on Δ^d , observe that Δ^d is a subset of a $(d - 1)$ -dimensional Euclidian subspace of \mathbb{R}^d ; so we can simply create the grid needed for Lemma 1 within this subspace. Recall that these two kernel have the form of $\exp(-\alpha d(x, y))$ where $d(x, y) = d_{\text{JS}}(x, y) = H(\frac{x+y}{2}) - \frac{H(x)+H(y)}{2}$ for Jensen-Shannon kernel and $d(x, y) = d_{\text{H}}(x, y) = \sum_{i=1}^d (\sqrt{x_i} - \sqrt{y_i})^2$ for Hellinger and note that $|K(x, z) - K(y, z)| \leq \alpha |d(x, z) - d(y, z)|$ for any $x, y, z \in \Delta^d$. It is easy to estimate that when x, y are sufficiently close, for JS kernel, $|d(x, z) - d(y, z)| \leq 2d \max_i |x_i - y_i| |\log |x_i - y_i|| \leq 2d \max_i \sqrt{|x_i - y_i|}$ and for Hellinger kernel, $|d(x, z) - d(y, z)| \leq 4d \max_i \sqrt{|x_i - y_i|}$. So even though these kernels are not Lipschitz, we can still modify the construction of the grid in Lemma 1 with width $\frac{1}{n^4}$ (assuming $d \leq n$) instead of $\frac{1}{\sqrt{dn}}$ such that if x, y lie in the same cell then $|K(x, z) - K(y, z)| = O(\frac{1}{n})$ for any $x, y, z \in \Delta^d$. Since all relevant points are in a bounded domain both kernels have c_K -bounded influence; setting $c_K = 1$ is sufficient.

► **Corollary 8.** *For Jensen-Shannon and Hellinger kernels, for any set $P \in \Delta^d$, there is a ε -KDE coreset of size $O((\sqrt{d}/\varepsilon)\sqrt{\log 1/\varepsilon})$.*

Note that these kernels are not rotation- and shift-invariant and therefore our lower bound result does not apply.

These kernels are based on widely-used information distances: the Jensen-Shannon distance $d_{\text{JS}}(x, p)$ and the Hellinger distance $d_{\text{H}}(x, p)$. These make sense when the input data

$x, p \in \Delta^d$ represent a "histogram," a discrete probability distribution over a d -variate domain. These are widely studied objects in information theory, and more commonly text analysis. For instance, a common text modeling approach is to represent each document v in a large corpus of documents V (e.g., a collection of tweets, or news articles, or wikipedia pages) as a set of word counts. That is, each coordinate v_j of v represents the number of times that word (indexed by) j occurs in that document. To remove length information from the documents (retaining only the topics), it is common to normalize each vector as $v \mapsto \frac{v}{\|v\|}$ so the j th coordinate represents the probability that a random word on the page is j . The most common modeling choice to measure distance between these distribution representations of documents are the Hellinger and Jensen-Shannon distances, and hence the most natural choice of similarity are the corresponding kernels we examine. In particular, with a very large corpus V of size n , Corollary 8 shows that we can approximate KDE_V , a kernel density estimate of V , with one described by a much smaller set $S \subset V$ so $\|\text{KDE}_V - \text{KDE}_S\| \leq \varepsilon$ and so $|S| = O(\sqrt{d}/\varepsilon \cdot \sqrt{\log 1/\varepsilon})$. Notably, when one has a fairly large d , and desires high accuracy (small ε), then our new result will provide the best possible ε -KDE coreset.

Exponential kernels. In order to apply our technique on \mathbb{S}^d , we can rewrite the kernel to be $K'(x, y) = K(\frac{x}{\|x\|}, \frac{y}{\|y\|})$ for all $x, y \in \mathbb{R}^d \setminus \{0\}$. We construct the grid in Lemma 1 on \mathbb{R}^d for K' and then only retain grid points which lie in the annulus $\mathbb{A}^d = \{x \in \mathbb{R}^d \mid \frac{1}{2} \leq \|x\| \leq \frac{3}{2}\}$. This annulus contains all grid points which could be the closest point of some point on \mathbb{S}^d , as required in Lemma 1. Moreover K' is C_K -Lipschitz on the annulus: it satisfies for any $x, y, z \in \mathbb{A}^d$ that $|K'(x, z) - K'(y, z)| \leq C_K \|x - y\|$, with $C_K = 4\alpha$. Since the domain is restricted to \mathbb{S}^d , similar to on the domain Δ^d , any kernel has c_K -bounded influence and setting $c_K = 1$ is sufficient.

► **Corollary 9.** *For the exponential kernel, for any set $P \in \mathbb{S}^d$, there is a ε -KDE coreset of size $O((\sqrt{d}/\varepsilon)\sqrt{\log 1/\varepsilon})$.*

The exponential kernel is not rotation- and shift-invariant and therefore our lower bound result does not apply.

Sinc kernel. Note that the sinc kernel is not everywhere positive, and as a result of its structure the VC-dimension is unbounded, so the approaches requiring those properties [16, 20] cannot be applied. It is also not characteristic, so the embedding-based results [13, 2] do not apply either. As a result, there is no non-trivial ε -KDE coreset for the sinc kernel. However, in our approach, the positivity of one single entry in the discrepancy matrix does not matter so long as the entire matrix is positive definite – which is the case for sinc. Therefore, our result could be applied to sinc kernel, with $c_K = 1$ (it has 1-bounded influence), $C_K = \alpha/\pi$ (it is (α/π) -Lipschitz) and $C_f = \alpha^2/2\pi^2$ (it is somewhere $(\alpha^2/2\pi^2)$ -steep).

► **Corollary 10.** *For sinc kernels, for any set $P \in \mathbb{R}^d$, there is a ε -KDE coreset of size $O((1/\varepsilon)\sqrt{\log 1/\varepsilon})$ (for $d = \{1, 2, 3\}$), and it cannot have a ε -KDE coreset of size $\Omega(1/\varepsilon)$.*

5 Conclusion

We proved that Gaussian kernel has a ε -KDE coreset of size $O(\frac{1}{\varepsilon}\sqrt{d \log \frac{1}{\varepsilon}})$ and the size must satisfy $\Omega(\frac{\sqrt{d}}{\varepsilon})$; both upper and lower bound result can be extended to a broad class of kernels. In particular the upper bound only requires that the kernel be positive definite (typically the same restriction needed for most machine learning techniques) and that it has a domain

which can be discretized over a bounded region without inducing too much error. This family of applicable kernels includes new options like the sinc kernel, which while positive definite in \mathbb{R}^d for $d = \{1, 2, 3\}$, it is not characteristic, is not always positive, and its super-level sets do not have bounded VC-dimension. This is the first non-trivial ε -KDE coreset result for these kernels.

By inspecting the new constructive algorithm for obtaining small discrepancy in the γ_2 -norm [4], the extra $\sqrt{\log}$ factor comes from the union bound over the randomness in the algorithm. Indeed, a previous result [21] showed that if $d = \frac{1}{\varepsilon^2}$ then the upper bound is $O(\frac{1}{\varepsilon^2})$, which is tight. This bound is deterministic and does not have an extra $\sqrt{\log}$ factor. Therefore, a natural conjecture is that the upper bound result can be further improved to $O(\sqrt{d}/\varepsilon)$, at least in a well-behaved setting like for the Gaussian kernel.

There are many other, even more diverse kernels which are positive definite, which operate on domains as diverse as graphs, time series, strings, and trees [15]. The heart of the upper bound construction which uses the decomposition of the associated positive definite matrix will work even for these kernels. However, it is less clear how to generate a finite gram or discrepancy matrix G , whose size depends polynomially on the data sets size for these discrete objects. Such constructions would further expand the pervasiveness of the ε -KDE coreset technique we present.

References

- 1 Ery Arias-Castro, David Mason, and Bruno Pelletier. On the estimation of the gradient lines of a density and the consistency of the mean-shift algorithm. *Journal of Machine Learning Research*, 17(43):1–28, 2016.
- 2 Francis Bach, Simon Lacoste-Julien, and Guillaume Obozinski. On the equivalence between herding and conditional gradient algorithms. In *ICML 2012 International Conference on Machine Learning*, 2012.
- 3 Wojciech Banaszczyk. Balancing vectors and gaussian measures of n-dimensional convex bodies. *Random Structures & Algorithms*, 12(4):351–360, 1998.
- 4 Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt walk: A cure for the Banaszczyk blues (to appear). *Proceedings of the fiftieth annual ACM symposium on Theory of computing*, 2018.
- 5 Jon Louis Bentley and James B. Saxe. Decomposable searching problems I: Static-to-dynamic transformations. *Journal of Algorithms*, 1(4), 1980.
- 6 Omer Bobrowski, Sayan Mukherjee, and Jonathan E. Taylor. Topological consistency via kernel estimation. *Bernoulli*, 23:288–328, 2017.
- 7 Bernard Chazelle. *The Discrepancy Method*. Cambridge, 2000.
- 8 Bernard Chazelle and Jiri Matousek. On linear-time deterministic algorithms for optimization problems in fixed dimensions. *J. Algorithms*, 21:579–597, 1996.
- 9 Luc Devroye and László Györfi. *Nonparametric Density Estimation: The L_1 View*. Wiley, 1984.
- 10 Petros Drineas and Michael W. Mahoney. On the Nyström method for approximating a Gram matrix for improved kernel-based learning. *JLMR*, 6:2153–2175, 2005.
- 11 Jianqing Fan and Irene Gijbels. *Local polynomial modelling and its applications: monographs on statistics and applied probability 66*, volume 66. CRC Press, 1996.
- 12 Brittany Terese Fasy, Fabrizio Lecci, Alessandro Rinaldo, Larry Wasserman, Sivaraman Balakrishnan, and Aarti Singh. Confidence sets for persistence diagrams. *The Annals of Statistics*, 42:2301–2339, 2014.

- 13 Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012.
- 14 Matthias Hein and Olivier Bousquet. Hilbertian metrics and positive definite kernels on probability measures. In *AISTATS*, pages 136–143, 2005.
- 15 Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. A review of kernel methods in machine learning. Technical Report 156, Max Planck Institute for Biological Cybernetics, 2006.
- 16 Sarang Joshi, Raj Varma Kommaraji, Jeff M Phillips, and Suresh Venkatasubramanian. Comparing distributions and shapes using the kernel distance. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 47–56. ACM, 2011.
- 17 Jiri Matousek. *Geometric Discrepancy; An Illustrated Guide, 2nd printing*. Springer-Verlag, 2010.
- 18 Jiri Matousek, Aleksandar Nikolov, and Kunal Talwar. Factorization norms and hereditary discrepancy. *arXiv preprint arXiv:1408.1376*, 2014.
- 19 Jeff M. Phillips. Algorithms for ε -approximations of terrains. In *ICALP*, 2008.
- 20 Jeff M Phillips. ε -samples for kernels. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1622–1632. SIAM, 2013.
- 21 Jeff M Phillips and Wai Ming Tai. Improved coresets for kernel density estimates. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2718–2727. SIAM, 2018.
- 22 Jeff M. Phillips, Bei Wang, and Yan Zheng. Geometric inference on kernel density estimates. In *SOCG*, 2015.
- 23 Alessandro Rinaldo and Larry Wasserman. Generalized density clustering. *The Annals of Statistics*, pages 2678–2722, 2010.
- 24 Isaac J Schoenberg. Metric spaces and completely monotone functions. *Annals of Mathematics*, pages 811–841, 1938.
- 25 Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.
- 26 Erich Schubert, Arthur Zimek, and Hans-Peter Kriegel. Generalized outlier detection with flexible kernel density estimates. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pages 542–550. SIAM, 2014.
- 27 David W. Scott. *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley, 1992.
- 28 Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, 1986.
- 29 Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *JMLR*, 11:1517–1561, 2010.
- 30 Yan Zheng and Jeff M. Phillips. l_∞ error and bandwidth selection for kernel density estimates of large data. In *KDD*, 2015.