

Comparing Distributions and Shapes using the Kernel Distance ^{*}

Sarang Joshi
University of Utah
sjoshi@sci.utah.edu

Jeff M. Phillips
University of Utah
jeffp@cs.utah.edu

Raj Varma Kommaraju
Microsoft
rajvarma@cs.utah.edu

Suresh Venkatasubramanian
University of Utah
suresh@cs.utah.edu

ABSTRACT

Starting with a similarity function between objects, it is possible to define a distance metric (the kernel distance) on pairs of objects, and more generally on probability distributions over them. These distance metrics have a deep basis in functional analysis and geometric measure theory, and have a rich structure that includes an isometric embedding into a Hilbert space. They have recently been applied to numerous problems in machine learning and shape analysis.

In this paper, we provide the first algorithmic analysis of these distance metrics. Our main contributions are as follows: (i) We present fast approximation algorithms for computing the kernel distance between two point sets \mathcal{P} and \mathcal{Q} that runs in near-linear time in the size of $\mathcal{P} \cup \mathcal{Q}$ (an explicit calculation would take quadratic time). (ii) We present polynomial-time algorithms for approximately minimizing the kernel distance under rigid transformation; they run in time $O(n + \text{poly}(1/\varepsilon, \log n))$. (iii) We provide several general techniques for reducing complex objects to convenient sparse representations (specifically to point sets or sets of points sets) which approximately preserve the kernel distance. In particular, this allows us to reduce problems of computing the kernel distance between various types of objects such as curves, surfaces, and distributions to computing the kernel distance between point sets.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms

General Terms

Algorithms, Theory

1. INTRODUCTION

Let $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow R$ be a kernel function, such as a Gaussian kernel; $K(p, q)$ describes how similar two points $p, q \in \mathbb{R}^d$ are. For point sets \mathcal{P}, \mathcal{Q} we can define a similarity function $\kappa(\mathcal{P}, \mathcal{Q}) = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} K(p, q)$. Then the *kernel distance* is defined as

$$D_K(\mathcal{P}, \mathcal{Q}) = \sqrt{\kappa(\mathcal{P}, \mathcal{P}) + \kappa(\mathcal{Q}, \mathcal{Q}) - 2\kappa(\mathcal{P}, \mathcal{Q})}. \quad (1.1)$$

By altering the kernel K , and the weighting of elements in κ , the kernel distance can capture distance between distributions, curves, surfaces, and even more general objects.

Motivation.

The earthmover distance (EMD) takes a metric space and two probability distributions over the space, and computes the amount of work needed to "transport" mass from one distribution to another. It has become a metric of choice in computer vision, where images are represented as intensity distributions over a Euclidean grid of pixels. It has also been applied to shape comparison [8], where shapes are represented by point clouds (discrete distributions) in space.

While the EMD is a popular way of comparing distributions over a metric space, it is also an expensive one. Computing the EMD requires solving an optimal transportation problem via the Hungarian algorithm, and while approximations exist for restricted cases like the Euclidean plane, they are still expensive, requiring either quadratic time in the size of the input or achieving at best a constant factor approximation when taking linear time [38, 2]. Further, it is hard to index structures using the EMD for performing near-neighbor, clustering and other data analysis operations. Indeed, there are lower bounds on our ability to embed the EMD into well-behaved normed spaces [4].

The kernel distance has thus become an effective alternative to comparing distributions on a metric space. In machine learning, the kernel distance has been used to build metrics on distributions [42, 23, 41, 39, 32] and to learn hidden Markov models [40]. In the realm of shape analysis [45, 18, 19, 17], the kernel distance (referred to there as the *current distance*) has been used to compare shapes, whether they be point sets, curves, or surfaces.

All of these methods utilize key structural features of the kernel distance. When constructed using a *positive definite*¹ similarity function K , the kernel distance can be interpreted

¹A positive definite function generalizes the idea of a positive definite matrix; see Section 2.

^{*}Full version available at arXiv:1001.0591. Research supported by NSF award CCF-0953066, and subawards to the University of Utah under NSF awards 0937060 and 1019343 to CRA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'11, June 13–15, 2011, Paris, France.

Copyright 2011 ACM 978-1-4503-0682-9/11/06 ...\$10.00.

through a lifting map $\phi : \mathbb{R}^d \rightarrow \mathcal{H}$ to a reproducing kernel Hilbert space (RKHS), \mathcal{H} . This lifting map ϕ is isometric; the kernel distance is precisely the distance induced by the Hilbert space ($D_K(\{p\}, \{q\}) = \|\phi(p) - \phi(q)\|_{\mathcal{H}}$). Furthermore, a point set \mathcal{P} has an isometric lifted representation $\Phi(\mathcal{P}) = \sum_{p \in \mathcal{P}} \phi(p)$ as a single vector in \mathcal{H} so $D_K(\mathcal{P}, \mathcal{Q}) = \|\Phi(\mathcal{P}) - \Phi(\mathcal{Q})\|_{\mathcal{H}}$. Moreover, by choosing an appropriately scaled basis, this becomes a simple ℓ_2 distance, so all algorithmic tools developed for comparing points and point sets under ℓ_2 can now be applied to distributions and shapes.

Dealing with uncertain data provides another reason to study the kernel distance. Rather than thinking of $K(\cdot, \cdot)$ as a similarity function, we can think of it as a way of capturing spatial uncertainty; $K(p, q)$ is the likelihood that the object claimed to be at p is actually at q . For example, setting $K(p, q) = \exp(-\|p - q\|^2 / \sigma) / (\sqrt{2\pi}\sigma)$ gives us a Gaussian blur function. In such settings, the kernel distance $D_K^2(\mathcal{P}, \mathcal{Q})$ computes the symmetric difference $|\mathcal{P} \Delta \mathcal{Q}|$ between shapes with uncertainty described by K .

Our work.

We present the first algorithmic analysis of the kernel distance. Our main contributions are as follows:

(i) We present fast approximation algorithms for computing the kernel distance between two point sets \mathcal{P} and \mathcal{Q} that runs in near-linear time in the size of $\mathcal{P} \cup \mathcal{Q}$ (note that an explicit calculation would take quadratic time). (ii) We present polynomial-time algorithms for approximately minimizing the kernel distance under rigid transformation; they run in time $O(n + \text{poly}(1/\varepsilon, \log n))$. (iii) We provide several general techniques for reducing complex objects to convenient sparse representations (specifically to point sets or sets of points sets) which approximately preserve the kernel distance. In particular, this allows us to reduce problems of computing the kernel distance between various types of objects such as curves, surfaces, and distributions to computing the kernel distance between point sets.

We build these results from two core technical tools. The first is a lifting map that maps objects into a finite-dimensional Euclidean space while approximately preserving the kernel distance. We believe that the analysis of lifting maps is of independent interest; indeed, these methods are popular in machine learning [41, 39, 40] but (in the realm of kernels) have received less attention in algorithms. Our second technical tool is a theorem relating ε -samples of range spaces defined with kernels to standard ε -samples of range spaces on $\{0, 1\}$ -valued functions. This gives a simpler algorithm than prior methods in learning theory that make use of the γ -fat shattering dimension, and it yields smaller ε -samples.

2. PRELIMINARIES

Definitions.

For the most general case of the kernel distance (that we will consider in this paper) we associate a unit vector $U(p)$ and a weighting $\mu(p)$ with every $p \in \mathcal{P}$. Similarly we associate a unit vector $V(q)$ and weighting $\nu(q)$ with every $q \in \mathcal{Q}$. Then we write

$$\kappa(\mathcal{P}, \mathcal{Q}) = \int_{p \in \mathcal{P}} \int_{q \in \mathcal{Q}} K(p, q) \langle U(p), V(q) \rangle d\mu(p) d\nu(q), \quad (2.1)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. This becomes a distance D_K , defined through (1.1).

When \mathcal{P} is a curve in \mathbb{R}^d we let $U(p)$ be the tangent vector at p and $\mu(p) = 1$. When \mathcal{P} is a surface in \mathbb{R}^3 we let $U(p)$ be the normal vector at p and $\mu(p) = 1$. This can be generalized to higher order surfaces through the machinery of k -forms and k -vectors [45, 35].

When \mathcal{P} is an arbitrary probability measure² in \mathbb{R}^d , then all $U(p)$ are identical unit vectors and $\mu(p)$ is the probability of p . For discrete probability measures, described by a point set, we replace the integral with a sum and $\mu(p)$ can be used as a weight $\kappa(\mathcal{P}, \mathcal{Q}) = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} K(p, q) \mu(p) \nu(q)$.

From distances to metrics.

When K is a symmetric similarity function (i.e. $K(p, p) = \max_{q \in \mathbb{R}^d} K(p, q)$, $K(p, q) = K(q, p)$, and $K(p, q)$ decreases as p and q become “less similar”) then D_K (defined through (2.1) and (1.1)) is a distance function, but may not be a metric. However, when K is positive definite, then this is sufficient for D_K to not only be a metric³, but also for D_K^2 to be of negative type [16].

We say that a symmetric function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a *symmetric positive definite kernel* if for any nonzero L_2 function f it satisfies $\int_{p \in \mathcal{P}} \int_{q \in \mathcal{Q}} f(q) K(p, q) f(p) dp dq > 0$. The proof of D_K being a metric follows by considering the reproducing kernel Hilbert space \mathcal{H} associated with such a K [5]. Moreover, D_K can be expressed very compactly in this space. The lifting map $\phi : \mathbb{R}^d \rightarrow \mathcal{H}$ associated with K has the “reproducing property” $K(p, q) = \langle \phi(p), \phi(q) \rangle_{\mathcal{H}}$. So by linearity of the inner product, $\Phi(\mathcal{P}) = \int_{p \in \mathcal{P}} \phi(p) d\mu(p)$ can be used to retrieve $D_K(\mathcal{P}, \mathcal{Q}) = \|\Phi(\mathcal{P}) - \Phi(\mathcal{Q})\|_{\mathcal{H}}$ using the induced norm $\|\cdot\|_{\mathcal{H}}$ of \mathcal{H} . Observe that this defines a norm $\|\Phi(\mathcal{P})\|_{\mathcal{H}} = \sqrt{\kappa(\mathcal{P}, \mathcal{P})}$ for a shape.

Examples.

If K is the “trivial” kernel, where $K(p, p) = 1$ and $K(p, q) = 0$ for $p \neq q$, then the distance between any two sets (without multiplicity) \mathcal{P}, \mathcal{Q} is $D_K^2(\mathcal{P}, \mathcal{Q}) = |\mathcal{P} \Delta \mathcal{Q}|$, where $\mathcal{P} \Delta \mathcal{Q} = \mathcal{P} \cup \mathcal{Q} \setminus (\mathcal{P} \cap \mathcal{Q})$ is the symmetric difference. In general for arbitrary probability measures, $D_K(\mathcal{P}, \mathcal{Q}) = \|\mu - \nu\|_2$. If $K(p, q) = \langle p, q \rangle$, the Euclidean dot product, then the resulting lifting map ϕ is the identity function, and the distance between two measures is the Euclidean distance between their means, which is a pseudometric but not a metric.

Gaussian properties.

To simplify much of the presentation of this paper we focus on the case where the kernel K is the Gaussian kernel; that is $K(p, q) = e^{-\|p - q\|^2 / h}$. Our techniques carry over to more general kernels, although the specific bounds will depend on the kernel being used. We now encapsulate some useful properties of Gaussian kernels in the following lemmata. When approximating $K(p, q)$, the first allows us to ignore pairs of points further than $\sqrt{h \ln(1/\gamma)}$ apart, the second allows us to approximate the kernel on a grid.

²We avoid the use of the term ‘probability distribution’ as this conflicts with the notion of a (Schwarz) distribution that itself plays an important role in the underlying theory.

³Technically this is not completely correct; there are a few special cases, as we will see, where it is a pseudometric [41].

LEMMA 2.1 (BOUNDED TAILS). *If $\|p-q\| > \sqrt{h \ln(1/\gamma)}$ then $K(p, q) < \gamma$.*

LEMMA 2.2 (LIPSCHITZ). *For $\delta \in \mathbb{R}^d$ where $\|\delta\| < \varepsilon$, for points $p, q \in \mathbb{R}^d$ we have $|K(p, q) - K(p, q + \delta)| \leq \varepsilon/\sqrt{h}$.*

2.1 Problem Transformations

In prior research employing the kernel distance, *ad hoc* discretizations are used to convert the input objects (whether they be distributions, point clouds, curves or surfaces) to weighted discrete point sets. This process introduces error in the distance computations that usually go unaccounted for. In this subsection, we provide algorithms and analysis for *rigorously* discretizing input objects with guarantees on the resulting error. These algorithms, as a side benefit, provide a formal error-preserving reduction from kernel distance computation on curves and surfaces to the corresponding computations on discrete point sets.

After this section, we will assume that all data sets considered \mathcal{P}, \mathcal{Q} are discrete point sets of size n with weight functions $\mu : \mathcal{P} \rightarrow \mathbb{R}$ and $\nu : \mathcal{Q} \rightarrow \mathbb{R}$. The weights need not sum to 1 (i.e. need not be *probability* measures), nor be the same for \mathcal{P} and \mathcal{Q} ; we will set $W = \max(\sum_{p \in \mathcal{P}} \mu(p), \sum_{q \in \mathcal{Q}} \nu(q))$ to denote the total measure. This implies (since $K(p, p) = 1$) that $\kappa(\mathcal{P}, \mathcal{P}) \leq W^2$. All our algorithms will provide approximations of $\kappa(\mathcal{P}, \mathcal{Q})$ within additive error εW^2 . Since without loss of generality we can always normalize so that $W = 1$, our algorithms all provide an additive error of ε . We also set $\Delta = (1/h) \max_{u, v \in \mathcal{P} \cup \mathcal{Q}} \|u - v\|$ to capture the normalized diameter of the data.

Reducing orientation to weights.

The kernel distance between two oriented curves or surfaces can be reduced to a set of distance computations on appropriately weighted point sets. We illustrate this in the case of surfaces in \mathbb{R}^3 . The same construction will also work for curves in \mathbb{R}^d .

For each point $p \in \mathcal{P}$ we can decompose $U(p) \triangleq (U_1(p), U_2(p), U_3(p))$ into three fixed orthogonal components such as the coordinate axes $\{e_1, e_2, e_3\}$. Now

$$\begin{aligned} \kappa(\mathcal{P}, \mathcal{Q}) &= \int_{p \in \mathcal{P}} \int_{q \in \mathcal{Q}} K(p, q) \langle U(p), V(q) \rangle d\mu(p) d\nu(q) \\ &= \int_{p \in \mathcal{P}} \int_{q \in \mathcal{Q}} K(p, q) \sum_{i=1}^3 (U_i(p) V_i(q)) d\mu(p) d\nu(q) \\ &= \sum_{i=1}^3 \int_{p \in \mathcal{P}} \int_{q \in \mathcal{Q}} K(p, q) (U_i(p) V_i(q)) d\mu(p) d\nu(q) \\ &= \sum_{i=1}^3 \kappa(\mathcal{P}_i, \mathcal{Q}_i), \end{aligned}$$

where each $p \in \mathcal{P}_i$ has measure $\mu_i(p) = \mu(p) \|U_i(p)\|$. When the problem specifies U as a unit vector in \mathbb{R}^d , this approach reduces to d independent problems without unit vectors.

Reducing continuous to discrete.

We now present two simple techniques (gridding and sampling) to reduce a continuous \mathcal{P} to a discrete point set, incurring at most εW^2 error.

We construct a grid G_ε (of size $O((\Delta/\varepsilon)^d)$) on a smooth shape \mathcal{P} , so no point⁴ $p \in \mathcal{P}$ is further than $\varepsilon\sqrt{h}$ from a

⁴For distributions with decaying but infinite tails, we can

point $g \in G_\varepsilon$. Let P_g be all points in \mathcal{P} closer to g than any other point in G_ε . Each point g is assigned a weight $\mu(g) = \int_{p \in P_g} 1 d\mu(p)$. The correctness of this technique follows by Lemma 2.2.

Alternatively, we can sample $n = O((1/\varepsilon^2)(d + \log(1/\delta)))$ points at random from \mathcal{P} . If we have not yet reduced the orientation information to weights, we can generate d points each with weight $U_i(p)$. This works with probability at least $1 - \delta$ by invoking a coresets technique summarized in Theorem 5.2.

For the remainder of the paper, we assume our input dataset \mathcal{P} is a weighted point set in \mathbb{R}^d of size n .

3. COMPUTING THE KERNEL DISTANCE I: WSPDS

The well-separated pair decomposition (WSPD) [9, 21] is a standard data structure to approximately compute pairwise sums of distances in near-linear time. A consequence of Lemma 2.2 is that we can upper bound the error of estimating $K(p, q)$ by a nearby pair $K(\tilde{p}, \tilde{q})$. Putting these observations together yields (with some work) an approximation to the kernel distance. Since $D_K^2(\mathcal{P}, \mathcal{Q}) = \kappa(\mathcal{P}, \mathcal{P}) + \kappa(\mathcal{Q}, \mathcal{Q}) - 2\kappa(\mathcal{P}, \mathcal{Q})$, the problem reduces to computing $\kappa(\mathcal{P}, \mathcal{Q})$ efficiently and with an error of at most $(\varepsilon/4)W^2$.

Two sets A and B are said to be α -separated [9] if

$$\max\{\text{diam}(A), \text{diam}(B)\} \leq \alpha \min_{a \in A, b \in B} \|a - b\|.$$

Let $A \otimes B = \{\{x, y\} \mid x \in A, y \in B\}$ denote the set of all unordered pairs of elements formed by A and B . An α -WSPD of a point set P is a set of pairs $\mathcal{W} = \{\{A_1, B_1\}, \dots, \{A_s, B_s\}\}$ such that (i) $A_i, B_i \subset P$ for all i , (ii) $A_i \cap B_i = \emptyset$ for all i , (iii) disjointly $\bigcup_{i=1}^s A_i \otimes B_i = P \otimes P$, and (iv) A_i and B_i are α -separated for all i . For a point set $P \subset \mathbb{R}^d$ of size n , we can construct an α -WSPD of size $O(n/\alpha^d)$ in time $O(n \log n + n/\alpha^d)$ [21, 14].

We can use the WSPD construction to compute $D_K^2(\mathcal{P}, \mathcal{Q})$ as follows. We first create an α -WSPD of $\mathcal{P} \cup \mathcal{Q}$ in $O(n \log n + n/\alpha^d)$ time. Then for each pair $\{A_i, B_i\}$ we also store four sets $A_{i,\mathcal{P}} = \mathcal{P} \cap A_i$, $A_{i,\mathcal{Q}} = \mathcal{Q} \cap A_i$, $B_{i,\mathcal{P}} = \mathcal{P} \cap B_i$, and $B_{i,\mathcal{Q}} = \mathcal{Q} \cap B_i$. Let $a_i \in A_i$ and $b_i \in B_i$ be arbitrary elements, and let $D_i = \|a_i - b_i\|$. By construction, D_i approximates the distance between any pair of elements in $A_i \times B_i$ with error at most $2\alpha D_i$.

In each pair $\{A_i, B_i\}$, we can compute the weight of the edges from \mathcal{P} to \mathcal{Q} :

$$W_i = \left(\sum_{p \in A_{i,\mathcal{P}}} \mu(p) \right) \left(\sum_{q \in B_{i,\mathcal{Q}}} \nu(q) \right) + \left(\sum_{q \in A_{i,\mathcal{Q}}} \nu(q) \right) \left(\sum_{p \in B_{i,\mathcal{P}}} \mu(p) \right).$$

We estimate the contribution of the edges in pair (A_i, B_i) to $\kappa(\mathcal{P}, \mathcal{Q})$ as

$$\begin{aligned} &\sum_{(a,b) \in A_{i,\mathcal{P}} \times B_{i,\mathcal{Q}}} \mu(a)\nu(b) e^{-D_i^2/h} \\ &+ \sum_{(a,b) \in A_{i,\mathcal{Q}} \times B_{i,\mathcal{P}}} \mu(b)\nu(a) e^{-D_i^2/h} = W_i e^{-D_i^2/h}. \end{aligned}$$

Since D_i has error at most $2\alpha D_i$ for each pair of points, Lemma 2.2 bounds the error as at most $W_i(2\alpha D_i/\sqrt{h})$.

truncate to ignore tails such that the integral of the ignored area is at most $(1 - \varepsilon/2)W^2$ and proceed with this approach using $\varepsilon/2$ instead of ε .

In order to bound the total error to $(\varepsilon/4)W^2$, we bound the error for each pair by $(\varepsilon/4)W_i$ since

$$\sum_i W_i = \sum_{p \in P} \sum_{q \in Q} \mu(p)\nu(q) = W^2.$$

By Lemma 2.1, if $D_i > \sqrt{h \ln(1/\gamma)}$, then $e^{-D_i^2/h} < \gamma$. So for any pair with $D_i > 2\sqrt{h \ln(4/\varepsilon)}$, (for $\alpha < 1/2$) we can ignore, because they cannot have an effect on $\kappa(\mathcal{P}, \mathcal{Q})$ of more than $(1/4)\varepsilon W_i$, and thus cannot have error more than $(1/4)\varepsilon W_i$.

Since we can ignore pairs with $D_i > 2\sqrt{h \ln(4/\varepsilon)}$, each pair will have error at most $W_i(2\alpha(2\sqrt{h \ln(4/\varepsilon)}/\sqrt{h}) = W_i(4\alpha\sqrt{\ln(4/\varepsilon)})$. We can set this equal to $(\varepsilon/4)W_i$ and solve for $\alpha < (1/4)\varepsilon/\sqrt{\ln(4/\varepsilon)}$. This ensures that each pair with $D_i \leq 2\sqrt{h \ln(4/\varepsilon)}$ will have error less than $(\varepsilon/4)W_i$, as desired.

THEOREM 3.1. *By building and using an $((\varepsilon/4)/\sqrt{\ln(4/\varepsilon)})$ -WSPD, we can compute a value U , such that $|U - D_K^2(\mathcal{P}, \mathcal{Q})| \leq \varepsilon W^2$, in time $O(n \log n + (n/\varepsilon^d) \log^{d/2}(1/\varepsilon))$.*

4. COMPUTING THE KERNEL DISTANCE II: APPROXIMATE FEATURE MAPS

In this section, we describe (approximate) feature representations $\Phi(\mathcal{P}) = \sum_{p \in \mathcal{P}} \phi(p)\mu(p)$ for shapes and distributions that reduce the kernel distance computation to an ℓ_2 distance calculation $\|\Phi(\mathcal{P}) - \Phi(\mathcal{Q})\|_{\mathcal{H}}$ in an RKHS, \mathcal{H} . This mapping immediately yields algorithms for a host of analysis problems on shapes and distributions, by simply applying Euclidean space algorithms to the resulting feature vectors.

The feature map ϕ allows us to translate the kernel distance (and norm) computations into operations in a RKHS that take time $O(n\rho)$ if \mathcal{H} has dimension ρ , rather than the brute force time $O(n^2)$. Unfortunately, \mathcal{H} is in general infinite dimensional, including the case of the Gaussian kernel. Thus, we use dimensionality reduction to find an approximate mapping $\tilde{\phi} : \mathbb{R}^d \rightarrow \mathbb{R}^\rho$ (where $\tilde{\Phi}(\mathcal{P}) = \sum_{p \in \mathcal{P}} \tilde{\phi}(p)$) that approximates $\kappa(\mathcal{P}, \mathcal{Q})$:

$$\left| \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} K(p, q)\mu(p)\nu(q) - \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \langle \tilde{\phi}(p), \tilde{\phi}(q) \rangle \right| \leq \varepsilon W^2.$$

The analysis in the existing literature on approximating feature space does not directly bound the dimension ρ required for a specific error bound⁵. We derive bounds from two known techniques: random projections [36] (for shift-invariant kernels, includes Gaussians) and the Fast Gauss Transform [48, 20] (for Gaussian kernel). We produce three different features maps, with different bounds on the number of dimensions ρ depending on $\log n$ (n is the number of points), ε (the error), δ (the probability of failure), Δ (the normalized diameter of the points), and/or d (the ambient dimension of the data *before* the map).

4.1 Random Projections Feature Space

Rahimi and Recht [36] proposed a feature mapping that essentially applies an implicit Johnson-Lindenstrauss projection from $\mathcal{H} \rightarrow \mathbb{R}^\rho$. The approach works for any shift

⁵Explicit matrix versions of the Johnson-Lindenstrauss lemma [24] cannot be directly applied because the source space is itself infinite dimensional, rather than \mathbb{R}^d .

invariant kernel (i.e one that can be written as $K(p, q) = k(p - q)$). For the Gaussian kernel, $k(z) = e^{-\|z\|^2/2}$, where $z \in \mathbb{R}^d$. Let the Fourier transform of $k : \mathbb{R}^d \rightarrow \mathbb{R}^+$ is $g(\omega) = (2\pi)^{-d/2} e^{-\|\omega\|^2/2}$. A basic result in harmonic analysis [37] is that k is a kernel if and only if g is a measure (and after scaling, is a probability distribution). Let ω be drawn randomly from the distribution defined by g :

$$k(x - y) = \int_{\omega \in \mathbb{R}^d} g(\omega) e^{i\langle \omega, x - y \rangle} d\omega = E_\omega[\langle \psi_\omega(x), \psi_\omega(y) \rangle],$$

where $\psi_\omega(z) = (\cos(\langle \omega, z \rangle), \sin(\langle \omega, z \rangle))$ are the real and imaginary components of $e^{i\langle \omega, z \rangle}$. This implies that $\langle \psi_\omega(x), \psi_\omega(y) \rangle$ is an unbiased estimator of $k(x - y)$.

We now consider a ρ -dimensional feature vector $\phi_\Upsilon : P \rightarrow \mathbb{R}^\rho$ where the $(2i - 1)$ th and $(2i)$ th coordinates are described by $\mu(p)\psi_{\omega_i}(p)2/\rho = (2\mu(p) \cos(\langle \omega_i, z \rangle))/\rho, 2\mu(p) \sin(\langle \omega_i, z \rangle)/\rho$ for some $\omega_i \in \Upsilon = \{\omega_1, \dots, \omega_{\rho/2}\}$ drawn randomly from g . Next we prove a bound on ρ using this construction.

LEMMA 4.1. *For $\phi_\Upsilon : \mathcal{P} \cup \mathcal{Q} \rightarrow \mathbb{R}^\rho$ with $\rho = O((1/\varepsilon^2) \log(n/\delta))$, with probability $\geq 1 - \delta$*

$$\left| \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} K(p, q)\mu(p)\nu(q) - \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle \right| \leq \varepsilon W^2.$$

PROOF. We make use of the following Chernoff-Hoeffding bound. Given a set $\{X_1, \dots, X_n\}$ of independent random variables, such that $|X_i - E[X_i]| \leq \Lambda$, then for $M = \sum_{i=1}^n X_i$ we can bound $\Pr[|M - E[M]| \geq \alpha] \leq 2e^{-2\alpha^2/(n\Lambda^2)}$. We can now bound the error of using ϕ_Υ for any pair $(p, q) \in \mathcal{P} \times \mathcal{Q}$ as follows:

$$\begin{aligned} & \Pr[|\langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle - \mu(p)\nu(q)k(p - q)| \geq \varepsilon\mu(p)\nu(q)] \\ &= \Pr[|\langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle - E_\Upsilon[\langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle]| \geq \varepsilon\mu(p)\nu(q)] \\ &\leq \Pr\left[\left| \sum_i \frac{2}{\rho} \mu(p)\nu(q) \langle \psi_{\omega_i}(p), \psi_{\omega_i}(q) \rangle - E_\Upsilon\left[\sum_i \frac{2}{\rho} \mu(p)\nu(q) \langle \psi_{\omega_i}(p), \psi_{\omega_i}(q) \rangle\right] \right| \geq \varepsilon\mu(p)\nu(q)\right] \\ &\leq 2e^{-2(\varepsilon\mu(p)\nu(q))^2/(\rho\Lambda^2/2)} \leq 2e^{-\rho\varepsilon^2/64}, \end{aligned}$$

where the last inequality follows by

$$\Lambda \leq 2 \max_{p, q} (2/\rho)\mu(p)\nu(q) \langle \psi_\omega(p), \psi_\omega(q) \rangle \leq 8(2/\rho)\mu(p)\nu(q)$$

since for each pair of coordinates $\|\psi_\omega(p)\| \leq 2$ for all $p \in \mathcal{P}$ (or $q \in \mathcal{Q}$). By the union bound, the probability that this holds for all pairs of points $(p, q) \in \mathcal{P} \times \mathcal{Q}$ is given by

$$\Pr[\forall (p, q) \in \mathcal{P} \times \mathcal{Q} \quad |\langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle - \mu(p)\nu(q)k(p - q)| \geq \varepsilon\mu(p)\nu(q)] \leq (n^2)2e^{-\rho\varepsilon^2/64}.$$

Setting $\delta \geq n^2 2e^{-\rho\varepsilon^2/64}$ and solving for ρ yields that for $\rho = O((1/\varepsilon^2) \log(n/\delta))$, with probability at least $1 - \delta$, for all $(p, q) \in \mathcal{P} \times \mathcal{Q}$ we have $|\mu(p)\nu(q)k(p - q) - \langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle| \leq \varepsilon\mu(p)\nu(q)$. It follows that with probability at least $1 - \delta$

$$\begin{aligned} & \left| \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p)\nu(q)K(p, q) - \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \langle \phi_\Upsilon(p), \phi_\Upsilon(q) \rangle \right| \\ & \leq \varepsilon \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p)\nu(q) \leq \varepsilon W^2. \quad \square \end{aligned}$$

Note that the analysis of Rahimi and Recht [36] is done for unweighted point sets (i.e. $\mu(p) = 1$) and actually goes

further, in that it yields a guarantee for any pair of points taken from a manifold \mathcal{M} having diameter Δ . They do this by building an ε -net over the domain and applying the above tail bounds to the ε -net. We can adapt this trick to replace the $(\log n)$ term in ρ by a $(d \log(\Delta/\varepsilon))$ term, recalling $\Delta = (1/h) \max_{p,p' \in \mathcal{P} \cup \Omega} \|p - p'\|$. This leads to the same guarantees as above with a dimension of $\rho = O((d/\varepsilon^2) \log(\Delta/\varepsilon\delta))$.

4.2 Fast Gauss Transform Feature Space

The above approach works by constructing features in the frequency domain. In what follows, we present an alternative approach that operates in the spatial domain directly. We base our analysis on the Improved Fast Gauss Transform (IFGT) [48], an improvement on the Fast Gauss Transform. We start with a brief review of the IFGT (see the original work [48] for full details).

IFGT feature space construction.

The goal of the IFGT is to approximate $\kappa(\mathcal{P}, \Omega)$. First we rewrite $\kappa(\mathcal{P}, \Omega)$ as the summation $\sum_{q \in \Omega} G(q)$ where $G(q) = \nu(q) \sum_{p \in \mathcal{P}} e^{-\|p-q\|^2/h^2} \mu(p)$. Next, we approximate $G(q)$ in two steps. First we rewrite

$$G(q) = \nu(q) \sum_{p \in \mathcal{P}} \mu(p) e^{-\frac{\|q-x_*\|^2}{h^2}} e^{-\frac{\|p-x_*\|^2}{h^2}} e^{\frac{2\|q-x_*\| \cdot \|p-x_*\|}{h^2}},$$

where the quantity x_* is a fixed vector that is usually the centroid of \mathcal{P} . The first two exponential terms can be computed for each p and q once. Second, we approximate the remaining exponential term by its (Hermite) Taylor expansion $e^v = \sum_{i \geq 0} \frac{v^i}{i!}$.

After a series of algebraic manipulations, the following expression emerges:

$$G(q) = \nu(q) e^{-\frac{\|q-x_*\|^2}{h^2}} \sum_{\alpha \geq 0} C_\alpha \left(\frac{q-x_*}{h} \right)^\alpha$$

where C_α is given by

$$C_\alpha = \frac{2^{|\alpha|}}{\alpha!} \sum_{p \in \mathcal{P}} \mu(p) e^{-\frac{\|p-x_*\|^2}{h^2}} \left(\frac{p-x_*}{h} \right)^\alpha.$$

The parameter α is a *multiindex*, and is actually a vector $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)$ of dimension d . The expression z^α , for $z \in \mathbb{R}^d$, denotes the monomial $z_1^{\alpha_1} z_2^{\alpha_2} \dots z_d^{\alpha_d}$, the quantity $|\alpha|$ is the *total degree* $\sum \alpha_i$, and the quantity $\alpha! = \prod_i (\alpha_i!)$. The multiindices are sorted in *graded lexicographic order*, which means that α comes before α' if $|\alpha| < |\alpha'|$, and two multiindices of the same degree are ordered lexicographically.

The above expression for $G(q)$ is an exact infinite sum, and is approximated by truncating the summation at multiindices of total degree $\tau - 1$. There are at most $\rho = \binom{\tau+d-1}{d} = O(\tau^d)$ such multiindices. Construct a mapping $\tilde{\phi} : \mathbb{R}^d \rightarrow \mathbb{R}^\rho$. Let $\tilde{\phi}(p)_\alpha = \sqrt{\frac{2^{|\alpha|}}{\alpha!}} \mu(p) e^{-\frac{\|p-x_*\|^2}{h^2}} \left(\frac{p-x_*}{h} \right)^\alpha$. Then

$$G(q) = \sum_{\alpha} \tilde{\phi}(q)_\alpha \sum_{p \in \mathcal{P}} \tilde{\phi}(p)_\alpha$$

and $S = \sum_{q \in \Omega} G(q)$ is then given by

$$S = \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} \sum_{\alpha} \tilde{\phi}(q)_\alpha \tilde{\phi}(p)_\alpha = \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} \langle \tilde{\phi}(q), \tilde{\phi}(p) \rangle.$$

IFGT error analysis.

The error incurred by truncating the sum at degree $\tau - 1$ is given by

$$\begin{aligned} \text{Err}(\tau) &= \left| \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} K(p, q) \mu(p) \nu(q) - \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} \langle \tilde{\phi}(p), \tilde{\phi}(q) \rangle \right| \\ &\leq \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} \mu(p) \nu(q) \frac{2^\tau}{\tau!} \Delta^{2\tau} = W^2 \frac{2^\tau}{\tau!} \Delta^{2\tau}. \end{aligned}$$

Set $\varepsilon W^2 = \text{Err}(\tau)$. Applying Stirling's approximation, we solve for τ in $\log(1/\varepsilon) \geq \tau \log(\tau/4\Delta^2)$. This yields the bounds $\tau = O(\Delta^2)$ and $\tau = O(\log(1/\varepsilon))$. Thus our error bound holds for $\tau = O(\Delta^2 + \log(1/\varepsilon))$. Using $\rho = O(\tau^d)$, we obtain the following result.

LEMMA 4.2. *There exists a mapping $\tilde{\phi} : \mathcal{P} \cup \Omega \rightarrow \mathbb{R}^\rho$ with $\rho = O(\Delta^{2d} + \log^d(1/\varepsilon))$ so*

$$\left| \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} K(p, q) \mu(p) \nu(q) - \sum_{p \in \mathcal{P}} \sum_{q \in \Omega} \langle \tilde{\phi}(p), \tilde{\phi}(q) \rangle \right| \leq \varepsilon W^2.$$

4.3 Summary of Feature Maps

We have developed three different bounds on the dimension required for feature maps that approximate $\kappa(\mathcal{P}, \Omega)$ to within εW^2 .
IFGT: $\rho = O(\Delta^{2d} + \log^d(1/\varepsilon))$. Lemma 4.2. Advantages: deterministic, independent of n , logarithmic dependence on $1/\varepsilon$. Disadvantages: polynomial dependence on Δ , exponential dependence on d .

Random-points: $\rho = O((1/\varepsilon^2) \log(n/\delta))$. Lemma 4.1. Advantages: independent of Δ and d . Disadvantages: randomized, dependent on n , polynomial dependence on $1/\varepsilon$.

Random-domain: $\rho = O((d/\varepsilon^2) \log(\Delta/\varepsilon\delta))$. (above) Advantages: independent of n , logarithmic dependence on Δ , polynomial dependence on d . Disadvantages: randomized, dependence on Δ and d , polynomial dependence on $1/\varepsilon$.

Feature-based computation of D_K .

As before, we can decompose $D_K^2(\mathcal{P}, \Omega) = \kappa(\mathcal{P}, \mathcal{P}) + \kappa(\Omega, \Omega) - 2\kappa(\mathcal{P}, \Omega)$ and use Lemma 4.1 to approximate each of $\kappa(\mathcal{P}, \mathcal{P})$, $\kappa(\Omega, \Omega)$, and $\kappa(\mathcal{P}, \Omega)$ with error $\varepsilon W^2/4$.

THEOREM 4.1. *We can compute a value U in time $O((n/\varepsilon^2) \log(n/\delta))$ such that $|U - D_K^2(\mathcal{P}, \Omega)| \leq \varepsilon W^2$, with probability at least $1 - \delta$.*

A nearest-neighbor algorithm.

The feature map does more than yield efficient algorithms for the kernel distance. As a representation for shapes and distributions, it allows us to solve other data analysis problems on shape spaces using off-the-shelf methods that apply to points in Euclidean space. As a simple example of this, we can combine the **Random-points** feature map with known results on approximate nearest-neighbor search in Euclidean space [3] to obtain the following result.

LEMMA 4.3. *Given a collection of m point sets $\mathcal{C} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m\}$, and a query surface Ω , we can compute the c -approximate nearest neighbor to Ω in \mathcal{C} under the kernel distance in time $O(\rho m^{1/c^2+o(1)})$ query time using $O(\rho m^{1+1/c^2+o(1)})$ space and preprocessing.*

5. CORESET FOR THE KERNEL DISTANCE

The kernel norm (and distance) can be approximated in near-linear time; however, this may be excessive for large data sets. Rather, we extract a small subset (a coreset) \mathcal{S} from the input \mathcal{P} such that the kernel distance between \mathcal{S} and \mathcal{P} is small. By triangle inequality, \mathcal{S} can be used as a proxy for \mathcal{P} . Specifically, we extend the notion of ε -samples for range spaces to handle non-binary range spaces defined by kernels.

Background on range spaces.

Let $\xi(P)$ denote the total weight of a set of points P , or cardinality if no weights are assigned. Let $P \subset \mathbb{R}^d$ be a set of points and let \mathcal{A} be a family of subsets of P . For examples of \mathcal{A} , let \mathcal{B} denote the set of all subsets defined by containment in a ball and let \mathcal{E} denote the set of all subsets defined by containment in an ellipse. We say (P, \mathcal{A}) is a *range space*. Let $\bar{\xi}_P(A) = \xi(A)/\xi(P)$. An ε -*sample* (or ε -*approximation*) of (P, \mathcal{A}) is a subset $Q \subset P$ such that

$$\max_{A \in \mathcal{A}} |\bar{\xi}_Q(Q \cap A) - \bar{\xi}_P(P \cap A)| \leq \varepsilon.$$

To create a coreset for the kernel norm, we want to generalize these notions of ε -samples to non-binary ($\{0, 1\}$ -valued instead of $\{0, 1\}$ -valued) functions, specifically to kernels. For two point sets P, Q , define

$$\bar{\kappa}(P, Q) = \frac{1}{\xi(P)} \frac{1}{\xi(Q)} \sum_{p \in P} \sum_{q \in Q} K(p, q),$$

and when we have a singleton set $Q = \{q\}$ and a subset $P' \subseteq P$ then we write $\bar{\kappa}_P(P', q) = (1/\xi(P)) \sum_{p \in P'} K(p, q)$. Let $K^+ = \max_{p, q \in P} K(p, q)$ be the maximum value a kernel can take on a dataset P , which can be normalized to $K^+ = 1$. We say a subset of $S \subset P$ is an ε -*sample of (P, K)* if

$$\max_q |\bar{\kappa}_P(P, q) - \bar{\kappa}_S(S, q)| \leq \varepsilon K^+.$$

The standard notion of VC-dimension [47] (and related notion of shattering dimension) is fundamentally tied to the binary ($\{0, 1\}$ -valued) nature of ranges, and as such, it does not directly apply to ε -samples of (P, K) . Other researchers have defined different combinatorial dimensions that can be applied to kernels [15, 25, 1, 46]. The best result is based on γ -fat shattering dimension F_γ [25], defined for a family of $(0, 1)$ -valued functions \mathcal{F} and a ground set P . A set $Y \subset P$ is γ -*fat shattered* by \mathcal{F} if there exists a function $\alpha : Y \rightarrow [0, 1]$ such that for all subsets $Z \subseteq Y$ there exists some $F_Z \in \mathcal{F}$ such that for every $x \in Z$ $F_Z(x) \geq \alpha(x) + \gamma$ and for every $x \in Y \setminus Z$ $F_Z(x) \leq \alpha(x) - \gamma$. Then $F_\gamma = \xi(Y)$ for the largest cardinality set $Y \subset P$ that can be γ -fat shattered. Bartlett *et al.* [7] show that a random sample of $O((1/\varepsilon^2)(F_\gamma \log^2(F_\gamma/\varepsilon) + \log(1/\delta)))$ elements creates an ε -sample (with probability at least $1 - \delta$) with respect to (P, \mathcal{F}) for $\gamma = \Omega(\varepsilon)$. Note that the γ -fat shattering dimension of Gaussian and other symmetric kernels in \mathbb{R}^d is $d + 1$ (by setting $\alpha(x) = .5$ for all x), the same as balls \mathcal{B} in \mathbb{R}^d , so this gives a random-sample construction for ε -samples of (P, K) of size $O((d/\varepsilon^2)(\log^2(1/\varepsilon) + \log(1/\delta)))$.

In this paper, we improve this result in two ways by directly relating a kernel range space (P, K) to a similar (binary) range space (P, \mathcal{A}) . First, this improves the random-sample bound because it uses sample-complexity results for binary range spaces that have been heavily optimized. Sec-

ond, this allows for all deterministic ε -sample constructions (which have no probability of failure) and can have much smaller size.

Constructions for ε -samples.

Vapnik and Chervonenkis [47] showed that the complexity of ε -samples is tied to the VC-dimension of the range space. That is, given a range space (X, \mathcal{A}) a subset $Y \subset X$ is said to be *shattered* by \mathcal{A} if all subsets of $Z \subset Y$ can be realized as $Z = Y \cap R$ for $R \in \mathcal{A}$. Then the *VC-dimension* of a range space (X, \mathcal{A}) is the cardinality of the largest subset $Y \subset X$ that can be shattered by \mathcal{A} . Vapnik and Chervonenkis [47] showed that if the VC-dimension of a range space (X, \mathcal{A}) is ν , then a random sample Y of $O((1/\varepsilon^2)(\nu \log(1/\varepsilon) + \log(1/\delta)))$ points from X is an ε -sample with probability at least $1 - \delta$. This bound was improved to $O((1/\varepsilon^2)(\nu + \log 1/\delta))$ by Talagrand [44, 26].

Alternatively, Matousek [28] showed that ε -samples of size $O((\nu/\varepsilon^2) \log(\nu/\varepsilon))$ could be constructed deterministically, that is there is no probability of failure. A simpler algorithm with more thorough runtime analysis is presented in Chazelle and Matousek [12], which runs in $O(d)^{3d} |X| (1/\varepsilon)^{2\nu} \log^\nu(1/\varepsilon)$ time. Smaller ε -samples exist; in particular Matousek, Welzl, and Wernisch [31], improved by Matousek [29], show that ε -samples exist of size $O((1/\varepsilon)^{2-2/(\nu+1)})$, based on a discrepancy result that says there exists a labeling $\chi : X \rightarrow \{-1, +1\}$ such that $\max_{R \in \mathcal{A}} \sum_{x \in R \cap X} \chi(x) \leq O(|X|^{1/2-1/2\nu})$. It is alluded by Chazelle [11] that if an efficient construction for such a labeling existed, then an algorithm for creating ε -samples of size $O((1/\varepsilon)^{2-2/(\nu+1)})$ would follow, see also Phillips [34]. Recently, Bansal [6] provided a randomized polynomial algorithm for the entropy method, which is central in proving these existential discrepancy bounds. This leads to an algorithm that runs in time $O(|X| \cdot \text{poly}(1/\varepsilon))$, as claimed by Charikar *et al.* [10, 33].

An alternate approach is through the VC-dimension of the *dual range space* $(\mathcal{A}, \bar{\mathcal{A}})$, of (primal) range space (X, \mathcal{A}) , where $\bar{\mathcal{A}}$ is the set of subsets of ranges \mathcal{A} defined by containing the same element of X . In our context, for range spaces defined by balls $(\mathbb{R}^d, \mathcal{B})$ and ellipses of fixed orientation $(\mathbb{R}^d, \mathcal{E})$ their dual range spaces have VC-dimension $\bar{\nu} = d$. Matousek [30] shows that a technique of matching with low-crossing number [13] along with Haussler's packing lemma [22] can be used to construct a low discrepancy coloring for range spaces where $\bar{\nu}$, the VC-dimension of the dual range space is bounded. This technique can be made deterministic and runs in $\text{poly}(|X|)$ time. Invoking this technique in the Chazelle and Matousek [12] framework yields an ε -sample of size $O((1/\varepsilon)^{2-2/(\bar{\nu}+1)} (\log(1/\varepsilon))^{2-1/(\bar{\nu}+1)})$ in $O(|X| \cdot \text{poly}(1/\varepsilon))$ time. Specifically, we attain the following result:

LEMMA 5.1. *For discrete ranges spaces (X, \mathcal{B}) and (X, \mathcal{E}) for $X \in \mathbb{R}^d$ of size n , we can construct an ε -sample of size $O((1/\varepsilon)^{2-2/(d+1)} (\log(1/\varepsilon))^{2-1/(d+1)})$ in $O(n \cdot \text{poly}(1/\varepsilon))$ time.*

For specific range spaces, the size of ε -samples can be improved beyond the VC-dimension bound. Phillips [34] showed for ranges \mathcal{R}_d consisting of axis-aligned boxes in \mathbb{R}^d , that ε -samples can be created of size $O((1/\varepsilon) \cdot \log^{2d}(1/\varepsilon))$. This can be generalized to ranges defined by k predefined normal directions of size $O((1/\varepsilon) \log^{2k}(1/\varepsilon))$. These algorithms run in time $O(|X|(1/\varepsilon^3) \text{poly} \log(1/\varepsilon))$. And for in-

tervals \mathcal{J} over \mathbb{R} , ε -samples of (X, \mathcal{J}) can be created of size $O(1/\varepsilon)$ by sorting points and retaining every $\varepsilon|X|$ th point in the sorted order [27].

ε -Samples for kernels.

The *super-level set* of a kernel given one input $q \in \mathbb{R}^d$ and a value $v \in \mathbb{R}^+$, is the set of all points $p \in \mathbb{R}^d$ such that $K(p, q) \geq v$. We say that a kernel is *linked* to a range space $(\mathbb{R}^d, \mathcal{A})$ if for every possible input point $q \in \mathbb{R}^d$ and any value $v \in \mathbb{R}^+$ that the super-level set of $K(\cdot, q)$ defined by v is equal to some $H \in \mathcal{A}$. For instance multi-variate Gaussian kernels with no skew are linked to $(\mathbb{R}^d, \mathcal{B})$ since all super-level sets are balls, and multi-variate Gaussian kernels with non-trivial covariance are linked to $(\mathbb{R}^d, \mathcal{E})$ since all super-level sets are ellipses.

THEOREM 5.1. *For any kernel $K : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}^+$ linked to a range space $(\mathcal{M}, \mathcal{A})$, an ε -sample S of (P, \mathcal{A}) for $S \subseteq \mathcal{M}$ is a ε -sample of (P, K) .*

A (flawed) attempt at a proof may proceed by considering a series of approximate level-sets, within which each point has about the same function value. Since S is an ε -sample of (P, \mathcal{A}) , we can guarantee the density of S and P in each level set is off by at most 2ε . However, the sum of absolute error over all approximate level-sets is approximately εK^+ times the number of approximate level sets. This analysis fails because it allows error to accumulate; however, a more careful application of the ε -sample property shows it cannot. A correct analysis follows using a charging scheme which prevents the error from accumulating.

PROOF. We can sort all $p_i \in P$ in similarity to q so that $p_i < p_j$ (and by notation $i < j$) if $K(p_i, q) > K(p_j, q)$. Thus any super-level set containing p_j also contains p_i for $i < j$. We can now consider the one-dimensional problem on this sorted order from q .

We now count the deviation $E(P, S, q) = \bar{\kappa}_P(P, q) - \bar{\kappa}_S(S, q)$ from p_1 to p_n using a charging scheme. That is each element $s_j \in S$ is charged to $\xi(P)/\xi(S)$ points in P . For simplicity we will assume that $k = \xi(P)/\xi(S)$ is an integer, otherwise we can allow fractional charges. We now construct a partition of P slightly differently, for positive and negative $E(P, S, q)$ values, corresponding to undercounts and overcounts, respectively.

Undercount of $\bar{\kappa}_S(S, q)$. For undercounts, we partition P into $2\xi(S)$ (possibly empty) sets $\{P'_1, P_1, P'_2, P_2, \dots, P'_{\xi(S)}, P_{\xi(S)}\}$ of consecutive points by the sorted order from q . Starting with p_1 (the closest point to q) we place points in sets P'_j or P_j following their sorted order. Recursively on j and i , starting at $j = 1$ and $i = 1$, we place each p_i in P'_j as long as $K(p_i, q) > K(s_j, q)$ (this may be empty). Then we place the next k points p_i into P_j . After k points are placed in P_j , we begin with P'_{j+1} , until all of P has been placed in some set. Let $t \leq \xi(S)$ be the index of the last set P_j such that $\xi(P_j) = k$. Note that for all $p_i \in P_j$ (for $j \leq t$) we have $K(s_j, q) \geq K(p_i, q)$, thus $\bar{\kappa}_S(\{s_j\}, q) \geq \bar{\kappa}_P(P_j, q)$.

We can now bound the undercount as

$$\begin{aligned} E(P, S, q) &= \sum_{j=1}^{\xi(S)} (\bar{\kappa}_P(P_j, q) - \bar{\kappa}_S(\{s_j\}, q)) + \sum_{j=1}^{\xi(S)} \bar{\kappa}_P(P'_j, q) \\ &\leq \sum_{j=1}^{t+1} \bar{\kappa}_P(P'_j, q) \end{aligned}$$

since the first term is at most 0 and since $\xi(P'_j) = 0$ for $j > t + 1$. Now consider a super-level set $H \in \mathcal{A}$ containing all points before s_{t+1} ; H is the smallest range that contains every non-empty P'_j . Because (for $j \leq t$) each set P_j can be charged to s_j , then $\sum_{j=1}^t \xi(P_j \cap H) = k \cdot \xi(S \cap H)$. And because S is an ε -sample of (P, \mathcal{A}) , then $\sum_{j=1}^{t+1} \xi(P'_j) = \left(\sum_{j=1}^{t+1} \xi(P'_j) + \sum_{j=1}^t \xi(P_j \cap H) \right) - k \cdot \xi(S \cap H) \leq \varepsilon \xi(P)$. We can now bound

$$\begin{aligned} E(P, S, q) &\leq \sum_{j=1}^{t+1} \bar{\kappa}_P(P'_j, q) = \sum_{j=1}^{t+1} \sum_{p \in P'_j} \frac{K(p, q)}{\xi(P)} \\ &\leq \frac{1}{\xi(P)} \sum_{j=1}^{t+1} \xi(P'_j) K^+ \leq \frac{1}{\xi(P)} (\varepsilon \xi(P)) K^+ = \varepsilon K^+. \end{aligned}$$

Overcount of $\bar{\kappa}_S(S, q)$: The analysis for overcounts is similar to undercounts, but we construct the partition in reverse and the leftover after the charging is not quite as clean to analyze. For overcounts, we partition P into $2\xi(S)$ (possibly empty) sets $\{P_1, P'_1, P_2, P'_2, \dots, P_{\xi(S)}, P'_{\xi(S)}\}$ of consecutive points by the sorted order from q . Starting with p_n (the furthest point from q) we place points in sets P'_j or P_j following their reverse-sorted order. Recursively on j and i , starting at $j = \xi(S)$ and $i = n$, we place each p_i in P'_j as long as $K(p_i, q) < K(s_j, q)$ (this may be empty). Then we place the next k points p_i into P_j . After k points are placed in P_j , we begin with P'_{j-1} , until all of P has been placed in some set. Let $t \leq \xi(S)$ be the index of the last set P_j such that $\xi(P_j) = k$ (the smallest such j). Note that for all $p_i \in P_j$ (for $j \geq t$) we have $K(s_j, q) \leq K(p_i, q)$, thus $\bar{\kappa}_S(\{s_j\}, q) \leq \bar{\kappa}_P(P_j, q)$.

We can now bound the (negative) undercount as

$$\begin{aligned} E(P, S, q) &= \sum_{j=\xi(S)}^t (\bar{\kappa}_P(P_j, q) - \bar{\kappa}_S(\{s_j\}, q)) \\ &\quad + \sum_{j=t-1}^1 (\bar{\kappa}_P(P_j, q) - \bar{\kappa}_S(\{s_j\}, q)) + \sum_{j=1}^{\xi(S)} \bar{\kappa}_P(P'_j, q) \\ &\geq (\bar{\kappa}_P(P_{t-1}, q) - \bar{\kappa}_S(\{s_{t-1}\}, q)) - \sum_{j=t-2}^1 \bar{\kappa}_S(\{s_j\}, q), \end{aligned}$$

since the first full term is at least 0, as is each $\bar{\kappa}_P(P_j, q)$ and $\bar{\kappa}_P(P'_j, q)$ term in the second and third terms. We will need the one term $\bar{\kappa}_P(P_{t-1}, q)$ related to P in the case when $1 \leq \xi(P_{t-1}) < k$.

Now, using that S is an ε -sample of (P, \mathcal{A}) , we will derive a bound on t , and more importantly $(t-2)$. We consider the maximal super-level set $H \in \mathcal{A}$ such that no points $H \cap P$ are in P'_j for any j . This is the largest set where each point $p \in P$ can be charged to a point $s \in S$ such that $K(p, q) > K(s, q)$, and thus presents the smallest (negative) undercount. In this case, $H \cap P = \cup_{j=1}^s P_j$ for some s and $H \cap S = \cup_{j=1}^s \{s_j\}$. Since $t \leq s$, then $\xi(H \cap P) = (s - t + 1)k + \xi(P_{t-1}) = (s - t + 1)\xi(P)/\xi(S) + \xi(P_{t-1})$ and $\xi(H \cap S) = s$. Thus

$$\begin{aligned} \varepsilon &\geq \bar{\xi}_S(H \cap S) - \bar{\xi}_P(H \cap P) \\ &= \frac{s}{\xi(S)} - \frac{(s - t + 1)\xi(P)/\xi(S)}{\xi(P)} - \frac{\xi(P_{t-1})}{\xi(P)} \\ &\geq \frac{t-1}{\xi(S)} - \frac{\xi(P_{t-1})}{\xi(P)}. \end{aligned}$$

Thus $(t-2) \leq \varepsilon \xi(S) + \xi(P_{t-1})(\xi(S)/\xi(P)) - 1$. Letting $p_i = \min_{i' \in P_{t-1}} K(p_{i'}, q)$ (note $K(p_i, q) \geq K(s_{t-1}, q)$)

$$\begin{aligned} & E(P, S, q) \\ & \geq \frac{\kappa(P_{t-1}, q)}{\xi(P)} - \frac{K(s_{t-1}, q)}{\xi(S)} - \left(\varepsilon \xi(S) + \xi(P_{t-1}) \frac{\xi(S)}{\xi(P)} - 1 \right) \frac{K^+}{\xi(S)} \\ & = -\varepsilon K^+ + K^+ \left(\frac{k - \xi(P_{t-1})}{\xi(P)} \right) - \frac{k \cdot K(s_{t-1}, q) - \kappa(P_{t-1}, q)}{\xi(P)} \\ & \geq -\varepsilon K^+ + K^+ \left(\frac{k - \xi(P_{t-1})}{\xi(P)} \right) - K(p_i, q) \left(\frac{k - \xi(P_{t-1})}{\xi(P)} \right) \\ & \geq -\varepsilon K^+. \quad \square \end{aligned}$$

COROLLARY 5.1. *For a Gaussian kernel, any ε -sample S for (P, \mathcal{B}) (or for (P, \mathcal{E}) if we consider covariance) guarantees that for any query $q \in \mathbb{R}^d$ that $|\bar{\kappa}_P(P, q) - \bar{\kappa}_S(S, q)| \leq \varepsilon K^+$.*

Coreset-based computation of kernel distance.

For convenience here we assume that our kernel has been normalized so $K^+ = 1$. Let P be an ε -sample of (\mathcal{P}, K) , and all points $p \in P$ have uniform weights so $\xi(P) = \xi(\mathcal{P}) = W$. Then for any $q \in \mathbb{R}^d$

$$\varepsilon \geq |\bar{\kappa}_P(P, q) - \bar{\kappa}_{\mathcal{P}}(\mathcal{P}, q)| = \left| \frac{\kappa(P, q)}{\xi(P)} - \frac{\kappa(\mathcal{P}, q)}{\xi(\mathcal{P})} \right|.$$

and hence

$$|\kappa(P, q) - \kappa(\mathcal{P}, q)| \leq \varepsilon \xi(\mathcal{P}) = \varepsilon W.$$

It follows that if Q is also an ε -sample for (\mathcal{Q}, K) , then $\|\kappa(P, Q) - \kappa(\mathcal{P}, \mathcal{Q})\| \leq 2\varepsilon W^2$. Hence, via known constructions of ε -samples randomized [47, 44] or deterministic [28, 12, 31, 29, 43, 34] (which can be applied to weighted point sets to create unweighted ones [28]) and Theorem 5.1 we have the following theorems. The first shows how to construct a coreset with respect to D_K .

THEOREM 5.2. *Consider any kernel K linked with $(\mathbb{R}^d, \mathcal{B})$ and objects $\mathcal{P}, \mathcal{Q} \subset \mathcal{M} \subset \mathbb{R}^d$, each with total weight W , and for constant d . We can construct sets $P \subset \mathcal{P}$ and $Q \subset \mathcal{Q}$ such that $|D_K(\mathcal{P}, \mathcal{Q}) - D_K(P, Q)| \leq \varepsilon W^2$ of size:*

- $O((1/\varepsilon^{2-1/(d+1)}) \log^{2-1/d+1}(1/\varepsilon))$, via Lemma 5.1; or
- $O(1/\varepsilon^2)(d + \log(1/\delta))$ via random sampling (correct with probability at least $(1 - \delta)$).

We present an alternative sampling condition to Theorem 5.2 in the full version. It has larger dependence on ε , and also has either dependence on Δ or on $\log n$ (and is independent of K^+). Also in the full version we show that it is NP-hard to optimize ε with a fixed subset size k .

The associated runtimes are captured in the following algorithmic theorem.

THEOREM 5.3. *When K is linked to $(\mathbb{R}^d, \mathcal{B})$, we can compute a number \hat{D} such that $|D_K(\mathcal{P}, \mathcal{Q}) - \hat{D}| \leq \varepsilon$ in time:*

- $O(n \cdot (1/\varepsilon^{2d+2}) \log^{d+1}(1/\varepsilon))$; or
- $O(n + \log n) \cdot ((1/\varepsilon^2) \log(1/\delta)) + (1/\varepsilon^4) \log^2(1/\delta)$ that is correct with probability at least $1 - \delta$.

Notice that these results automatically work for any kernel linked with $(\mathbb{R}^d, \mathcal{B})$ (or more generally with $(\mathbb{R}^d, \mathcal{E})$) with no extra work; this includes not only Gaussians (with non-trivial covariance), but any other standard kernel such as triangle, ball, or Epanechnikov.

6. MINIMIZING THE KERNEL DISTANCE UNDER TRANSLATION AND ROTATION

We attack the problem of minimizing the kernel distance between \mathcal{P} and \mathcal{Q} under a set of transformations: translations or translations and rotations. A *translation* $T \in \mathbb{R}^d$ is a vector so $\mathcal{Q} \oplus T = \{q + T \mid q \in \mathcal{Q}\}$. The translation $T^* = \arg \min_{T \in \mathbb{R}^d} D_K(\mathcal{P}, \mathcal{Q} \oplus T)$, applied to \mathcal{Q} minimizes the kernel norm. A *rotation* $R \in \text{SO}(d)$ can be represented as a special orthogonal matrix. We can write $R \circ \mathcal{Q} = \{R(q) \mid q \in \mathcal{Q}\}$, where $R(q)$ rotates q about the origin, preserving its norm. The set of a translation and rotation $(T^*, R^*) = \arg \min_{(T, R) \in \mathbb{R}^d \times \text{SO}(d)} D_K(\mathcal{P}, R \circ (\mathcal{Q} \oplus T))$ applied to \mathcal{Q} minimizes the kernel norm.

Decomposition.

In minimizing $D_K(\mathcal{P}, R \circ (\mathcal{Q} \oplus T))$ under all translations and rotations, we can reduce this to a simpler problem. The first term $\kappa(\mathcal{P}, \mathcal{P}) = \sum_{p_1 \in \mathcal{P}} \sum_{p_2 \in \mathcal{P}} \mu(p_1) \mu(p_2) K(p_1, p_2)$ has no dependence on T or R , so it can be ignored. And the second term $\kappa(\mathcal{Q}, \mathcal{Q}) = \sum_{q_1 \in \mathcal{Q}} \sum_{q_2 \in \mathcal{Q}} \nu(q_1) \nu(q_2) K(R(q_1 + T), R(q_2 + T))$ can also be ignored because it is invariant under the choice of T and R . Each subterm $K(R(q_1 + T), R(q_2 + T))$ only depends on $\|R(q_1 + T) - R(q_2 + T)\| = \|q_1 - q_2\|$, which is also independent of T and R . Thus we can rephrase the objective as finding

$$\begin{aligned} (T^*, R^*) &= \arg \max_{(T, R) \in \mathbb{R}^d \times \text{SO}(d)} \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p) \nu(q) K(p, R(q + T)) \\ &= \arg \max_{(T, R) \in \mathbb{R}^d \times \text{SO}(d)} \kappa(\mathcal{P}, R \circ (\mathcal{Q} \oplus T)). \end{aligned}$$

We start by providing an approximately optimal translation. Then we describe the adapted result for handling both translations and rotations, leaving the details for the full version.

6.1 Approximately Optimal Translations

We describe, for any parameter $\varepsilon > 0$, an algorithm for a translation \hat{T} such that $D_K^2(\mathcal{P}, \mathcal{Q} \oplus \hat{T}) - D_K^2(\mathcal{P}, \mathcal{Q} \oplus T^*) \leq \varepsilon W^2$. We begin with a key lemma providing analytic structure to our problem.

LEMMA 6.1. $\kappa(\mathcal{P}, \mathcal{Q} \oplus T^*) \geq W^2/n^2$.

PROOF. When $T \in \mathbb{R}^d$ aligns $q \in \mathcal{Q}$ so $q + T = p$ for $p \in \mathcal{P}$ it ensures that $K(p, q) = 1$. We can choose the points p and q such that $\mu(p)$ and $\nu(q)$ are as large as possible. They must each be at least W/n , so $K(p, q) \mu(p) \nu(q) \geq W^2/n^2$. All other subterms in $\kappa(\mathcal{P}, \mathcal{Q} \oplus T)$ are at least 0. Thus $\kappa(\mathcal{P}, \mathcal{Q} \oplus T) \geq W^2/n^2$. \square

Thus, if $\kappa(\mathcal{P}, \mathcal{Q} \oplus T^*) \geq W^2/n^2$, then for some pair of points $p \in \mathcal{P}$ and $q \in \mathcal{Q}$ we must have $\mu(p) \nu(q) K(p, q + T^*) \geq \mu(p) \nu(q)/n^2$, i.e. $K(p, q + T^*) \geq 1/n^2$. Otherwise, if all n^2 pairs (p, q) satisfy $\mu(p) \nu(q) K(p, q + T^*) < \mu(p) \nu(q)/n^2$, then

$$\begin{aligned} \kappa(\mathcal{P}, \mathcal{Q} \oplus T^*) &= \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p) \nu(q) K(p, q + T^*) \\ &< \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p) \nu(q)/n^2 = W^2/n^2. \end{aligned}$$

Thus some pair $p \in \mathcal{P}$ and $q \in \mathcal{Q}$ must satisfy $\|p - (q + T^*)\| \leq \sqrt{\ln(n^2)}$, via Lemma 2.1 with $\gamma = 1/(n^2)$.

Let G_ε be a grid on \mathbb{R}^d so that when any point $p \in \mathbb{R}^d$ is snapped to the nearest grid point $g \in G_\varepsilon$, we guarantee

that $\|g - p\| \leq \varepsilon$. We can define an orthogonal grid $G_\varepsilon = \{(\varepsilon/\sqrt{d})z \mid z \in \mathbb{Z}^d\}$, where \mathbb{Z}^d is the d -dimensional lattice of integers. Let $\mathcal{G}[\varepsilon, p, \Lambda]$ represent the subset of the grid G_ε that is within a distance Λ of the point p . In other words, $\mathcal{G}[\varepsilon, p, \Lambda] = \{g \in G_\varepsilon \mid \|g - p\| \leq \Lambda\}$.

Algorithm.

These results imply the following algorithm. For each point $p \in \mathcal{P}$, for each $q \in \mathcal{Q}$, and for each $g \in \mathcal{G}[\varepsilon/2, p, \sqrt{\ln(n^2)}]$ we consider the translation $T_{p,q,g}$ such that $q + T_{p,q,g} = g$. We return the translation $T_{p,q,g}$ which maximizes $\kappa(\mathcal{P}, \mathcal{Q} \oplus T_{p,q,g})$, by evaluating κ at each such translation of \mathcal{Q} .

THEOREM 6.1. *The above algorithm runs in time $O((1/\varepsilon)^d \cdot n^4 \log^{d/2} n)$, for fixed d , and is guaranteed to find a translation \hat{T} such that $D_K^2(\mathcal{P}, \mathcal{Q} \oplus \hat{T}) - D_K^2(\mathcal{P}, \mathcal{Q} \oplus T^*) \leq \varepsilon W^2$.*

PROOF. We know that the optimal translation T^* must result in some pair of points $p \in \mathcal{P}$ and $q \in \mathcal{Q}$ such that $\|p - (q + T^*)\| \leq \sqrt{\ln(n^2)}$ by Lemma 6.1. So checking all pairs $p \in \mathcal{P}$ and $q \in \mathcal{Q}$, one must have $\|p - q\| \leq \sqrt{\ln(n^2)}$. Assuming we have found this closest pair, $p \in \mathcal{P}$ and $q \in \mathcal{Q}$, we only need to search in the neighborhood of translations $T = p - q$.

Furthermore, for some translation $T_{p,q,g} = g - q$ we can claim that $\kappa(\mathcal{P}, \mathcal{Q} \oplus T^*) - \kappa(\mathcal{P}, \mathcal{Q} \oplus T_{p,q,g}) \leq \varepsilon$. Since $\|T^* - T_{p,q,g}\| \leq \varepsilon/2$, we have the bound on subterm $|K(p, q + T^*) - K(p, q + T_{p,q,g})| \leq \varepsilon/2$, by Lemma 2.2. In fact, for every other pair $p' \in \mathcal{P}$ and $q' \in \mathcal{Q}$, we also know $|K(p', q' + T^*) - K(p', q' + T_{p,q,g})| \leq \varepsilon/2$. Thus the sum of these subterms has error at most $(\varepsilon/2) \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{Q}} \mu(p)\nu(q) = (\varepsilon/2)W^2$.

Since, the first two terms of $D_K^2(\mathcal{P}, \mathcal{Q} \oplus T)$ are unaffected by the choice of T , this provides an ε -approximation for $D_K^2(\mathcal{P}, \mathcal{Q} \oplus T)$ because all error is in the $(-2)\kappa(\mathcal{P}, \mathcal{Q} \oplus T)$ term.

For the runtime we need to bound the number of pairs from \mathcal{P} and \mathcal{Q} (i.e. $O(n^2)$), the time to calculate $\kappa(\mathcal{P}, \mathcal{Q} \oplus T)$ (i.e. $O(n^2)$), and finally the number of grid points in $\mathcal{G}[\varepsilon/2, p, \sqrt{\ln(n^2)}]$. The last term requires $O((1/\varepsilon)^d)$ points per unit volume, and a ball of radius $\sqrt{\ln(n^2)}$ has volume $O(\log^{d/2} n)$, resulting in $O((1/\varepsilon)^d \log^{d/2} n)$ points. This product produces a total runtime of $O((1/\varepsilon)^d n^4 \log^{d/2} n)$. \square

For a constant dimension d , using Theorem 5.2 to construct a coresset, we can first set $n = O((1/\varepsilon^2) \log(1/\delta))$ and now the time to calculate $\kappa(\mathcal{P}, \mathcal{Q} \oplus T)$ is $O((1/\varepsilon^4) \log^2(1/\delta))$ after spending $O(n + (1/\varepsilon^2) \log(1/\delta) \log n)$ time to construct the coressets. Hence the total runtime is $O(n + \log n(1/\varepsilon^2) \cdot (\log(1/\delta)) + (1/\varepsilon^{d+8}) \cdot \log^{d/2}((1/\varepsilon) \log(1/\delta)) \log^4(1/\delta))$, and is correct with probability at least $1 - \delta$.

THEOREM 6.2. *For fixed d , in $O(n + \log n(1/\varepsilon^2)(\log(1/\delta)) + (1/\varepsilon^{d+8}) \cdot \log^{d/2}((1/\varepsilon) \log(1/\delta)) \log^4(1/\delta))$ time we can find a translation \hat{T} such that $D_K^2(\mathcal{P}, \mathcal{Q} \oplus \hat{T}) - D_K^2(\mathcal{P}, \mathcal{Q} \oplus T^*) \leq \varepsilon W^2$, with probability at least $1 - \delta$.*

Translations and rotations.

Using a similar, but more involved, approach (detailed in the full version) we can find the approximately optimal translation and rotation (\hat{T}, \hat{R}) to minimize D_K in $O(n + \text{poly}(1/\varepsilon, \log n)^d)$.

THEOREM 6.3. *For fixed d , in $O(n + \log n(1/\varepsilon^2)(\log(1/\delta)) + (1/\varepsilon)^{(d^2+7d+6)/2}(\log(1/\varepsilon\delta))^{(d^2+7d+10)/4})$, time we can find a translation and rotation pair (\hat{T}, \hat{R}) , such that*

$$D_K^2(\mathcal{P}, \hat{R} \circ (\mathcal{Q} \oplus \hat{T})) - D_K^2(\mathcal{P}, R^* \circ \mathcal{Q} \oplus T^*) \leq \varepsilon W^2,$$

with probability at least $1 - \delta$.

7. REFERENCES

- [1] Noga Alon, Shai Ben-David, Nocolò Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of ACM*, 44:615–631, 1997.
- [2] Alexandr Andoni, Khanh Do Ba, Piotr Indyk, and David Woodruff. Efficient sketches for Earth-mover distance, with applications. In *50th Annual Symp. Foundations of Computer Science*, 2009.
- [3] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *47th Annual Symp. Foundations of Computer Science*, 2006.
- [4] Alexandr Andoni, Piotr Indyk, and Robert Krauthgamer. Earth mover distance over high-dimensional spaces. In *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms*, 2008.
- [5] N. Aronszajn. Theory of reproducing kernels. *Transactions AMS*, 68:337–404, 1950.
- [6] Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *Proc. 51st IEEE Symp. Foundations of Computer Science*, 2010.
- [7] Peter L. Bartlett, Philip M. Long, and Robert C. Williamson. Fat-shattering and the learnability of real-valued functions. *Journal of Computer and System Sciences*, 52(3):434–452, 1996.
- [8] Sergio Cabello, Panos Giannopoulos, Christian Knauer, and Günter Rote. Matching point sets with respect to the earth mover’s distance. *Computational Geometry*, 39:118–133, 2008.
- [9] Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest neighbors and n -body potential fields. *Journal of ACM*, 42:67–90, 1995.
- [10] Moses Charikar, Alantha Newman, and Aleksandar Nikolov. Tight hardness results for minimizing discrepancy. In *Proceedings 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, 2011.
- [11] Bernard Chazelle. *The Discrepancy Method*. Cambridge University Press, 2000.
- [12] Bernard Chazelle and Jiri Matousek. On linear-time deterministic algorithms for optimization problems in fixed dimensions. *Journal of Algorithms*, 21:579–597, 1996.
- [13] Bernard Chazelle and Emo Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete and Computational Geometry*, 4:467–489, 1989.
- [14] Kenneth L. Clarkson. Fast algorithms for the all nearest neighbors problem. In *Proceedings 34th Annual Symposium on Foundations of Computer Science*, 1983.
- [15] Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag, 1996.

- [16] Michel Marie Deza and Monique Laurent. *Geometry of Cuts and Metrics*. Springer, 1997.
- [17] Stanley Durrleman, Xavier Pennec, Alain Trounev, and Nicholas Ayache. Sparse approximation of currents for statistics on curves and surfaces. In *11th International Conference on Medical Image Computing and Computer Assisted Intervention*, 2008.
- [18] Joan Glaunès. *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique*. PhD thesis, Université Paris 13, 2005.
- [19] Joan Glaunès and Sarang Joshi. Template estimation form unlabeled point set data and surfaces for computational anatomy. In *Math. Found. Comp. Anatomy*, 2006.
- [20] Leslie Greengard and John Strain. The fast Gauss transform. *Journal Scientific and Statistical Computing*, 12, 1991.
- [21] Sarel Har-Peled. *Approximation Algorithms in Geometry*. <http://valis.cs.uiuc.edu/~sarel/teach/notes/aprx/>.
- [22] David Haussler. Sphere packing numbers for subsets of the boolean n -cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69:217–232, 1995.
- [23] Matriel Hein and Olivier Bousquet. Hilbertian metrics and positive definite kernels on probability measures. In *Proceedings 10th International Workshop on Artificial Intelligence and Statistics*, 2005.
- [24] William B. Johnson and Joram Lindenstrauss. Extensions of Lipschitz maps into a Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.
- [25] Michael Kerns and Robert E. Shapire. Efficient distribution-free learning of probabilistic concepts. *Journal of Computer and System Sciences*, 48:464–497, 1994.
- [26] Yi Li, Philip M. Long, and Aravind Srinivasan. Improved bounds on the samples complexity of learning. *Journal of Computer and System Science*, 62:516–527, 2001.
- [27] Maarten Löffler and Jeff M. Phillips. Shape fitting on point sets with probability distributions. In *Proceedings 17th Annual European Symposium on Algorithms*, 2009.
- [28] Jiri Matoušek. Approximations and optimal geometric divide-and-conquer. In *Proceedings 23rd Symposium on Theory of Computing*, pages 505–511, 1991.
- [29] Jiri Matoušek. Tight upper bounds for the discrepancy of halfspaces. *Discrete and Computational Geometry*, 13:593–601, 1995.
- [30] Jiri Matoušek. *Geometric Discrepancy*. Springer, 1999.
- [31] Jiri Matoušek, Emo Welzl, and Lorenz Wernisch. Discrepancy and approximations for bounded VC-dimension. *Combinatorica*, 13:455–466, 1993.
- [32] Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29:429–443, 1997.
- [33] Aleksandar Nikolov. Personal communication, 2011.
- [34] Jeff M. Phillips. Algorithms for ε -approximations of terrains. In *Proceedings 35th International Colloquium on Automata, Languages, and Programming*, 2008.
- [35] Jeff M. Phillips and Suresh Venkatasubramanian. A gentle introduction to the kernel distance. Technical report, arXiv:1103.1625, 2011.
- [36] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Neural Information Processing Systems*, 2007.
- [37] W. Rudin. *Fourier analysis on groups*. Wiley-Interscience, 1962.
- [38] Sameer Shirdhonkar and David W. Jacobs. Approximate Earth mover's distance in linear time. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [39] Alex J. Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A Hilbert space embedding for distributions. In *Proceedings of Algorithmic Learning Theory*, 2007.
- [40] Le Song, Byron Boots, Sajid Siddiqi, Geoffrey Gordon, and Alex J. Smola. Hilbert space embeddings of hidden markov models. In *Proceedings of International Conference on Machine Learning*, 2010.
- [41] Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.
- [42] Ch. Suquet. Distances Euclidiennes sur les mesures signées et application à des théorèmes de Berry-Esséen. *Bulletin Belgium Mathematics Society*, 2:161–181, 1995.
- [43] Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Range counting over multidimensional data streams. In *Proceedings 20th Symposium on Computational Geometry*, pages 160–169, 2004.
- [44] M. Talagrand. Sharper bounds for Gaussian and empirical processes. *Annals of Probability*, 22:76, 1994.
- [45] Marc Vaillant and Joan Glaunès. Surface matching via currents. In *Proceedings Information Processing in Medical Imaging*, volume 19, pages 381–92, 2005.
- [46] Vladimir Vapnik. Inductive principles of the search for empirical dependencies. In *Proceedings of the Second Annual Workshop on Computational Learning Theory*, pages 3–21, 1989.
- [47] Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- [48] C. Yang, R. Duraiswami, N.A. Gumerov, and L. Davis. Improved fast Gauss transform and efficient kernel density estimation. In *Proceedings 9th International Conference on Computer Vision*, pages 664–671, 2003.