Subsampling in Smoothed Range Spaces *

Jeff M. Phillips and Yan Zheng

University of Utah

Abstract. We consider smoothed versions of geometric range spaces, so an element of the ground set (e.g. a point) can be contained in a range with a non-binary value in $[0, 1]$. Similar notions have been considered for kernels; we extend them to more general types of ranges. We then consider approximations of these range spaces through ε-nets and ε-samples (aka ε-approximations). We characterize when size bounds for ε-samples on kernels can be extended to these more general smoothed range spaces. We also describe new generalizations for ε-nets to these range spaces and show when results from binary range spaces can carry over to these smoothed ones.

1 Introduction

This paper considers traditional sample complexity problems but adapted to when the range space (or function space) smoothes out its boundary. This is important in various scenarios where either the data points or the measuring function is noisy. Similar problems have been considered in specific contexts of functions classes with a $[0, 1]$ range or kernel density estimates. We extend and generalize various of these results, motivated by scenarios like the following.

(S1) Consider maintaining a random sample of noisy spatial data points (say twitter users with geo-coordinates), and we want this sample to include a witness to every large enough event. However, because the data coordinates are noisy we use a kernel density estimate to represent the density. And moreover, we do not want to consider regions with a single or constant number of data points which only occurred due to random variations. In this scenario, how many samples do we need to maintain?

(S2) Next consider a large approximate (say high-dimensional image feature [1]) dataset, where we want to build a linear classifier. Because the features are approximate (say due to feature hashing techniques), we model the classifier boundary to be randomly shifted using Gaussian noise. How many samples from this dataset do we need to obtain a desired generalization bound?

(S3) Finally, consider one of these scenarios in which we are trying to create an informative subset of the enormous full dataset, but have the opportunity to do so in ways more intelligent than randomly sampling. On such a reduced dataset one may want to train several types of classifiers, or to estimate the density of various subsets. Can we generate a smaller dataset compared to what would be required by random sampling?

* Thanks to supported by NSF CCF-1350888, IIS-1251019, and ACI-1443046.

The traditional way to study related sample complexity problems is through range spaces (a ground set X, and family of subsets $\mathcal{A}$) and their associated dimension (e.g., VC-dimension [25]). We focus on a smooth extension of range spaces defined on a geometric ground set. Specifically, consider the ground set P to be a subset of points in $\mathbb{R}^d$, and let $\mathcal{A}$ describe subsets defined by some geometric objects, for instance a halfspace or a ball. Points $p \in \mathbb{R}^d$ that are inside the object (e.g., halfspace or ball) are typically assigned a value 1, and those outside a value 0. In our smoothed setting points near the boundary are given a value between 0 and 1, instead of discretely switching from 0 to 1.

In learning theory these smooth range spaces can be characterized by more general notions called P-dimension [22] (or Pseudo dimension) or V-dimension [24] (or "fat" versions of these [2]) and can be used to learn real-valued functions for regression or density estimation, respectively.

In geometry and data structures, these smoothed range spaces are of interest in studying noisy data. Our work extends some recent work [12, 21] which examines a special case of our setting that maps to kernel density estimates, and matches or improves on related bounds for non-smoothed versions.

Main contributions. We next summarize the main contributions in this paper.

- We define a general class of *smoothed range spaces* (Sec 3.1), with application to density estimation and noisy agnostic learning, and we show that these can inherit sample complexity results from *linked* non-smooth range spaces (Corollary 1).
- We define an (ε, τ)-net for a smoothed range space (Sec 3.3). We show how this can inherit sampling complexity bounds from *linked* non-smooth range spaces (Theorem 2), and we relate this to non-agnostic density estimation and hitting set problems.
- We provide discrepancy-based bounds and constructions for ε-samples on smooth range spaces requiring significantly fewer points than uniform sampling approaches (Theorems 4 and 5), and also smaller than discrepancy-based bounds on the linked binary range spaces.

2 Definitions and Background

Recall that we will focus on geometric range spaces $(P, \mathcal{A})$ where the ground set $P \subset \mathbb{R}^d$ and the family of ranges $\mathcal{A}$ are defined by geometric objects. It is common to approximate a range space in one of two ways, as an ε-sample (aka ε-approximation) or an ε-net. An *ε-sample* for a range space $(P, \mathcal{A})$ is a subset $Q \subset P$ such that

$$\max_{A \in \mathcal{A}} \left| \frac{|A \cap P|}{|P|} - \frac{|Q \cap A|}{|Q|} \right| \leq \varepsilon.$$

An *ε-net* of a range space $(P, \mathcal{A})$ is a subset $Q \subset P$ such that

$$\text{for all } A \in \mathcal{A} \text{ such that } \frac{|P \cap A|}{|P|} \geq \varepsilon \text{ then } A \cap Q \neq \emptyset.$$

Given a range space $(P, \mathcal{A})$ where $|P| = m$, then $\pi_{\mathcal{A}}(m)$ describes the maximum number of possible distinct subsets of P defined by some $A \in \mathcal{A}$. If we can bound, $\pi_{\mathcal{A}}(m) \leq Cm^{\nu}$ for absolute constant C, then $(P, \mathcal{A})$ is said to have *shatter dimension* ν. For instance the shatter dimension of $\mathcal{H}$ halfspaces in $\mathbb{R}^d$ is d, and for $\mathcal{B}$ balls in $\mathbb{R}^d$ is $d+1$. For a range space with shatter dimension ν, a random sample of size $O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ is an ε-sample with probability at least $1-\delta$ [25, 14], and a random sample of size $O((\nu/\varepsilon)\log(1/\varepsilon\delta))$ is an ε-net with probability at least $1-\delta$ [11, 18].

An ε-sample Q is sufficient for agnostic learning with generalization error ε, where the best classifier might misclassify some points. An ε-net Q is sufficient for non-agnostic learning with generalization error ε, where the best classifier is assumed to have no error on P.

The size bounds can be made deterministic and slightly improved for certain cases. An ε-sample Q can be made of size $O(1/\varepsilon^{2\nu/(\nu+1)})$ [15] and this bound can be no smaller [16] in the general case. For balls $\mathcal{B}$ in $\mathbb{R}^d$ which have shatter-dimension $\nu = d+1$, this can be improved to $O(1/\varepsilon^{2d/(d+1)}\log^{d/(d+1)}(1/\varepsilon))$ [4, 16], and the best known lower bound is $O(1/\varepsilon^{2d/(d+1)})$. For axis-aligned rectangles $\mathcal{R}$ in $\mathbb{R}^d$ which have shatter-dimension $\nu = 2d$, this can be improved to $O((1/\varepsilon)\log^{d+1/2}(1/\varepsilon))$ [13].

For ε-nets, the general bound of $O((\nu/\varepsilon)\log(1/\varepsilon))$ can also be made deterministic [15], and for halfspaces in $\mathbb{R}^4$ the size must be at least $\Omega((1/\varepsilon)\log(1/\varepsilon))$ [19]. But for halfspaces in $\mathbb{R}^3$ the size can be $O(1/\varepsilon)$ [17, 10], which is tight. By a simple lifting, this also applies for balls in $\mathbb{R}^2$. For other range spaces, such as axis-aligned rectangles in $\mathbb{R}^2$, the size bound is $\Theta((1/\varepsilon)\log\log(1/\varepsilon))$ [3, 19].

2.1 Kernels

A *kernel* is a bivariate similarity function $K : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^+$, which can be normalized so $K(x,x) = 1$ (which we assume through this paper). Examples include ball kernels ($K(x,p) = \{1$ if $\|x-p\| \leq 1$ and 0 otherwise$\}$), triangle kernels ($K(x,p) = \max\{0, 1-\|x-p\|\}$), Epanechnikov kernels ($K(x,p) = \max\{0, 1-\|x-p\|^2\}$), and Gaussian kernels ($K(x,p) = \exp(-\|x-p\|^2)$, which is reproducing). In this paper we focus on symmetric, shift invariant kernels which depend only on $z = \|x-p\|$, and can be written as a single parameter function $K(x,p) = k(z)$; these can be parameterized by a single bandwidth (or just width) parameter w so $K(x,p) = k_w(\|x-p\|/w)$.

Given a point set $P \subset \mathbb{R}^d$ and a kernel, a *kernel density estimate* KDE_P is the convolution of that point set with K. For any $x \in \mathbb{R}^d$ we define $\text{KDE}_P(x) = \frac{1}{|P|}\sum_{p \in P} K(x,p)$.

A *kernel range space* [12, 21] $(P, \mathcal{K})$ is an extension of the combinatorial concept of a range space $(P, \mathcal{A})$ (or to distinguish it we refer to the classic notion as a *binary range space*). It is defined by a point set $P \subset \mathbb{R}^d$ and a kernel K. An element K_x of $\mathcal{K}$ is a kernel $K(x, \cdot)$ applied at point $x \in \mathbb{R}^d$; it assigns a value in $[0,1]$ to each point $p \in P$ as $K(x,p)$. If we use a ball kernel, then each value is exactly $\{0,1\}$ and we recover exactly the notion of a binary range space for geometric ranges defined by balls.

The notion of an ε-kernel sample [12] extends the definition of ε-sample. It is a subset $Q \subset P$ such that

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

A binary range space $(P, \mathcal{A})$ is *linked* to a kernel range space $(P, \mathcal{K})$ if the set $\{p \in P \mid K(x,p) \geq \tau\}$ is equal to $P \cap A$ for some $A \in \mathcal{A}$, for any threshold value τ. [12] showed that an ε-sample of a linked range space $(P, \mathcal{A})$ is also an ε-kernel sample of a corresponding kernel range space $(P, \mathcal{K})$. Since all range spaces defined by symmetric, shift-invariant kernels are linked to range spaces defined by balls, they inherit all ε-sample bounds, including that random samples of size $O((1/\varepsilon^2)(d+\log(1/\delta))$ provide an ε-kernel sample with probability at least $1-\delta$. Then [21] showed that these bounds can be improved through discrepancy-based methods to $O(((1/\varepsilon)\sqrt{\log(1/\varepsilon\delta)})^{2d/(d+2)})$, which is $O((1/\varepsilon)\sqrt{\log(1/\varepsilon\delta)})$ in $\mathbb{R}^2$.

A more general concept has been studied in learning theory on real-valued functions, where a function f as a member of a function class $\mathcal{F}$ describes a mapping from $\mathbb{R}^d$ to $[0,1]$ (or more generally $\mathbb{R}$). A kernel range space where the linked binary range space has bounded shatter-dimension ν is said to have bounded V-dimension [24] (see [2]) of ν. Given a ground set X, then for $(X, \mathcal{F})$ this describes the largest subset Y of X which can be shattered in the following sense. Choose any value $s \in [0,1]$ for all points $y \in Y$, and then for each subset of $Z \subset Y$ there exists a function $f \in \mathcal{F}$ so $f(y) > s$ if $y \in Z$ and $f(y) < s$ if $y \notin Z$. The best sample complexity bounds for ensuring Q is an ε-sample of P based on V-dimension are derived from a more general sort of dimension (called a P-dimension [22] where in the shattering definition, each y may have a distinct $s(y)$ value) requires $|Q| = O((1/\varepsilon^2)(\nu + \log(1/\delta)))$ [14]. As we will see, these V-dimension based results are also general enough to apply to the to-be-defined smooth range spaces.

3 New Definitions

In this paper we extend the notion of a kernel range spaces to other *smoothed range spaces* that are "linked" with common range spaces, e.g., halfspaces. These inherent the construction bounds through the linking result of [12], and we show cases where these bounds can also be improved. We also extend the notion of ε-nets to kernels and smoothed range spaces, and showing linking results for these as well.

3.1 Smoothed Range Spaces

Here we will define the primary smoothed combinatorial object we will examine, starting with halfspaces, and then generalizing. Let $\mathcal{H}_w$ denote the family of *smoothed halfspaces with width parameter* w, and let $(P, \mathcal{H}_w)$ be the associated smoothed range space where $P \subset \mathbb{R}^d$. Given a point $p \in P$, then smoothed halfspace $h \in \mathcal{H}_w$ maps p to a value $v_h(p) \in [0,1]$ (rather than the traditional $\{0,1\}$ in a binary range space).

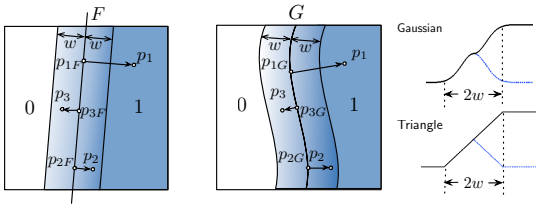

	$\|p - p_F\|$	$?(p \in h)$	$v_h(p)$
p_1	$3w/2$	TRUE	1
p_2	$3w/4$	TRUE	$7/8$
p_3	$w/2$	FALSE	$1/4$

Fig. 1. Illustration of the smoothed halfspace, and smoothed polynomial surface, with function value of three points $\{p_1, p_2, p_3\}$ defined using a triangle kernel.

We first describe a specific mapping to the function value $v_h(p)$ that will be sufficient for the development of most of our techniques. Let F be the $(d-1)$-flat defining the boundary of halfspace h. Given a point $p \in \mathbb{R}^d$, let $p_F = \arg\min_{q \in F} \|p - q\|$ describe the point on F closest to p. Now we define

$$v_{h,w}(p) = \begin{cases} 1 & p \in h \text{ and } \|p - p_F\| \geq w \\ \frac{1}{2} + \frac{1}{2}\frac{\|p-p_F\|}{w} & p \in h \text{ and } \|p - p_F\| < w \\ \frac{1}{2} - \frac{1}{2}\frac{\|p-p_F\|}{w} & p \notin h \text{ and } \|p - p_F\| < w \\ 0 & p \notin h \text{ and } \|p - p_F\| \geq w. \end{cases}$$

These points within a slab of width $2w$ surrounding F can take on a value between 0 and 1, where points outside of this slab revert back to the binary values of either 0 or 1.

We can make this more general using a shift-invariant kernel $k(\|p - x\|) = K(p, x)$, where $k_w(\|p - x\|) = k(\|p - x\|/w)$ allows us to parameterize by w. Define $v_{h,w}(p)$ as follows.

$$v_{h,w}(p) = \begin{cases} 1 - \frac{1}{2}k_w(\|p - p_F\|) & p \in h \\ \frac{1}{2}k_w(\|p - p_F\|) & p \notin h. \end{cases}$$

For brevity, we will omit the w and just use $v_h(p)$ when clear. These definitions are equivalent when using the triangle kernel. But for instance we could also use a Epanechnikov kernel or Gaussian kernel. Although the Gaussian kernel does not satisfy the restriction that only points in the width $2w$ slab take non $\{0, 1\}$ values, we can use techniques from [21] to extend to this case as well. This is illustrated in Figure 1. Another property held by this definition which we will exploit is that the slope ς of these kernels is bounded by $\varsigma = O(1/w) = c/w$, for some constant c; the constant $c = 1/2$ for triangle and Gaussian, and $c = 1$ for Epanechnikov.

Finally, we can further generalize this by replacing the flat F at the boundary of h with a polynomial surface G. The point $p_G = \arg\min_{q \in G} \|p - q\|$ replaces p_F in the above definitions. Then the slab of width $2w$ is replaced with a curved volume in $\mathbb{R}^d$; see Figure 1. For instance, if G defines a circle in $\mathbb{R}^d$, then v_h

defines a disc of value 1, then an annulus of width $2w$ where the function value decreases to 0. Alternatively, if G is a single point, then we essentially recover the kernel range space, except that the maximum height is $1/2$ instead of 1. We will prove the key structural results for polynomial curves in Section 5, but otherwise focus on halfspaces to keep the discussion cleaner. The most challenging elements of our results are all contained in the case with F as a $(d-1)$-flat.

3.2 ε-Sample in a Smoothed Range Space

It will be convenient to extend the notion of a kernel density estimate to these smoothed range space. A *smoothed density estimate* SDE_P is defined for any $h \in \mathcal{H}_w$ as

$$\text{SDE}_P(h) = \frac{1}{|P|} \sum_{p \in P} v_h(p).$$

An *ε-sample Q of a smoothed range space* $(P, \mathcal{H}_w)$ is a subset $Q \subset P$ such that

$$\max_{h \in \mathcal{H}_w} |\text{SDE}_P(h) - \text{SDE}_Q(h)| \leq \varepsilon.$$

Given such an ε-sample Q, we can then consider a subset $\bar{\mathcal{H}}_w$ of $\mathcal{H}_w$ with bounded integral (perhaps restricted to some domain like a unit cube that contains all of the data P). If we can learn the smooth range $\hat{h} = \arg\max_{h \in \bar{\mathcal{H}}_w} \text{SDE}_Q(h)$, then we know $\text{SDE}_P(h^*) - \text{SDE}_Q(\hat{h}) \leq \varepsilon$, where $h^* = \arg\max_{h \in \bar{\mathcal{H}}_w} \text{SDE}_P(h)$, since $\text{SDE}_Q(\hat{h}) \geq \text{SDE}_Q(h^*) \geq \text{SDE}_P(h^*) - \varepsilon$. Thus, such a set Q allows us to learn these more general density estimates with generalization error ε.

We can also learn smoothed classifiers, like scenario (S2) in the introduction, with generalization error ε, by giving points in the negative class a weight of -1; this requires separate $(\varepsilon/2)$-samples for the negative and positive classes.

3.3 (ε, τ)-Net in a Smoothed Range Space

We now generalize the definition of an ε-net. Recall that it is a subset $Q \subset P$ such that Q "hits" all large enough ranges ($|P \cap A|/|P| \geq \varepsilon$). However, the notion of "hitting" is now less well-defined since a point $q \in Q$ may be in a range but with value very close to 0; if a smoothed range space is defined with a Gaussian or other kernel with infinite support, any point q will have a non-zero value for *all* ranges! Hence, we need to introduce another parameter $\tau \in (0, \varepsilon)$, to make the notion of hitting more interesting in this case.

A subset $Q \subset P$ is an *(ε, τ)-net of smoothed range space* $(P, \mathcal{H}_w)$ if for any smoothed range $h \in \mathcal{H}_w$ such that $\text{SDE}_P(h) \geq \varepsilon$, then there exists a point $q \in Q$ such that $v_h(q) \geq \tau$.

The notion of ε-net is closely related to that of hitting sets. A *hitting set* of a binary range space $(P, \mathcal{A})$ is a subset $Q \subset P$ so every $A \in \mathcal{A}$ (not just the large enough ones) contains some $q \in Q$. To extend these notions to the smoothed setting, we again need an extra parameter $\tau \in (0, \varepsilon)$, and also need to only consider large enough smoothed ranges, since there are now an infinite number

even if P is finite. A subset $Q \subset P$ is an (ε, τ)-*hitting set of smoothed range space* $(P, \mathcal{H}_w)$ if for any $h \in \mathcal{H}_w$ such that $\text{SDE}_P(h) \geq \varepsilon$, then $\text{SDE}_Q(h) \geq \tau$.

In the binary range space setting, an ε-net Q of a range space $(P, \mathcal{A})$ is sufficient to learn the best classifier on P with generalization error ε in the non-agnostic learning setting, that is assuming a perfect classifier exists on P from $\mathcal{A}$. In the density estimation setting, there is not a notion of a perfect classifier, but if we assume some other properties of the data, the (ε, τ)-net will be sufficient to recover them. For instance, consider (like scenario (S1) in the introduction) that P is a discrete distribution so for some "event" points $p \in P$, there is at least an ε-fraction of the probability distribution describing P at p (e.g., there are more than $\varepsilon|P|$ points very close to p). In this setting, we can recover the location of these points since they will have probability at least τ in the (ε, τ)-net Q.

4 Linking and Properties of (ε, τ)-Nets

First we establish some basic connections between ε-sample, (ε, τ)-net, and (ε, τ)-hitting set in smoothed range spaces. In binary range spaces an ε-sample Q is also an ε-net, and a hitting set is also an ε-net; we show a similar result here up to the covering constant τ.

Lemma 1. *For a smoothed range space $(P, \mathcal{H}_w)$ and $0 < \tau < \varepsilon < 1$, an (ε, τ)-hitting set Q is also an (ε, τ)-net of $(P, \mathcal{H}_w)$.*

Proof. The (ε, τ)-hitting set property establishes for all $h \in \mathcal{H}_w$ with $\text{SDE}_P(h) \geq \varepsilon$, then also $\text{SDE}_Q(h) \geq \tau$. Since $\text{SDE}_Q(h) = \frac{1}{|Q|} \sum_{q \in Q} v_h(q)$ is the average over all points $q \in Q$, then it implies that at least one point also satisfies $v_h(q) \geq \tau$. Thus Q is also an (ε, τ)-net. □

In the other direction an (ε, τ)-net is not necessarily an (ε, τ)-hitting set since the (ε, τ)-net Q may satisfy a smoothed range $h \in \mathcal{H}_w$ with a single point $q \in Q$ such that $v_h(q) \geq \tau$, but all others $q' \in Q \setminus \{q\}$ having $v_h(q') \ll \tau$, and thus $\text{SDE}_Q(h) < \tau$.

Theorem 1. *For $0 < \tau < \varepsilon < 1$, an $(\varepsilon - \tau)$-sample Q in smoothed range space $(P, \mathcal{H}_w)$ is an (ε, τ)-hitting set in $(P, \mathcal{H}_w)$, and thus also an (ε, τ)-net of $(P, \mathcal{H}_w)$.*

Proof. Since Q is the $(\varepsilon - \tau)$-sample in the smoothed range space, for any smoothed range $h \in \mathcal{H}_w$ we have $|\text{SDE}_P(h) - \text{SDE}_Q(h)| \leq \varepsilon - \tau$. We consider the upper and lower bound separately.

If $\text{SDE}_P(h) \geq \varepsilon$, when $\text{SDE}_P(h) \geq \text{SDE}_Q(h)$, we have

$$\text{SDE}_Q(h) \geq \text{SDE}_P(h) - (\varepsilon - \tau) \geq \varepsilon - (\varepsilon - \tau) = \tau.$$

And more simply when $\text{SDE}_Q(h) \geq \text{SDE}_P(h)$ and $\text{SDE}_P(h) \geq \varepsilon \geq \tau$, then $\text{SDE}_Q(h) \geq \tau$. Thus in both situations, Q is an (ε, τ)-hitting set of $(P, \mathcal{H}_w)$. And then by Lemma 1 Q is also an (ε, τ)-net of $(P, \mathcal{H}_w)$. □

4.1 Relations between Smoothed Range Spaces and Linked Binary Range Spaces

Consider a smoothed range space $(P, \mathcal{H}_w)$, and for one smoothed range $h \in \mathcal{H}_w$, examine the range boundary F (e.g. a $(d-1)$-flat, or polynomial surface) along with a symmetric, shift invariant kernel K that describes v_h. The *superlevel* set $(v_h)^\tau$ is all points $x \in \mathbb{R}^d$ such that $v_h(x) \geq \tau$. Then recall a smoothed range space $(P, \mathcal{H}_w)$ is *linked* to a binary range space $(P, \mathcal{A})$ if every set $\{p \in P \mid v_h(p) \geq \tau\}$ for any $h \in \mathcal{H}_w$ and any $\tau > 0$, is exactly the same as some range $A \cap P$ for $A \in \mathcal{A}$. For smoothed range spaces defined by halfspaces, then the linked binary range space is also defined by halfspaces. For smoothed range spaces defined by points, mapping to kernel range spaces, then the linked binary range spaces are defined by balls.

Joshi *et al.* [12] established that given a kernel range space $(P, \mathcal{K})$, a linked binary range space $(P, \mathcal{A})$, and an ε-sample Q of $(P, \mathcal{A})$, then Q is also an ε-kernel sample of $(P, \mathcal{K})$. An inspection of the proof reveals the same property holds directly for smoothed range spaces, as the only structural property needed is that all points $p \in P$, as well as all points $q \in Q$, can be sorted in decreasing function value $K(p, x)$, where x is the center of the kernel. For smoothed range space, this can be replaced with sorting by $v_h(p)$.

Corollary 1 ([12]). *Consider a smoothed range space $(P, \mathcal{H}_w)$, a linked binary range space $(P, \mathcal{A})$, and an ε-sample Q of $(P, \mathcal{A})$ with $\varepsilon \in (0, 1)$. Then Q is an ε-sample of $(P, \mathcal{H}_w)$.*

We now establish a similar relationship to (ε, τ)-nets of smoothed range spaces from $(\varepsilon - \tau)$-nets of linked binary range spaces.

Theorem 2. *Consider a smoothed range space $(P, \mathcal{H}_w)$, a linked binary range space $(P, \mathcal{A})$, and an $(\varepsilon - \tau)$-net Q of $(P, \mathcal{A})$ for $0 < \tau < \varepsilon < 1$. Then Q is an (ε, τ)-net of $(P, \mathcal{H}_w)$.*

Proof. Let $|P| = n$. Then since Q is an $(\varepsilon - \tau)$-net of $(P, \mathcal{A})$, for any range $A \in \mathcal{A}$, if $|P \cap A| \geq (\varepsilon - \tau)n$, then $Q \cap A \neq \emptyset$.

Suppose $h \in \mathcal{H}_w$ has $\text{SDE}_P(h) \geq \varepsilon$ and we want to establish that $\text{SDE}_Q(h) \geq \tau$. Let $A \in \mathcal{A}$ be the range such that $(\varepsilon - \tau)n$ points with largest $v_h(p_i)$ values are exactly the points in A. We now partition P into three parts (1) let P_1 be the $(\varepsilon - \tau)n - 1$ points with largest v_h values, (2) let y be the point in P with $(\varepsilon - \tau)n$th largest v_h value, and (3) let P_2 be the remaining $n - n(\varepsilon - \tau)$ points. Thus for every $p_1 \in P_1$ and every $p_2 \in P_2$ we have $v_h(p_2) \leq v_h(y) \leq v_h(p_1) \leq 1$.

Now using our assumption $n \cdot \text{SDE}_P(h) \geq n\varepsilon$ we can decompose the sum

$$n \cdot \text{SDE}_P(h) = \sum_{p_1 \in P_1} v_h(p_1) + v_h(y) + \sum_{p_2 \in P_2} v_h(p_2) \geq n\varepsilon,$$

and hence using upper bounds $v_h(p_1) \leq 1$ and $v_h(p_2) \leq v_h(y)$,

$$\begin{aligned} v_h(y) &\geq n\varepsilon - \sum_{p_1 \in P_1} v_h(p_1) - \sum_{p_2 \in P_2} v_h(p_2) \\ &\geq n\varepsilon - (n(\varepsilon - \tau) - 1) \cdot 1 - (n - n(\varepsilon - \tau))v_h(y). \end{aligned}$$

Solving for $v_h(y)$ we obtain

$$v_h(y) \geq \frac{n\tau + 1}{n - n(\varepsilon - \tau) + 1} \geq \frac{n\tau}{n - n(\varepsilon - \tau)} \geq \frac{n\tau}{n} = \tau.$$

Since $(P, \mathcal{A})$ is linked to $(P, \mathcal{H}_w)$, there exists a range $A \in \mathcal{A}$ that includes precisely $P_1 \cup y$ (or more points with the same $v_h(y)$ value as y). Because Q is an $(\varepsilon - \tau)$-net of $(P, \mathcal{A})$, Q contains at least one of these points, lets call it q. Since all of these points have function value $v_h(p) \geq v_h(y) \geq \tau$, then $v_h(q) \geq \tau$. Hence Q is also an (ε, τ)-net of $(P, \mathcal{H}_w)$, as desired. □

This implies that if $\tau \leq c\varepsilon$ for any constant $c < 1$, then creating an (ε, τ)-net of a smoothed range space, with a known linked binary range space, reduces to computing an ε-net for the linked binary range space. For instance any linked binary range space with shatter-dimension ν has an ε-net of size $O(\frac{\nu}{\varepsilon} \log \frac{1}{\varepsilon})$, including halfspaces in $\mathbb{R}^d$ with $\nu = d$ and balls in $\mathbb{R}^d$ with $\nu = d + 1$; hence there exists $(\varepsilon, \varepsilon/2)$-nets of the same size. For halfspaces in $\mathbb{R}^2$ or $\mathbb{R}^3$ (linked to smoothed halfspaces) and balls in $\mathbb{R}^2$ (linked to kernels), the size can be reduced to $O(1/\varepsilon)$ [17, 10, 23].

5 Min-Cost Matchings within Cubes

Before we proceed with our construction for smaller ε-samples for smoothed range spaces, we need to prepare some structural results about min-cost matchings. Following some basic ideas from [21], these matchings will be used for discrepancy bounds on smoothed range spaces in Section 6.

In particular, we analyze some properties of the interaction of a min-cost matching M and some basic shapes ([21] considered only balls). Let $P \subset \mathbb{R}^d$ be a set of $2n$ points. A *matching* $M(P)$ is a decomposition of P into n pairs $\{p_i, q_i\}$ where $p_i, q_i \in P$ and each p_i (and q_i) is in exactly one pair. A *min-cost matching* is the matching M that minimizes $\mathsf{cost}_1(M, P) = \sum_{i=1}^{n} \|p_i - q_i\|$. The min-cost matching can be computed in $O(n^3)$ time by [9] (using an extension of the Hungarian algorithm from the bipartite case). In $\mathbb{R}^2$ it can be calculated in $O(n^{3/2} \log^5 n)$ time [26].

Following [21], again we will base our analysis on a result of [5] which says that if $P \subset [0, 1]^d$ (a unit cube) then for d a constant, $\mathsf{cost}_d(M, P) = \sum_{i=1}^{n} \|p_i - q_i\|^d = O(1)$, where M is the min-cost matching. We make no attempt to optimize constants, and assume d is constant.

One simple consequence, is that if P is contained in a d-dimensional cube of side length ℓ, then $\mathsf{cost}_d(M, P) = \sum_{i=1}^{n} \|p_i - q_i\|^d = O(\ell^d)$.

We are now interested in interactions with a matching M for P in a d-dimensional cube of side length ℓ $\mathsf{C}_{\ell,d}$ (call this shape an (ℓ, d)-*cube*), and more general objects; in particular C_w a (w, d)-cube and, S_w a slab of width $2w$, both restricted to be within $\mathsf{C}_{\ell,d}$. Now for such an object O_w (which will either be C_w or S_w) and an edge $\{p, q\}$ where line segment $\overline{pq}$ intersects O_w define point p_B (resp. q_B) as the point on segment $\overline{pq}$ inside O_w closest to p (resp. q). Note if p

(resp. q) is inside O then $p_B = p$ (resp. $q_B = q$), otherwise it is on the boundary of O_w. For instance, see C_{20w} in Fig 2.

Define the *length* of a matching M restricted to an object $O_w \subset \mathbb{R}^d$ as

$$\rho(O_w, M) = \sum_{(q,p)\in M} \min\left\{(2w)^d, \|p_B - q_B\|^d\right\}.$$

Note this differs from a similar definition by [21] since that case did not need to consider when both p and q were both outside of O_w, and did not need the $\min\{(2w)^d, \ldots\}$ term because all objects had diameter 2.

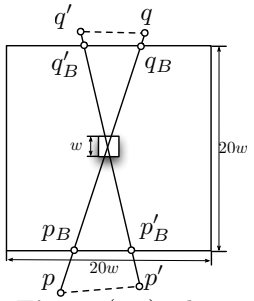

Fig. 2. (T3) edges

Lemma 2. *Let $P \subset \mathsf{C}_{\ell,d}$, where d is constant, and M be its min-cost matching. For any (w,d)-cube $C_w \subset \mathsf{C}_{\ell,d}$ we have $\rho(C_w, M) = O(w^d)$.*

Proof. We cannot simply apply the result of [5] since we do not restrict that $P \subset C_w$. We need to consider cases where either p or q or both are outside of C_w. As such, we have three types of edges we consider, based on a cube C_{20w} of side length $20w$ and with center the same as C_w.

(T1) Both endpoints are within C_{20w} of edge length at most $\sqrt{d}20w$.
(T2) One endpoint is in C_w, the other is outside C_{20w}.
(T3) Both endpoints are outside C_{20w}.

For all (T1) edges, the result of Bern and Eppstein can directly bound their contribution to $\rho(C_w, M)$ as $O(w^d)$ (scale to a unit cube, and rescale). For all (T2) edges, we can also bound their contribution to $\rho(C_w, M)$ as $O(w^d)$, by extending an analysis of [21] when both C_w and C_{20w} are similarly proportioned balls. This analysis shows there are $O(1)$ such edges.

We now consider the case of (T3) edges, restricting to those that also intersect C_w. We argue there can be at most $O(1)$ of them. In particular consider two such edges $\{p, q\}$ and $\{p', q'\}$, and their mappings to the boundary of C_{20w} as p_B, q_B, p'_B, q'_B; see Figure 2. If $\|p_B - p'_B\| \leq 10w$ and $\|q_B - q'_B\| \leq 10w$, then we argue next that this cannot be part of a min-cost matching since $\|p - p'\| + \|q - q'\| < \|p - q\| + \|p' - q'\|$, and it would be better to swap the pairing. Then it follows from the straight-forward net argument below that there can be at most $O(1)$ such pairs.

We first observe that $\|p_B - p'_B\| + \|q_B - q'_B\| \leq 10w + 10w < 20w + 20w \leq \|p_B - q_B\| + \|p'_B - q'_B\|$. Now we can obtain our desired inequality using that $\|p - q\| = \|p - p_B\| + \|p_B - q_B\| + \|q_B - q\|$ (and similar for $\|p' - q'\|$) and that $\|p - p'\| \leq \|p - p_B\| + \|p_B - p'_B\| + \|p'_B - p'\|$ by triangle inequality (and similar for $\|q - q'\|$).

Next we describe the net argument that there can be at most $O(d^2 \cdot 2^{2d}) = O(1)$ such pairs with $\|p_B - p'_B\| > 10w$ and $\|q_B - q'_B\| > 10w$. First place a $5w$-net $\mathcal{N}_f$ on each $(d-1)$-dimensional face f of C_{20w} so that any point $x \in f$ is within $5w$ of some point $\eta \in \mathcal{N}_f$. We can construct $\mathcal{N}_f$ of size $O(2^d)$ with a simple grid.

Then let $\mathcal{N} = \bigcup_f \mathcal{N}_f$ as the union of the nets on each face; its size is $O(d \cdot 2^d)$. Now for any point $p \notin C_{20w}$ let $\eta(p) = \arg\min_{\eta \in \mathcal{N}} \|p_B - \eta\|$ be the closest point in $\mathcal{N}$ to p_B. If two points p and p' have $\eta(p) = \eta(p')$ then $\|p - p'\| \leq 10w$. Hence there can be at most $O((d \cdot 2^d)^2)$ edges with $\{p, q\}$ mapping to unique $\eta(p)$ and $\eta(q)$ if no other edge $\{p', q'\}$ has $\|p_B - p'_B\| \leq 10w$ and $\|q_B - q'_B\| \leq 10w$.

Concluding, there can be at most $O(d^2 \cdot 2^{2d}) = O(1)$ edges in M of type (T3), and the sum of their contribution to $\rho(C_w, M)$ is at most $O(w^d)$, completing the proof. □

Lemma 3. *Let $P \subset \mathsf{C}_{\ell,d}$, where d is constant, and let M be its min-cost matching. For any width $2w$ slab S_w restricted to $\mathsf{C}_{\ell,d}$ we have $\rho(S_w, M) = O(\ell^{d-1} w)$.*

Proof. We can cover the slab S_w with $O((\ell/w)^{d-1})$ (w, d)-cubes. To make this concrete, we cover $\mathsf{C}_{\ell,d}$ with $\lceil \ell/w \rceil^d$ cubes on a regular grid. Then in at least one basis direction (the one closest to orthogonal to the normal of F) any column of cubes can intersect S_w in at most 4 cubes. Since there are $\lceil \ell/w \rceil^{d-1}$ such columns, the bound holds. Let $\mathcal{C}_w$ be the set of these cubes covering S_w.

Restricted to any one such cube C_w, the contribution of those edges to $\rho(S_w, M)$ is at most $O(w^d)$ by Lemma 2. Now we need to argue that we can just sum the effect of all covering cubes. The concern is that an edge goes through many cubes, only contributing a small amount to each $\rho(C_w, M)$ term, but when the total length is taken to the dth power it is much more. However, since each edge's contribution is capped at $(2w)^2$, we can say that if any edge goes through more than $O(1)$ cubes, its length must be at least w, and its contribution in one such cube is already $\Omega(w)$, so we can simply inflate the effect of each cube towards $\rho(S_w, M)$ by a constant.

In particular, consider any edge $\overline{pq}$ that has $p \in C_w$. Each cube has $3^d - 1$ neighboring cubes, including through vertex incidence. Thus if edge $\overline{pq}$ passes through more than 3^d cubes, q must be in a cube that is not one of C'_w's neighbors. Thus it must have length at least w; and hence its length in at least one cube C'_w must be at least $w/3^d$, with its contribution to $\rho(C'_w, M) > w^d/(3^{d^2})$. Thus we can multiply the effect of each edge in $\rho(C_w, M)$ by $3^{d^2} 2^d = O(1)$ and be sure it is at least as large as the effect of that edge in $\rho(S_w, M)$. Hence

$$\begin{aligned}\rho(S_w, M) &\leq 3^{d^2} 2^d \sum_{C_w \in \mathcal{C}_w} \rho(C_w, M) \leq O(1) \sum_{C_w \in \mathcal{C}_w} O(w^d) \\ &= O((\ell/w)^{d-1}) \cdot O(w^d) = O(\ell^{d-1} w).\end{aligned}$$ □

We can apply the same decomposition as used to prove Lemma 3 to also prove a result for a w-expanded volume G_w around a degree g polynomial surface G. A degree g polynomial surface can intersect a line at most g times, so for some $\mathsf{C}_{\ell,d}$ the expanded surface $G_w \cap \mathsf{C}_{\ell,d}$ can be intersected by $O(g(\ell/w)^{d-1})$ (w, d)-cubes. Hence we can achieve the following bound.

Corollary 2. *Let $P \subset \mathsf{C}_{\ell,d}$, where d is constant, and let M be its min-cost matching. For any volume G_w defined by a polynomial surface of degree g expanded by a width w, restricted to $\mathsf{C}_{\ell,d}$ we have $\rho(G_w, M) = O(g\ell^{d-1} w)$.*

6 Constructing ε-Samples for Smoothed Range Spaces

In this section we build on the ideas from [21] and the new min-cost matching results in Section 5 to produce new discrepancy-based ε-sample bounds for smoothed range spaces. The basic construction is as follows. We create a min-cost matching M on P, then for each pair $(p, q) \in M$, we retain one of the two points at random, halving the point set. We repeat this until we reach our desired size. This should not be unfamiliar to readers familiar with discrepancy-based techniques for creating ε-samples of binary range spaces [16, 6]. In that literature similar methods exist for creating matchings "with low-crossing number". Each such matching formulation is specific to the particular combinatorial range space one is concerned with. However, in the case of smoothed range spaces, we show that the min-cost matching approach is a universal algorithm. It means that an ε-sample Q for one smoothed range space $(P, \mathcal{H}_w)$ is also an ε-sample for any other smoothed range space $(P, \mathcal{H}'_w)$, perhaps up to some constant factors. We also show how these bounds can sometimes improve upon ε-sample bounds derived from linked range spaces; herein the parameter w will play a critical role.

6.1 Discrepancy for Smoothed Halfspaces

To simplify arguments, we first consider $P \subset \mathbb{R}^2$ extending to $\mathbb{R}^d$ in Section 6.4.

Let $\chi : P \to \{-1, +1\}$ be a coloring of P, and define the *discrepancy* of $(P, \mathcal{H}_w)$ with coloring χ as $\mathsf{disc}_\chi(P, \mathcal{H}_w) = \max_{h \in \mathcal{H}_w} |\sum_{p \in P} \chi(p) v_h(p)|$. Restricted to one smoothed range $h \in \mathcal{H}_w$ this is $\mathsf{disc}_\chi(P, h) = |\sum_{p \in P} \chi(p) v_h(p)|$. We construct a coloring χ using the min-cost matching M of P; for each $\{p_i, q_i\} \in M$ we randomly select one of p_i or q_i to have $\chi(p_i) = +1$, and the other $\chi(q_i) = -1$. We next establish bounds on the discrepancy of this coloring for a ς-bounded smoothed range space $(P, \mathcal{H}_w)$, i.e., where the gradient of v_h is bounded by $\varsigma \leq c_1/w$ for a constant c_1 (see Section 3.1).

For any smoothed range $h \in \mathcal{H}_w$, we can now define a random variable $X_j = \chi(p_j) v_h(p_j) + \chi(q_j) v_h(q_j)$ for each pair $\{p_j, q_j\}$ in the matching M. This allows us to rewrite $\mathsf{disc}_\chi(P, h) = |\sum_j X_j|$. We can also define a variable $\Delta_j = 2|v_h(p_j) - v_h(q_j)|$ such that $X_j \in \{-\Delta_j/2, \Delta_j/2\}$. Now following the key insight from [21] we can bound $\sum_j \Delta_j^2$ using results from Section 5, which shows up in the following Chernoff bound from [8]: Let $\{X_1, X_2, \ldots\}$ be independent random variables with $\mathbf{E}[X_j] = 0$ and $X_j = \{-\Delta_j/2, \Delta_j/2\}$ then

$$\Pr\left[\mathsf{disc}_\chi(P, h) \geq \alpha\right] = \Pr\left[\left|\sum_j X_j\right| \geq \alpha\right] \leq 2\exp\left(\frac{-2\alpha^2}{\sum_j \Delta_j^2}\right). \tag{1}$$

Lemma 4. *Assume $P \subset \mathbb{R}^2$ is contained in some cube $\mathsf{C}_{\ell,2}$ and with min-cost matching M defining χ, and consider a ς-bounded smoothed halfspace $h \in \mathcal{H}_w$ associated with slab S_w. Let $\rho(S_w, M) \leq c_2(\ell w)$ for constant c_2 (see definition of ρ in Section 5). Then $\Pr\left[\mathsf{disc}_\chi(P, h) > C\sqrt{\frac{\ell}{w}\log(2/\delta)}\right] \leq \delta$ for any $\delta > 0$ and constant $C = c_1\sqrt{2c_2}$.*

Proof. Using the gradient of v_h is at most $\varsigma = c_1/w$ and $|v_h(p_j) - v_h(q_j)| \leq \varsigma \max\{2w, \|p_j - q_j\|\}$ we have

$$\sum_j \Delta_j^2 = \sum_j 4(v_h(p_j) - v_h(q_j))^2 \leq 4\varsigma^2 \rho(S_w, M) \leq 4c_1^2/w^2 \cdot c_2 \ell w = 4c_1^2 c_2 \ell / w,$$

where the second inequality follows by Lemma 3 which shows that $\rho(S_w, M) = \sum_j \max\{(2w)^2, \|p_j - q_j\|^2\} \leq c_2(\ell w)$.

We now study the random variable $\mathsf{disc}_\chi(P, h) = |\sum_i X_i|$ for a single $h \in \mathcal{H}_w$. Invoking (1) we can bound $\Pr[\mathsf{disc}_\chi(P, h) > \alpha] \leq 2\exp(-\alpha^2/(2c_1^2 c_2 \ell/w))$. Setting $C = c_1\sqrt{2c_2}$ and $\alpha = C\sqrt{\frac{\ell}{w}\log(2/\delta)}$ reveals $\Pr\left[\mathsf{disc}_\chi(P, h) > C\sqrt{\frac{\ell}{w}\log(2/\delta)}\right] \leq \delta$. □

6.2 From a Single Smoothed Halfspace to a Smoothed Range Space

The above theorems imply small discrepancy for a single smoothed halfspace $h \in \mathcal{H}_w$, but this does not yet imply small discrepancy $\mathsf{disc}_\chi(P, \mathcal{H}_w)$, for all choices of smoothed halfspaces simultaneously. And in a smoothed range space, the family $\mathcal{H}_w$ is not finite, since even if the same set of points have $v_h(p) = 1$, $v_h(p) = 0$, or are in the slab S_w, infinitesimal changes of h will change $\textsc{sde}_P(h)$. So in order to bound $\mathsf{disc}_\chi(P, \mathcal{H}_w)$, we will show that there are polynomial in n number of smoothed halfspaces that need to be considered, and then apply a union bound across this set. The proof is deferred to the full version.

Theorem 3. *For $P \subset \mathbb{R}^2$ of size n, for $\mathcal{H}_w$, and value $\Psi(n, \delta) = O\big(\sqrt{\frac{\ell}{w}\log\frac{n}{\delta}}\big)$ for $\delta > 0$, we can choose a coloring χ such that $\Pr[\mathsf{disc}_\chi(P, \mathcal{H}_w) > \Psi(n, \delta)] \leq \delta$.*

6.3 ε-Samples for Smoothed Halfspaces

To transform this discrepancy algorithm to ε-samples, let $f(n) = \mathsf{disc}_\chi(P, \mathcal{H}_w)/n$ be the value of ε in the ε-samples generated by a single coloring of a set of size n. Solving for n in terms of ε, the sample size is $s(\varepsilon) = O(\frac{1}{\varepsilon}\sqrt{\frac{\ell}{w}\log\frac{\ell}{w\varepsilon\delta}})$. We can then apply the *MergeReduce* framework [7]; iteratively apply this random coloring in $O(\log n)$ rounds on disjoint subsets of size $O(s(\varepsilon))$. Using a generalized analysis (c.f., Theorem 3.1 in [20]), we have the same ε-sample size bound.

Theorem 4. *For $P \subset \mathsf{C}_{\ell,2} \subset \mathbb{R}^2$, with probability at least $1-\delta$, we can construct an ε-sample of $(P, \mathcal{H}_w)$ of size $O(\frac{1}{\varepsilon}\sqrt{\frac{\ell}{w}\log\frac{\ell}{w\varepsilon\delta}})$.*

To see that these bounds make rough sense, consider a random point set P in a unit square. Then setting $w = 1/n$ will yield roughly $O(1)$ points in the slab (and should roughly revert to the non-smoothed setting); this leads to $\mathsf{disc}_\chi(P, \mathcal{H}_w) = O(\sqrt{n}\sqrt{\log(n/\delta)})$ and an ε-sample of size $O((1/\varepsilon^2)\sqrt{\log(1/\varepsilon\delta)})$, basically the random sampling bound. But setting $w = \varepsilon$ so about εn points

are in the slab (the same amount of error we allow in an ε-sample) yields $\mathsf{disc}_\chi(P, \mathcal{H}_w) = O((1/\sqrt{\varepsilon n}) \cdot \sqrt{\log(n/\delta)})$ and the size of the ε-sample to be $O(\frac{1}{\varepsilon}\sqrt{\log(1/\varepsilon\delta)})$, which is a large improvement over $O(1/\varepsilon^{4/3})$, and the best bound known for non-smoothed range spaces [16].

However, the assumption that $P \subset \mathsf{C}_{\ell,2}$ (although not uncommon [16]) can be restrictive. In the full version we relax the condition for well clustered data.

6.4 Generalization to d Dimensions

We now extend from $\mathbb{R}^2$ to $\mathbb{R}^d$ for $d > 2$. Using results from Section 5 we implicitly get a bound on $\sum_j \Delta_j^d$, but the Chernoff bound we use requires a bound on $\sum_j \Delta_j^2$. As in [21], we can attain a weaker bound using Jensen's inequality over at most n terms

$$\left(\sum_j \frac{1}{n}\Delta_j^2\right)^{d/2} \leq \sum_j \frac{1}{n}\left(\Delta_j^2\right)^{d/2} \quad \text{so} \quad \sum_j \Delta_j^2 \leq n^{1-2/d}\left(\sum_j \Delta_j^d\right)^{2/d}. \tag{2}$$

Replacing this bound and using $\rho(S_w, M) \leq O(\ell^{d-1}w)$ in Lemma 4 and considering $\varsigma = c_1/w$ for some constant c_1 results in the next lemma. Its proof is deferred to the full version.

Lemma 5. *Assume $P \subset \mathbb{R}^d$ is contained in some cube $\mathsf{C}_{\ell,d}$ and with min-cost matching M, and consider a ς-bounded smoothed halfspace $h \in \mathcal{H}_w$ associated with slab S_w. Let $\rho(S_w, M) \leq c_2(\ell^{d-1}w)$ for constant c_2. Then $\Pr\left[\mathsf{disc}_\chi(P,h) > Cn^{1/2-1/d}(\ell/w)^{1-1/d}\sqrt{\log(2/\delta)}\right] \leq \delta$ for any $\delta > 0$ and $C = \sqrt{2}c_1(c_2)^{1/d}$.*

For all choices of smoothed halfspaces, applying the union bound, the discrepancy is increased by a $\sqrt{\log n}$ factor, with the following probabilistic guarantee,

$$\Pr[\mathsf{disc}_\chi(P, \mathcal{H}_w) > Cn^{1/2-1/d}(\ell/w)^{1-1/d}\sqrt{\log(n/\delta)}] \leq \delta.$$

Ultimately, we can extend Theorem 4 to the following.

Theorem 5. *For $P \subset \mathsf{C}_{\ell,d} \subset \mathbb{R}^d$, where d is constant, with probability at least $1-\delta$, we can construct an ε-sample of $(P, \mathcal{H}_w)$ of size $O\Big((\ell/w)^{2(d-1)/(d+2)} \cdot \left(\frac{1}{\varepsilon}\sqrt{\log\frac{\ell}{w\varepsilon\delta}}\right)^{2d/(d+2)}\Big)$.*

Note this result addresses scenario (S3) from the introduction where we want to find a small set (the ε-sample) so that it could be much smaller than the d/ε^2 random sampling bound, and allows generalization error $O(\varepsilon)$ for agnostic learning as described in Section 3.2. When ℓ/w is constant, the exponents on $1/\varepsilon$ are also better than those for binary ranges spaces (see Section 2).

References

1. A. F. Alessandro Bergamo, Lorenzo Torresani. Picodes: Learning a compact code for novel-category recognition. In *NIPS*, 2011.
2. N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of ACM*, 44:615–631, 1997.
3. B. Aronov, E. Ezra, and M. Sharir. Small size ε-nets for axis-parallel rectangles and boxes. *Siam Journal of Computing*, 39:3248–3282, 2010.
4. J. Beck. Irregularities of distribution I. *Acta Mathematica*, 159:1–49, 1987.
5. M. Bern and D. Eppstein. Worst-case bounds for subadditive geometric graphs. *SOCG*, 1993.
6. B. Chazelle. *The Discrepancy Method*. Cambridge, 2000.
7. B. Chazelle and J. Matousek. On linear-time deterministic algorithms for optimization problems in fixed dimensions. *J. Algorithms*, 21:579–597, 1996.
8. D. P. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge, 2009.
9. J. Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
10. S. Har-Peled, H. Kaplan, M. Sharir, and S. Smorodinksy. ε-nets for halfspaces revisited. Technical report, arXiv:1410.3154, 2014.
11. D. Haussler and E. Welzl. epsilon-nets and simplex range queries. *Disc. & Comp. Geom.*, 2:127–151, 1987.
12. S. Joshi, R. V. Kommaraju, J. M. Phillips, and S. Venkatasubramanian. Comparing distributions and shapes using the kernel distance. *SOCG*, 2011.
13. K. G. Larsen. On range searching in the group model and combinatorial discrepancy. In *FOCS*, 2011.
14. Y. Li, P. M. Long, and A. Srinivasan. Improved bounds on the samples complexity of learning. *J. Comp. and Sys. Sci.*, 62:516–527, 2001.
15. J. Matoušek. Tight upper bounds for the discrepancy of halfspaces. *Discrete & Computational Geometry*, 13:593–601, 1995.
16. J. Matoušek. *Geometric Discrepancy*. Springer, 1999.
17. J. Matoušek, R. Seidel, and E. Welzl. How to net a lot with little: Small ε-nets for disks and halfspaces. In *SOCG*, 1990.
18. J. Pach and P. K. Agarwal. *Combinatorial geometry*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1995.
19. J. Pach and G. Tardos. Tight lower bounds for the size of epsilon-nets. *Journal of American Mathematical Society*, 26:645–658, 2013.
20. J. M. Phillips. Algorithms for ε-approximations of terrains. In *Automata, Languages and Programming*, pages 447–458. Springer, 2008.
21. J. M. Phillips. eps-samples for kernels. *SODA*, 2013.
22. D. Pollard. *Emperical Processes: Theory and Applications*. NSF-CBMS REgional Confernece Series in Probability and Statistics, 1990.
23. E. Pyrga and S. Ray. New existence proofs ε-nets. In *SOCG*, 2008.
24. V. Vapnik. Inductive principles of the search for empirical dependencies. In *COLT*, 1989.
25. V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
26. K. R. Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *FOCS*, 1998.