

Distribution-Sensitive Bounds on Relative Approximations of Geometric Ranges

Yufei Tao Yu Wang

Chinese University of Hong Kong

Hong Kong

`{taoyf, yuwang}@cse.cuhk.edu.hk`

March 16, 2019

Abstract

A family \mathcal{R} of ranges and a set X of points, all in \mathbb{R}^d , together define a range space $(X, \mathcal{R}|_X)$, where $\mathcal{R}|_X = \{X \cap h \mid h \in \mathcal{R}\}$. We want to find a structure to estimate the quantity $|X \cap h|/|X|$ for any range $h \in \mathcal{R}$ with the (ρ, ϵ) -*guarantee*: (i) if $|X \cap h|/|X| > \rho$, the estimate must have a relative error ϵ ; (ii) otherwise, the estimate must have an absolute error $\rho\epsilon$. The objective is to minimize the size of the structure. Currently, the dominant solution is to compute a relative (ρ, ϵ) -approximation, which is a subset of X with $\tilde{O}(\lambda/(\rho\epsilon^2))$ points, where λ is the VC-dimension of $(X, \mathcal{R}|_X)$, and \tilde{O} hides polylog factors.

This paper shows a more general bound sensitive to the content of X . We give a structure that stores $O(\log(1/\rho))$ integers plus $\tilde{O}(\theta \cdot (\lambda/\epsilon^2))$ points of X , where θ — called the *disagreement coefficient* — measures how much the ranges differ from each other in their intersections with X . The value of θ is between 1 and $1/\rho$, such that our space bound is never worse than that of relative (ρ, ϵ) -approximations, but we improve the latter's $1/\rho$ term whenever $\theta = o(\frac{1}{\rho \log(1/\rho)})$. We also prove that, in the worst case, summaries with the $(\rho, 1/2)$ -guarantee must consume $\Omega(\theta)$ words even for $d = 2$ and $\lambda \leq 3$.

We then constrain \mathcal{R} to be the set of halfspaces in \mathbb{R}^d for a constant d , and prove the existence of structures with $o(1/(\rho\epsilon^2))$ size offering (ρ, ϵ) -guarantees, when X is generated from various stochastic distributions. This is the first formal justification on why the term $1/\rho$ is not compulsory for "realistic" inputs.

Acknowledgements: This work was partially supported by a direct grant (Project Number: 4055079) from CUHK and by a Faculty Research Award from Google.

1 Introduction

A (data) *summary*, in general, refers to a structure that captures certain information up to a specified precision about a set of objects, but using space significantly smaller than the size of the set. These summaries have become important tools in algorithm design, especially in distributed/parallel computing where the main performance goal is to minimize the communication across different servers.

In this paper, we revisit the problem of finding a small-space summary to perform range estimation in \mathbb{R}^d with relative-error guarantees. Let \mathcal{R} be a family of geometric ranges in \mathbb{R}^d (e.g., a “halfspace family” \mathcal{R} is the set of all halfspaces in \mathbb{R}^d), and X be a set of points in \mathbb{R}^d . \mathcal{R} and X together define a range space $(X, \mathcal{R}|_X)$, where $\mathcal{R}|_X = \{X \cap h \mid h \in \mathcal{R}\}$. Denote by λ the VC-dimension of $(X, \mathcal{R}|_X)$.

Following the notations of [6, 11], define

$$\overline{X}(h) = |X \cap h|/|X|$$

for each $h \in \mathcal{R}$, namely, $\overline{X}(h)$ is the fraction of points in X that are covered by h . Given real-valued parameters $0 < \rho, \epsilon < 1$, we need to produce a structure — called a (ρ, ϵ) -*summary* henceforth — that allows us to obtain, for every range $h \in \mathcal{R}$, a real-valued estimate τ satisfying the following (ρ, ϵ) -*guarantee*:

$$|\overline{X}(h) - \tau| \leq \epsilon \cdot \max\{\rho, \overline{X}(h)\}. \quad (1)$$

Phrased differently, the guarantee says that (i) if $\overline{X}(h) > \rho$, τ must have a relative error at most ϵ ; (ii) otherwise, τ must have an absolute error at most $\rho\epsilon$. The main challenge is to minimize the size of the structure.

1.1 Previous Results

Throughout the paper, all logarithms have base 2 by default.

Sample-Based (ρ, ϵ) -Summaries. We say that a (ρ, ϵ) -summary of $(X, \mathcal{R}|_X)$ is *sample-based* if it meets the requirements below: it stores a subset $Z \subseteq X$ such that, for any range $h \in \mathcal{R}$ with $Z \cap h = \emptyset$, it returns an estimate 0 for $\overline{X}(h)$.

A *relative (ρ, ϵ) -approximation* [11, 17] is a subset $Z \subseteq X$ such that $|\overline{X}(h) - \overline{Z}(h)| \leq \epsilon \cdot \max\{\rho, \overline{X}(h)\}$ holds for all ranges $h \in \mathcal{R}$. Hence, the (ρ, ϵ) -guarantee can be fulfilled by simply setting τ to $\overline{Z}(h)$, rendering Z a legal (sample-based) (ρ, ϵ) -summary. Strengthening earlier results [3, 12, 21], Li et al. [17] proved that a random sample of X with size $O(\frac{1}{\rho} \cdot \frac{1}{\epsilon^2} (\lambda \log \frac{1}{\rho} + \log \frac{1}{\delta}))$ is a relative (ρ, ϵ) -approximation with probability at least $1 - \delta$. This implies the existence of a (ρ, ϵ) -summary of size $O(\frac{1}{\rho} \cdot \frac{\lambda}{\epsilon^2} \log \frac{1}{\rho})$.

A range space $(X, \mathcal{R}|_X)$ of a constant VC-dimension is said to be *well-behaved*, if $\mathcal{R}|_X$ contains at most $O(|X|) \cdot k^{O(1)}$ sets of size not exceeding k , for any integer k from 1 to $|X|$. Ezra [6] showed that such a range space admits a sample-based (ρ, ϵ) -summary of size $O(\frac{1}{\rho} \cdot \frac{1}{\epsilon^2} (\log \frac{1}{\epsilon} + \log \log \frac{1}{\rho}))$; note that this is smaller than the corresponding result $O(\frac{1}{\rho} \cdot \frac{1}{\epsilon^2} \log \frac{1}{\rho})$ of [17] when $\rho \ll \epsilon$. It is worth mentioning that, when $d \leq 3$ and \mathcal{R} is the halfspace family, any $(X, \mathcal{R}|_X)$ is well-behaved; this, however, is not true when $d \geq 4$.

As opposed to the above “generic” bounds, Har-Peled and Sharir [11] proved specific bounds on the halfspace family \mathcal{R} . For $d = 2$, they showed that any (X, \mathcal{R}) has a relative (ρ, ϵ) -approximation

of size $O(\frac{1}{\rho} \cdot \frac{1}{\epsilon^{4/3}} \log^{4/3} \frac{1}{\rho\epsilon})$; similarly, for $d = 3$, the bound becomes $O(\frac{1}{\rho} \cdot \frac{1}{\epsilon^{3/2}} \log^{3/2} \frac{1}{\rho\epsilon})$. Combining these results and those of [6] gives the currently best bounds for these range spaces.

A Lower Bound of $\Omega(1/\rho)$. Notice that all the above bounds contain a term $1/\rho$. This is not a coincidence, but instead is due to a connection to “ ϵ -nets”. Given a range space $(X, \mathcal{R}|_X)$, an ϵ -net [13] is a subset $Z \subseteq X$ such that $\overline{Z}(h) > 0$ holds for any range $h \in \mathcal{R}$ satisfying $\overline{X}(h) \geq \epsilon$. As can be verified easily, any sample-based $(\rho, 1/2)$ -summary of $(X, \mathcal{R}|_X)$ must also be a ρ -net. This implies that the smallest size of sample-based $(\rho, 1/2)$ -summaries must be at least that of ρ -nets.

Regarding the sizes of ϵ -nets, a lower bound of $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ is known for many range families \mathcal{R} (see [14, 15, 20] and the references therein). More precisely, this means that, for each such family \mathcal{R} , one cannot hope to obtain an ϵ -net of size $o(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ for *every* possible X . It thus follows that, for these families \mathcal{R} , $\Omega(\frac{1}{\rho} \log \frac{1}{\rho})$ is a lower bound on the sizes of sample-based $(\rho, 1/2)$ -summaries. This, in turn, indicates that range spaces $(X, \mathcal{R}|_X)$ defined by such an \mathcal{R} cannot always be well-behaved.

Coming back to the sizes of ϵ -nets, a weaker lower bound of $\Omega(1/\epsilon)$ holds quite commonly even on the range families that evade the $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ bound. Consider, for example, the halfspace family \mathcal{R} in \mathbb{R}^2 . For any X , the range space $(X, \mathcal{R}|_X)$ definitely has an ϵ -net of size $O(1/\epsilon)$ [10, 19]. This is tight: place a set X of points on the boundary of a circle; and it is easy to show that any ϵ -net of $(X, \mathcal{R}|_X)$ must have a size of at least $1/\epsilon$. This means that the size of any sample-based $(\rho, 1/2)$ -summary of $(X, \mathcal{R}|_X)$ must be $\Omega(1/\rho)$.

1.2 Our Results

On One Input: Moving Beyond $\Omega(1/\rho)$. The $\Omega(1/\rho)$ lower bound discussed earlier holds only in the *worst case*, i.e., it is determined by the “hardest” X . For other X , the range space $(X, \mathcal{R}|_X)$ may admit much smaller (ρ, ϵ) -summaries. For example, let \mathcal{R} be again the set of halfplanes in \mathbb{R}^2 . When all the points of X lie on a line, (X, \mathcal{R}) has a sample-based (ρ, ϵ) -approximation of size only $O(\frac{1}{\epsilon} \log \frac{1}{\rho})$; also, it would be interesting to note that (X, \mathcal{R}) has an ϵ -net that contains only 2 points! In general, the existing bounds on (ρ, ϵ) -summaries can be excessively loose on individual inputs X . This calls for an alternative, *distribution-sensitive*, analytical framework that is able to prove tighter bounds using extra complexity parameters that depend on the *content* of X .

The first contribution of this paper is to establish such a framework by resorting to the concept of *disagreement coefficient* [8] from active learning. This notion was originally defined in a context different from ours; and we will adapt it to range spaces in the next section. At this moment, it suffices to understand that the disagreement coefficient θ is a real value satisfying: $1 \leq \theta \leq 1/\rho$. The coefficient quantifies the differences among the sets in $\mathcal{R}|_X$ (a larger θ indicates greater differences). Even under the same \mathcal{R} , θ may vary considerably depending on X .

We will show that, for any range space $(X, \mathcal{R}|_X)$ of VC-dimension λ , there is a (ρ, ϵ) -summary that keeps $O(\log(1/\rho))$ integers plus

$$O\left(\min\left\{\frac{1}{\rho}, \theta \log \frac{1}{\rho}\right\} \cdot \frac{\lambda}{\epsilon^2} \log \frac{1}{\rho}\right) \quad (2)$$

points of X . The above is never worse than the general bound $O(\frac{1}{\rho} \cdot \frac{\lambda}{\epsilon^2} \log \frac{1}{\rho})$ of relative (ρ, ϵ) -approximations.

We will also prove that $\Omega(\theta)$ is a lower bound on the number of words needed to encode a $(\rho, 1/2)$ -summary even when $d = 2$ and $\lambda \leq 3$. This generalizes the $\Omega(1/\rho)$ lower bound in Section 1.1

because θ is at most, but can reach, $1/\rho$. Thus, our result in (2) reflects the hardness of the input, and is tight within polylog factors for constant ϵ . Our lower bound is information-theoretic, and does not require the summary to be sample-based.

On a Distribution of Inputs: Small Summaries for Halfspaces. Our framework allows us to explain — for the first time we believe — why $\Omega(1/\rho) \cdot \text{poly}(1/\epsilon)$ is too pessimistic a bound on the sizes of (ρ, ϵ) -summaries for inputs encountered in practice. For this purpose, we *must not* allow arbitrary inputs because of the prevalent $\Omega(1/\rho)$ lower bound; instead, we will examine a class of inputs following a certain distribution.

In this paper, we demonstrate the above by concentrating on the family \mathcal{R} of halfspaces in \mathbb{R}^d where the dimensionality d is a constant; this is arguably the “most-studied” family in the literature of relative (ρ, ϵ) -approximations. The core of our solutions concerns two stochastic distributions that have drastically different behavior:

- **(Box Uniform)** Suppose that X is obtained by drawing n points uniformly at random from the *unit box* $[0, 1]^d$. When $\rho = \Omega(\frac{\log n}{n})$, we will prove that $\theta = O(\text{polylog } \frac{1}{\rho})$ with high probability (i.e., at least $1 - 1/n^2$). Accordingly, (2) becomes $O(\frac{1}{\epsilon^2} \text{polylog } \frac{1}{\rho})$, improving the general bound of relative (ρ, ϵ) -approximations by almost a factor of $O(1/\rho)$.
- **(Ball Uniform)** Consider instead that the n points are drawn uniformly at random from the *unit ball*: $\{x \in \mathbb{R}^d \mid \sum_{i=1}^d x[i]^2 \leq 1\}$, where $x[i]$ represents the i -th coordinate of point x . This time, we will prove that $\theta = O((\frac{1}{\rho})^{\frac{d-1}{d+1}})$ with high probability for $\rho = \Omega(\frac{\log n}{n})$. (2) indicates the existence of a (ρ, ϵ) -summary with size $\tilde{O}((\frac{1}{\rho})^{\frac{d-1}{d+1}} \cdot \frac{1}{\epsilon^2})$ for $\rho = \Omega(\frac{\log n}{n})$, again circumventing the $\Omega(1/\rho)$ lower bound.

The very same bounds can also be obtained in *non-uniform* settings. Suppose that X is obtained by drawing n points in an *iid* manner, according to a distribution that can be described by a probabilistic density function (pdf) $\pi(x)$ over \mathbb{R}^d where $d = O(1)$. Define the *support region* of π as $\text{supp}(\pi) = \{x \in \mathbb{R}^d \mid \pi(x) > 0\}$. When π satisfies:

- *C1*: $\text{supp}(\pi)$ is the unit box (or unit ball, resp.);
- *C2*: for every point $x \in \text{supp}(\pi)$, it holds that $\pi(x) = \Omega(1)$;

we will show that $(X, \mathcal{R}|_X)$ has a (ρ, ϵ) -summary whose size is asymptotically the same as the aforementioned bound for box uniform (or ball uniform, resp.). Conditions *C1* and *C2* are satisfied by many distributions encountered in practice (e.g., the truncated versions of the Gaussian, Elliptical, and Laplace distributions, etc.), suggesting that real-world datasets may have much smaller (ρ, ϵ) -summaries than previously thought.

Even better, the linearity of halfspaces implies that, the same bounds still hold even when the shape of $\text{supp}(\pi)$ in Condition *C1* is obtained from the unit box/ball by an affine transformation. Call a distribution *atomic* if it satisfies *C1* (perhaps after an affine transformation) and *C2*. Our results hold also on “composite distributions” synthesized from a constant z number of atomic distributions whose support regions may overlap *arbitrarily*. Specifically, let $\pi_1, \pi_2, \dots, \pi_z$ be the pdfs of atomic distributions; and define $\pi(x) = \sum_{i=1}^z \gamma_i \cdot \pi_i(x)$, for arbitrary positive constants $\gamma_1, \gamma_2, \dots, \gamma_z$ that sum up to 1; see Figure 1 for an example. Then, the (ρ, ϵ) -summary bound on π is asymptotically determined by the highest of the (ρ, ϵ) -summary bounds on π_1, \dots, π_z .




Figure 1: Mixture of 3 truncated Gaussian distributions in 2D space

2 Disagreement Coefficients

Existing Definitions on Distributions. Disagreement coefficient was introduced by Hanneke [8] to analyze active learning algorithms (although a similar concept had been coined earlier [1] in statistics).

Let \mathcal{D} be a distribution over \mathbb{R}^d . For any region $A \subseteq \mathbb{R}^d$, we denote by $\mathbf{Pr}_{\mathcal{D}}[A]$ the probability of $x \in A$ when x is drawn from \mathcal{D} . Let \mathcal{R} be a family of geometric ranges. Given a subset $\mathcal{R}' \subseteq \mathcal{R}$, define the *disagreement region* $DIS(\mathcal{R}')$ of \mathcal{R}' as

$$DIS(\mathcal{R}') = \{x \in \mathbb{R}^d \mid \exists h_1, h_2 \in \mathcal{R}' \text{ s.t. } x \in h_1 \text{ and } x \notin h_2\}.$$

That is, $DIS(\mathcal{R}')$ includes every such point $x \in \mathbb{R}^d$ that does not fall in all the ranges in \mathcal{R}' , and in the meantime, does not fall outside all the ranges in \mathcal{R}' , either. Given a range $h \in \mathcal{R}$ and a real value $r > 0$, define its r -ball $B_{\mathcal{D}}(h, r)$ as the set of all ranges $h' \in \mathcal{R}$ satisfying $\mathbf{Pr}_{\mathcal{D}}[DIS(\{h, h'\})] \leq r$. It is worth mentioning that $DIS(\{h, h'\})$ is simply the symmetric difference between h and h' .

Now, fix a range h , and consider increasing r continuously; this can only expand the set $B_{\mathcal{D}}(h, r)$, and hence, also $DIS(B_{\mathcal{D}}(h, r))$. Interestingly, even though $\mathbf{Pr}_{\mathcal{D}}[DIS(B_{\mathcal{D}}(h, r))]$ is monotonically increasing, the ratio $\mathbf{Pr}_{\mathcal{D}}[DIS(B_{\mathcal{D}}(h, r))]/r$ may remain bounded by a certain quantity. Given a real value $\sigma \geq 0$, the *disagreement coefficient* $\theta_{\mathcal{D}}^h(\sigma)$ of h measures this quantity with respect to all $r > \sigma$:

$$\theta_{\mathcal{D}}^h(\sigma) = \max \left\{ 1, \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{D}}[DIS(B_{\mathcal{D}}(h, r))]}{r} \right\}. \quad (3)$$

The function $\theta_{\mathcal{D}}^h(\sigma)$ has several useful properties:

1. By definition, $\theta_{\mathcal{D}}^h(\sigma)$ is between 1 and $1/\sigma$, regardless of \mathcal{D} and h .
2. The supremum in (3) ensures that $\theta_{\mathcal{D}}^h(\sigma)$ is monotonically decreasing.
3. For any $c \geq 1$, it holds that $\theta_{\mathcal{D}}^h(\sigma) \leq c \cdot \theta_{\mathcal{D}}^h(c\sigma)$ (see Corollary 7.2 of [9]).

New Definitions on Range Spaces. The above definitions rely on \mathcal{D} , and are not suitable for our problem settings where the input X is a finite set. Next, we present a way to adapt the definitions to a range space $(X, \mathcal{R}|_X)$ for analyzing geometric algorithms.

We impose a uniform distribution over X : let $\mathcal{U}(X)$ be the distribution of a random point drawn uniformly from X . By replacing \mathcal{D} with $\mathcal{U}(X)$ in (3), we rewrite (3) into the following for any $\sigma \geq 0$:

$$\theta_{\mathcal{U}(X)}^h(\sigma) = \max \left\{ 1, \sup_{r > \sigma} \frac{\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]}{r} \right\}. \quad (4)$$

Set

$$\sigma_{min} = \frac{\min_{h \in \mathcal{R}} |X \cap h|}{n} \quad (5)$$

We define the *disagreement coefficient of the range space* $(X, \mathcal{R}|_X)$ as a function $\theta_X(\sigma) : [\sigma_{min}, \infty) \rightarrow \mathbb{R}$ where

$$\theta_X(\sigma) = \min_{h \in \mathcal{R} \text{ s.t. } \overline{X}(h) \leq \sigma} \left\{ \theta_{\mathcal{U}(X)}^h(\sigma) \right\}. \quad (6)$$

It is clear from the above discussion that $1 \leq \theta_X(\sigma) \leq 1/\sigma$ and $\theta_X(\sigma)$ is monotonically decreasing.

As a remark, the finiteness of X gives a simpler interpretation of the r -ball $B_{\mathcal{U}(X)}(h, r)$: it is the set of ranges $h' \in \mathcal{R}$ such that $DIS(\{h, h'\})$ covers no more than $r|X|$ points in X . Also, $\Pr_{\mathcal{U}(X)}[A]$ for any region $A \subseteq \mathbb{R}^d$ is simply $|X \cap A|/|X|$.

3 Small (ρ, ϵ) -Summaries Based on Disagreement Coefficients

Given a range space (X, \mathcal{R}) with VC-dimension λ , we will show how to find a (ρ, ϵ) -summary whose size can be bounded using disagreement coefficients. Our algorithm is randomized, and succeeds with probability at least $1 - \delta$ for a real-valued parameter $0 < \delta < 1$. Set $n = |X|$. We require that $\rho \geq \sigma_{min}$; otherwise, manually increasing ρ to σ_{min} achieves the same approximation guarantee.

3.1 Algorithms

3.1.1 Computing a (ρ, ϵ) -Summary

We will shrink \mathcal{R} progressively by removing a range h from \mathcal{R} once we are sure we can provide an accurate estimate for $\overline{X}(h)$. Define $\mathcal{R}_0 = \mathcal{R}$. We perform at most $\lceil \log(1/\rho) \rceil$ rounds. Given \mathcal{R}_{i-1} , Round $i \geq 1$ is executed as follows:

1. $m_i \leftarrow$ the number of points $x \in X$ such that x falls in *all* the ranges in \mathcal{R}_{i-1}
2. $X_i \leftarrow X \cap DIS(\mathcal{R}_{i-1})$
3. draw a set S_i of points uniformly at random from X_i with

$$|S_i| = O \left(\frac{|X_i|}{n} \cdot \frac{2^i}{\epsilon^2} \left(\lambda \log \frac{1}{\rho} + \log \frac{\log(1/\rho)}{\delta} \right) \right) \quad (7)$$

4. $\mathcal{R}_i = \{h \in \mathcal{R}_{i-1} \mid \overline{S_i}(h) \cdot |X_i| + m_i < n/2^i\}$

The algorithm terminates when either $i = \lceil \log(1/\rho) \rceil$ or $\mathcal{R}_i = \emptyset$. Suppose that in total t rounds are performed. The final (ρ, ϵ) -summary consists of sets S_1, S_2, \dots, S_t , and $2t + 1$ integers $n, m_1, m_2, \dots, m_t, |X_1|, |X_2|, \dots, |X_t|$.

3.1.2 Performing Estimation

Given a range $h \in \mathcal{R}$, we deploy the summary to estimate $\overline{X}(h)$ in two steps:

1. $j \leftarrow$ the largest $i \in [1, t]$ such that $h \in \mathcal{R}_i$
2. return $\overline{S}_j(h) \cdot \frac{|X_j|}{n} + \frac{m_j}{n}$ as the estimate

Regarding Step 1, whether $h \in \mathcal{R}_i$ can be detected as follows. First, if $h \notin \mathcal{R}_{i'}$ for any $i' < i$, then immediately $h \notin \mathcal{R}_i$. Otherwise, compute $\overline{S}_i(h)$, and declare $h \in \mathcal{R}_i$ if and only if $\overline{S}_i(h) \cdot |X_i| + m_i < n/2^i$.

3.2 Analysis

We now proceed to prove the correctness of our algorithms, and bound the size of the produced summary. It suffices to consider $\epsilon \leq 1/3$ (otherwise, lower ϵ to $1/3$ and then apply the argument below).

The subsequent discussion is carried out under the event that, for every $i \in [1, t]$, S_i is a relative $(\rho_i, \epsilon/4)$ -approximation of X_i with respect to the ranges in \mathcal{R} where

$$\rho_i = \frac{n(1+\epsilon)}{2^i \cdot |X_i|}.$$

By the result of [17] (reviewed in Section 1.1), with $|S_i|$ shown in (7), the event happens with a probability at least $1 - \delta \cdot \frac{t}{\lceil \log(1/\rho) \rceil} \geq 1 - \delta$.

3.2.1 Correctness

To show that our algorithm indeed outputs a (ρ, ϵ) -summary, the key step is to prove:

Lemma 1. *The following are true for all $i \in [1, t]$: (i) for every range $h \in \mathcal{R}_i$, $\overline{X}(h) < (1 + \epsilon)/2^i$; (ii) for every range $h \notin \mathcal{R}_i$, $\overline{X}(h) \geq (1 - \epsilon)/2^i$.*

Proof. See Appendix A. □

Now consider the estimation algorithm in Section 3.1.2. Given the value j obtained at Step 1 for the input range $h \in \mathcal{R}$, the above lemma suggests that

$$(1 - \epsilon)/2^{j+1} \leq \overline{X}(h) < (1 + \epsilon)/2^j.$$

This, together with S_j being a $(\rho_j, \epsilon/4)$ -approximation of X_j , ensures that our estimate satisfies the (ρ, ϵ) -guarantee for h . The details can be found in Appendix B.

3.2.2 Bounding the Size

To bound the size of our (ρ, ϵ) -summary, we will focus on bounding $\sum_{i=1}^t |S_i|$, because the rest of the summary clearly needs $O(t) = O(\log(1/\rho))$ extra integers. Let us start with a trivial bound that follows directly from $|X_i| \leq n$:

$$\begin{aligned} \sum_{i=1}^t |S_i| &= O\left(\sum_{i=1}^t \frac{2^i}{\epsilon^2} \left(\lambda \log \frac{1}{\rho} + \log \frac{\log(1/\rho)}{\delta}\right)\right) \\ &= O\left(\frac{1}{\rho \epsilon^2} \left(\lambda \log \frac{1}{\rho} + \log \frac{\log(1/\rho)}{\delta}\right)\right). \end{aligned} \tag{8}$$

Next, we use disagreement coefficients to prove a tighter bound. Fix $h \in \mathcal{R}$ to be an *arbitrary* range such that $\bar{X}(h) \leq \rho$ (h definitely exists because $\rho \geq \sigma_{\min}$).

Lemma 2. $\mathcal{R}_i \subseteq B(h, \rho + (1 + \epsilon)/2^i)$.

Proof. It suffices to prove that, for any $h' \in \mathcal{R}_i$, $\Pr_{\mathcal{U}(X)}[DIS(\{h, h'\})] \leq \rho + (1 + \epsilon)/2^i$, or equivalently, $|X \cap DIS(\{h, h'\})| \leq n(\rho + (1 + \epsilon)/2^i)$.

This holds because $|X \cap DIS(\{h, h'\})| \leq |(X \cap h) \cup (X \cap h')|$. By definition of h , we know $|X \cap h| \leq n\rho$, while By Lemma 1, we know $|X \cap h'| \leq n(1 + \epsilon)/2^i$. Therefore, $|X \cap DIS(\{h, h'\})| \leq n(\rho + (1 + \epsilon)/2^i)$. \square

Lemma 3. $|X_i|/n \leq \theta_{\mathcal{U}(X)}^h(2\rho) \cdot (\rho + \frac{1+\epsilon}{2^{i-1}})$.

Proof. Lemma 2 tells us that $DIS(\mathcal{R}_{i-1}) \subseteq DIS(B_{\mathcal{U}(X)}(h, \rho + \frac{1+\epsilon}{2^{i-1}}))$. Thus:

$$\begin{aligned} |X_i|/n &= \Pr_{\mathcal{U}(X)}(DIS(\mathcal{R}_{i-1})) \\ &\leq \Pr_{\mathcal{U}(X)}\left(DIS\left(B_{\mathcal{U}(X)}\left(h, \rho + \frac{1+\epsilon}{2^{i-1}}\right)\right)\right) \\ (\text{by (4)}) &\leq \theta_{\mathcal{U}(X)}^h\left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) \cdot \left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) \end{aligned}$$

By $1/2^{i-1} > \rho$, and the fact that $\theta_{\mathcal{U}(X)}^h$ is monotonically decreasing, the above leads to

$$\begin{aligned} \theta_{\mathcal{U}(X)}^h\left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) \cdot \left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) &\leq \theta_{\mathcal{U}(X)}^h(\rho + \rho) \cdot \left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) \\ &= \theta_{\mathcal{U}(X)}^h(2\rho) \cdot \left(\rho + \frac{1+\epsilon}{2^{i-1}}\right). \end{aligned}$$

\square

Therefore:

$$\begin{aligned} \sum_{i=1}^t \frac{|X_i| \cdot 2^i}{n} &\leq \theta_{\mathcal{U}(X)}^h(2\rho) \cdot \sum_{i=1}^t 2^i \cdot \left(\rho + \frac{1+\epsilon}{2^{i-1}}\right) \\ &= \theta_{\mathcal{U}(X)}^h(2\rho) \cdot \sum_{i=1}^t (2^i \cdot \rho + O(1)) \\ (\text{by } 1/2^i = \Omega(\rho)) &= \theta_{\mathcal{U}(X)}^h(2\rho) \cdot O(t) \\ &= \theta_{\mathcal{U}(X)}^h(2\rho) \cdot O(\log(1/\rho)) \\ &= \theta_{\mathcal{U}(X)}^h(\rho) \cdot O(\log(1/\rho)) \end{aligned} \tag{9}$$

where the last equality used the fact that $\theta_{\mathcal{U}(X)}^h(2\rho) \leq 2 \cdot \theta_{\mathcal{U}(X)}^h(\rho)$.

Remember that the above holds for *all* $h \in \mathcal{R}$ satisfying $\bar{X}(h) \leq \rho$. By the definition in (6), we can improve the bound of (9) to

$$\sum_{i=1}^t \frac{|X_i| \cdot 2^i}{n} = \theta_X(\rho) \cdot O(\log(1/\rho)). \tag{10}$$

Combining the above with (7) gives $\sum_{i=1}^t |S_i| = O(\frac{1}{\epsilon^2} \cdot \theta_X(\rho) \log(1/\rho) \cdot (\lambda \log(1/\rho) + \log \frac{\log(1/\rho)}{\delta}))$. Putting this together with (8) and setting δ to a constant gives:

Theorem 1. *For any $\rho \geq \sigma_{\min}$ and any $0 < \epsilon < 1$, a range space $(X, \mathcal{R}|_X)$ of VC-dimension λ has a (ρ, ϵ) -summary which keeps $O(\log(1/\rho))$ integers and $O(\min\{\frac{1}{\rho}, \theta_X(\rho) \cdot \log \frac{1}{\rho}\} \cdot \frac{\lambda}{\epsilon^2} \log \frac{1}{\rho})$ points of X . Here, σ_{\min} is defined in (5), and θ_X is the disagreement coefficient function defined in (6).*

3.2.3 A Remark

Our (ρ, ϵ) -summary is currently not sample-based, but this can be fixed by keeping — at Step 1 of the computation algorithm in Section 3.1 — an arbitrary point counted by m_i .

The (ρ, ϵ) -summary after the fix also serves as a ρ -net. Thus, by setting ϵ to a constant in Theorem 1, we know that for any $\rho \geq \sigma_{\min}$, the range space $(X, \mathcal{R}|_X)$ in Theorem 1 has an ρ -net of size $O(\min\{\frac{1}{\rho}, \theta_X(\rho) \cdot \log \frac{1}{\rho}\} \cdot \lambda \log \frac{1}{\rho})$. However, it should be pointed out that this bound on ρ -nets can be slightly improved, as is implied by Theorem 5.1 of [9] and made explicit in [16].

4 Bridging Distribution and Finite-Set Disagreement Coefficients

This section will establish another theorem which will be used together with Theorem 1 to explain why we are able to obtain (ρ, ϵ) -summarizes of $o(1/\rho)$ size on practical datasets. Suppose that the input X has been generated by taking n points independently following the same distribution \mathcal{D} over \mathbb{R}^d . The learning literature (see, e.g., [9]) has developed a solid understanding on when the quantity $\theta_{\mathcal{D}}^h(\sigma)$ is small. Unfortunately, those findings can rarely be applied to $\theta_{\mathcal{U}(X)}^h(\sigma)$ because they are conditioned on requirements that must be met by \mathcal{D} , e.g., one common requirement is continuity. $\mathcal{U}(X)$, due to its discrete nature, seldom meets the requirements.

On the other hand, clearly $\mathcal{U}(X)$ approximates \mathcal{D} increasingly better as n grows. Thus, we ask the question:

How large n needs to be for $\theta_{\mathcal{U}(X)}^h(\sigma)$ to be asymptotically the same as $\theta_{\mathcal{D}}^h(\sigma)$?

We partially answer the question in the next theorem:

Theorem 2 (The Bridging Theorem). *Let \mathcal{D} be a distribution over \mathbb{R}^d , and \mathcal{R} be a family of ranges. Denote by λ the VC-dimension of the range space $(\mathbb{R}^d, \mathcal{R})$.*

Fix an arbitrary range $h \in \mathcal{R}$, an arbitrary integer n , a real value $0 < \delta < 1$, a real value σ satisfying $n \geq \frac{c}{\sigma}(\log \frac{n}{\delta} + \lambda \log \frac{1}{\sigma})$ for some universal constant c . If we draw a set X of n points independently from \mathcal{D} , then with probability at least $1 - \delta$, it holds that $\theta_{\mathcal{U}(X)}^h(\sigma) \leq 8 \cdot \theta_{\mathcal{D}}^h(2\sigma)$.

The rest of the section serves as a proof of the theorem. Let us first get rid of two easy cases:

- If $\sigma \geq 1$, $\theta_{\mathcal{U}(X)}^h(\sigma) = \theta_{\mathcal{D}}^h(2\sigma) = 1$ by definition of (4); and the theorem obviously holds.
- If $\sigma < 1/n$, observe that every range $h' \in B_{\mathcal{U}(X)}(h, \sigma)$ covers exactly the same set of points in X as h . Hence, $\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))] = 0$. It follows from (4) that $\theta_{\mathcal{U}(X)}^h(\sigma) = 1$. The theorem again obviously holds because $\theta_{\mathcal{D}}^h(2\sigma) \geq 1$, by definition.

Hence, it suffices to consider $1/n \leq \sigma < 1$. Define $S = \{i/n \mid i \text{ is an integer in } [\sigma n, n]\}$. For $\sigma \geq 1/n$, (4) implies

$$\theta_{\mathcal{U}(X)}^h(\sigma) \leq \max \left\{ 1, 2 \cdot \max_{r \in S} \frac{\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]}{r} \right\}. \quad (11)$$

Consider an arbitrary $r \in S$. We will show that, when n satisfies the condition in the theorem, with probability at least $1 - \delta/n$, it holds that

$$\frac{\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]}{r} \leq 4 \cdot \theta_{\mathcal{D}}^h(2\sigma). \quad (12)$$

Once this is done, applying the union bound on all the $r \in S$ will prove that (11) is at most $8 \cdot \theta_{\mathcal{D}}^h(2\sigma)$ with probability at least $1 - \delta$, as claimed in the theorem.

We aim to establish the following equivalent form of (12):

$$\frac{|X \cap DIS(B_{\mathcal{U}(X)}(h, r))|}{nr} \leq 4 \cdot \theta_{\mathcal{D}}^h(2\sigma). \quad (13)$$

For the above purpose, the most crucial step is to prove:

Lemma 4. *When $n \geq \frac{c_1}{r}(\lambda \log \frac{1}{r} + \log \frac{n}{\delta})$ for some universal constant c_1 , it holds with probability at least $1 - \delta/(2n)$ that $B_{\mathcal{U}(X)}(h, r) \subseteq B_{\mathcal{D}}(h, 2r)$.*

Proof. The rationale of our proof is that any $h' \notin B_{\mathcal{D}}(h, 2r)$ is unlikely to appear in $B_{\mathcal{U}(X)}(h, r)$ when n is large. Indeed, $h' \notin B_{\mathcal{D}}(h, 2r)$ indicates that a point x drawn from D has probability over $2r$ to fall in $DIS(\{h, h'\})$. Hence, $|X \cap DIS(\{h, h'\})|$ should be sharply concentrated around $2r \cdot n$, rendering $h' \notin B_{\mathcal{U}(X)}(h, r)$. The challenge, however, is that there can be an infinite number of ranges h' to consider. To tackle the challenge, we need to bring down the number of ranges somehow to $n^{O(\lambda)}$. We achieve the purpose by observing that we can define another range space with VC-dimension $O(\lambda)$ to capture the disagreement regions of range pairs from \mathcal{R} , as shown below.

Define $\mathcal{R}^{dis} = \{DIS(\{h, h'\}) \mid h, h' \in \mathcal{R}\}$. We observe that the range space $(\mathbb{R}^d, \mathcal{R}^{dis})$ has VC-dimension $O(\lambda)$. To explain why, for any $h \in \mathcal{R}$, define $\bar{h} = \mathbb{R}^d \setminus h$. Accordingly, define $\bar{\mathcal{R}} = \{\bar{h} \mid h \in \mathcal{R}\}$. The two range spaces $(\mathbb{R}^d, \mathcal{R})$ and $(\mathbb{R}^d, \bar{\mathcal{R}})$ have the same VC-dimension λ . Therefore, the range space $(\mathbb{R}^d, \mathcal{R} \cup \bar{\mathcal{R}})$ has VC-dimension at most $2\lambda + 1$. Now apply a 2-fold intersection on $(\mathbb{R}^d, \mathcal{R} \cup \bar{\mathcal{R}})$ to create $(\mathbb{R}^d, \mathcal{R}_1)$ where $\mathcal{R}_1 = \{h \cap h' \mid h, h' \in \mathcal{R} \cup \bar{\mathcal{R}}\}$. By a result of [2], the VC dimension of $(\mathbb{R}^d, \mathcal{R}_1)$ is bounded by $O(\lambda)$. Finally, apply a 2-fold union on $(\mathbb{R}^d, \mathcal{R}_1)$ to create $(\mathbb{R}^d, \mathcal{R}_2)$ where $\mathcal{R}_2 = \{h \cup h' \mid h, h' \in \mathcal{R}_1\}$. By another result of [2], the VC dimension of $(\mathbb{R}^d, \mathcal{R}_2)$ is bounded by $O(\lambda)$. Notice that \mathcal{R}^{dis} is a subset of \mathcal{R}_2 . It thus follows that the VC-dimension of $(\mathbb{R}^d, \mathcal{R}^{dis})$ must be $O(\lambda)$.

Essentially, now the task is to draw a sufficiently large set X of points from \mathcal{D} to guarantee with probability at least $1 - \delta/(2n)$: for every range $h \in \mathcal{R}^{dis}$ with $\Pr_{\mathcal{D}}(h) > 2r$, we ensure $|X \cap h|/|X| > r$. By applying a result of [17] on general range spaces, we know that $|X|$ only needs to be $\frac{c_1}{r}(\lambda \log \frac{1}{r} + \log \frac{n}{\delta})$ for some constant c_1 which does not depend on r, δ , and n . \square

Set $r' = \Pr_{\mathcal{D}}(DIS(B_{\mathcal{D}}(h, 2r)))$; notice that, by definition of $\theta_{\mathcal{D}}^h(2r)$, $r' \leq 2r \cdot \theta_{\mathcal{D}}^h(2r)$. We want to draw a sufficiently large set X of points from \mathcal{D} to guarantee, with probability at least $1 - \delta/(2n)$, $|X \cap DIS(B_{\mathcal{D}}(h, 2r))| \leq 2n \cdot \max\{r, r'\}$. By Chernoff bounds, n only needs to be at least $\frac{c_2}{r} \log \frac{n}{\delta}$ for some universal constant c_2 .

Now, set $c = \max\{c_1, c_2\}$ and $n = \frac{c}{r}(\lambda \log \frac{1}{r} + \log \frac{n}{\delta})$. With probability at least $1 - \delta/n$, we can

derive (13) from the above discussion as follows:

$$\begin{aligned}
\frac{|X \cap \text{DIS}(B_{\mathcal{U}(X)}(h, r))|}{nr} &\leq \frac{|X \cap \text{DIS}(B_{\mathcal{D}}(h, 2r))|}{nr} \quad (\text{by Lemma 4}) \\
&\leq \frac{2n \cdot \max\{r, r'\}}{nr} = 2 \cdot \max\{1, r'/r\} \\
&\leq 2 \cdot \max\{1, 2 \cdot \theta_{\mathcal{D}}^h(2r)\} = 4 \cdot \theta_{\mathcal{D}}^h(2r) \leq 4 \cdot \theta_{\mathcal{D}}^h(2\sigma)
\end{aligned}$$

where the last inequality used $r \geq \sigma$ and the fact that $\theta_{\mathcal{D}}^h$ is monotonically decreasing. This establishes (13) and hence completes the proof of Theorem 2.

5 $o(1/\rho)$ -Size Summaries for Halfspace Ranges

We are ready to explain why a set of points generated from a stochastic distribution often admits (ρ, ϵ) -summaries of $o(1/\rho)$ size for fixed ϵ . This requires specializing \mathcal{R} into a concrete range family. We will do so by constraining \mathcal{R} to be the set of halfspaces in \mathbb{R}^d , because this family has received considerable attention (as reviewed in Section 1.1).

We prove in the appendix the next two technical lemmas regarding the disagreement coefficients on box-uniform and ball-uniform distributions:

Lemma 5. *Let \mathcal{U} be the distribution where a point is drawn uniformly at random from the unit box $[0, 1]^d$ with $d = O(1)$. For any halfspace h disjoint with the box, it holds that $\theta_{\mathcal{U}}^h(\sigma) = O(\log^{d-1} \frac{1}{\sigma})$ for all $\sigma > 0$.*

Lemma 6. *Let \mathcal{U} be the distribution where a point is drawn uniformly at random from the unit ball $\{x \in \mathbb{R}^d \mid \sum_{i=1}^d x[i]^2 \leq 1\}$ with $d = O(1)$. For any halfspace h disjoint with the ball, it holds that $\theta_{\mathcal{U}}^h(\sigma) = O((\frac{1}{\sigma})^{\frac{d-1}{d+1}})$ for all $\sigma > 0$.*

Next, we establish our main result for *non-uniform* distributions:

Theorem 3. *Let \mathcal{R} be the family of halfspaces in \mathbb{R}^d with a constant dimensionality d . Let \mathcal{D} be a distribution over \mathbb{R}^d such that the pdf π of \mathcal{D} satisfies Conditions C1 and C2 as prescribed in Section 1.2. Suppose that we draw a set X of n points independently from \mathcal{D} . Both of the following hold with probability at least $1 - 1/n^2$:*

- When $\text{supp}(\pi)$ is the unit box, for any $0 < \epsilon < 1$ and any $\rho \geq \frac{c \log n}{n}$ where $c > 0$ is a constant, X has a (ρ, ϵ) -summary that keeps $O(\log(1/\rho))$ integers and $O(\frac{1}{\epsilon^2} \log^{d+1} \frac{1}{\rho})$ points of X .
- When $\text{supp}(\pi)$ is the unit ball, for any $0 < \epsilon < 1$ and any $\rho \geq \frac{c \log n}{n}$ where $c > 0$ is a constant, X has a (ρ, ϵ) -summary that keeps $O(\log(1/\rho))$ integers and $O(\frac{1}{\epsilon^2} \cdot (\frac{1}{\rho})^{\frac{d-1}{d+2}} \cdot \log^2 \frac{1}{\rho})$ points of X .

The constant c in the above does not depend on \mathcal{D} , n , ρ , and ϵ .

Proof. We will prove only the case where $\text{supp}(\pi)$ is the unit box because the unit-ball case is similar. Set $\sigma^* = \frac{c \log n}{n}$ where c is some constant to be determined later. Thanks to Theorem 1, it suffices to prove that with probability at least $1 - 1/n^2$, $\theta_X(\rho) = O(\log^{d-1} \frac{1}{\rho})$ at every $\rho \geq \sigma^*$. We will argue that, with probability at least $1 - 1/n^2$, there exists a halfspace $h \in \mathcal{R}$ such that $\bar{X}(h) \leq \rho$ and $\theta_{\mathcal{U}(X)}^h(\rho) = O(\log^{d-1} \frac{1}{\rho})$. Once this is done, we know $\theta_X(\rho) = O(\log^{d-1} \frac{1}{\rho})$ from (6).

Condition C_2 says that the pdf π satisfies $\pi(x) \geq \gamma$ for any point x in $\text{supp}(\pi)$ (i.e., the unit box), where γ is a positive constant. Remember that, by definition of $\text{supp}(\pi)$, $\pi(x) = 0$ for any x outside $\text{supp}(\pi)$.

Simply set h to a halfspace as stated in Lemma 5, i.e., $\theta_{\mathcal{U}}^h(\sigma) = O(\log^{d-1} \frac{1}{\sigma})$. Let $\pi_{\mathcal{U}}$ be the pdf of \mathcal{U} : $\pi_{\mathcal{U}}(x)$ equals 1 if $x \in [0, 1]^d$, or 0 otherwise. Define α as any constant such that $\alpha \leq \gamma$. We have $\alpha \cdot \pi_{\mathcal{U}}(x) \leq \pi(x) \leq 1 \leq \frac{1}{\alpha} \cdot \pi_{\mathcal{U}}(x)$ for all $x \in \mathbb{R}^d$. Given this, Theorem 7.6 of [9] tells us that $\theta_{\mathcal{D}}^h(\sigma) = O(\theta_{\mathcal{U}}^h(\sigma/\alpha))$. It thus follows that $\theta_{\mathcal{D}}^h(\sigma) = O(\log^{d-1} \frac{1}{\sigma})$ for all $\sigma > 0$.

Now, apply Theorem 2 on h by setting $\delta = 1/n^2$ and $\lambda = O(1)$. The theorem shows that, when $n \geq \frac{\beta \cdot \log n}{\rho}$ for some constant β , $\theta_{\mathcal{U}(X)}^h(\rho) \leq 8 \cdot \theta_{\mathcal{D}}^h(\rho) = O(\log^{d-1} \frac{1}{\rho})$ with probability at least $1 - 1/n^2$. We set $c \geq \beta$ to ensure $n \geq \frac{\alpha \cdot \log n}{\rho}$. Note also that the choice of h guarantees $\bar{X}(h) = 0 < \rho$. This makes h a halfspace we are looking for, and concludes the proof. \square

Remark 1 (Composite Distributions). Let \mathcal{D}_1 and \mathcal{D}_2 be two distributions over \mathbb{R}^d with pdfs π_1 and π_2 , respectively (the support regions of π_1 and π_2 may overlap). Define a distribution \mathcal{D} with pdf $\pi(x) = \gamma \cdot \pi_1(x) + (1 - \gamma) \cdot \pi_2(x)$, for some constant $0 < \gamma < 1$. Theorem 7.7 of [9] tells us that, for any halfspace $h \in \mathcal{R}$ and any $\sigma > 0$, $\theta_{\mathcal{D}}^h(\sigma) \leq \theta_{\mathcal{D}_1}^h(\frac{\sigma}{\gamma}) + \theta_{\mathcal{D}_2}^h(\frac{\sigma}{1-\gamma})$.

It thus follows from Lemma 5 that, when \mathcal{D}_1 and \mathcal{D}_2 are atomic distributions with support regions obtainable from the unit box through affine transformations, $\theta_{\mathcal{D}}^h(\sigma) = O(\log^{d-2} \frac{1}{\sigma})$ for any h disjoint with $\text{supp}(\mathcal{D}_1) \cup \text{supp}(\mathcal{D}_2)$. The unit-box result of Theorem 3 can be easily shown to hold on this \mathcal{D} as well, by adapting the proof in a straightforward manner. The same is true for the unit-ball result of Theorem 3. All these results can now be extended to a composite distribution synthesized from a constant number of atomic distributions (see Section 1.2).

Remark 2 (More Distribution with Near-Constant θ). What is given in Lemma 5 is only one scenario where $\theta_{\mathcal{D}}^h(\sigma)$ is nearly a constant. There are other combinations of \mathcal{D} and \mathcal{R} where $\theta_{\mathcal{D}}^h(\sigma) = \tilde{O}(1)$ for all $h \in \mathcal{R}$; see [5, 7, 8, 9, 22] (in some of those combinations, \mathcal{R} may not contain all the halfspaces in \mathbb{R}^d ; e.g., a result of [8] concerns only the halfspaces whose boundary planes pass the origin). The proof of Theorem 3 can be adapted to show that X has a (ρ, ϵ) -summary of size $\tilde{O}(1/\epsilon^2)$ with high probability when $\rho = \Omega(\max\{\frac{\log n}{n}, \min_{h \in \mathcal{R}} \Pr_{\mathcal{D}}(h)\})$.

Remark 3 (Time Complexity). In general, for any X , a (ρ, ϵ) -summary (for the halfspace family \mathcal{R} in constant-dimensional space) can be found in polynomial time even by implementing the algorithm of Section 3.1 naively. The time can be improved to $O(n \text{ polylog } n) + n^{1-\Omega(1)} \cdot s^{O(1)}$, where s is the size of the returned summary, by utilizing specialized data structures [4, 18].

6 A Lower Bound with Disagreement Coefficients

In this section, we will prove a lower bound on the sizes of $(\rho, 1/2)$ -summaries in relation to disagreement coefficients. Our core result is:

Theorem 4. *Let \mathcal{R} be the family of all halfplanes in \mathbb{R}^2 . Fix any integer w as the number of bits in a word. Choose arbitrary integers η , q , and k such that $\eta \geq 4$, q is a multiple of η , and $1 \leq k \leq q/(4\eta)$. There must exist a set \mathcal{C} of range spaces $(X, \mathcal{R}|_X)$, each satisfying the following conditions:*

- X is a set of $q + k$ points in \mathbb{R}^2 .
- The disagreement coefficient of $(X, \mathcal{R}|_X)$ satisfies $\theta_X(\frac{k}{q+k}) = \frac{k+q}{k+q/\eta}$.




Figure 2: A set of points in \mathcal{X} ($\eta = 4$)

- Any encoding, which encodes a $(\frac{k}{q+k}, 1/2)$ -summary for each range space in \mathcal{C} , must use at least $\eta \cdot w$ bits on at least one range space in \mathcal{C} .

Therefore, for $\rho = \frac{k}{q+k}$, if one wishes to store a $(\rho, 1/2)$ -summary for each range space in \mathcal{C} , at least $\eta \cdot w$ bits (namely, η words) are needed on at least one range space. Since $\theta_X(\rho) = \frac{k+q}{k+q/\eta} \leq \eta$, this establishes $\theta_X(\rho)$ as a space lower bound for $(\rho, 1/2)$ -summaries. In the theorem, any X has dimensionality $d = 2$ and any (X, \mathcal{R}) has VC-dimension at most 3; hence, Theorem 1 is tight up to polylog factors on constant λ and ϵ .

The flexibility of η , q , and k allows the lower bound to hold in a variety of more concrete settings. For example, by adjusting k and q , one sees that $\theta_X(\rho)$ is a lower bound for the whole range of $\rho \in (0, O(1)]$. On the other hand, by focusing on any specific $\rho \in (0, O(1)]$ but adjusting η , one sees that $\theta_X(\rho)$ remains as a lower bound when $\theta_X(\rho)$ goes from $O(1)$ to $\Omega(1/\rho)$.¹

Proof of Theorem 4. Fix integers η , q , and k as stated in Theorem 4. Define $n = q + k$. Next, we construct a class \mathcal{X} of point sets, each consisting of n points in \mathbb{R}^2 . First, place k points at coordinates $(0, \infty)$. Call them the *outer* points; they belong to all the sets in \mathcal{X} .

For each set $X \in \mathcal{X}$, we generate q extra *inner* points. For this purpose, place an arbitrary polygon Ψ with η vertices, making sure that all the vertices fall on the upper arc of the unit circle (i.e., the arc is $\{(x[1], x[2]) \mid x[1]^2 + x[2]^2 = 1 \text{ and } x[2] \geq 0\}$). Then, given each vertex of Ψ , we add q/η inner points to X , all of which are located at that vertex. See Figure 2 for an example with $\eta = 4$. This finishes the construction of X . It is important to note that a different Ψ is used for each X . Thus, \mathcal{X} includes an infinite number of inputs, each corresponding to a possible Ψ . Our construction ensures a nice property:

Lemma 7. *Fix any $X \in \mathcal{X}$. Given any $(k/n, 1/2)$ -summary of X , we are able to infer all the vertices of Ψ used to construct X .*

Proof. We say that a halfplane in \mathcal{R} is *upward* if it covers the point $(0, \infty)$. Our aim is to prove that, the summary allows us to determine whether an arbitrary upward halfplane covers any inner point of X . This implies that we can reconstruct all the vertices of Ψ using the summary.²

Given an upward halfplane h , we use the summary to obtain an estimate — denoted as τ — of $\overline{X}(h)$. If $\tau \geq 2k/n$, we return “yes” (i.e., h covers at least one inner point); otherwise, we

¹This rules out, for example, a claim of the form: “when $\theta_X(\rho) \geq \sqrt{1/\rho}$, there is a $(\rho, 1/2)$ -summary of size $O(\sqrt{\theta_X(\rho)})$ ”.

²To see this, consider any vertex v of Ψ , and use the summary to distinguish the line ℓ tangent to the arc at v and a line that is parallel to ℓ , but moves slightly away from the arc.

return “no”. To see that this is correct, first note that $\overline{X}(h)$ must be at least k/n , and hence $0.5\overline{X}(h) \leq \tau \leq 1.5\overline{X}(h)$. Therefore, if h covers no inner points, $\overline{X}(h) = k/n$, indicating $\tau < 1.5k/n$. Otherwise, $\overline{X}(h) \geq \frac{k+q/\eta}{n} \geq 5k/n$, indicating $\tau \geq 2.5k/n$. \square

We prove in the appendix:

Lemma 8. *For each $X \in \mathcal{X}$, the disagreement coefficient of $(X, \mathcal{R}|_X)$ satisfies $\theta_X(k/n) = \frac{n}{k+q/\eta}$.*

The set \mathcal{C} is simply the set $\{(X, \mathcal{R}) \mid X \in \mathcal{X}\}$. Recall that each $X \in \mathcal{X}$ corresponds to a distinct η -vertex polygon Ψ . Hence, by Lemma 7, the $(k/n, 1/2)$ -summaries associated with the range spaces in \mathcal{C} serve as an encoding of all such Ψ 's.

So far the number of Ψ 's is infinite, which does not fit the purpose of arguing for a space lower bound in RAM with a finite word length w . This can be easily fixed by creating 2^w choices for each vertex of Ψ , such that each of the η vertices can independently take a choice of its own. This generates $2^{\eta w}$ polygons for Ψ , and hence, the same number of inputs in \mathcal{C} . Lemmas 7 and 8 are still valid. Therefore, any encoding, which encodes a $(k/n, 1/2)$ -summary for each range space in \mathcal{C} , can be used to distinguish all those $2^{\eta w}$ choices of Ψ . The encoding, therefore, must use $\eta \cdot w$ bits for at least one range space. This completes the proof of Theorem 4.

Appendix

A Proof of Lemma 1

Proof of the First Statement. We will prove that, if a range $h \in \mathcal{R}_{i-1}$ satisfies $\overline{X}(h) \geq (1+\epsilon)/2^i$, then $\overline{S}_i(h) \cdot |X_i| + m_i \geq n/2^i$; hence, h cannot belong to \mathcal{R}_i .

Consider first the case where $\overline{X}_i(h) < \rho_i$. Since S_i is a $(\rho_i, \epsilon/4)$ -approximation of X_i , we know

$$\begin{aligned} \overline{S}_i(h) &\geq \overline{X}_i(h) - \rho_i \cdot \epsilon/4 = (\overline{X}(h) \cdot n - m_i)/|X_i| - \rho_i \cdot \epsilon/4 \\ &\geq \left(\frac{n(1+\epsilon)}{2^i} - m_i \right) \frac{1}{|X_i|} - \frac{n(1+\epsilon) \cdot \epsilon/4}{2^i \cdot |X_i|} \\ &= \frac{n(1+\epsilon)(1-\epsilon/4)}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} > \frac{n}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} \end{aligned}$$

as desired. Similarly, for the other case where $\overline{X}_i(h) \geq \rho_i$:

$$\begin{aligned} \overline{S}_i(h) &\geq \overline{X}_i(h)(1-\epsilon/4) = (\overline{X}(h) \cdot n - m_i)(1-\epsilon/4)/|X_i| \\ &\geq \left(\frac{n(1+\epsilon)}{2^i} - m_i \right) \frac{1-\epsilon/4}{|X_i|} \\ &> \frac{n(1+\epsilon)(1-\epsilon/4)}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} > \frac{n}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|}. \end{aligned}$$

Proof of the Second Statement. We will prove that every range $h \in \mathcal{R}_{i-1} \setminus \mathcal{R}_i$ satisfies $\overline{X}(h) \geq (1-\epsilon)/2^i$. Combining this with the fact that $\mathcal{R}_i \subseteq \mathcal{R}_{i-1}$ will establish the second bullet.

Given an $h \in \mathcal{R}_{i-1} \setminus \mathcal{R}_i$ but $\overline{X}(h) < (1-\epsilon)/2^i$, we will show that $\overline{S}_i(h) \cdot |X_i| + m_i < n/2^i$. Consider first the case where $\overline{X}_i(h) < \rho_i$. As S_i is a $(\rho_i, \epsilon/4)$ -approximation of X_i , we know

$$\begin{aligned} \overline{S}_i(h) &\leq \overline{X}_i(h) + \rho_i \cdot \epsilon/4 = (\overline{X}(h) \cdot n - m_i)/|X_i| + \rho_i \cdot \epsilon/4 \\ &\leq \left(\frac{n(1-\epsilon)}{2^i} - m_i \right) \frac{1}{|X_i|} + \frac{n(1-\epsilon) \cdot \epsilon/4}{2^i \cdot |X_i|} \\ &= \frac{n(1-\epsilon+\epsilon/4+\epsilon^2/4)}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} < \frac{n}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} \end{aligned}$$

as desired. Similarly, for the other case where $\overline{X}_i(h) \geq \rho_i$:

$$\begin{aligned} \overline{S}_i(h) &\leq \overline{X}_i(h)(1+\epsilon/4) = (\overline{X}(h) \cdot n - m_i)(1+\epsilon/4)/|X_i| \\ &\leq \left(\frac{n(1-\epsilon)}{2^i} - m_i \right) \frac{1+\epsilon/4}{|X_i|} \\ &= \frac{n(1-\epsilon)(1+\epsilon/4)}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|} < \frac{n}{2^i \cdot |X_i|} - \frac{m_i}{|X_i|}. \end{aligned}$$

B Correctness of the Estimation Algorithm in Section 3.1.2

Proposition 1. For any $h \in \mathcal{R}$ with $\overline{X}(h) \geq \rho$, our estimate $\overline{S}_j(h) \frac{|X_j|}{n} + \frac{m_j}{n}$ falls between $\overline{X}(h)(1-\epsilon)$ and $\overline{X}(h)(1+\epsilon)$.

Proof. It is equivalent to show:

$$\begin{aligned} \left| \overline{S_j}(h) - \frac{\overline{X}(h) \cdot n - m_j}{|X_j|} \right| &\leq \frac{n \cdot \overline{X}(h) \cdot \epsilon}{|X_j|} \\ \Leftrightarrow |\overline{S_j}(h) - \overline{X_j}(h)| &\leq \frac{n \cdot \overline{X}(h) \cdot \epsilon}{|X_j|} \end{aligned} \quad (14)$$

Consider first the case where $\overline{X_j}(h) < \rho_j$. Since S_j is a $(\rho_j, \epsilon/4)$ -approximation of X_j , we know

$$|\overline{S_j}(h) - \overline{X_j}(h)| \leq \rho_j \cdot \epsilon/4 = \frac{n(1+\epsilon) \cdot \epsilon/4}{2^j \cdot |X_j|} \quad (15)$$

By comparing (14) and (15), we can see that it suffices to prove

$$\overline{X}(h) \geq \frac{1+\epsilon}{4 \cdot 2^j}. \quad (16)$$

For this purpose, we distinguish two cases:

- If $j = t$, then $\overline{X}(h) \geq \rho > 1/2^{t-1}$ (by definition of t), giving $\overline{X}(h) > (1+\epsilon)/(4 \cdot 2^t)$.
- Otherwise, by Lemma 1, we know $\overline{X}(h) \geq (1-\epsilon)/2^{j-1}$ which is at least $(1+\epsilon)/(4 \cdot 2^j)$ for $\epsilon \leq 1/3$.

For the other case where $\overline{X_j}(h) \geq \rho_j$, we know

$$|\overline{S_j}(h) - \overline{X_j}(h)| \leq \overline{X_j}(h) \cdot \epsilon/4 \quad (17)$$

By comparing (14) and (17), we can see that it suffices to prove

$$\overline{X_j}(h) \cdot |X_j| \leq 4n \cdot \overline{X}(h)$$

which is obviously true because $\overline{X_j}(h) \cdot |X_j| = |h \cap X_j| \leq |h \cap X| = n \cdot \overline{X}(h)$. \square

Proposition 2. For any $h \in \mathcal{R}$ with $\overline{X}(h) < \rho$, our estimate $\overline{S_j}(h) \frac{|X_j|}{n} + \frac{m_j}{n}$ falls between $\overline{X}(h) - \epsilon\rho$ and $\overline{X}(h) + \epsilon\rho$.

Proof. Similar to the proof of the previous proposition, it is equivalent to show:

$$|\overline{S_j}(h) - \overline{X_j}(h)| \leq \frac{\epsilon\rho \cdot n}{|X_j|} \quad (18)$$

Consider first the case where $\overline{X_j}(h) < \rho_j$. Since S_j is a $(\rho_j, \epsilon/4)$ -approximation of X_j , we know

$$|\overline{S_j}(h) - \overline{X_j}(h)| \leq \rho_j \cdot \epsilon/4 = \frac{n(1+\epsilon) \cdot \epsilon/4}{2^j \cdot |X_j|} \quad (19)$$

By comparing (18) and (19), we can see that it suffices to prove

$$\rho \geq \frac{1+\epsilon}{4 \cdot 2^j}.$$

The argument for (16) in the proof of Proposition 1 still holds here. Thus, the above follows from $\rho > \overline{X}(h)$.

For the other case where $\overline{X_j}(h) \geq \rho_j$, we know

$$|\overline{S_j}(h) - \overline{X_j}(h)| \leq \overline{X_j}(h) \cdot \epsilon/4 \quad (20)$$

By comparing (18) and (20), we can see that it suffices to prove

$$\overline{X_j}(h) \cdot |X_j| \leq 4n\rho$$

which is true because $\overline{X_j}(h) \cdot |X_j| = |h \cap X_j| \leq |h \cap X| = n \cdot \overline{X}(h) < n\rho$. \square

C Proof of Lemma 5

For any $0 < r < 1$, define $I(r)$ as the intersection between the region $\{x \in \mathbb{R}^d \mid \prod_{i=1}^d x[i] \leq r\}$ and the unit box. We will start by proving a geometric fact:

Proposition 3. *For any $0 < r < 1$, $I(r)$ has a volume bounded by $O(r \cdot \log^{d-1} \frac{1}{r})$.*

Proof. We partition $I(r)$ into two parts:

1. $[r, 1]^d \cap \{x \in \mathbb{R}^d \mid \prod_{i=1}^d x[i] \leq r\}$
2. $([0, 1]^d \setminus [r, 1]^d) \cap \{x \in \mathbb{R}^d \mid \prod_{i=1}^d x[i] \leq r\}$

We will prove that both parts have volume $O(r \cdot \log^{d-1} \frac{1}{r})$. In fact, this is obvious for the second part, because its volume is bounded by the volume of $[0, 1]^d \setminus [r, 1]^d$, which is $1^d - (1 - r)^d = O(r)$.

The volume of the first part can be calculated explicitly as

$$\begin{aligned}
& \int_r^1 \dots \int_r^1 \left(\int_r^{r / \prod_{i=1}^{d-1} x[i]} \mathbf{d}x[d] \right) \mathbf{d}x[d-1] \dots \mathbf{d}x[1] \\
&= \int_r^1 \dots \int_r^1 \left(r / \prod_{i=1}^{d-1} x[i] - r \right) \mathbf{d}x[d-1] \dots \mathbf{d}x[1] \\
&\leq r \int_r^1 \dots \int_r^1 \frac{1}{\prod_{i=1}^{d-1} x[i]} \mathbf{d}x[d-1] \dots \mathbf{d}x[1] \\
&= r \int_r^1 \dots \int_r^1 \left(\frac{1}{\prod_{i=1}^{d-2} x[i]} \int_r^1 \frac{1}{x[d-1]} \mathbf{d}x[d-1] \right) \mathbf{d}x[d-2] \dots \mathbf{d}x[1] \\
&= r \int_r^1 \dots \int_r^1 \left(\frac{1}{\prod_{i=1}^{d-2} x[i]} \cdot \ln \frac{1}{r} \right) \mathbf{d}x[d-2] \dots \mathbf{d}x[1] \\
&= r \ln \frac{1}{r} \int_r^1 \dots \int_r^1 \frac{1}{\prod_{i=1}^{d-2} x[i]} \mathbf{d}x[d-2] \dots \mathbf{d}x[1] \\
&\dots \\
&= r \ln^{d-1} \frac{1}{r}.
\end{aligned}$$

□

Let p be one of the 2^d corners of the unit box. We say that a halfspace $h \in \mathcal{R}$ is *monotone* with respect to p if h satisfies the following conditions: (i) h has a non-empty intersection with the unit box, and (ii) for every point x covered by h in the unit box, h must contain the entire axis-parallel rectangle with p and x as opposite corners. The fact below is rudimentary:

Proposition 4. *If h has a non-empty intersection with the unit box, then h must be monotone to at least one corner of the unit box.* □

Define \mathcal{R}_1 to be the subset of the halfspaces in \mathcal{R} that are monotone to the origin of \mathbb{R}^d . With respect to the uniform distribution \mathcal{U} inside the unit box, we have:

Proposition 5. For any $0 < r < 1$,

$$[0, 1]^d \cap \left(\bigcup_{h \in \mathcal{R}_1: \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \quad (21)$$

is a subset of $I(r)$.

Proof. We will prove that if a point $x \in [0, 1]^d$ is outside $I(r)$, it does not belong to (21). As $x \notin I(r)$, x satisfies $\prod_{i=1}^d x[i] > r$. By monotonicity, any $h \in \mathcal{R}_1$ containing x must cover the entire axis-parallel rectangle that has x and the origin as opposite corners. Since this rectangle has volume $\prod_{i=1}^d x[i]$, we know $\mathbf{Pr}_{\mathcal{U}}(h) \geq \prod_{i=1}^d x[i] > r$. Hence, if $\mathbf{Pr}_{\mathcal{U}}(h) \leq r$, h must not contain x . It thus follows that x is not in (21). \square

We are now ready to prove Lemma 5. Consider any halfspace h^* that is disjoint with the unit box. We need to show that $\theta_{\mathcal{U}}^{h^*}(\sigma) = O(\log^{d-1} \frac{1}{\sigma})$ for all $\sigma > 0$.

For each $i \in [1, 2^d]$, define \mathcal{R}_i as the subset of the halfspaces in \mathcal{R} that are monotone to the i -th corner of the unit box (with the 1st corner being the origin). For any halfspace h with a non-empty intersection with the unit box, $\mathbf{Pr}_{\mathcal{U}}[DIS(\{h^*, h\})] = \mathbf{Pr}_{\mathcal{U}}(h)$. Hence, $B_{\mathcal{U}}(h^*, r) = \{h \in \mathcal{R} \mid 0 \leq \mathbf{Pr}_{\mathcal{U}}(h) \leq r\}$. This further indicates that

$$\begin{aligned} [0, 1]^d \cap DIS(B_{\mathcal{U}}(h^*, r)) &= [0, 1]^d \cap DIS(\{h \in \mathcal{R} \mid 0 \leq \mathbf{Pr}_{\mathcal{U}}(h) \leq r\}) \\ &= [0, 1]^d \cap \left(\bigcup_{h \in \mathcal{R} : 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right). \end{aligned}$$

By (3), we can compute $\theta_{\mathcal{U}}^{h^*}(\sigma)$ as

$$\begin{aligned} \theta_{\mathcal{U}}^{h^*}(\sigma) &= \max \left\{ 1, \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{U}} \left[[0, 1]^d \cap \left(\bigcup_{h \in \mathcal{R} : 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \right]}{r} \right\} \\ \text{(by Proposition 4)} &\leq \max \left\{ 1, \sum_{i=1}^{2^d} \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{U}} \left[[0, 1]^d \cap \left(\bigcup_{h \in \mathcal{R}_i : \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \right]}{r} \right\} \\ \text{(by symmetry)} &= \max \left\{ 1, 2^d \cdot \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{U}} \left[[0, 1]^d \cap \left(\bigcup_{h \in \mathcal{R}_1 : \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \right]}{r} \right\} \\ \text{(by Proposition 5)} &\leq \max \left\{ 1, 2^d \cdot \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{U}}[I(r)]}{r} \right\} \\ \text{(by Proposition 3)} &\leq \max \left\{ 1, O \left(\sup_{r > \sigma} \frac{r \log^{d-1} \frac{1}{r}}{r} \right) \right\} = O \left(\log^{d-1} \frac{1}{\sigma} \right). \end{aligned}$$

D Proof of Lemma 6

For any $0 < r < 1$, define $b(r)$ as the ball $\{x \in \mathbb{R}^d \mid \sum_{i=1}^d x[i]^2 \leq r^2\}$; notice that $b(1)$ is simply the unit ball. We will establish the following fact at the end of the proof:




Figure 3: A 2D spherical cap

Proposition 6. *For any $0 < r < 1$, the volume of*

$$b(1) \cap \left(\bigcup_{h \in \mathcal{R} : 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \quad (22)$$

is $O(r^{\frac{2}{d+1}})$.

Consider any halfspace h^* disjoint with $b(1)$. We need to show that $\theta_{\mathcal{U}}^{h^*}(\sigma) = O((\frac{1}{\sigma})^{\frac{d-1}{d+1}})$ for all $\sigma > 0$. For any halfspace h with a non-empty intersection with $b(1)$, $\mathbf{Pr}_{\mathcal{U}}[DIS(\{h^*, h\})] = \mathbf{Pr}_{\mathcal{U}}(h)$. Hence, $B_{\mathcal{U}}(h^*, r) = \{h \in \mathcal{R} \mid 0 \leq \mathbf{Pr}_{\mathcal{U}}(h) \leq r\}$, leading to

$$\begin{aligned} b(1) \cap DIS(B_{\mathcal{U}}(h^*, r)) &= b(1) \cap DIS(\{h \in \mathcal{R} \mid 0 \leq \mathbf{Pr}_{\mathcal{U}}(h) \leq r\}) \\ &= b(1) \cap \left(\bigcup_{h \in \mathcal{R} : 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right). \end{aligned}$$

By (3), $\theta_{\mathcal{U}}^{h^*}(\sigma)$ can now be derived as

$$\begin{aligned} \theta_{\mathcal{U}}^{h^*}(\sigma) &= \max \left\{ 1, \sup_{r > \sigma} \frac{\mathbf{Pr}_{\mathcal{U}} \left[b(1) \cap \left(\bigcup_{h \in \mathcal{R} : 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h \right) \right]}{r} \right\} \\ \text{(by Proposition 6)} &= \max \left\{ 1, O \left(\sup_{r > \sigma} \frac{r^{\frac{2}{d+1}}}{r} \right) \right\} = O \left(\left(\frac{1}{\sigma} \right)^{\frac{d-1}{d+1}} \right). \end{aligned}$$

Proof of Proposition 6. Given a halfspace h , define $\mathcal{C}(h)$ as the “spherical cap” which is the intersection between h and the unit ball. Furthermore, define the *height* of $\mathcal{C}(h)$ as the maximum Euclidean distance between a point in $\mathcal{C}(h)$ and the boundary plane of h . See Figure 3 for an illustration.

Proposition 7. *For any halfspace h , the volume of $\mathcal{C}(h)$ is $\Omega(r^{\frac{d+1}{2}})$, where r is the height of $\mathcal{C}(h)$.*

Proof. Due to symmetry, we will prove the statement only for the halfspace $h = \{x \in \mathbb{R}^d \mid x[d] \geq 1 - r\}$. Note that $\mathcal{C}(h)$ has height r .

Recall that, given a $(d-1)$ -dimensional ball b' in \mathbb{R}^d (called the “base”) and a point $p \in \mathbb{R}^d$ (called the “apex vertex”), a cone in \mathbb{R}^d is defined as the union of the line segments connecting p to all the points in b' . Consider the cone whose (i) base is the $(d-1)$ -dimensional ball corresponding to the intersection of $b(1)$ and the plane $\{x \in \mathbb{R}^d \mid x[d] = 1-r\}$, and (ii) its apex vertex is $(0, \dots, 0, 1)$. The cone has a volume $\Theta(r^{\frac{d+1}{2}})$. As the cone is contained by $\mathcal{C}(r)$, the volume of $\mathcal{C}(r)$ is $\Omega(r^{\frac{d+1}{2}})$. \square

Let h be a halfspace satisfying $0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r$. The spherical cap $\mathcal{C}(h)$ has a volume that equals $\mathbf{Pr}_{\mathcal{U}}(h)$ multiplied by the volume of the unit ball. By Proposition 7, the height of $\mathcal{C}(h)$ is at most $c \cdot \mathbf{Pr}_{\mathcal{U}}(h)^{\frac{2}{d+1}}$ for some constant $c > 0$. Hence, every point in h has a distance at least $1 - c \cdot \mathbf{Pr}_{\mathcal{U}}(h)^{\frac{2}{d+1}} \geq 1 - c \cdot r^{\frac{2}{d+1}}$ from the origin.

We now know that $\bigcup_{h: 0 < \mathbf{Pr}_{\mathcal{U}}(h) \leq r} h$ — the region given in (22) — is fully contained in the d -dimensional annulus $b(1) \setminus b(1 - c \cdot r^{\frac{2}{d+1}})$. The annulus has volume $O(r^{\frac{2}{d+1}})$, thus completing the proof of Proposition 6.

E Proof of Lemma 8

Let Ψ be the polygon used to generate X . By Definition 5, to compute $\theta_X(k/n)$, we should focus on the halfplanes $h \in \mathcal{R}$ with $\overline{X}(h) \leq k/n$. Such halfplanes belong to one of the following categories:

1. h covers the k outer points, but not any inner points of X .
2. h does not cover any points of X at all.

Category 1. Consider an h of the first category. By (4), to derive $\theta_{\mathcal{U}(X)}^h(k/n)$, we need to calculate the ratio $\frac{1}{r} \cdot \mathbf{Pr}_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ at each $r > k/n$.

- Case 1: $r \geq \frac{q/\eta}{n}$. Let h' be any upward halfplane covering exactly one vertex v of Ψ . $DIS(\{h, h'\}) \cap X$ is the set of q/η inner points at v . Therefore, $h' \in B_{\mathcal{U}(X)}(h, r)$, which indicates that $DIS(B_{\mathcal{U}(X)}(h, r))$ covers all the inner points of X .

Consider now a halfplane $h'' \in \mathcal{R}$ that does not cover any points of X . $DIS(\{h, h''\}) \cap X$ is the set of k outer points. Since $k < q/\eta$, $h'' \in B_{\mathcal{U}(X)}(h, r)$, which indicates that $DIS(B_{\mathcal{U}(X)}(h, r))$ covers all the outer points of X as well.

Therefore, $\mathbf{Pr}_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))] = 1$. This means that $\frac{1}{r} \cdot \mathbf{Pr}_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ is maximized at $\frac{n}{q/\eta}$ for $r \geq \frac{q/\eta}{n}$.

- Case 2: $\frac{k}{n} < r < \frac{q/\eta}{n}$. Consider an arbitrary halfplane $h' \in B_{\mathcal{U}(X)}(h, r)$. Since h' cannot cover any vertex of Ψ ,³ $DIS(B_{\mathcal{U}(X)}(h, r))$ contains no inner points. On the other hand, h' may be a halfplane that covers no points of X at all. Hence, $DIS(B_{\mathcal{U}(X)}(h, r))$ is precisely the set of k outer points. Therefore, the supremum of $\frac{1}{r} \cdot \mathbf{Pr}_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ is 1 when $\frac{k}{n} < r < \frac{q/\eta}{n}$.

We conclude that $\theta_{\mathcal{U}(X)}^h(k/n) = \frac{n}{q/\eta}$ for any h in Category 1.

Category 2. Consider an h of the secondary category. To compute $\theta_{\mathcal{U}(X)}^h(k/n)$, we distinguish three cases:

³Otherwise, h' contains at least q/η inner points, giving $\mathbf{Pr}_{\mathcal{U}(X)}[DIS(\{h, h'\})] \geq \frac{q/\eta}{n}$ and contradicting $h' \in B_{\mathcal{U}(X)}(h, r)$.

- Case 1: $r \geq \frac{k}{n} + \frac{q/\eta}{n}$. Let h' be any upward halfplane covering exactly one vertex v of Ψ . $DIS(\{h, h'\}) \cap X$ is the set of q/η inner points at v , plus the k outer points. Therefore, $h' \in B_{\mathcal{U}(X)}(h, r)$. This indicates that $DIS(B_{\mathcal{U}(X)}(h, r))$ is the entire X . Therefore, $\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))] = 1$; and $\frac{1}{r} \cdot \Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ is maximized at $\frac{n}{k+q/\eta}$ for $r \geq \frac{k}{n} + \frac{q/\eta}{n}$.
- Case 2: $\frac{q/\eta}{n} \leq r < \frac{k}{n} + \frac{q/\eta}{n}$. Consider an arbitrary halfplane $h' \in B_{\mathcal{U}(X)}(h, r)$. Since h' may be a halfplane that covers all the k outer points but no inner points, $DIS(B_{\mathcal{U}(X)}(h, r))$ contains all the k outer points.

Note that h' can cover at most one vertex of Ψ . Furthermore, if it covers a vertex of Ψ , it cannot contain any outer points. Recall that all the vertices of Ψ are on the upper arc of the unit circle. Hence, if h' covers a single vertex of Ψ , the vertex must either be the leftmost vertex v_{left} of Ψ , or the rightmost one v_{right} .⁴ This indicates that $DIS(B_{\mathcal{U}(X)}(h, r))$ contains $2q/\eta$ inner points: q/η ones placed at v_{left} and v_{right} , respectively.

Therefore, $\Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))] = \frac{k+2q/\eta}{n}$. Hence, $\frac{1}{r} \cdot \Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ is maximized at $\frac{k+2q/\eta}{q/\eta}$ when $\frac{q/\eta}{n} \leq r < \frac{k}{n} + \frac{q/\eta}{n}$.

- Case 3: $\frac{k}{n} < r < \frac{q/\eta}{n}$. Consider an arbitrary halfplane $h' \in B_{\mathcal{U}(X)}(h, r)$. Since h' cannot cover any inner points, $DIS(B_{\mathcal{U}(X)}(h, r))$ contains no inner points. On the other hand, h' may be a halfplane that covers all the k outer points, indicating that $DIS(B_{\mathcal{U}(X)}(h, r))$ is the set of k outer points. Therefore, the supremum of $\frac{1}{r} \cdot \Pr_{\mathcal{U}(X)}[DIS(B_{\mathcal{U}(X)}(h, r))]$ is 1 when $\frac{k}{n} < r < \frac{q/\eta}{n}$.

When k, q , and η satisfy the requirements of Theorem 4, $\frac{n}{k+q/\eta} > \frac{k+2q/\eta}{q/\eta} > 1$. We conclude that $\theta_{\mathcal{U}(X)}^h(k/n) = \frac{n}{k+q/\eta}$ for any h in Category 2.

Combining the above discussion of both categories shows that $\theta_X(k/n) = \frac{n}{k+q/\eta}$, as claimed in the lemma.

References

- [1] Kenneth S. Alexander. Rates of growth and sample moduli for weighted empirical processes indexed by sets. *Probability Theory and Related Fields*, 75:379–423, 1987.
- [2] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 36(4):929–965, 1989.
- [3] Hervé Brönnimann, Bernard Chazelle, and Jiří Matoušek. Product range spaces, sensitive sampling, and derandomization. *SIAM Journal on Computing*, 28(5):1552–1575, 1999.
- [4] Timothy M. Chan. Optimal partition trees. *Discrete & Computational Geometry*, 47(4):661–690, 2012.
- [5] Ran El-Yaniv and Yair Wiener. Active learning via perfect selective classification. *Journal of Machine Learning Research*, 13:255–279, 2012.

⁴Suppose that the (only) vertex v of Ψ covered by h' is neither v_{left} nor v_{right} . Then, the boundary of h' must intersect both the segment connecting v_{left} , v , and the segment connecting v_{right} , v . This suggests that h' must be upward, but this contradicts the fact that h' covers no outer points.

- [6] Esther Ezra. Small-size relative (p, ϵ) -approximations for well-behaved range spaces. In *Proceedings of Symposium on Computational Geometry (SoCG)*, pages 233–242, 2013.
- [7] Wayne A. Fuller. *Sampling Statistics*. Wiley, 2009.
- [8] Steve Hanneke. A bound on the label complexity of agnostic active learning. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 353–360, 2007.
- [9] Steve Hanneke. Theory of active learning. *Manuscript downloadable at <http://www.stevehanneke.com>*, 2014.
- [10] Sarel Har-Peled, Haim Kaplan, Micha Sharir, and Shakhar Smorodinsky. Epsilon-nets for halfspaces revisited. *CoRR*, abs/1410.3154, 2014.
- [11] Sarel Har-Peled and Micha Sharir. Relative (p, ϵ) -approximations in geometry. *Discrete & Computational Geometry*, 45(3):462–496, 2011.
- [12] David Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Inf. Comput.*, 100(1):78–150, 1992.
- [13] David Haussler and Emo Welzl. Epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.
- [14] János Komlós, János Pach, and Gerhard J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.
- [15] Andrey Kupavskii, Nabil H. Mustafa, and János Pach. New lower bounds for epsilon-nets. In *Proceedings of Symposium on Computational Geometry (SoCG)*, pages 54:1–54:16, 2016.
- [16] Andrey Kupavskii and Nikita Zhivotovskiy. When are epsilon-nets small? *CoRR*, abs/1711.10414, 2017.
- [17] Yi Li, Philip M. Long, and Aravind Srinivasan. Improved bounds on the sample complexity of learning. *Journal of Computer and System Sciences (JCSS)*, 62(3):516–527, 2001.
- [18] Jiří Matoušek. Efficient partition trees. *Discrete & Computational Geometry*, 8:315–334, 1992.
- [19] Jiří Matoušek, Raimund Seidel, and Emo Welzl. How to net a lot with little: Small epsilon-nets for disks and halfspaces. In *Proceedings of Symposium on Computational Geometry (SoCG)*, pages 16–22, 1990.
- [20] János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Proceedings of Symposium on Computational Geometry (SoCG)*, pages 458–463, 2011.
- [21] D. Pollard. Rates of uniform almost-sure convergence for empirical processes indexed by unbounded classes of functions. *Manuscript*, 1986.
- [22] Liwei Wang. Smoothness, disagreement coefficient, and the label complexity of agnostic active learning. *Journal of Machine Learning Research*, 12:2269–2292, 2011.