

Learning convex bodies is hard

Navin Goyal*
Microsoft Research India
navingo@microsoft.com

Luis Rademacher†
Georgia Tech
lrademac@cc.gatech.edu

Abstract

We show that learning a convex body in \mathbb{R}^d , given random samples from the body, requires $2^{\Omega(\sqrt{d/\epsilon})}$ samples. By learning a convex body we mean finding a set having at most ϵ relative symmetric difference with the input body. To prove the lower bound we construct a hard to learn family of convex bodies. Our construction of this family is very simple and based on error correcting codes.

1 Introduction

We consider the following problem: Given uniformly random points from a convex body in \mathbb{R}^d , we would like to approximately learn the body with as few samples as possible. In this question, and throughout this paper, we are interested in the number of samples but not in the computational requirements for constructing such an approximation. Our main result will show that this needs about $2^{\Omega(\sqrt{d})}$ samples. This problem is a special case of the statistical problem of inferring information about a probability distribution from samples. For example, one can approximate the centroid of the body with a sample of size roughly linear in d . On the other hand, a sample of size polynomial in d is not enough to approximate the volume of a convex body within a constant factor ([3], and see Section 5 here for a discussion). Note that known approximation algorithms for the volume (e.g., [2]) do not work in this setting as they need a membership oracle and random points from various carefully chosen subsets of the input body.

Our problem also relates to work in learning theory (e.g., [12, 8]), where one is given samples generated according to (say) the Gaussian distribution

*Supported by ARC ThinkTank.

†Supported by ARC ThinkTank and NSF.

and each sample is labeled “positive” or “negative” depending on whether it belongs to the body. Aside from different distributions, another difference between the learning setting of [8] and ours is that in ours one gets only positive examples. Klivans et al. [8] give an algorithm and a nearly matching lower bound for learning convex bodies with labeled samples chosen according to the Gaussian distribution. Their algorithm takes time $2^{\tilde{O}(\sqrt{d})}$ and they also show a lower bound of $2^{\Omega(\sqrt{d})}$. Note that our lower bound is not just a simplification of theirs: Neither lower bound seems to imply the other in a simple way.

The problem of learning convex sets from uniformly random samples from them was raised by Frieze et al. [4]. They gave a polynomial time algorithm for learning parallelepipeds. Another somewhat related direction is the work on the learnability of discrete distributions by Kearns et al. [7]. We remark that in the traditional distribution-independent PAC-learning setting, it is not possible to learn convex bodies efficiently because their VC-dimension is infinite.

Our lower bound result (like that of [8]) also allows for membership oracle queries. Note that it is known that estimating the volume of convex bodies requires an exponential number of membership queries if the algorithm is *deterministic* [1], which implies that learning bodies requires an exponential number of membership queries because if an algorithm can learn the body then it can also estimate its volume.

To formally define the notion of learning we need to specify a distance $d(\cdot, \cdot)$ between bodies. A natural choice in our setting is to consider the total variation distance of the uniform distribution on each body (see Section 2).

We will use the term *random oracle of a convex body K* for a black box that when queried outputs a uniformly random point from K .

Theorem 1. *There exists a distribution \mathcal{D} on the set of convex bodies in \mathbb{R}^d satisfying the following: Let ALG be a randomized algorithm that, given a random convex body K according to \mathcal{D} , makes at most q total queries to random and membership oracles of K and outputs a set C such that, for $8/d \leq \epsilon \leq 1/8$,*

$$\Pr(d(C, K) \leq \epsilon) \geq 1/2$$

where the probability is over K , the random sample and any randomization by ALG. Then

$$q \geq 2^{\Omega(\sqrt{d/\epsilon})}.$$

Remarkably, the lower bound of Klivans et al. [8] is numerically essentially identical to ours ($2^{\Omega(\sqrt{d})}$ for ϵ constant). Constructions similar to

theirs are possible for our particular scenario [5]. We believe that our argument is considerably simpler, and elementary compared to that of [8]. Furthermore, our construction of the hard to learn family is explicit. Our construction makes use of error correcting codes. To our knowledge, this connection with error correcting codes is new in such contexts and may find further applications. See Section 5 for some further comparison.

An informal outline of the proof. The idea of the proof is to find a large family of convex bodies in \mathbb{R}^n satisfying two conflicting goals: (1) Any two bodies in the family are almost disjoint; (2) and yet they look alike in the sense that a small sample of random points from any such body is insufficient for determining which one it is because there are many bodies in the family consistent with the sample. Since any two bodies are almost disjoint, even approximating a body would allow one to determine it exactly. This will imply that it is also hard to approximate.

We first construct a family of bodies that although not almost disjoint, have sufficiently large symmetric difference. We will then be able to construct a family with almost disjoint bodies by taking products of bodies in the first family.

The first family is quite natural (it is described formally in Sec. 3.1). Consider the cross polytope O_n in \mathbb{R}^n (generalization of the octahedron to n dimensions: convex hull of the vectors $\{\pm e_i : i \in [n]\}$, where e_i is the unit vector in \mathbb{R}^n with the i th coordinate 1 and the rest 0). A *peak* attached to a facet F of O_n is a pyramid that has F as its base and has its other vertex outside O_n on the normal to F going through its centroid. If the height of the peak is sufficiently small then attaching peaks to any subset of the 2^n facets will result in a convex polytope. We will show later that we can choose the height so that the volume of all the 2^n peaks is $\Omega(1/n)$ fraction of the volume of O_n . We call this family of bodies \mathcal{P} . [We remark that our construction of cross-polytopes with peaks has resemblance to a construction in [10] with different parameters, but there does not seem to be any connection between the problem studied there and the problem we are interested in.]

Intuitively, a random point in a body from this family tells one that if the point is in one of the peaks then that peak is present, otherwise one learns nothing. Therefore if the number of queries is at most a polynomial in n , then one learns nothing about most of the peaks and so the algorithm cannot tell which body it got.

But these bodies do not have large symmetric difference (can be as small

as a $O(1/(n2^n))$ fraction of the cross polytope if the two bodies differ in just one peak) but we can pick a subfamily of them having pairwise symmetric difference at least $\Omega(1/n)$ by picking a large random subfamily. We will do it slightly differently which will be more convenient for the proof: Bodies in \mathcal{P} have one-to-one correspondence with binary strings of length 2^n : each facet corresponds to a coordinate of the string which takes value 1 if that facet has a peak attached, else it has value 0. To ensure that any two bodies in our family differ in many peaks it suffices to ensure that their corresponding strings have large Hamming distance. Large sets of such strings are of course furnished by good error correcting codes.

From this family we can obtain another family of almost disjoint bodies by taking products, while preserving the property that polynomially many random samples do not tell the bodies apart. This product trick (also known as tensoring) has been used many times before, in particular for amplifying hardness, but we are not aware of its use in a setting similar to ours. Our construction of the product family also resembles the operation of concatenation in coding theory.

Acknowledgments. We are grateful to Adam Kalai and Santosh Vempala for useful discussions.

2 Preliminaries

Let $K, L \subseteq \mathbb{R}^n$ be bounded and measurable. We define a distance $\text{dist}(K, L)$ as the total variation distance between the uniform distributions in K and L , that is,

$$\text{dist}(K, L) = \begin{cases} \frac{|K \setminus L|}{|K|} & \text{if } |K| \geq |L| \\ \frac{|L \setminus K|}{|L|} & \text{if } |L| > |K|. \end{cases}$$

We will use $|A|$ to denote the volume of sets $A \subset \mathbb{R}^n$, and also to denote the cardinality of finite sets A ; which one is meant in a particular case will be clear from the context.

Let $\mathbf{1}$ denote the vector $(1, \dots, 1)$. “log” denotes logarithm with base 2.

We will need some basic definitions and facts from coding theory; see, e.g., [11]. For a finite alphabet Σ , and word length n , a *code* C is a subset of Σ^n . For any two codewords $x, y \in C$, distance $\text{dist}(x, y)$ between them is defined by $\text{dist}(x, y) := |\{i \in [n] : x_i \neq y_i\}|$. The *relative minimum distance* for code C is $\min_{x, y \in C, x \neq y} \text{dist}(x, y)/n$. For $\Sigma = \{0, 1\}$, the *weight* of a codeword x is $|\{i \in [n] : x_i \neq 0\}|$. Define $V_q(n, r) := \sum_{i=0}^r \binom{n}{i} (q-1)^i$. The following is well-known and easy to prove:

Theorem 2 (Gilbert–Varshamov). *For alphabet size q , code length n , and minimum distance d , there exists a code of size at least $q^n/V_q(n, d - 1)$.*

When the alphabet is $\Sigma = \{0, 1\}$, we define the *complement* \bar{c} of a codeword $c \in C$ as $\bar{c}_i := 1 - c_i$.

3 A hard to learn family of convex bodies

The construction proceeds in two steps. In the first step we construct a large subfamily of \mathcal{P} such that the relative pairwise symmetric difference between the bodies is $\Omega(1/n)$. This symmetric difference is however not sufficiently large for our lower bound. The second step of the construction amplifies the symmetric difference by considering products of the bodies from the first family.

3.1 The inner family: Cross-polytope with peaks

We first construct a family with slightly weaker properties. The family consists of what we call “cross polytope with peaks”. The n -dimensional cross polytope O_n is the convex hull of the $2n$ points $\{\pm e_i, i \in [n]\}$. Let F be a facet of O_n , and let c_F be the center of F . The *peak* associated to F is the convex hull of F and the point αc_F , where $\alpha > 1$ is a positive scalar defined as follows: α is picked as large as possible so that the union of the cross polytope and all 2^n peaks is a convex body. A *cross polytope with peaks* will then be the union of the cross polytope and any subfamily of the 2^n possible peaks. The set of all 2^{2^n} bodies of this type will be denoted \mathcal{P} . By fixing of an ordering of the facets of the cross polytope, there is a one-to-one correspondence between the cross polytope with peaks and 0–1 vectors with 2^n coordinates.

Let P denote the cross polytope with all 2^n peaks. We will initially choose α as large as possible so that the following condition—necessary for convexity of P but not clearly sufficient—is satisfied: for every pair of adjacent facets F, G of O_n , the vertex of each peak is in the following half-space: the halfspace containing O_n and whose boundary is the hyperplane orthogonal to the (vector connecting the origin to the) center of $F \cap G$, and containing $F \cap G$. A straightforward computation shows that $\alpha = n/(n - 1)$ for this condition. This implies by another easy computation that the volume of all the peaks is $|O_n|/(n - 1)$. We will now show that this weaker condition on α is actually sufficient for the convexity of P and any cross polytope with peaks.

Proposition 3. *Every set in \mathcal{P} is convex.*

Proof. Let Q be the intersection of all halfspaces of the form

$$\{x \in \mathbb{R}^n : a \cdot x \leq 1\}$$

where $a \in \mathbb{R}^n$ is a vector having entries in $\{-1, 0, 1\}$ and exactly one zero entry. Equivalently, the boundary of each such halfspace is a hyperplane orthogonal to the center of some $(n-2)$ -dimensional face of O_n and containing that face.

In the rest of the proof we will show that $P = Q$, which gives the convexity of P . This equality implies that a cross polytope with only some peaks is also convex: any such body can be obtained from P by intersecting P with the halfspaces induced by the facets of O_n associated to the missing peaks, and it is easy to see from the definition of the peaks that each such intersection removes exactly one peak.

It is clear that $P \subseteq Q$. For the other inclusion, let $x \in Q$. By symmetry we can assume $x \geq 0$. If $\sum x_i \leq 1$, then $x \in O_n \subseteq P$. If $\sum x_i > 1$, we will show that x is in the peak of the positive orthant. We would like to write x as a convex combination of e_1, \dots, e_n and the extra vertex of the peak, $v = \mathbf{1}/(n-1)$. Let $\mu = (n-1)((\sum x_i) - 1) > 0$. We want a vector $\lambda = (\lambda_1, \dots, \lambda_n)$ such that x is a convex combination of the vertices of the peak:

$$x = \mu v + \sum \lambda_i e_i = \mu v + \lambda,$$

that is, $\lambda = x + \mathbf{1} - \mathbf{1} \sum x_i$. It satisfies

$$\mu + \sum \lambda_i = 1$$

and $\lambda_i = x_i + 1 - \sum x_i$, and this is non-negative: By definition of Q we have for all $j \in [n]$

$$\sum x_i \leq 1 + x_j.$$

This shows that x belongs to the peak in the positive orthant. \square

For notational convenience we let $N := 2^n$. Recall that we identify bodies in \mathcal{P} with binary strings in $\{0, 1\}^N$. Let $C \subset \{0, 1\}^N$ be a code with relative minimum distance at least $1/4$. To simplify computations involving distance “dist” between bodies, it will be convenient to have the property that all codewords in C have weight $N/2$. We can ensure this easily as follows. Let $\tilde{C} \subset \{0, 1\}^{N/2}$ be a code with relative minimum distance at least $1/4$, then set $C := \{(c, \bar{c}) : c \in \tilde{C}\}$. Clearly $|C| = |\tilde{C}|$. By Theorem 2

we can choose \tilde{C} such that $|\tilde{C}| \geq 2^{c_1 N}$, for a positive constant c_1 . We fix C to be such a code, i.e. a code with relative minimum distance at least $1/4$, size $2^{c_1 N}$, and all codewords with weight $N/2$.

We define the family \mathcal{P}_C as the family consisting of bodies in \mathcal{P} corresponding to codewords in C . As all codewords in C have the same weight, we have that all bodies in \mathcal{P}_C have the same volume. Recall that the volume of each peak is $\frac{|O_n|}{2^n \binom{n-1}{i}}$. Therefore for distinct $P, Q \in \mathcal{P}_C$ the volume of the symmetric difference of P and Q is at least $\frac{|O_n|}{4 \binom{n-1}{i}}$.

3.2 The outer family: The product construction

Let C' be a code with codewords of length k and minimum distance at least $k/2$ on the alphabet \mathcal{P}_C . That is, codewords in C' can be represented as (B_1, \dots, B_k) , where $B_i \in \mathcal{P}_C$ for $i = 1, \dots, k$. The product family $\mathcal{P}_{C'}$ corresponding to code C' , has $|C'|$ bodies in \mathbb{R}^{kn} , one for each codeword. The body for codeword $(B_1, \dots, B_k) \in C'$ is simply $B_1 \times \dots \times B_k$.

Clearly $|\mathcal{P}_{C'}| = |C'|$. Using Theorem 2 we can choose C' such that $|C'| \geq q^k / V_q(k, k/2)$. Now note that

$$\begin{aligned} V_q(k, k/2) &= \sum_{i=0}^{k/2} \binom{k}{i} (q-1)^i \leq (q-1)^{k/2} \sum_{i=0}^{k/2} \binom{k}{i} \\ &< 2^k (q-1)^{k/2} < (4q)^{k/2}. \end{aligned}$$

Therefore $q^k / V_q(k, k/2) > q^k / (4q)^{k/2} = (q/4)^{k/2}$. Setting $q = 2^{c_1 N}$, we get $|C'| > 2^{(c_1 N - 2)k/2} > 2^{c_2 k N}$, for constant $c_2 > 0$, assuming N is sufficiently large. We just showed:

Lemma 4. $|\mathcal{P}_{C'}| > 2^{c_2 k 2^n}$.

The following lemma shows that the bodies in $\mathcal{P}_{C'}$ are almost pairwise disjoint.

Lemma 5. For distinct $A, B \in \mathcal{P}_{C'}$ we have

$$\text{dist}(A, B) = 1 - \frac{|A \cap B|}{|A|} > 1 - e^{-k/(16n)}.$$

Proof. We constructed $\mathcal{P}_{C'}$ so that all bodies in it have the same volume. This implies

$$\text{dist}(A, B) = 1 - \frac{|A \cap B|}{|A|}.$$

Let $A = A_1 \times \dots \times A_k$ and $B = B_1 \times \dots \times B_k$. Then

$$\frac{|A \cap B|}{|A|} = \frac{|A_1 \cap B_1| \times \dots \times |A_k \cap B_k|}{|A_1| \times \dots \times |A_k|}.$$

Since the minimum relative distance in C is at least $1/4$ and the weight of each codeword is $N/2$, we have that for $A_i \neq B_i$ the number of peaks in $A_i \cap B_i$ is at most $2^n \cdot 3/8$. Hence

$$\frac{|A_i \cap B_i|}{|A_i|} \leq \frac{1 + 3/(8(n-1))}{1 + 1/(2(n-1))}.$$

Since the minimum distance of C' is at least $k/2$, we have $A_i \neq B_i$ for at least $k/2$ values of i in $[k]$. Therefore we get

$$\begin{aligned} \frac{|A \cap B|}{|A|} &\leq \left(\frac{1 + 3/(8(n-1))}{1 + 1/(2(n-1))} \right)^{k/2} \\ &\leq \left(1 - \frac{1}{8n} \right)^{k/2} < e^{-k/(16n)}. \end{aligned}$$

□

4 Proof of the lower bound

Proof of Theorem 1. We will make use of the family $\mathcal{P}_C^{C'}$ that we constructed in the previous section. Recall that the bodies in this family live in \mathbb{R}^d , for $d := kn$. For this proof we will think of d as fixed and we will choose n appropriately for the lower bound proof. By a straightforward but tedious argument it is enough to prove the theorem assuming that d is a power of 2.

We will use Yao's principle (see, e.g., [9]). To this end, we will first show that the interaction between an algorithm and the oracles can be assumed to be "discrete", which in turn will imply that effectively there is only a finite number of deterministic algorithms that make at most q queries. The discretization of the oracles also serves a second purpose: that we can see deterministic algorithms as finite decision trees and use counting arguments to show a lower bound on the query complexity.

Fix a body K from $\mathcal{P}_C^{C'}$. Suppose that a randomized algorithm has access to the following discretizations of the oracles:

- A discrete random oracle that generates a random point $X = (X_1, \dots, X_k)$ from $K = \prod_i K_i$ and, for each $i \in [k]$ outputs whether X_i lies in the corresponding cross-polytope or in which peak it lies.

- A discrete membership oracle that when given a sequence of indices of peaks $I = (i_1, \dots, i_k)$ outputs, for each $i \in [k]$, whether peak i is present in K_i .

Claim: A randomized algorithm with access to discrete versions of the oracles can simulate a randomized algorithm with access to continuous oracles with the same number of queries.

Proof of claim: We will show it for bodies in \mathcal{P} , i.e. cross polytopes with peaks; the generalization of this argument to product bodies in $\mathcal{P}_C^{C'}$ is straightforward. Let A and B be the algorithm with access to the continuous and discrete oracles respectively. Algorithm B acts as A , except when A invokes an oracle, where B will do as follows: When A makes a query p to the continuous membership oracle, B will query the peak that contains p (we can assume that p lies in a peak, as otherwise the query provided no new information). Now suppose that A makes a query to the continuous random oracle and gets a point p . Then B makes a query to the discrete random oracle. B then generates a uniformly random point p' in the region that it got from the oracle. Clearly p' has the same distribution as p , namely uniform distribution on the body.

If we see deterministic algorithms as decision trees, it is clear that there are only a finite number of deterministic algorithms that make at most q queries to the discrete oracles of K . Thus, by Yao's principle, for any distribution \mathcal{D} on inputs, the probability of error of any randomized algorithm against \mathcal{D} is at least the probability of error of the best deterministic algorithm against \mathcal{D} .

Our hard input distribution \mathcal{D} is the uniform distribution over $\mathcal{P}_C^{C'}$. Now, in the decision tree associated to a deterministic algorithm, each node associated to a membership query has two children (either the query point is in the body or not), while a node associated to a random sample has at most $(2^n + 1)^k$ children (the random sample can lie in one of the 2^n peaks or in O_n , for each factor in the product body). Thus, if the algorithm makes at most q queries in total, then the decision tree has at most $(2^n + 1)^{kq}$ leaves. These leaves induce a partition of the family of inputs $\mathcal{P}_C^{C'}$. By Lemma 5, the distance between any pair of bodies is at least $1 - e^{-k/16n} = 1 - e^{-d/16n^2}$, where n is chosen so that the output of the algorithm can be within ϵ of at most one body in each part of the partition. That is,

$$2\epsilon < 1 - e^{-k/16n},$$

which implies that we should take

$$n < 4\sqrt{\frac{d}{\ln \frac{1}{1-2\epsilon}}}.$$

As $d = kn$ is a power of 2, we can satisfy the previous inequality and the integrality constraints of k and n by using our assumption that $8/d \leq \epsilon \leq 1/8$ and letting n be a power of 2 such that

$$2 \leq \sqrt{\frac{d}{\ln \frac{1}{1-2\epsilon}}} \leq n < 4\sqrt{\frac{d}{\ln \frac{1}{1-2\epsilon}}} \leq d.$$

By Lemma 4, the total number of bodies is

$$|\mathcal{P}_C^{C'}| \geq 2^{c_2 k 2^n}.$$

This implies that the probability of error is at least

$$1 - \frac{(2^n + 1)^{kq}}{|\mathcal{P}_C^{C'}|} \geq 1 - \frac{(2^n + 1)^{kq}}{2^{c_2 k 2^n}}.$$

If we want this error to be less than a given δ , then for some $c_3, c_4 > 0$ we need

$$\begin{aligned} q &\geq c_3 \left(\frac{\log(1 - \delta)}{kn} + \frac{2^n}{n} \right) \\ &\geq c_3 \left(\frac{\log(1 - \delta)}{d} + \frac{1}{\sqrt{d}} \sqrt{\log \frac{1}{1 - 2\epsilon}} \cdot 2^{\sqrt{\log \frac{d}{1 - 2\epsilon}}} \right) \\ &\geq c_3 \left(\frac{\log(1 - \delta)}{d} + \sqrt{\frac{\epsilon}{d}} \cdot 2^{\sqrt{\log \frac{d}{1 - 2\epsilon}}} \right). \end{aligned}$$

For $\delta = 1/2$ and $\epsilon \leq 1/4$ this implies

$$q \geq 2^{\Omega(\sqrt{d/\epsilon})}.$$

□

5 Discussion

Informally, our construction of $\mathcal{P}_C^{C'}$ can be thought of as “codes” in \mathbb{R}^d , namely sets in \mathbb{R}^d that are far from each other; the difficulty in the construction of such codes comes from the requirements of convexity and that

the distributions of polynomially many random samples look alike. By using slightly more involved arguments we can handle ϵ arbitrarily close to 1 and prove a similar lower bound. It is not clear if such a lower bound is possible for other learning settings that have been studied in the past, e.g. labeled samples from Gaussian distribution. Unlike that setting, we do not know a matching upper bound for learning convex bodies in our model.

Our construction of the hard family is more “explicit” than that of [8]: The hard family they construct is obtained by a probabilistic argument; our construction can be made explicit by using good error correcting codes.

We mention here a somewhat surprising corollary of our result, without detailed proof or precise numerical constants. Informally, it shows instability of the reconstruction of a convex body as a function of the volume of its intersection with halfspaces, relative to its volume. It is an exercise to see that knowledge of $|K \cap H|/|K|$ for every halfspace H uniquely determines K . Moreover, given d^c (for some fixed constant $c > 0$) random samples from a convex body $K \subseteq \mathbb{R}^d$, with high probability we can estimate $|K \cap H|/|K|$ for *all* half-spaces H within additive error of $O(1/d^{c'})$, where c' is a positive constant depending on c . This can be proved using standard arguments about ϵ -approximations and the fact that the VC-dimension of halfspaces in \mathbb{R}^d is $d + 1$.

We say that two convex bodies K and L are α -*halfspace-far* if there is a halfspace H such that $||K \cap H|/|K| - |L \cap H|/|L|| > \alpha$. Thus if we choose some $t < c'$ and K and L are $1/d^t$ -halfspace-far, then we can detect this using d^c random points, with high probability. Now, we claim that there is a pair of bodies in $\mathcal{P}_C^{C'}$ that is not far. For otherwise, all pairs would be far and we would be able to distinguish every body in $\mathcal{P}_C^{C'}$ from every other body in $\mathcal{P}_C^{C'}$ with a sample of size d^c , and thus learn it. But as we have proved, this is impossible. So we can conclude that there are two bodies in $\mathcal{P}_C^{C'}$ that are not $1/d^t$ -halfspace-far, i.e. they are $1/d^t$ -halfspace-close. This gives:

Corollary 6. *For any constant $t > 0$ and sufficiently large d there exist two convex bodies $K, L \subseteq \mathbb{R}^d$ such that K, L are $1/d^t$ -halfspace-close: for every halfspace H*

$$\left| \frac{|K \cap H|}{|K|} - \frac{|L \cap H|}{|L|} \right| \leq \frac{1}{d^t},$$

but $\text{dist}(K, L) > 1/8$.

An earlier version of this manuscript mentioned the problem of whether the volume of a convex body in \mathbb{R}^d can be estimated from $\text{poly}(d)$ uniformly

random samples. Very recently, Ronen Eldan [3] has answered this in the negative. His result provides a probabilistic construction of a family of convex bodies such that the volume a random body from this family is hard to estimate from random samples. His result does not supersede ours in the sense that our lower bound of $2^{\Omega(\sqrt{d})}$ is stronger, and perhaps optimal, and our construction of the hard family is explicit.

It is known that if the convex body is a polytope with $\text{poly}(d)$ facets, then it can be learned with $\text{poly}(d)$ uniformly random samples [6] in an information-theoretical sense. However, whether this can be done efficiently remains open:

Problem. *Can one learn polytopes with $\text{poly}(d)$ facets from $\text{poly}(d)$ uniformly random (over the polytope) samples in $\text{poly}(d)$ time?*

References

- [1] I. Bárány and Z. Füredi. Computing the volume is difficult. *Discrete & Computational Geometry*, 2:319–326, 1987.
- [2] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991.
- [3] R. Eldan. A polynomial number of random points does not determine the volume of a convex body. *Manuscript*, 2009.
- [4] A. M. Frieze, M. Jerrum, and R. Kannan. Learning linear transformations. In *FOCS*, pages 359–368, 1996.
- [5] A. Kalai. Personal communication. 2006.
- [6] A. Kalai and S. Vempala. Personal communication. 2007.
- [7] M. J. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. E. Schapire, and L. Sellie. On the learnability of discrete distributions. In *STOC*, pages 273–282, 1994.
- [8] A. Klivans, R. O’Donnell, and R. Servedio. Learning geometric concepts via gaussian surface area. In *FOCS*, 2008.
- [9] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

- [10] L. Rademacher and S. Vempala. Testing geometric convexity. In *FSTTCS*, pages 469–480, 2004.
- [11] J. van Lint. *Introduction to Coding Theory*. Springer, third edition, 1998.
- [12] S. Vempala. *The Random Projection Method*. American Mathematical Society, DIMACS, 2004.