

On Hardness of Learning Intersection of Two Halfspaces

Subhash Khot^{*}
NYU and Georgia Tech
khot@cs.nyu.edu

Rishi Saket[†]
Georgia Tech
saket@cc.gatech.edu

ABSTRACT

We show that unless $NP = RP$, it is hard to (even) weakly PAC-learn intersection of two halfspaces in \mathbf{R}^n using a hypothesis which is a function of up to ℓ linear threshold functions for any integer ℓ . Specifically, we show that for every integer ℓ and an arbitrarily small constant $\varepsilon > 0$, unless $NP = RP$, no polynomial time algorithm can distinguish whether there is an intersection of two halfspaces that correctly classifies a given set of labeled points in \mathbf{R}^n , or whether any function of ℓ linear threshold functions can correctly classify at most $\frac{1}{2} + \varepsilon$ fraction of the points.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Theory, Algorithms

Keywords

Learning, hardness, approximation, halfspaces

1. INTRODUCTION

A halfspace in \mathbf{R}^n is given by the set $\{\mathbf{x} \mid \mathbf{r} \cdot \mathbf{x} \leq c\}$ for some non-zero vector \mathbf{r} and a number c . The problem of learning a halfspace or an intersection of a small number of halfspaces is an extremely well-studied problem in machine learning, with several applications to computer vision [20], artificial intelligence [21] and data mining [23]. It is well-known that a single halfspace can be PAC-learned efficiently by sampling

^{*}Supported by NSF CAREER award CCF-0643626, Microsoft New Faculty Fellowship, and Sloan Foundation fellowship.

[†]Supported by the same awards as of Subhash Khot.

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

STOC'08, May 17–20, 2008, Victoria, British Columbia, Canada.
Copyright 2008 ACM 978-1-60558-047-0/08/05 ...\$5.00.

a polynomial number of data points and finding a separating hyperplane via linear programming [9]. Blum, Frieze, Kannan, and Vempala [6] showed how to learn a single halfspace even in presence of random classification noise, whereas Kalai, Klivans, Mansour and Servedio [14] gave polynomial time algorithm for learning a single halfspace in presence of adversarial noise under certain distributional assumptions.

For learning intersection of halfspaces, algorithms are known for various special cases. When the data points are drawn from the uniform distribution over the unit ball, Blum and Kannan [7] and Vempala [26] gave algorithms to PAC-learn intersection of a constant number of halfspaces. For the uniform distribution over the boolean hypercube, Klivans, O'Donnell and Servedio [17] obtained an algorithm for learning intersection of constant number of halfspaces. Arriaga and Vempala [4] and Klivans and Servedio [18] gave algorithms for learning intersection of halfspaces when no data point is too close to any separating hyperplane (i.e. the problem instance has a good *margin*). However the general problem of learning intersection of halfspaces remains an open problem.

In this paper we study the hardness of learning intersection of *two* halfspaces with a hypothesis that is a function of a constant (but arbitrarily large) number linear thresholds. We prove that this problem is hard to even weakly PAC-learn unless $NP = RP$. In particular our result holds for hypothesis class of an intersection of a constant number of halfspaces. We state our result formally after reviewing previous work on the hardness side.

1.1 Previous Work

For the problem of proper learning a single halfspace with adversarial noise, Feldman, Gopalan, Khot and Ponnuswami [10] and Guruswami and Raghavendra [12] independently proved that the problem is hard to even weakly PAC-learn¹. More specifically, they proved:

THEOREM 1. *Let $\delta, \varepsilon > 0$ be arbitrarily small constants. Then, given a set of labeled points in \mathbf{R}^n with a guarantee that there is a halfspace that classifies $1 - \delta$ fraction of points correctly, there is no polynomial time algorithm to find a halfspace that classifies $\frac{1}{2} + \varepsilon$ fraction of points correctly, unless $P = NP$.*

The above theorem is really a hardness of approximation result for an optimization problem, where the goal is to find

¹The result of Guruswami and Raghavendra [12] holds even when the points are from a boolean hypercube.

a halfspace that maximizes the number of correctly classified points. By considering a distribution that is uniform on all the data points, the theorem implies hardness of weakly PAC-learning a halfspace when the hypothesis is also required to be a halfspace (known as *proper learning*). Note that the imperfect completeness is necessary in the above theorem, since via linear programming, one can always efficiently find a halfspace that correctly classifies *all* the points, if one exists. The theorem is optimal, since one can easily classify $\frac{1}{2}$ fraction of the data points correctly, by taking an arbitrary halfspace or its complement as a hypothesis. From the learning theory perspective, such an optimal hardness result is especially satisfying, since if one could efficiently find $(\frac{1}{2} + \varepsilon)$ -consistent hypothesis (i.e. weakly PAC-learn), then one can use boosting techniques [24] to efficiently find a $(1 - \varepsilon)$ -consistent hypothesis (i.e. PAC-learn)².

However, an optimal result such as Theorem 1 was not known for learning intersection of (two) halfspaces. Blum and Rivest [8] showed that it is NP-hard to learn the intersection of two halfspaces with intersection of two halfspaces, and Alekhovich, Braverman, Feldman, Klivans and Pitassi [1] proved a similar result even when the hypothesis is an intersection of ℓ halfspaces for any constant ℓ . Both the results are only NP-hardness results and do not prove APX-hardness for the underlying optimization problem.

In a different line of work, under cryptographic assumptions, Klivans and Sherstov [19] showed that there is no polynomial time algorithm to PAC-learn intersection of n^ϵ halfspaces, and the result holds without any restriction on the hypothesis class.

1.2 Our Results

The following theorem is the main result in this paper.

THEOREM 2. *Let ℓ be any fixed integer and $\varepsilon > 0$ be an arbitrarily small constant. Then, given a set of labeled points in \mathbf{R}^n with a guarantee that there is an intersection of two halfspaces that classifies all the points correctly, there is no polynomial time algorithm to find a function f of up to ℓ linear threshold functions that classifies $\frac{1}{2} + \varepsilon$ fraction of points correctly, unless $NP = RP$.*

We state the above result in terms of functions of ℓ linear threshold functions. This encompasses hypotheses such as intersection of ℓ halfspaces since a linear threshold function $\text{sgn}(c - \mathbf{r} \cdot \mathbf{x})$ is same as the halfspace $\{\mathbf{x} \mid \mathbf{r} \cdot \mathbf{x} \leq c\}$ for any (unit) vector \mathbf{r} in \mathbf{R}^n and real number c . Note that the result holds with perfect completeness, i.e. the problem is hard even when an intersection of two halfspaces is guaranteed to classify *every* point correctly. The result is optimal since an arbitrary halfspace or its complement has a success rate of $\frac{1}{2}$ on any given data set. It provides evidence that the approach of weak learning intersection of two halfspaces with a function of a constant number of halfspaces followed by boosting may not work.

Overview of the reduction. Our reduction starts with an instance of the Label Cover problem on n vertices and label set $[k]$. The reduction produces an instance in nk dimensional space with a block of k dimensions for each vertex. In

²After applying the boosting algorithm, the hypothesis is a majority over a set of hypotheses used for the weak PAC-learning.

the first step we create a set of points for each vertex which are $\{-1, 1\}$ combinations in the k dimensions corresponding to the vertex and 0 everywhere else. These points simulate a junta test in the following manner. Suppose that a hyperplane $\mathbf{r} \cdot \mathbf{x} = c$ passes close to many points corresponding to a particular vertex v , then \mathbf{r} has the property that whatever mass it has in the k coordinates corresponding to v is concentrated in a few coordinates. However, for this to work, one has to ensure that \mathbf{r} has *some* non-negligible mass in those k coordinates to begin with.

To ensure this, in the second step we replace each point created in the first step with two small spheres of random points, where points in one of the spheres are labeled ‘+’ and those in the other are labeled ‘-’. If these spheres are sufficiently ‘dense’, one can show that any hyperplane that separates a pair of such spheres must have non negligible mass in the coordinates corresponding to the vertex v . We use the fact that the class of halfspaces has polynomially bounded VC dimension and therefore with high probability a polynomially large set of random points on a sphere is an ϵ -sample for all halfspaces.

In order to enforce consistency between the labels of different vertices of Label Cover instance, the third step involves a folding procedure over a subspace defined by the constraints of the Label Cover instance. A similar folding technique was used in [11]. The construction is such that the existence of a ‘good’ function of ℓ halfspaces gives us one single halfspace which can be used to extract a good labeling to the Label Cover instance.

Remark. The gap instance that we construct is such that in the YES case, the set of points is classified correctly by the intersection of two parallel halfspaces of the form $\mathbf{r} \cdot \mathbf{x} \geq -c$ and $\mathbf{r} \cdot \mathbf{x} \leq c$, for some $c > 0$. This implies that the points are correctly classified by the degree 2 polynomial given by $(\mathbf{r} \cdot \mathbf{x})^2 \leq c^2$. So, using linear programming a degree 2 polynomial can be efficiently found that classifies the points in our instance. However, the polynomial obtained may not be factorizable into two parallel hyperplanes.

2. PRELIMINARIES

We start by formally defining the problem.

DEFINITION 1. *An instance of INTERSECTION-HALFSPACE $_\ell$ is a set of points in \mathbf{R}^n each labeled either ‘+’ or ‘-’ and the goal is to find a function of at most ℓ linear threshold functions (halfspaces) which correctly classifies the maximum number of points, where a ‘+’ point is classified correctly if it lies inside the intersection and a ‘-’ point is classified correctly if it lies outside of it.*

We show that the problem INTERSECTION-HALFSPACE $_\ell$ is hard by giving a gap-preserving reduction from the Label Cover problem, a canonical problem used to prove many hardness of approximation results. For the purposes of our reduction we require the Label Cover instance to satisfy a certain ‘smoothness’ property. Similar ‘smooth’ versions of Label Cover have been used in earlier hardness reductions [15, 13, 16]. In addition to smoothness we also require that sufficiently large induced subgraphs of the label cover instance have a large number of edges. We define the following version of the Label Cover problem that captures both these additional properties.

DEFINITION 2. An instance \mathcal{L} of $\text{SMOOTH LABEL COVER}(t, \mu, \nu, k, m)$ consists of a (multi)graph $G(V, E)$ and mappings $\{\pi_{u,e}\}_{e \in E, u \in e}$ where $\pi_{u,e} : [k] \mapsto [m]$. A labeling $\sigma : V \mapsto [k]$ is said to satisfy an edge e between u and v if $\pi_{u,e}(\sigma(u)) = \pi_{v,e}(\sigma(v))$. The instance satisfies:

- (Smoothness:) For any vertex $u \in V$ and any set $S \subseteq [k]$ of size at most t ,

$$\Pr_{e:u \in e} [\exists i, j \in S, i \neq j, \text{ s.t. } \pi_{u,e}(i) = \pi_{u,e}(j)] \leq \mu,$$

where the probability is taken over a random edge incident on u .

- For any $V' \subseteq V$ such that $|V'| = \xi|V|$, the induced subgraph on V' has at least $(\xi^2/2)|E|$ edges, for any $\xi \geq \nu$.

The following theorem is proved using the PCP Theorem [3, 2] combined with Raz's Parallel Repetition Theorem [22]. We give a proof of the theorem in Appendix B based on the smooth version of label cover constructed in [16].

THEOREM 3. For any constant t and arbitrarily small constants $\mu, \nu, \eta > 0$, there exist constants k and m such that given an instance \mathcal{L} of $\text{SMOOTH LABEL COVER}(t, \mu, \nu, k, m)$ it is NP-hard to distinguish between the following two cases:

- YES Case/Completeness: There is a labeling to the vertices of \mathcal{L} which satisfies all the edges.
- NO Case/Soundness: No labeling to the vertices of \mathcal{L} satisfies more than η fraction of the edges.

We give a gap-preserving reduction from Label Cover to $\text{INTERSECTION-HALFSPACE}_\ell$ stated in the following theorem, which along with Theorem 3 implies Theorem 2.

THEOREM 4. For any constant $\varepsilon > 0$ and integer $\ell > 0$, there is a randomized polynomial time reduction from an instance \mathcal{L} of $\text{SMOOTH LABEL COVER}(t, \mu, \nu, k, m)$ to an instance \mathcal{I} of $\text{INTERSECTION-HALFSPACE}_\ell$ for appropriately chosen parameters t, μ, ν and soundness η , such that,

- YES Case/Completeness: If \mathcal{L} is a YES instance, then there is an intersection of two halfspaces which correctly classifies all the points in instance \mathcal{I} .
- NO Case/Soundness: If \mathcal{L} is a NO instance, then with probability at least $\frac{9}{10}$, there is no function of up to ℓ linear threshold functions that correctly classifies more than $\frac{1}{2} + \varepsilon$ fraction of points in instance \mathcal{I} .

3. REDUCTION

The reduction proceeds in three steps. In the first step we construct an initial set of unlabeled points, from the instance of Label Cover. In the second step we replace each initial point with two small spheres of points, with points in one sphere labeled '+' and points in the other labeled '-'. The third step consists of reducing the problem into a lower dimensional space via a *folding* over the subspace induced by the consistency constraints of the Label Cover instance. In the end we obtain a set of points each labeled either '+' or '-' as an instance of $\text{INTERSECTION-HALFSPACE}_\ell$ for a given constant ℓ .

3.1 Step 1

We start with an instance \mathcal{L} of $\text{SMOOTH LABEL COVER}(t, \mu, \nu, k, m)$, where we will fix t, μ and ν later. Let $|V| = n$. First we define the space in which the points lie. For every vertex $v \in V$, we have a set of k coordinates labeled by $M(v) := \{v(i)\}_{i=1}^k$. The complete set of coordinates is the union of these sets over all vertices, say $\mathcal{M} := \bigcup_{v \in V} M(v)$. Therefore, the points we construct lie in nk dimensional real space $\mathbf{R}^{\mathcal{M}}$. The construction of the points is as follows.

1. For every vertex v , define,

$$s(v) := \{\mathbf{x} \in \mathbf{R}^{\mathcal{M}} \mid \forall i \in [k], \mathbf{x}(v(i)) \in \{-1, 1\} \text{ and } \mathbf{x}(u(i)) = 0, \forall u \neq v\}.$$

Thus, $s(v)$ is the set of all vectors that are $\{-1, 1\}$ combinations on the coordinates $M(v)$ and 0 on all other coordinates. Note that $|s(v)| = 2^k$ for all $v \in V$.

2. Let $S = \bigcup_{v \in V} s(v)$. Clearly $|S| = n \cdot 2^k$, since the sets $s(v)$ are disjoint for all $v \in V$.

We output S , as the initial set of points at the end of Step 1. One would like to say that if any hyperplane say $\mathbf{r} \cdot \mathbf{x} = a$, passes through a significant fraction of the points in S , then \mathbf{r} should be close to a 'junta' i.e. most of the mass of the vector \mathbf{r} should be concentrated in a few coordinates of $M(v)$ for a significant fraction of vertices $v \in V$. However, it is possible that \mathbf{r} has zero mass in the coordinates $M(v)$ for almost all $v \in V$ and yet the hyperplane $\mathbf{r} \cdot \mathbf{x} = 0$ passes through most of the points in S . To overcome this problem, we replace each point in S with two spheres of points as described in Step 2.

3.2 Step 2

We start with the set of points $S = \bigcup_{v \in V} s(v)$ constructed in Step 1. For every point in S , we create two 'spheres' of points separated by a small distance. This step is randomized as it requires sampling a suitable number of points from a unit sphere. The following lemma is proved in Appendix A.

LEMMA 5. Let $\varepsilon' > 0$ be any constant and n be a sufficiently large integer. Let R be a set of $(nk)^2$ unit vectors chosen uniformly at random in nk -dimensional real space. Then with probability at least $1 - 1/n$, the set R satisfies the following property: for any subset $T \subseteq R$ such that $|T| = \varepsilon'|R|$ and for any unit vector \mathbf{r} , there exist $\mathbf{z}', \mathbf{z}'' \in T$, such that $|\mathbf{r} \cdot \mathbf{z}' - \mathbf{r} \cdot \mathbf{z}''| \geq \varepsilon'/(100\sqrt{nk})$.

Now we describe our construction in Step 2.

1. Set parameters $\delta = 2^{-(nk)^{100}}$ and $\gamma = 1/(100\sqrt{n})$.
2. Let R be the set of $(nk)^2$ unit vectors in $\mathbf{R}^{\mathcal{M}}$ as in Lemma 5 (ε' will be chosen later and is related to the soundness parameter of the reduction).
3. Let \mathbf{x} be a point in S . Construct two sets $\alpha(\mathbf{x})$ and $\beta(\mathbf{x})$ as follows,

$$\alpha(\mathbf{x}) := \{(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z} \mid \mathbf{z} \in R\}$$

³The parameter δ here is essentially the *margin* in the YES case (upto a polynomial factor). We set δ to be exponentially small in n and k . However, our reduction goes through even if δ is taken to be $O(\frac{1}{n^2})$.

and,

$$\beta(\mathbf{x}) := \{(1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z} \mid \mathbf{z} \in R\}.$$

4. For every vertex $v \in V$, let $A(v) := \bigcup_{\mathbf{x} \in s(v)} \alpha(\mathbf{x})$ and $B(v) := \bigcup_{\mathbf{x} \in s(v)} \beta(\mathbf{x})$.
5. Output the sets $A := \bigcup_{v \in V} A(v)$ and $B := \bigcup_{v \in V} B(v)$.

The points created have the property that any hyperplane that separates the sets $\alpha(\mathbf{x})$ from $\beta(\mathbf{x})$ for a significant fraction of points $\mathbf{x} \in S$, must essentially be a ‘junta’ in the coordinates $M(v)$ for a significant fraction of vertices $v \in V$. This property will be formally stated and used in the soundness analysis to decode a labeling for the Label Cover instance. In conjunction with this property, one needs to enforce the consistency constraints of the Label Cover instance. We achieve this in the third step of the reduction, by folding over a subspace defined by these constraints.

3.3 Step 3

For the sake of convenience, let $<$ be any arbitrary total order on V . Let e be an edge between u and v in G , with $u < v$. Let $\mathbf{h}_e^j \in \mathbf{R}^M$, for $j \in [m]$, be defined in the following manner: set $\mathbf{h}_e^j(u(i)) = 1$ for all $i \in \pi_{u,e}^{-1}(j)$, set $\mathbf{h}_e^j(v(i)) = -1$ for all $i \in \pi_{v,e}^{-1}(j)$ and set all the other coordinates to 0. Note that for any vector $\mathbf{r} \in \mathbf{R}^M$,

$$\begin{aligned} \forall u, v \in e, e \in E, \forall j \in [m], \\ \mathbf{r} \cdot \mathbf{h}_e^j = 0 \iff \\ \sum_{i \in \pi_{u,e}^{-1}(j)} \mathbf{r}(u(i)) = \sum_{i \in \pi_{v,e}^{-1}(j)} \mathbf{r}(v(i)). \end{aligned} \quad (1)$$

The folding is done as follows.

1. Let $T = \bigcup_{e \in E, j \in [m]} \{\mathbf{h}_e^j\}$. Let $H \subset \mathbf{R}^M$, where $H = \text{span}(T)$, and F be the subspace of \mathbf{R}^M orthogonal to H such that $\mathbf{R}^M = F \oplus H$ and $F \perp H$.
2. Let $\{\lambda_i\}_{i=1}^{nk}$ be an orthonormal basis for \mathbf{R}^M such that $F = \text{span}(\{\lambda_i\}_{i=1}^g)$ for some $g \leq nk$.
3. Write down all the points in the sets A and B in the basis $\{\lambda_i\}_{i=1}^{nk}$ and only consider the coordinates corresponding to the basis $\{\lambda_i\}_{i=1}^g$ of the g -dimensional space F . Ignoring the rest of the coordinates, obtain sets A' and B' , which are essentially projections (with multiplicities) of sets A and B respectively onto the subspace F .

We label all the points in A' as ‘+’, and all the points in B' as ‘-’ and output these points as an instance of INTERSECTION-HALFSPACE $_\ell$.

4. ANALYSIS

4.1 YES Case

If the instance \mathcal{L} of SMOOTHLABELCOVER(t, μ, ν, k, m) is a YES instance, then there is a labeling, say σ to the vertices in V that satisfies all the edges in E . We need to exhibit two halfspaces such that the points in A' lie inside their intersection while the points in B' lie outside their intersection, where A' and B' are the sets obtained through the reduction given above.

Let us consider the vector $\mathbf{r}^* \in \mathbf{R}^M$, where $\mathbf{r}^*(v(\sigma(v))) = 1/\sqrt{n}$ for all $v \in V$, and all other coordinates are set to 0. So, \mathbf{r}^* has exactly n non-zero coordinates, each set to $1/\sqrt{n}$. Clearly $\|\mathbf{r}^*\| = 1$. We prove the following two lemmas.

LEMMA 6. *For every point $\mathbf{y} \in A$, $|\mathbf{r}^* \cdot \mathbf{y}| < 1/\sqrt{n}$, and for every point $\mathbf{w} \in B$, $|\mathbf{r}^* \cdot \mathbf{w}| > 1/\sqrt{n}$.*

PROOF. Since \mathbf{r}^* is a unit vector, for any unit vector \mathbf{z} , we have,

$$|\mathbf{r}^* \cdot \gamma\delta\mathbf{z}| \leq \gamma\delta = \delta \left(\frac{1}{100\sqrt{n}} \right). \quad (2)$$

Consider a point $\mathbf{y} \in A$, such that $\mathbf{y} = (1 - \delta)\mathbf{x} + \gamma\delta\mathbf{z}$, where \mathbf{x} is a $\{-1, 1\}$ vector in the coordinates $M(v)$ for some $v \in V$, and 0 on all other coordinates, and \mathbf{z} is a unit vector. Now since \mathbf{r}^* has a single non-zero coordinate set to $1/\sqrt{n}$ in each set $M(v)$, clearly $|\mathbf{r}^* \cdot \mathbf{x}| = 1/\sqrt{n}$ and therefore $|\mathbf{r}^* \cdot (1 - \delta)\mathbf{x}| = (1 - \delta)(1/\sqrt{n})$. Combining with (2), we get $|\mathbf{r}^* \cdot ((1 - \delta)\mathbf{x} + \gamma\delta\mathbf{z})| = |\mathbf{r}^* \cdot \mathbf{y}| < 1/\sqrt{n}$. Similarly, for any $\mathbf{w} \in B$, we obtain $|\mathbf{r}^* \cdot \mathbf{w}| > 1/\sqrt{n}$. \square

Consider any edge $e \in E$ between two vertices u and v in V . Since \mathbf{r}^* is $1/\sqrt{n}$ in exactly one coordinate $u(\sigma(u))$ in $M(u)$ and 0 on all others, for any $j \in [m]$, we have that the quantity $\sum_{i \in \pi_{u,e}^{-1}(j)} \mathbf{r}^*(u(i))$ is $1/\sqrt{n}$ iff $\sigma(u) \in \pi_{u,e}^{-1}(j)$ and 0 otherwise. And similarly, $\sum_{i \in \pi_{v,e}^{-1}(j)} \mathbf{r}^*(v(i))$ is $1/\sqrt{n}$ iff $\sigma(v) \in \pi_{v,e}^{-1}(j)$ and 0 otherwise. Since σ is a satisfying assignment, $\exists j' \in [m]$ such that $\sigma(u) \in \pi_{u,e}^{-1}(j')$ and $\sigma(v) \in \pi_{v,e}^{-1}(j')$. Therefore we have,

$$\begin{aligned} \sum_{i \in \pi_{u,e}^{-1}(j)} \mathbf{r}^*(u(i)) = \sum_{i \in \pi_{v,e}^{-1}(j)} \mathbf{r}^*(v(i)), \\ \forall u, v \in e, e \in E, \forall j \in [m]. \end{aligned} \quad (3)$$

Combining the above with (1) we obtain,

$$\mathbf{r}^* \cdot \mathbf{h}_e^j = 0, \quad \forall e \in E, \forall j \in [m]. \quad (4)$$

Since H was defined to be the span of $\{\mathbf{h}_e^j\}_{e \in E, j \in [m]}$, equation (4) implies that $\mathbf{r}^* \perp H$. Let $\bar{\mathbf{r}}^*$ be the projection of \mathbf{r}^* onto F , where $F \perp H$ and $\mathbf{R}^M = F \oplus H$. For any $\mathbf{y} \in \mathbf{R}^M$, let $\bar{\mathbf{y}}$ be the projection of \mathbf{y} onto F . Then, since \mathbf{r}^* lies entirely in F we have $\mathbf{r}^* \cdot \mathbf{y} = \bar{\mathbf{r}}^* \cdot \bar{\mathbf{y}}$. Combined with Lemma 6 this implies that for every point $\bar{\mathbf{y}} \in A'$, $|\bar{\mathbf{r}}^* \cdot \bar{\mathbf{y}}| < 1/\sqrt{n}$ and for every point $\bar{\mathbf{w}} \in B'$, $|\bar{\mathbf{r}}^* \cdot \bar{\mathbf{w}}| > 1/\sqrt{n}$. Therefore the intersection of the two halfspaces in F , namely $\{\mathbf{y} \mid \bar{\mathbf{r}}^* \cdot \mathbf{y} \leq 1/\sqrt{n}\}$ and $\{\mathbf{y} \mid \bar{\mathbf{r}}^* \cdot \mathbf{y} \geq -1/\sqrt{n}\}$, classifies all the points in A' and B' correctly. Note that the intersection of halfspaces that we obtain is the region between two parallel hyperplanes.

4.2 NO Case

In this case, we assume that the instance \mathcal{L} of SMOOTHLABELCOVER(t, μ, ν, k, m) has soundness η . For a contradiction, we assume that we have a function f of ℓ linear threshold functions in F that classifies $\frac{1}{2} + \varepsilon$ fraction of the points in A' and B' correctly, where A' is the set of ‘+’ points and B' is the set of ‘-’ points. We will henceforth refer to linear threshold functions as halfspaces, since they are exactly the same. Let the halfspaces on which f depends be given by the equations,

$$\bar{\mathbf{r}}_i \cdot \bar{\mathbf{y}} \leq c_i \quad \text{for } i = 1, \dots, \ell,$$

where $\bar{\mathbf{r}}_i \in F$ and $\|\bar{\mathbf{r}}_i\| = 1$ for all $i = 1, \dots, \ell$. Let \mathbf{r}_i be the vector in \mathbf{R}^M obtained from $\bar{\mathbf{r}}_i \in F$, by adding zeros on the

coordinates corresponding to the basis of H , and rewriting it in the coordinates $\mathcal{M} = \bigcup_{v \in V} M(v)$. Clearly $\|\mathbf{r}_i\| = 1$ for $i = 1, \dots, \ell$. Let f' be the function in $\mathbf{R}^{\mathcal{M}}$ given by the predicate of f applied on the halfspaces $\{\mathbf{r}_i \cdot \mathbf{y} \leq c_i\}$ for $i = 1, \dots, \ell$, where the halfspace $\{\bar{\mathbf{r}}_i \cdot \bar{\mathbf{y}} \leq c_i\}$ in F is replaced by the halfspace $\{\mathbf{r}_i \cdot \mathbf{y} \leq c_i\}$ in $\mathbf{R}^{\mathcal{M}}$ in the predicate of f . Note that f' is exactly the function f applied on points in $\mathbf{R}^{\mathcal{M}}$ after projection onto F . We have the following simple lemma.

LEMMA 7. *The the function f' of the halfspaces $\{\mathbf{r}_i \cdot \mathbf{y} \leq c_i\}$ for $i = 1, \dots, \ell$ classifies $\frac{1}{2} + \varepsilon$ fraction of the points in $A \cup B$ correctly, where a point in A is classified correctly if it lies inside the intersection and a point in B is classified correctly if it lies outside.*

PROOF. We observe that since $\bar{\mathbf{r}}_i \in F$, if $\mathbf{y} \in \mathbf{R}^{\mathcal{M}}$ has a projection $\bar{\mathbf{y}} \in F$, then $\mathbf{r}_i \cdot \mathbf{y} = \bar{\mathbf{r}}_i \cdot \bar{\mathbf{y}}$. Now, since A' and B' are (multi)sets of points in F and are projections of the sets A and B respectively of points in $\mathbf{R}^{\mathcal{M}}$, the lemma follows. \square

For the rest of the analysis, we will consider only the sets of points A and B and the halfspaces $\{\mathbf{r}_i \cdot \mathbf{y} \leq c_i\}$ for $i = 1, \dots, \ell$ in $\mathbf{R}^{\mathcal{M}}$. For every vertex v , and every $\mathbf{x} \in s(v)$, there are $|R|$ pairs of points given by the sets $\{(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z}\}$. In total there are $|V|2^k|R|$ such pairs, which partition the set $A \cup B$, where each pair has one point from A and one point from B . We say that a pair $\{\mathbf{y}_1, \mathbf{y}_2\}$ where $\mathbf{y}_1 \in A$ and $\mathbf{y}_2 \in B$ is correctly classified by f' if both \mathbf{y}_1 and \mathbf{y}_2 are correctly classified by f' . Since the function f' of halfspaces $\{\mathbf{r}_i \cdot \mathbf{y} \leq c_i\}$ for $i = 1, \dots, \ell$ correctly classifies $\frac{1}{2} + \varepsilon$ fraction of the points in $A \cup B$, it follows that it correctly classifies $\varepsilon/2$ fraction of pairs $\{(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z}\}$. For a pair to be classified correctly by f' , it must be separated by at least one of the ℓ halfspaces on which f' depends. Thus, there must be at least one out of the ℓ halfspaces that separates $\varepsilon/(2\ell)$ fraction of the pairs. Without loss of generality, we can assume that the halfspace $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$ separates $\varepsilon/(2\ell)$ fraction of the pairs. The rest of the analysis uses \mathbf{r}_1 to deduce a labeling to the vertices of Label Cover instance that satisfies a significant fraction of the edges.

Let $\varepsilon' = \varepsilon/(32\ell)$. By an averaging argument we have that for ε' fraction of the vertices $v \in V$, for ε' fraction of vectors $\mathbf{x} \in s(v)$, for $2\varepsilon'$ fraction of $\mathbf{z} \in R$, the pair $\{(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z}\}$ is separated by the halfspace $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$. Call such vertices ‘good’. Let u be one such vertex. We will show that the vector \mathbf{r}_1 must have a significant mass in the coordinates $M(u)$, and moreover that the mass is concentrated in a few of the coordinates in $M(u)$.

We know from above that for u , there is a vector $\mathbf{x}' \in s(u)$ such that for $2\varepsilon'$ fraction of $\mathbf{z} \in R$ the pair $\{(1 - \delta)\mathbf{x}' + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x}' + \delta\gamma\mathbf{z}\}$ is separated by $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$. Let us fix this particular $\mathbf{x}' \in s(u)$. Let us say that a pair is separated ‘correctly’ by the halfspace $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$ if the ‘+’ point is inside the halfspace and the ‘-’ point is outside, and otherwise we say that the pair is separated ‘incorrectly’. Based on the above, for our choice of $\mathbf{x}' \in s(u)$ we have the following two cases.

Case 1. The halfspace $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$ separates ‘correctly’ the pair $\{(1 - \delta)\mathbf{x}' + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x}' + \delta\gamma\mathbf{z}\}$ for ε' fraction of $\mathbf{z} \in R$. Let T be this set of ‘good’ vectors $\mathbf{z} \in R$, for which the

corresponding pairs are separated correctly. and so $|T| = \varepsilon'|R|$. Since R satisfies the property stated in Lemma 5, there exist $\mathbf{z}', \mathbf{z}'' \in T$ such that,

$$|\mathbf{r}_1 \cdot \mathbf{z}' - \mathbf{r}_1 \cdot \mathbf{z}''| \geq \varepsilon'/(100\sqrt{nk}). \quad (5)$$

Moreover, since the pairs are separated ‘correctly’ we have that,

$$\mathbf{r}_1 \cdot ((1 - \delta)\mathbf{x}' + \gamma\delta\mathbf{z}') - c_1 \leq 0 \quad (6)$$

$$\mathbf{r}_1 \cdot ((1 + \delta)\mathbf{x}' + \gamma\delta\mathbf{z}') - c_1 \geq 0, \quad (7)$$

$$\mathbf{r}_1 \cdot ((1 - \delta)\mathbf{x}' + \gamma\delta\mathbf{z}'') - c_1 \leq 0 \quad (8)$$

$$\mathbf{r}_1 \cdot ((1 + \delta)\mathbf{x}' + \gamma\delta\mathbf{z}'') - c_1 \geq 0. \quad (9)$$

Subtracting equation (6) from (9), and (8) from (7), we obtain,

$$2\delta\mathbf{r}_1 \cdot \mathbf{x}' - \delta\gamma\mathbf{r}_1 \cdot (\mathbf{z}' - \mathbf{z}'') \geq 0$$

$$2\delta\mathbf{r}_1 \cdot \mathbf{x}' + \delta\gamma\mathbf{r}_1 \cdot (\mathbf{z}' - \mathbf{z}'') \geq 0.$$

Combining the above with equation (5) we get that $|2\delta\mathbf{r}_1 \cdot \mathbf{x}'| \geq \gamma\delta\varepsilon'/(100\sqrt{nk})$. Substituting the value of γ and simplifying we have $|\mathbf{r}_1 \cdot \mathbf{x}'| \geq \varepsilon'/(2 \cdot 10^4 n\sqrt{k})$. Since \mathbf{x}' takes values 1 or -1 on coordinates in $M(u)$ and is 0 on all other coordinates, this implies,

$$\sum_{i \in k} |\mathbf{r}_1(u(i))| \geq \frac{\varepsilon'}{2 \cdot 10^4 n\sqrt{k}} \quad (10)$$

Case 2. In this case we have that the halfspace $\{\mathbf{r}_1 \cdot \mathbf{y} \leq c_1\}$ separates ‘incorrectly’ the pair $\{(1 - \delta)\mathbf{x}' + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x}' + \delta\gamma\mathbf{z}\}$ for ε' fraction of $\mathbf{z} \in R$. The analysis continues along the same line as Case 1, taking the set T to be the set of ‘good’ vectors $\mathbf{z} \in R$, for which the corresponding pairs are separated ‘incorrectly’. Since the pairs are now separated ‘incorrectly’, the inequalities (6), (7), (8) and (9) are reversed. We omit the rest of the argument which remains essentially the same as Case 1 and we obtain the same bound given by equation (10).

The above analysis shows that the vector \mathbf{r}_1 has significant mass in the coordinates $M(u)$. To show it is concentrated in a small number of coordinates, we need the following lemma.

LEMMA 8. *There is a set $Q \subseteq s(u)$, s.t. $|Q| \geq \varepsilon'|s(u)|$ and for every $\mathbf{x} \in Q$, $\mathbf{r}_1 \cdot \mathbf{x} \in [c_1 - 2\delta\sqrt{k}, c_1 + 2\delta\sqrt{k}]$.*

PROOF. We consider the set Q of points $\mathbf{x} \in s(u)$, such that a pair $\{(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z}, (1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z}\}$ for some $\mathbf{z} \in R$ is separated by $\mathbf{r}_1 \cdot \mathbf{y} \leq c_1$. Clearly, $|Q| \geq \varepsilon'|s(u)| = \varepsilon'2^k$. Now, for any given $\mathbf{x} \in s(u)$, all the points of the form $(1 - \delta)\mathbf{x} + \delta\gamma\mathbf{z}$ and $(1 + \delta)\mathbf{x} + \delta\gamma\mathbf{z}$ for any $\mathbf{z} \in R$ lie in a ball of radius $2\delta\sqrt{k}$ around \mathbf{x} . Therefore, for all $\mathbf{x} \in Q$, the hyperplane $\mathbf{r}_1 \cdot \mathbf{y} = c_1$ passes at a perpendicular distance of at most $2\delta\sqrt{k}$ from \mathbf{x} . The lemma follows. \square

The following is a well known lemma (see lemma 7.3 of [12]). We state a version based on lemma 3.5 proved in [5].

LEMMA 9. *Let X_1, \dots, X_p be i.i.d $\{-1, 1\}$ valued Bernoulli random variables, with $\Pr[1] = \frac{1}{2}$, and let $\omega_1, \dots, \omega_p$ be positive real numbers. Then there is a universal constant b such that, for any $c \in \mathbf{R}$ and $\zeta > 0$, if,*

$$\Pr \left[\sum_{i=1}^p \omega_i X_i \in [c - \zeta, c + \zeta] \right] \geq \frac{b}{p^{\frac{1}{2}}}$$

then, $\exists i \in [p]$ such that $\omega_i \leq \zeta$.

Let X_1, \dots, X_k be i.i.d $\{-1, 1\}$ valued Bernoulli random variables with $\Pr[1] = \frac{1}{2}$. For convenience, we let $\delta' = \delta\sqrt{k}$. Observe that Lemma 8 implies,

$$\Pr \left[\sum_{i=1}^k |\mathbf{r}_1(u(i))| X_i \in [c_1 - 2\delta', c_1 + 2\delta'] \right] \geq \varepsilon'. \quad (11)$$

Suppose we apply Lemma 9 to the above. Then it gives us a coordinate in $M(u)$ such that on that coordinate \mathbf{r}_1 has very small mass. Removing that coordinate, we can again apply the lemma to the remaining coordinates and do this until a small number of coordinates remain. If we ensure that at each step we remove a coordinate from $M(u)$ on which \mathbf{r}_1 has small mass, then the total mass of the coordinates removed is small. Combining this with the lower bound given by equation (10), this would imply that most of the mass of \mathbf{r}_1 in $M(u)$ is concentrated in a small number of coordinates. This would enable us to select a labeling for the vertex u from among those ‘large’ coordinates.

Therefore, we apply Lemma 9 iteratively, until the set of coordinates is of size at most b^2/ε'^2 , in the following manner. Initialize $I_0 = [k]$.

1. At step j , we have a set of indices I_j of size $k - j$, with the following inequality satisfied,

$$\Pr \left[\sum_{i \in I_j} |\mathbf{r}_1(u(i))| X_i \in [c_1 - 2^{j+1}\delta', c_1 + 2^{j+1}\delta'] \right] \geq \varepsilon' \quad (12)$$

2. If $k - j < b^2/\varepsilon'^2$, then we stop and obtain a set $I^u = I_j$ of indices, such that $|I^u| < b^4/\varepsilon'^4$.
3. If $k - j \geq b^2/\varepsilon'^2$, then we apply Lemma 9 to obtain $i' \in I_j$ such that $|\mathbf{r}_1(u(i'))| \leq 2^{j+1}\delta'$. Now for any setting $x_i \in \{-1, 1\}$ of variables X_i for $i \in I_j$,

$$\begin{aligned} \sum_{i \in I_j} |\mathbf{r}_1(u(i))| x_i \in [c_1 - 2^{j+1}\delta', c_1 + 2^{j+1}\delta'] &\implies \\ \sum_{i \in I_j \setminus \{i'\}} |\mathbf{r}_1(u(i))| x_i \in [c_1 - 2^{j+2}\delta', c_1 + 2^{j+2}\delta']. \end{aligned}$$

Therefore we have,

$$\Pr \left[\sum_{i \in I_j \setminus \{i'\}} |\mathbf{r}_1(u(i))| X_i \in [c_1 - 2^{j+2}\delta', c_1 + 2^{j+2}\delta'] \right] \geq \varepsilon'.$$

So, we set $I_{j+1} = I_j \setminus \{i'\}$ and proceed to step $j + 1$.

At the j th step, an index corresponding to a coordinate of mass at most $2^{j+1}\delta'$ is removed. There are at most k steps for $j = 0, \dots, k - 1$. Therefore, the total mass of the coordinates removed is at most $2^{k+1}\delta'$. Combining this with (10) and with a small enough choice of δ , we have a set $I^u \subseteq [k]$ such that $|I^u| \leq b^2/(\varepsilon')^2$ and

$$\sum_{i \in I^u} |\mathbf{r}_1(u(i))| \geq \frac{\varepsilon'}{2 \cdot 10^4 n \sqrt{k}} - 2^{k+1}\delta' \geq \frac{\varepsilon'}{4 \cdot 10^4 n \sqrt{k}} \quad (13)$$

and,

$$\sum_{i \in [k] \setminus I^u} |\mathbf{r}_1(u(i))| \leq 2^{k+1}\delta' \leq \frac{\varepsilon'}{16 \cdot 10^4 n \sqrt{k}} \quad (14)$$

Since u was one of the ε' fraction of the vertices of V that are ‘good’, we can obtain such sets of indices I^v satisfying the above properties for all ‘good’ vertices v . Construct the labeling σ^* to these vertices by choosing a label for every ‘good’ vertex $v \in V$ randomly from I^v . From the properties of the instance \mathcal{L} of LABELCOVER(t, μ, ν, k, m), if we choose $\nu \ll \varepsilon'$, then the set of ‘good’ vertices induces $(\varepsilon')^2/2$ fraction of edges in E . Let e be a random edge in E , and say e is between vertices v_1 and v_2 in V . Then with probability $(\varepsilon')^2/2$, both v_1 and v_2 are ‘good’. Now, suppose that we have chosen $t \gg b^2/(\varepsilon')^2$, then except with probability 2μ , π_{e, v_1} maps the elements of I^{v_1} to distinct elements $J^{v_1} \subseteq [m]$ and π_{e, v_2} maps the elements of I^{v_2} to distinct elements in $J^{v_2} \subseteq [m]$.

Suppose for a contradiction that J^{v_1} and J^{v_2} are disjoint. This implies that $\pi_{e, v_2}^{-1}(J^{v_1})$ and I^{v_2} are disjoint. Since \mathbf{r}_1 is orthogonal to the subspace H , from (1) we have that for every $j \in J^{v_1}$,

$$\sum_{i \in \pi_{e, v_1}^{-1}(j)} \mathbf{r}_1(v_1(i)) = \sum_{i \in \pi_{e, v_2}^{-1}(j)} \mathbf{r}_1(v_2(i))$$

and taking the absolute values and summing over all $j \in J^{v_1}$, we have,

$$\begin{aligned} \sum_{j \in J^{v_1}} \left| \sum_{i \in \pi_{e, v_1}^{-1}(j)} \mathbf{r}_1(v_1(i)) \right| \\ = \sum_{j \in J^{v_1}} \left| \sum_{i \in \pi_{e, v_2}^{-1}(j)} \mathbf{r}_1(v_2(i)) \right| \end{aligned} \quad (15)$$

Now, since π_{e, v_1} maps elements of I^{v_1} to distinct elements $J^{v_1} \subseteq [m]$,

$$\begin{aligned} \sum_{j \in J^{v_1}} \left| \sum_{i \in \pi_{e, v_1}^{-1}(j)} \mathbf{r}_1(v_1(i)) \right| \\ \geq \sum_{i \in I^{v_1}} |\mathbf{r}_1(v_1(i))| - \sum_{i \in [k] \setminus I^{v_1}} |\mathbf{r}_1(v_1(i))|. \end{aligned} \quad (16)$$

From equation (15) and the above we have,

$$\begin{aligned} \sum_{j \in J^{v_1}} \left| \sum_{i \in \pi_{e, v_2}^{-1}(j)} \mathbf{r}_1(v_2(i)) \right| \\ \geq \sum_{i \in I^{v_1}} |\mathbf{r}_1(v_1(i))| - \sum_{i \in [k] \setminus I^{v_1}} |\mathbf{r}_1(v_1(i))| \\ \geq \frac{\varepsilon'}{4 \cdot 10^4 n \sqrt{k}} - \frac{\varepsilon'}{16 \cdot 10^4 n \sqrt{k}} \\ = \frac{3\varepsilon'}{16 \cdot 10^4 n \sqrt{k}} \end{aligned} \quad (17)$$

where we used (13) and (14) applied to v_1 . But since, $\pi_{e, v_2}^{-1}(J^{v_1}) \subseteq [k] \setminus I^{v_2}$, equation (17) is a contradiction to equation (14) applied to v_2 . Therefore, J^{v_1} and J^{v_2} are not disjoint. So, with probability $1/(|I^{v_1}| |I^{v_2}|) = (\varepsilon')^4/b^4$, the labeling σ^* satisfies the edge e . Combining everything, we obtain that there is a labeling to the vertices of V that satisfies $((\varepsilon')^2/2 - 2\mu)((\varepsilon')^4/b^4)$ fraction of the edges in E . By choosing the smoothness parameter μ and the soundness parameter η of the Label Cover instance to be arbitrarily

small, we obtain a contradiction. Thus, if the instance \mathcal{L} of $\text{SMOOTH LABEL COVER}(t, \mu, \nu, k, m)$ is a NO instance, then with high probability, there is no function of up to ℓ halfspaces that correctly classifies $\frac{1}{2} + \varepsilon$ fraction of the points in $A' \cup B'$. This, along with the analysis of the YES case proves Theorem 4, and hence Theorem 2.

5. CONCLUSION

We proved a tight hardness result for learning intersection of two halfspaces using functions of up to ℓ linear threshold functions (halfspaces) for any constant ℓ . An interesting open question is whether a similar hardness result holds for learning intersection of halfspaces by more general classes of hypotheses such as (functions of) low degree polynomials. As noted in the remark in section 1.2 our reduction does not extend even to degree 2 polynomials.

6. ACKNOWLEDGEMENT

We thank the attentive audience of IAS, Princeton for pointing out that our hardness result also holds for any arbitrary function of linear threshold functions, not just intersection of halfspaces as stated in an earlier version of the paper.

7. REFERENCES

- [1] M. Alekhnovich, M. Braverman, V. Feldman, A. Klivans, and T. Pitassi. Learnability and automatizability. In *Proc. 45th IEEE FOCS*, pages 621–630, 2004.
- [2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [3] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [4] R. Arriaga and S. Vempala. An algorithmic theory of learning: Robust concepts and random projection. In *Proc. 40th IEEE FOCS*, pages 616–623, 1999.
- [5] I. Benjamin, G. Kalai, and O. Schramm. Noise sensitivity of boolean functions and applications to percolation. *Inst. Hautes Études Sci. Publ. Math.*, (90):5–43, 1999.
- [6] A. Blum, A. Frieze, R. Kannan, and S. Vempala. A polynomial-time algorithm for learning noisy linear threshold functions. In *Proc. 37th IEEE FOCS*, pages 330–338, 1996.
- [7] A. Blum and R. Kannan. Learning an intersection of a constant number of halfspaces over a uniform distribution. *J. Comput. Syst. Sci.*, 54(2):371–380, 1997.
- [8] A. Blum and R. Rivest. Training a 3-node neural network is NP-complete. In *Proc. Machine Learning: From Theory to Applications*, pages 9–28, 1993.
- [9] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. Warmuth. Learnability and the Vapnik Chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989.
- [10] V. Feldman, P. Gopalan, S. Khot, and A. Ponnuswami. New results for learning noisy parities and halfspaces. In *Proc. 47th IEEE FOCS*, pages 563–574, 2006.

- [11] P. Gopalan, S. Khot, and R. Saket. Hardness of reconstructing multivariate polynomials over finite fields. In *Proc. 48th IEEE FOCS*, pages 349–359, 2007.
- [12] V. Guruswami and P. Raghavendra. Hardness of learning halfspaces with noise. In *Proc. 47th IEEE FOCS*, pages 543–552, 2006.
- [13] J. Holmerin and S. Khot. A new PCP outer verifier with applications to homogeneous linear equations and max-bisection. In *Proc. 36th ACM STOC*, pages 11–20, 2004.
- [14] A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. Agnostically learning halfspaces. In *Proc. 46th IEEE FOCS*, pages 11–20, 2005.
- [15] S. Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. In *Proc. 43rd IEEE FOCS*, pages 23–32, 2002.
- [16] S. Khot and R. Saket. A 3-query non-adaptive PCP with perfect completeness. In *Proc. IEEE CCC*, pages 159–169, 2006.
- [17] A. Klivans, R. O’Donnell, and R. Servedio. Learning intersections and thresholds of halfspaces. In *Proc. 43rd IEEE FOCS*, pages 177–186, 2002.
- [18] A. Klivans and R. Servedio. Learning intersections of halfspaces with a margin. In *Proc. COLT*, pages 348–362, 2004.
- [19] A. Klivans and A. Sherstov. Cryptographic hardness for learning intersections of halfspaces. In *Proc. 47th IEEE FOCS*, pages 553–562, 2006.
- [20] A. Levin and A. Shashua. Principal component analysis over continuous subspaces and intersection of half-spaces. In *Proc. ECCV(3)*, pages 635–650, 2002.
- [21] D. Murphy. Nearest neighbor pattern classification perceptrons. *Proceedings of the IEEE*, 78(10):1595–1598, 1990.
- [22] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [23] U. Rückert, L. Richter, and S. Kramer. Quantitative association rules based on half-spaces: An optimization approach. In *Proc. ICDM*, pages 507–510, 2004.
- [24] R. Schapire. The strength of weak learnability. *Machine Learning*, 5:197–227, 1990.
- [25] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [26] S. Vempala. A random sampling based algorithm for learning the intersection of half-spaces. In *Proc. 38th IEEE FOCS*, pages 508–513, 1997.

APPENDIX

A. PROOF OF LEMMA 5

First we need some definitions.

DEFINITION 3. A range space is a pair (X, \mathcal{F}) , where X is a set and \mathcal{F} is a family of subsets of X , i.e. $\mathcal{F} \subseteq 2^X$.

DEFINITION 4. For any set $A \subseteq X$, define $P_{\mathcal{F}}(A)$ the projection of \mathcal{F} onto A , as $P_{\mathcal{F}}(A) := \{F \cap A : F \in \mathcal{F}\}$.

DEFINITION 5. We say that a set $A \subseteq X$ is shattered by (X, \mathcal{F}) if $P_{\mathcal{F}}(A) = 2^A$.

The VC dimension of a range space is defined as follows.

DEFINITION 6. *The VC dimension of (X, \mathcal{F}) is the cardinality of the maximum set it shatters, i.e. $VC \dim = \sup\{|A| : A \text{ is shattered}\}$. It may be infinite.*

We use the following theorem regarding sampling from range spaces of bounded VC dimension [25].

THEOREM 10. *Let (X, \mathcal{F}) a range space of VC dimension d , and let ϕ be a uniform measure on X . There is a universal constant C_{VC} such that with probability at least $1 - \delta$, a random set $S \subseteq X$ of size,*

$$C_{VC} \left(\frac{d}{\tau^2} \log \left(\frac{d}{\tau} \right) + \frac{1}{\tau^2} \log \left(\frac{1}{\delta} \right) \right)$$

satisfies,

$$\left| \frac{|S \cap F|}{|S|} - \phi(F) \right| \leq \tau,$$

for all $F \in \mathcal{F}$.

Let $N = nk$, let \mathbf{S}^{N-1} be the unit sphere in N dimensions, and let ϕ be a uniform measure over \mathbf{S}^{N-1} . Define $P(\mathbf{r}, [a, b]) := \{\mathbf{z} \in \mathbf{S}^{N-1} \mid \mathbf{r} \cdot \mathbf{z} \in [a, b]\}$. The set $P(\mathbf{r}, [a, b])$ is exactly the set of unit vectors whose dot product with \mathbf{r} lies in the interval $[a, b]$. Let $\mathcal{P} := \{P(\mathbf{r}, [a, b]) \mid \|\mathbf{r}\| = 1, b - a = \varepsilon'/(100\sqrt{N})\}$. By Stirling's approximation, we have that for large enough N , the surface area of the \mathbf{S}^{N-1} is at least $1/\sqrt{N}$ times the surface area of \mathbf{S}^{N-2} . As a result the following fact is easy to derive.

FACT 11. *For large enough N , for any $P \in \mathcal{P}$, $\phi(P) \leq \varepsilon'/10$.*

Moreover, we observe that every element $P \in \mathcal{P}$ is a set of points in \mathbf{S}^{N-1} that lie in an intersection of two halfspaces. Since, in \mathbf{R}^N , the VC dimension of the class of all N dimensional halfspaces is $N + 1$, the VC dimension of the class of N dimensional halfspaces for the set \mathbf{S}^{N-1} is at most $N + 1$. Using this we have the following bound which we state without proof.

LEMMA 12. *The VC dimension of the range space $(\mathbf{S}^{N-1}, \mathcal{P})$ is at most $4(N + 1) \log(2(N + 1))$.*

Now if R is a set of N^2 , random unit vectors from \mathbf{S}^{N-1} , then for large enough N , Theorem 10 along with the above lemma implies that with probability at least $1 - 1/n$,

$$\left| \frac{|R \cap P|}{|R|} - \phi(P) \right| \leq \varepsilon'/10,$$

for any $P \in \mathcal{P}$. The above, coupled with Fact 11 implies that with probability $1 - 1/n$ over the choice of R ,

$$\frac{|R \cap P|}{|R|} \leq \varepsilon'/5,$$

for all $P \in \mathcal{P}$. In other words, with probability $1 - 1/n$ over the choice of R , for any unit vector \mathbf{r} , at most $\varepsilon'/5$ fraction of points in R are contained in the set $\{\mathbf{z} \mid \mathbf{r} \cdot \mathbf{z} \in [a, b]\}$ for any a, b s.t. $b - a = \varepsilon'/(100\sqrt{N})$. This implies that with probability $1 - 1/n$ over the choice of R , for any set $T \subseteq R$, such that $|T| = \varepsilon'|R|$, for any unit vector \mathbf{r} , there exist $\mathbf{z}', \mathbf{z}'' \in T$, such that $|\mathbf{r} \cdot \mathbf{z}' - \mathbf{r} \cdot \mathbf{z}''| \geq \varepsilon'/(100\sqrt{N})$. This proves Lemma 5.

B. PROOF OF THEOREM 3

Let us first define the 'smooth' version of the bipartite label cover problem.

DEFINITION 7. *An instance of SMOOTHLABELCOVERBIPARTITE(k, m, T) consists of a bipartite graph $G(U, V, E)$, where the vertices in U have the same degree, and a set of projections $\pi^{vu} : [k] \mapsto [m]$ for all $\{u, v\} \in E$ such that $u \in U, v \in V$. A labeling σ to the vertices in G satisfies and edge $e = \{u, v\}$ s.t. $u \in U, v \in V$ iff $\pi^{vu}(\sigma(v)) = \sigma(u)$. Moreover, for any vertex $v \in V$ for any $i, j \in [k], i \neq j$,*

$$\Pr_{e=\{u,v\} \in E} [\pi^{vu}(i) = \pi^{vu}(j)] \leq \frac{1}{T}, \quad (18)$$

where the probability is taken over a random edge incident on v .

The following theorem was proved in [16], using the PCP Theorem [2, 3] and Raz's Parallel Repetition Theorem [22].

THEOREM 13. *For any constant $\delta > 0$, for any constant $T > 0$, there exist k and m such that given an instance \mathcal{L}' of SMOOTHLABELCOVERBIPARTITE(k, m, T) it is NP-hard to distinguish between the following two cases, YES Case: There is a labeling to the vertices of \mathcal{L}' that satisfies all the edges. NO Case: No labeling to the vertices of \mathcal{L}' satisfies more than δ fraction of the edges.*

The construction of an instance \mathcal{L} of SMOOTHLABELCOVER(t, μ, ν, k, m) is as follows. We start with an instance \mathcal{L}' of SMOOTHLABELCOVERBIPARTITE(k, m, T) where we will fix T later. The vertex set of \mathcal{L} is the V side of \mathcal{L}' . An edge of \mathcal{L} is constructed as follows: select a vertex u from U and for every two neighbors v_1 and v_2 of u in \mathcal{L}' , add an edge e between them in \mathcal{L} . Set $\pi_{e,v_1} = \pi^{v_1 u}$ and $\pi_{e,v_2} = \pi^{v_2 u}$ in E . Let $E(u)$ be the set of such edges added in \mathcal{L} corresponding to a vertex $u \in U$. Note that we are constructing a multi-graph, since two vertices v_1 and v_2 in V might share two different neighbors in U , in which case there will be multiple edges between v_1 and v_2 in \mathcal{L} . Clearly the sets $E(u)$ for $u \in U$ are a partition of edges in \mathcal{L} , and since U side is regular, the sets $E(u)$ are of equal size. Essentially, we are adding a clique of edges $E(u)$ corresponding to u on its neighborhood $N(u) \subseteq V$ for every $u \in U$. Let v be a vertex in V and let $S \subseteq [k]$ be a set of size t , then applying equation (18) to all pairs in S and taking union bound, we have,

$$\Pr_{e:v \in e} [\exists x, y \in S, x \neq y : \pi_{e,v}(x) = \pi_{e,v}(y)] \leq \frac{t^2}{T}$$

where the probability is over the edges incident on v in \mathcal{L} . Note that we have used the fact the vertices in U have the same degree. Now, taking T to be large enough, we can reduce this probability to at most μ . To verify the second property, let V' be a subset of V such that $|V'| = \xi|V|$, for some $\xi > 0$. Now, consider a vertex u in U , and let p_u be the probability that a random neighbor of u falls in V' . From the proof of Theorem 13 one can see that vertices on U side have the same degree, say d which can be increased to any arbitrary constant by parallel repetition. Therefore, $E_u[p_u] = \xi$. Moreover, the fraction of edges in \mathcal{L} that lies inside V' is the probability for a random $u \in U$,

a random pair of its neighbors lies in V' . For a particular u this is $p_u^2 - 1/d$, where $1/d$ is the probability of selecting the same vertex twice out of d neighbors of u . Hence, we have that the fraction of edges induced by V' in \mathcal{L} is $\mathbb{E}_u[p_u^2 - 1/d] \geq (\mathbb{E}_u[p_u])^2 - 1/d = \xi^2 - 1/d$. This fraction is at least $\xi^2/2$ if $\xi \geq \sqrt{\frac{2}{d}}$, and so we are done by taking $\nu = \sqrt{\frac{2}{d}}$ which can be made arbitrarily small by taking d to be large enough.

Now, if \mathcal{L}' was a YES instance, then there is a labeling σ to vertices $U \cup V$ that satisfies all the edges of \mathcal{L}' . This implies,

$$\pi^{v_1 u}(\sigma(v_1)) = \sigma(u) = \pi^{v_2 u}(\sigma(v_2)),$$

for all edges $e_1 = \{u, v_1\}, e_2 = \{u, v_2\}$ of \mathcal{L}' , $u \in U, v_1, v_2 \in N(u) \subseteq V$, where $N(u)$ is the neighborhood of $u \in U$ in \mathcal{L}' . Consider the edge $e \in E(u)$ between v_1 and v_2 in \mathcal{L} . Clearly, $\pi_{e, v_1}(\sigma(v_1)) = \pi_{e, v_2}(\sigma(v_2))$. Therefore, the labeling σ restricted to V satisfies all the edges of \mathcal{L} .

Now consider a labeling σ' to V that satisfies ε fraction of the edges in \mathcal{L} . Consider any vertex $u \in U$. For $j \in [m]$, let $S_u^j \subseteq N(u)$ the set of vertices $v \in N(u)$ such that $\pi^{v u}(\sigma'(v)) = j$. It can be seen that the sets S_u^j ($j \in [m]$) form a partition of $N(u)$ and the disjoint union of edges (corresponding to u) induced by each S_u^j in $E(u)$ is exactly the subset of edges of $E(u)$ that are satisfied by σ' . Let $l_u = \operatorname{argmax}_j |S_u^j|$ for each $u \in U$. Observe that seen that any subset S of $N(u)$ containing c ($c < 1$) fraction of vertices of $N(u)$ induces in $E(u)$ at most c^2 fraction of the total edges of $E(u)$. Suppose σ' satisfies ε_u fraction of the edges of $E(u)$, then a simple argument shows that $S_u^{l_u}$ must contain at least ε_u fraction of vertices in $N(u)$. Now σ' satisfies ε fraction of all the edges of \mathcal{L} , and since the U side is regular in \mathcal{L}' , we have that $\mathbb{E}_u[|S_u^{l_u}|/|N(u)|] \geq \mathbb{E}_u[\varepsilon_u] \geq \varepsilon$. Therefore, by extending the labeling σ' to U by setting $\sigma'(u) = l_u$ for $u \in U$, we can satisfy the edges of \mathcal{L}' between vertices of $S_u^{l_u}$ and u for all $u \in U$. This would satisfy ε fraction of the edges in \mathcal{L}' . So, if the instance of \mathcal{L}' is a NO instance with soundness η then there is no labeling to the vertices of \mathcal{L} which satisfies more than η fraction of the edges of \mathcal{L} . This completes the proof of Theorem 3.