

Hardness of Learning DNFs using Halfspaces

Suprovat Ghoshal*

Rishi Saket†

Abstract

The problem of learning t -term DNF formulas (for $t = O(1)$) has been studied extensively in the PAC model since its introduction by Valiant (STOC 1984). A t -term DNF can be efficiently learnt using a t -term DNF only if $t = 1$ i.e., when it is an AND, while even *weakly* learning a 2-term DNF using a constant term DNF was shown to be NP-hard by Khot and Saket (FOCS 2008). On the other hand, Feldman et al. (FOCS 2009) showed the hardness of weakly learning a noisy AND using a *halfspace* – the latter being a generalization of an AND, while Khot and Saket (STOC 2008) showed that an intersection of two halfspaces is hard to weakly learn using any function of constantly many halfspaces. The question of whether a 2-term DNF is efficiently learnable using 2 or constantly many halfspaces remained open. In this work we answer this question in the negative by showing the hardness of weakly learning a 2-term DNF as well as a noisy AND using any function of a constant number of halfspaces. In particular we prove the following.

For any constants $\nu, \zeta > 0$ and $\ell \in \mathbb{N}$, given a distribution over point-value pairs $\{0, 1\}^n \times \{0, 1\}$, it is NP-hard to decide whether,

- YES Case. There is a 2-term DNF that classifies all the points of the distribution, and an AND that classifies at least $1 - \zeta$ fraction of the points correctly.
- NO Case. Any boolean function depending on at most ℓ halfspaces classifies at most $1/2 + \nu$ fraction of the points of the distribution correctly.

Our result generalizes and strengthens the previous best results mentioned above on the hardness of learning a 2-term DNF, learning an intersection of two halfspaces, and learning a noisy AND.

*Indian Institute of Science. E-mail: suprovat@iisc.ac.in

†IBM Research, Bangalore. E-mail: rissaket@in.ibm.com

1 Introduction

A boolean function $f : \{0, 1\}^n \mapsto \{0, 1\}$ can always be represented in a *disjunctive normal form* (DNF) consisting of an OR over ANDs of boolean literals. DNFs are widely studied in several fields in computer science such as complexity theory [SBI04] [Ale05][AMO15], learning theory [KLW10] [Lee10], pseudorandomness [Baz09] [GMR13] and combinatorics [LLZ18] [LZ19] to name a few. In many scenarios, such as circuit design or neural network training, the *size* - given by the number of terms - of a DNF is an important consideration, and much work has studied the properties of small DNFs. Contributing to this line of research, our work investigates the learnability of bounded term DNFs in the *probably approximate correct* (PAC) model [Val84].

Formally, a *concept class* \mathcal{C} of boolean functions is efficiently learnable if, given access to random samples from a distribution over point-value pairs realizable by an (unknown) member of \mathcal{C} , there exists a polynomial time randomized algorithm which outputs with high probability a hypothesis (from a *hypothesis class* \mathcal{H}) which is consistent with the distribution's samples with probability (accuracy) arbitrarily close to 1. A few useful variants are *proper learning* when the hypothesis is restricted to be from the concept class ($\mathcal{H} = \mathcal{C}$), and *weak learning* wherein the accuracy of learning is allowed to be any constant greater than the trivially achievable random threshold of half.

The question of efficient learnability of DNFs was raised in Valiant's celebrated work [Val84] which introduced PAC learning. While a 1-term DNF (i.e., an AND), is properly learnable, it is NP-hard to properly learn a t -term DNF [PV88], for any $t \geq 2$. This is not true when allowed more general hypotheses (improper learning): Valiant [Val84] showed that a t -term DNF can be efficiently learnt using a *conjunctive normal form* (CNF) which is an AND of ORs (clauses), where each clause has at most t literals.

For learning unbounded term DNFs, Bshouty [Bsh96] first gave a $2^{O(\sqrt{n \log t})}$ -time algorithm which was later improved by Klivans and Servedio [KS04]. They showed that any DNF can be expressed as a polynomial threshold of degree $O(n^{1/3} \log t)$, and therefore can be learnt using linear programming in time $n^{\tilde{O}(n^{1/3})}$ which is state of the art for learning unrestricted DNF.

On the complexity side, assuming $\text{NP} \not\subseteq \text{ZPP}$, Nock, Jappy and Salantin [NJS98] showed that any n^c -sized DNFs cannot efficiently be learnt using $n^{\gamma c + \eta}$ -term DNFs, for certain ranges γ, η , which was later strengthened by Alekhovich et al. [ABF⁺08] to any constant value of γ , assuming $\text{NP} \neq \text{RP}$ thus ruling out proper learning of arbitrary DNFs in polynomial time. Subsequently, Feldman [Fel09] showed this to hold even with access to membership queries. While these results do not rule out properly learning restricted DNFs, Alekhovich et al. [ABF⁺08] showed that it is NP-hard to learn a 2-term DNFs using a t -term DNF for any constant t . Subsequently, this was strengthened by Khot and Saket [KS08a] who showed the NP-hardness of even weakly learning 2-term DNFs using constant term DNFs. In related work, Feldman [Fel06] and Feldman et al. [FGKP09] showed the hardness of weak agnostic learning *noisy* AND with an AND i.e., in the presence of a small fraction of adversarially perturbed sample labels, whereas Khot and Saket [KS08a] showed the same when using a CNF with bounded clause width as hypothesis¹, complementing Valiant's [Val84] algorithmic work mentioned above in the non-noise setting.

Linear threshold functions (LTFs), a.k.a. *halfspaces*, given by $\{\mathbf{x} : \langle \mathbf{c}, \mathbf{x} \rangle + \theta > 0\}$, are a natural and well studied class of hypothesis used in machine learning, especially as the output of linear kernel models such as Perceptron and SVMs. They generalize DNFs: a t -term DNF is can be represented as an OR of t halfspaces. While Valiant's work [Val84] showed that a t -term DNF can be learnt

¹This also rules out as hypothesis AND of $O(1)$ -arity functions, as the latter can be represented as $O(1)$ -width CNFs.

using an intersection of *unbounded* number of halfspaces, using linear programming a 1-term DNF can be efficiently learnt using one halfspace. A natural question is whether the latter holds for $t = 2$ as well:

$$\text{Can a 2-term DNF be efficiently learnt using two halfspaces?} \quad (1)$$

In this context Khot and Saket [KS08b] showed that an intersection of two halfspaces – which captures (the negation of) a 2-term DNF – cannot be efficiently weakly learnt using any function of constant number of halfspaces, unless $\text{NP} \neq \text{RP}$. Feldman et al. [FGRW12] showed the hardness of weakly learning a noisy AND using a single halfspace, generalizing previous works of Guruswami and Raghavendra [GR09] and Feldman et al. [FGKP09], which proved the same hardness for learning a noisy halfspace, and results of [Fel06, FGKP09] on the hardness of learning a noisy AND with an AND.

In our work we answer (1) in the negative, proving the NP-hardness of learning 2-term DNF (as well as noisy AND) using functions of any constant number of halfspaces. In particular we prove:

Theorem 1.1 *For any constants $\nu, \zeta > 0$ and $\ell \in \mathbb{N}$, the following holds. Given a distribution² \mathcal{D} over $\{0, 1\}^n \times \{0, 1\}$, it is NP-hard to distinguish between the two following cases:*

- YES Case: *There exist two ANDs \mathcal{A}_1 and \mathcal{A}_2 such that*

$$\Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [(\mathcal{A}_1(\mathbf{x}) \vee \mathcal{A}_2(\mathbf{x})) = a] = 1, \quad (2) \quad \Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [\mathcal{A}_1(\mathbf{x}) = a] \geq 1 - \zeta. \quad (3)$$

- NO Case: *Any boolean function f depending on at most ℓ halfspaces satisfies,*

$$\Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [f(\mathbf{x}) = a] \leq \frac{1}{2} + \nu.$$

Our result strengthens and generalizes the hitherto best results on the hardness of learning (i) constant term DNFs using constant term DNFs, (ii) noisy ANDs using halfspaces, and (iii) intersections of two halfspaces using constantly many halfspaces, given in the works of [KS08a], [FGRW12] and [KS08b]. An interesting feature of our result is that the hardness results for 2-term DNF as well as noisy AND and *simultaneously* shown through the same hard instance satisfying (2) and (3) in Theorem 1.1. We note that our result regarding 2-term DNF is tight in terms of the degree of the hypothesis threshold function: any 2-term DNF can be efficiently learnt using a single quadratic threshold. More generally, by representing the negation of a t -term DNF as the sign of a product of t linear forms, a t -term DNF is learnable using a degree- t *polynomial threshold function* (PTF) in time $n^{O(t)}$.

1.1 Other Related Work

The problem of learning DNFs has also been studied in distributional settings. In particular, Verbeugt [Ver90] gave a quasi-polynomial time algorithm for learning DNFs under the uniform distribution on the hypercube. Blum et al. [BFJ⁺94] gave a polynomial time algorithm for weakly learning DNFs under the uniform distribution using membership queries. This was subsequently improved on by Jackson [Jac97] who gave a polynomial time algorithm for strongly learning DNFs in the same setting. There have been several other works which have studied the problem of learning DNFs under a wide range of alternative restricted settings [BMOS03][Ser04][KST09][Fel12].

²The distribution in our reduction is explicitly described by a polynomial sized set of point-label pairs.

There has been extensive work on algorithms which learn concepts using halfspaces as the output hypothesis class such as Perceptron [Ros58], SVMs [V+98] etc., as well as on algorithmic results for learning more general concepts. Klivans et al. [KOS04] showed that one can learn arbitrary functions of constant number of halfspaces in quasi-polynomial time under the uniform distribution. Diakonikolas et al. [DHK+10] gave an algorithm that learns degree- d PTFs in time $2^{O_\varepsilon(d^2)}$ under the uniform distribution over the hypercube. Recently Gottlieb et al. [GKKN18] gave an algorithm for PAC learning intersections of t -halfspaces with intersections of $t \log t$ halfspaces with margin γ in time $n^{O(1/\gamma^2)} \text{poly}(t)$. On the hardness side, Diakonikolas et al. [DOSW11] improved on the works of [GR09, FGKP09] showing the intractability of weakly learning a noisy halfspace using degree-2 PTFs, which was generalized to rule out all constant degree PTFs as hypotheses by Bhattacharyya et al. [BGS18].

There are also hypothesis independent hardness of learning results for learning (relatively) small DNFs. In [DSS16], Daniely showed that unless there exist efficient algorithm for strongly refuting K -XOR formulas for certain clause densities, there does not exist any polynomial time algorithm for learning t -term DNFs (with $t = \omega(\log n)$), using any polynomial time evaluable hypothesis class. Earlier, Klivans and Sherstov [KS09] had shown that there is no efficient algorithm for learning intersection of n^ε -halfspaces using unrestricted hypothesis classes, under different cryptographic hardness assumptions. Such a result is not possible for $O(1)$ -term DNF as it is efficiently learnable using a CNF (by [Val84]) and by an $O(1)$ -degree PTF as observed earlier in this section.

On this note, we remark that the result of Applebaum, Barak and Xiao [ABX08] shows that hypothesis independent hardness of learning results assuming $P \neq NP$ are unlikely without making significant breakthroughs in complexity theory. Therefore, any complexity theoretic hardness of learning result for noisy AND will probably require some restriction of the hypothesis. While our result rules out functions of constantly many halfspaces as hypotheses, it remains an important open problem to show the same for (functions of) polynomial thresholds.

1.2 Overview of the Reduction

Our reduction is from a variant of Label Cover and uses the standard template of a suitably constructed *dictatorship test* defined over a set of coordinates. This test gives a distribution over point-value pairs when applied to the coordinates corresponding to each edge (or a local collections of edges) of the Label Cover, and the union of these distributions constitutes the instance of the learning problem. Roughly speaking, the distribution for an edge should satisfy (i) (completeness) any matching labeling should yield a good classifier from the concept class, and (ii) (soundness) any good enough hypothesis should yield a matching labeling to the vertices of the edge. For ease of exposition in this section, we shall not elaborate on our hardness for learning noisy AND, and focus on the 2-term DNF concept class. Furthermore, instead of 2-term DNF we shall consider 2-clause CNF (these are essentially equivalent: negating one yields the other).

We begin by describing in Figure 1 an elementary dictatorship test – a generalized version of which was used by [KS08a] to rule out weakly learning a 2-clause CNF using constant clause CNF (equivalently a 2-term DNF with constant term DNF). Let us denote by \mathcal{I} the generated distribution, and by \mathcal{I}_0 and \mathcal{I}_1 its restrictions to the 0 and 1 valued points respectively.

For the completeness, observe that for every $i \in [M]$, the CNF: $X_i \wedge (\bigvee_{r \in [k]} Y_{r,i})$ classifies all the points of \mathcal{D} . On the other hand, it can be shown that any t -clause CNF that classifies $1/2 + \Omega(1)$ fraction of the points correctly yields two different \mathbf{Y}_r and $\mathbf{Y}_{r'}$ that can be *decoded* into a common element of $[M]$. The analysis in [KS08a] used fairly simple structural and probabilistic arguments

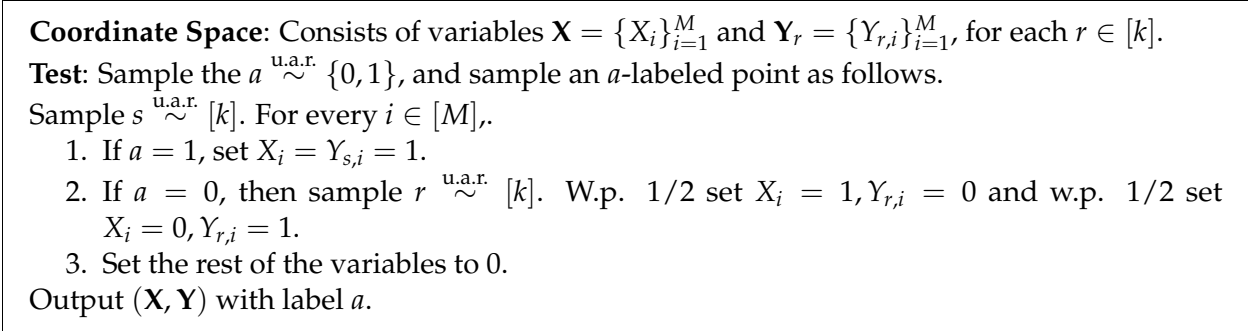


Figure 1: Distribution \mathcal{I}

and we refer the reader to it for the detailed proof.

Unfortunately, it is easy to see that the above distribution fails with halfspaces as hypotheses because of the following simple observation: the first moment of each of the variables $\{X_i\}_{i=1}^M$ under \mathcal{I}_1 is 1, while it is $1/2$ under \mathcal{I}_0 . Since $\{X_i\}_{i=1}^M$ are independent Bernoulli variables in both cases, by straightforward concentration this difference in first moments can be leveraged by a non-dictatorial halfspace (whose linear form simply sums up the X_i s) to classify the points with high probability.

Another issue that impairs \mathcal{I} from working for halfspaces is the lack of *noise*. Consider a linear form whose coefficient for each X_i is 2^i and for each $\{Y_{r,i}\}_{r=1}^k$ is (-2^i) . Such a linear form always evaluates to 0 on \mathcal{I}_1 while it is always non-zero on a \mathcal{I}_0 , and therefore this instance admits a good classifier using an intersection of two halfspaces. This pathological example also illustrates the complications arising out of coefficient vectors that are *not regular* i.e., they contain large sequences of geometrically decreasing coefficients.

Our test – represented by the distribution \mathcal{D} and its restrictions \mathcal{D}_0 and \mathcal{D}_1 – is designed to fix the above shortcomings. Firstly, we ensure by having k blocks of the \mathbf{X} variables as well, that the individual marginal distribution of any variable is the same for both, \mathcal{D}_0 and \mathcal{D}_1 . In order to introduce noise, a large common randomly chosen set of variables are sampled i.i.d Bernoulli, under both \mathcal{D}_0 and \mathcal{D}_1 . Restricted to each $i \in [M]$, this set resides either in \mathbf{X} or \mathbf{Y} depending on which side holds non-zero variables under \mathcal{D}_0 . Furthermore, the test also incorporates Label Cover projections $[M] \rightarrow [m]$, so that all the variables corresponding to those labels projecting to some $j \in [m]$ are sampled together, and independently of those projecting to $j' \neq j$. In order to bound the variance of the difference of samples from \mathcal{D}_0 and \mathcal{D}_1 , after fixing the common noise set, we replace each variable Z with a collection of Q variables $\{Z_q\}_{q=1}^Q$, sampling them u.a.r. from $\{0, 1\}^Q$ if the Z was part of the noise set, u.a.r. from³ $\{\mathbf{e}_1, \dots, \mathbf{e}_Q\}$ if Z contributed to the difference between \mathcal{D}_0 and \mathcal{D}_1 , and are each set to 0 if Z was set to 0. The parameter Q is taken to be much larger than the pre-image size of the Label Cover projections.

The analysis of a coefficient vector utilizes the notion of its *critical index* which was introduced by Servedio [Ser06], and also used in the above mentioned work of [FGRW12]. As part of their analysis, [FGRW12] proved the invariance of halfspaces with regular coefficients under distributions with matching first to third moments. They apply this to \mathcal{D}_0 and \mathcal{D}_1 ensuring that their appropriate moments are matched. In our case however, this is impossible for points classified by a 2-clause

³Here \mathbf{e}_q is the Q -length indicator vector of the q th coordinate, $q \in [Q]$.

CNF, since the product of the two linear forms representing the CNF is always zero under \mathcal{D}_0 , and positive under \mathcal{D}_1 . The invariance shown by [FGRW12] only bounded the deviation of the biases of halfspaces under \mathcal{D}_0 and \mathcal{D}_1 , and thus their analysis only rules out single halfspaces as hypothesis. In contrast, our work bounds the expected point-wise deviation allowing us to extend our result to functions of constantly many halfspaces. For technical considerations, as in [FGRW12] we combine our test with *Smooth Label Cover* whose (smoothness) property guarantees that the projections for most its edges behave as bijections when restricted to fixed small sets of labels for its vertices.

At a high level our analysis proceeds as follows: we first truncate the coefficient vector to a small irregular part. Subsequently, our analysis bounds the variance of the difference between \mathcal{D}_0 and \mathcal{D}_1 given a noise set which we show picks up enough mass with high probability. Finally, we directly apply anti-concentration on the noise variables to bound the deviation between \mathcal{D}_0 and \mathcal{D}_1 . In the rest of this section we present a simplified version of the our test distribution, and informally describe its analysis. A slightly modified version (with differently biased random choices) yields the desired hardness result for noisy AND as well.

1.3 Simplified Distribution and Sketch of Analysis

Consider the distribution \mathcal{D} given in Figure 2. The completeness case is straightforward. Indeed, fix any $i^* \in [M]$ and the following 2-clause CNF: $f^*(\mathbf{X}, \mathbf{Y}) := \left(\bigvee_{r \in [k]} \bigvee_{q \in [Q]} X_{r,i^*,q} \right) \wedge \left(\bigvee_{r \in [k]} \bigvee_{q \in [Q]} Y_{r,i^*,q} \right)$. Let j^* be such that $i \in B_{j^*}$. If $a = 0$, then $b_{j^*} = 0$ which implies $\mathbf{Y}^{(j^*)} = \mathbf{0}$, and $b_{j^*} = 1$ which implies $\mathbf{X}^{(j^*)} = \mathbf{0}$. Therefore, either $\bigvee_{r \in [k]} \bigvee_{q \in [Q]} X_{r,i^*,q} = 0$, or $\bigvee_{r \in [k]} \bigvee_{q \in [Q]} Y_{r,i^*,q} = 0$, which means that $f^* = 0$. On the other hand if $a = 1$, then there are r, r' s.t. $\mathbf{X}_{r,i^*}, \mathbf{Y}_{r',i^*} \in \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$, implying that $f^* = 1$.

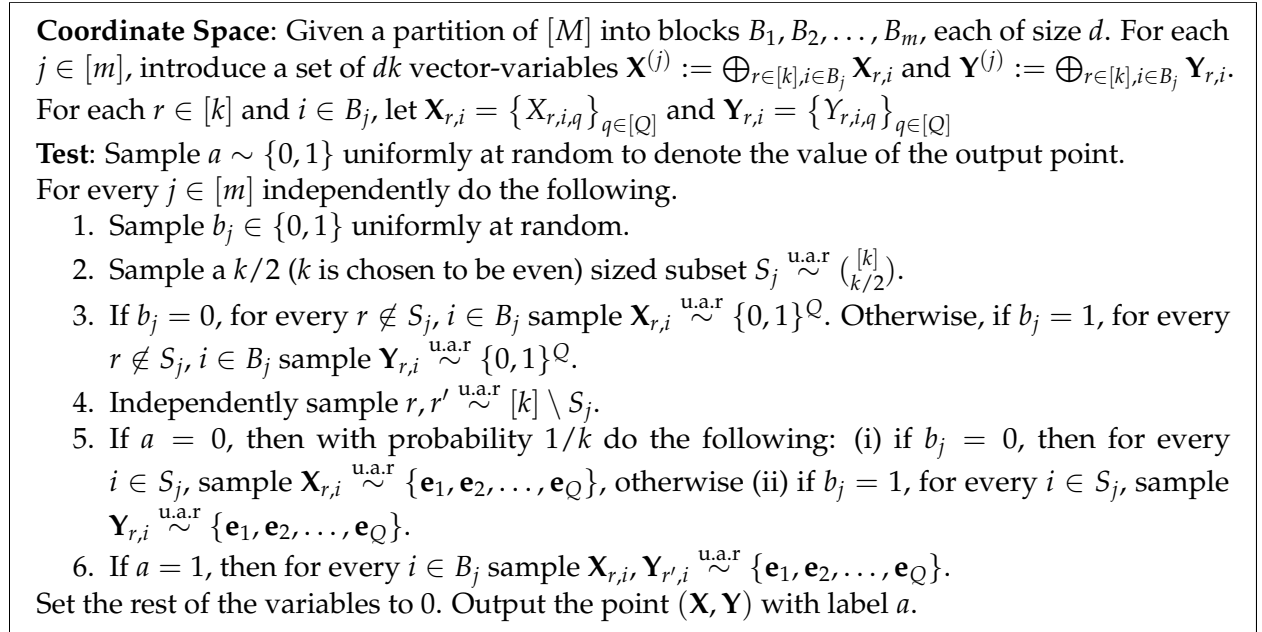


Figure 2: Distribution \mathcal{D}

For the soundness assume that there is some function f of ℓ halfspaces $\text{pos}(h_s(\mathbf{X}, \mathbf{Y}))$, $s = 1, \dots, \ell$ that classifies $1/2 + \nu$ fraction of the points of \mathcal{D} correctly, where $\text{pos}(\cdot)$ is the sign function.

Using the definitions related to the critical index [Ser06] (see Section 2.1) we define the sets $C_{X,s,r}$, $C_{\bar{X},s,r}^{\leq K}$ and $I_{X,s,r}$, letting \mathbf{c} to be the coefficient vector of h_s corresponding to the variables \mathbf{X}_r , as follows:

- $C_{X,s,r}$ is the set $C_\tau(\mathbf{c})$, and $C_{\bar{X},s,r}^{\leq K} \subseteq C_{X,s,r}$ are the top K elements of $C_{X,s,r}$ given by $C_{\bar{\tau}}^{\leq K}(\mathbf{c})$
- $I_{X,s,r} := C_{\bar{X},s,r}^{\leq K} \cup \{i \in [M] \setminus C_{X,s,r} : \|\mathbf{c}_i\|^2 > (1/d^8) \sum_{i' \in [M] \setminus C_{X,s,r}} \|\mathbf{c}_{i'}\|^2\}$,

and similarly $C_{Y,s,r}$, $C_{\bar{Y},s,r}^{\leq K}$ and $I_{Y,s,r}$, for $s \in [\ell]$, $r \in [k]$, and some parameters τ and K that we choose appropriately. Clearly the size of any $I_{X,s,r}$ or $I_{Y,s,r}$ is at most $2(K + d^8)$. Using the smoothness property of the Label Cover instance we can assume that for each r

$$\left| \left(\bigcup_{s \in [\ell]} I_{X,s,r} \right) \cap B_j \right| \leq 1, \text{ and } \left| \left(\bigcup_{s \in [\ell]} I_{Y,s,r} \right) \cap B_j \right| \leq 1 \text{ for all } j \in [m]. \quad (4)$$

Further we may assume that no two subsets from $\{C_{\bar{X},s,r}^{\leq K}, C_{\bar{Y},s,r}^{\leq K} : s \in [\ell], r \in [k]\}$ can have indices from the same block B_j . Otherwise using (4), there are two distinct tuples (Z_1, r) and (Z_2, r') where $Z_1, Z_2 \in \{X, Y\}$ such that $C_{\bar{Z}_1,s,r}^{\leq K}$ and $C_{\bar{Z}_2,s',r'}^{\leq K}$ share indices from a common block. This yields a good labeling. This, along with (4) implies that any B_j can have at most one element from $\bigcup_{s \in [\ell], r \in [k]} (C_{\bar{X},s,r}^{\leq K} \cup C_{\bar{Y},s,r}^{\leq K})$.

Pairing Distribution. Using the description of \mathcal{D} along with the above observations one can construct a distribution $\widehat{\mathcal{D}}$ over $((\mathbf{X}^0, \mathbf{Y}^0), (\mathbf{X}^1, \mathbf{Y}^1))$ such that its marginals are \mathcal{D}_0 and \mathcal{D}_1 , as follows.

- Sample the subsets $\{S_j\}_{j \in [m]}$, and for all $j \in [m]$ such that B_j contains some element (at most one from above) from $\bigcup_{s \in [\ell], r \in [k]} (C_{\bar{X},s,r}^{\leq K} \cup C_{\bar{Y},s,r}^{\leq K})$, sample b_j .
- Sample the values of all $\mathbf{X}_{r,i}$ s.t. $i \in C_{\bar{X},s,r}^{\leq K}$ for some $s \in [\ell]$, and all $\mathbf{Y}_{r,i}$ s.t. $i \in C_{\bar{Y},s,r}^{\leq K}$ for some $s \in [\ell]$. Now sample the rest of the b_j s.
- Sample values of the rest of (*noise variables*) $\mathbf{X}_{r,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q$ s.t. $r \notin S_j, b_j = 0$ and $\mathbf{Y}_{r,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q$ s.t. $r \notin S_j, b_j = 1$.
- Finally, conditioned on all these fixings, sample $(\mathbf{X}^0, \mathbf{Y}^0) \leftarrow \mathcal{D}_0$ and $(\mathbf{X}^1, \mathbf{Y}^1) \leftarrow \mathcal{D}_1$.

Note that this distribution is independent of the choice of any specific halfspace $\text{pos}(h_s(\mathbf{X}, \mathbf{Y}))$, which is crucial to the argument. From the goodness of the classifier f and an averaging argument we obtain that there exists one out of the ℓ halfspaces, $\text{pos}(h_{s^*}(\mathbf{X}, \mathbf{Y}))$ such that,

$$\mathbb{E}_{\widehat{\mathcal{D}}} \left[\left| \text{pos}(h_{s^*}(\mathbf{X}^1, \mathbf{Y}^1)) - \text{pos}(h_{s^*}(\mathbf{X}^0, \mathbf{Y}^0)) \right| \right] \geq \nu/2\ell. \quad (5)$$

Truncation Step. This step uses the property that squared mass corresponding to the indices in $C_{\bar{X},s^*,r}^{\leq K}$ is much larger than that in $C_{X,s^*,r} \setminus C_{\bar{X},s^*,r}^{\leq K}$ and similarly for $C_{\bar{Y},s^*,r}^{\leq K}$ and $C_{Y,s^*,r}$, $r \in [k]$. Using this we show that in the coefficient vector of h_{s^*} corresponding to some \mathbf{X}_r can be truncated by zeroing out those coefficients corresponding to $C_{X,s^*,r} \setminus C_{\bar{X},s^*,r}^{\leq K}$, while not disturbing (5) appreciably. This is obtained by bounding the deviation in the value of h_{s^*} due to this zeroing out, and showing that this is overwhelmed by the anti-concentration of the noisy variables corresponding to $C_{\bar{X},s^*,r}^{\leq K/4}$. Our anti-concentration bound leverages the *Littlewood-Offord-Erdős* lemma (see Lemma 2.2), as well as the standard Berry-Esseen bound. Doing this truncation for coefficients corresponding to each \mathbf{X}_r and \mathbf{Y}_r , we obtain a linear form h which satisfies (5) with $\nu/4\ell$ on the RHS. Let the coefficients of h be given by \mathbf{c}_X and \mathbf{c}_Y corresponding to \mathbf{X} and \mathbf{Y} respectively. Note that truncation implies,

$C_{\bar{\tau}}^{\leq K}(\mathbf{c}_{X,r}) = C_{\tau}(\mathbf{c}_{X,r})$ and $C_{\bar{\tau}}^{\leq K}(\mathbf{c}_{Y,r}) = C_{\tau}(\mathbf{c}_{Y,r})$, for $r \in [k]$. Call the blocks B_j which are disjoint from all $C_{\tau}(\mathbf{c}_{X,r}), C_{\tau}(\mathbf{c}_{Y,r})$ as *regular blocks*.

Structural Lemma. In this we show that unless the truncated coefficient vectors of h satisfy a certain structural property (which we leave unstated in this sketch), the LHS of (5) for h is very small, leading to a contradiction. On the other hand, this property leads to a good labeling. The main idea is to show that (assuming this property is not satisfied) the truncation of $(\mathbf{c}_X, \mathbf{c}_Y)$ implies with high probability over \widehat{D} that the squared mass corresponding to the noise variables in the regular blocks is a significant fraction of the total squared mass in the regular blocks. The regularity of the blocks is leveraged to apply the Chernoff-Hoeffding inequality to obtain this. On the other hand, the variance of the difference of h between $(\mathbf{X}^0, \mathbf{Y}^0)$ and $(\mathbf{X}^1, \mathbf{Y}^1)$ is nearly always much smaller than the squared mass of the noise variables. This uses the fact that our choice of Q is large (in fact we need to set it to $\text{poly}(d)$). Another application of Berry-Esseen along with Chebyshev's inequality completes the proof.

1.4 Organization

Section 2 provides the necessary technical preliminaries required for the paper and describes the Smooth Label Cover problem. In Section 3 we give the theorem which states the guarantees of our reduction from Smooth Label Cover, followed by a description of the reduction. We prove the completeness and soundness of the reduction in Sections 4 and 5. In Section 6 we prove the guarantee of the truncation step, and in Section 7 we prove the main structural lemma used in our soundness analysis. Finally in Appendix A we prove the anti-concentration bound used in the truncation step, using the Littlewood-Offord-Erdős lemma as well as the Berry-Esseen theorem.

2 Preliminaries

Throughout the paper we use $\text{pos}(x) = \mathbb{1}\{x \geq 0\}$ to denote the sign function. The following is a well known quantitative form of the central limit theorem.

Theorem 2.1 (Berry-Esseen Theorem [O'D14]) *Let X_1, \dots, X_n be independent random variables with $E[X_i] = 0$ and $\text{Var}(X_i) = \sigma_i^2$, and assume that $\sum_{i \in [n]} \sigma_i^2 = 1$. Let $\gamma := \sum_{i \in [n]} \|X_i\|_{L_3}^3$. Then*

$$\sup_{t \in \mathbb{R}} \left| \Pr_{X_1, \dots, X_n} \left[\sum_{i \in [n]} X_i \leq t \right] - \Phi(t) \right| \leq c\gamma \quad (6)$$

where c is a universal constant, and $\Phi(\cdot)$ is the standard Gaussian CDF. Note that it follows from above that (6) also holds when taking $\gamma := \max_{i \in [n]} (\|X_i\|_{L_3}^3 / \sigma_i^2)$.

Apart from the above, our anti-concentration bounds also use the following classical result for Bernoulli sums.

Lemma 2.2 (Littlewood-Offord-Erdős Lemma [Erd45]) *Let X_1, X_2, \dots, X_n be i.i.d $\{0, 1\}$ -Bernoulli Random variables, and let $a_1, a_2, \dots, a_n \in \mathbb{R}$ such that for all $i \in [n]$ we have $|a_i| \geq 1$. Then there exists a constant $C > 0$ such that*

$$\Pr_{X_1, \dots, X_n} \left[\left| \sum_{i \in [n]} a_i X_i + \theta \right| \leq 1 \right] \leq \frac{C}{\sqrt{n}}$$

for any constant θ .

We also use the following well known bounds of Chernoff-Hoeffding and Chebyshev.

Theorem 2.3 (Chernoff-Hoeffding) *Let X_1, \dots, X_n be independent random variables, each bounded as $a_i \leq X_i \leq b_i$ with $\Delta_i = b_i - a_i$ for $i = 1, \dots, n$. Then, for any $t > 0$,*

$$\Pr \left[\left| \sum_{i=1}^n X_i - \sum_{i=1}^n \mathbb{E}[X_i] \right| > t \right] \leq 2 \cdot \exp \left(-\frac{2t^2}{\sum_{i=1}^n \Delta_i^2} \right).$$

Chebyshev's Inequality. For any random variable X and $t > 0$, $\Pr[|X| > t] \leq \mathbb{E}[X^2]/t^2$.

2.1 Critical Index

Consider a vector of coefficients $\mathbf{c} \in \mathbb{R}^{[M] \times [q]}$ such that $\mathbf{c} = (\mathbf{c}_i)_{i=1}^M$ and $\mathbf{c}_i \in \mathbb{R}^{[q]}$ for each $i \in [M]$. Let $\sigma : [M] \rightarrow [M]$ be an ordering such that $\|\mathbf{c}_{\sigma(1)}\|_2 \geq \|\mathbf{c}_{\sigma(2)}\|_2 \geq \dots \geq \|\mathbf{c}_{\sigma(M)}\|_2$. For a given $\tau \in (0, 1)$, the τ -critical index $i_\tau(\mathbf{c})$ to be the minimum index $i \in [M]$ such that $\|\mathbf{c}_{\sigma(i)}\|_2^2 \leq \tau \sum_{i' \geq i} \|\mathbf{c}_{\sigma(i')}\|_2^2$. If $i_\tau(\mathbf{c}) = 1$ then \mathbf{c} is said to be τ -regular. Define $C_\tau(\mathbf{c}) := \{\sigma(i) : i < i_\tau(\mathbf{c})\}$, and for any integer $K \in \mathbb{N}$, $C_\tau^{\leq K}(\mathbf{c}) = \{\sigma(i) : 1 \leq i \leq K, i < i_\tau(\mathbf{c})\}$ to denote the set of the first K indices in $C_\tau(\mathbf{c})$. The following proposition summarizes well known properties of critical indices:

Proposition 2.4 ([Ser06]) *For the above setting the following condition hold:*

- For any $1 \leq i_1 < i_2 \leq i_\tau(\mathbf{c})$, we have $\|\mathbf{c}_{\sigma(i_2)}\|_2^2 \leq \frac{1}{\tau}(1 - \tau)^{i_1 - i_2} \|\mathbf{c}_{\sigma(i_1)}\|_2^2$.
- The vector $(\mathbf{c}_i)_{i \in [M] \setminus C_\tau(\mathbf{c})}$ is τ -regular.

2.2 Smooth Label Cover

Our reduction is from the following hypergraph variant of Smooth Label Cover.

Definition 2.5 (SMOOTH-2k-LABEL COVER) *A SMOOTH-2k-LABEL COVER instance $\mathcal{L}((V_{\mathcal{L}}, E_{\mathcal{L}}), M, m, \{\pi_{e,v}\}_{e \in E_{\mathcal{L}}, v \in e})$ consists of a regular, $2k$ -uniform, connected hypergraph with a vertex set $V_{\mathcal{L}}$, a hyperedge set $E_{\mathcal{L}}$ and a set of projections $\{\pi_{e,v} : [M] \mapsto [m]\}_{e \in E_{\mathcal{L}}, v \in e}$.*

A labeling $\sigma : V_{\mathcal{L}} \mapsto [M]$ is said to strongly satisfy a hyperedge e if for every $v, v' \in e$, we have $\pi_{e,v}(\sigma(v)) = \pi_{e,v'}(\sigma(v'))$. It is said to weakly satisfy a hyperedge e if there exists a distinct pair of vertices $v, v' \in e$ such that $\pi_{e,v}(\sigma(v)) = \pi_{e,v'}(\sigma(v'))$.

The following theorem gives the hardness of SMOOTH-2k-LABEL COVER.

Theorem 2.6 ([FGRW12]) *There exists an absolute constant $\gamma_0 > 0$ such that for all integer parameters z and J , it NP-Hard to distinguish whether an instance \mathcal{L} of SMOOTH-2k-LABEL COVER with $M = 7^{(J+1)^z}$ and $N = 2^{z7^z}$, satisfies,*

- (YES): *There exists a labeling $\sigma : V \mapsto [M]$ which strongly satisfies all the hyperedges.*
- (NO): *There is no labeling $\sigma : V \mapsto [M]$ which weakly satisfies more than $2k^2 2^{-\gamma_0 z}$ -fraction of hyperedges.*

Additionally, \mathcal{L} satisfies the following properties:

- (Smoothness) *For every vertex $v \in V$ and for a randomly sampled hyperedge incident on v*

$$\Pr_{e \sim v} [\pi_{e,v}(i) = \pi_{e,v}(j)] \leq \frac{1}{J}$$

for any fixed pair of distinct labels $i, j \in [M]$.

- For any hyperedge $e \in E$ and vertex $v \in e$, and any label on the smaller side $\alpha \in [m]$, $|\pi_{e,v}^{-1}(\alpha)| \leq d$, where $d = 4^z$.

3 Hardness Reduction

In this section we prove the following theorem.

Theorem 3.1 *For any constants $\nu, \zeta \in (0, 1/2)$ and $\ell \in \mathbb{N}$, there exists a choice of z, J and k in Theorem 2.6 and a polynomial time reduction from the corresponding SMOOTH-2k-LABEL COVER instance \mathcal{L} to a distribution \mathcal{D} supported on point label pairs in $\{0, 1\}^n \times \{0, 1\}$ such that*

- **Completeness:** *If \mathcal{L} is a YES instance, then there exists two ORs C_1 and C_2 such that*

$$\Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [(C_1(\mathbf{x}) \wedge C_2(\mathbf{x})) = a] = 1, \quad \text{and} \quad \Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [C_1(\mathbf{x}) = a] \geq 1 - \zeta.$$

- **Soundness:** *If \mathcal{L} is a NO instance, then for any function $f : \{0, 1\}^n \mapsto \{0, 1\}$ of ℓ -halfspaces*

$$\Pr_{(\mathbf{x}, a) \leftarrow \mathcal{D}} [f(\mathbf{x}) = a] \leq \frac{1}{2} + \nu.$$

Since the negation of an s -term DNF is an s -clause CNF, the above theorem along with Theorem 2.6 implies Theorem 1.1. In the remainder of the section we provide a detailed construction of the distribution towards proving the above theorem.

The hardness reduction is from an instance \mathcal{L} of SMOOTH-2k-LABEL COVER as given in Theorem 2.6 instantiated with

$$J = \frac{10^3 \cdot \ell^2 \cdot \log^2(dk) \cdot d^{20}}{\nu(\zeta(1-\zeta))^2} \quad (7)$$

where γ_0 and d are as given in Theorem 2.6, and we shall fix the parameter z later. Additionally, we also define the parameters

$$Q = 16dk, \quad \tau = (10k \log Q)^{-2}, \quad K = \frac{20}{\tau} \log \frac{Q}{\tau}, \quad k = 10/(\zeta(1-\zeta))^2, \quad t = k/4. \quad (8)$$

For every hyperedge e , we define an arbitrary partition of its $2k$ vertices into two subsets of size k each, given by $e = e_X \cup e_Y$ (this notation shall become clear below).

Coordinate Set: For every vertex $v \in V_{\mathcal{L}}$ of the SMOOTH-2k-LABEL COVER instance and every label $i \in [M]$ we introduce two vectors of Q boolean valued variables (coordinates) each: $\mathbf{X}_{v,i} := (X_{v,i,1}, \dots, X_{v,i,Q})$ and $\mathbf{Y}_{v,i} := (Y_{v,i,1}, \dots, Y_{v,i,Q})$. Let $\mathbf{X}_v := \bigoplus_{i=1}^M \mathbf{X}_{v,i}$, and $\mathbf{Y}_v := \bigoplus_{i=1}^M \mathbf{Y}_{v,i}$. Let \mathbf{X} and \mathbf{Y} denote $\{\mathbf{X}_v\}_{v \in V_{\mathcal{L}}}$ and $\{\mathbf{Y}_v\}_{v \in V_{\mathcal{L}}}$. In particular, the points of the instance lie in the $(2qM|V_{\mathcal{L}}|)$ -dimensional boolean space.

Point-label distribution: The point-label distribution $\mathcal{D}_{\text{global}}$ is given in Figure 3, with t being a parameter to be decided later.

The restriction of $\mathcal{D}_{\text{global}}$ to a hyperedge e is denoted by \mathcal{D}^e . The restrictions of \mathcal{D}^e to the 0 and 1 points are given by \mathcal{D}_0^e and \mathcal{D}_1^e respectively. Fix a hyperedge e , and some $v \in e_X$.

Consider the distributions \mathcal{D}_0^e and \mathcal{D}_1^e conditioned on a fixation of $\{S_j, S'_j\}_{j \in [m]}$. Under both distributions, if $b_j = 0$ then for each $v \in S_j$ independently each of $\{\mathbf{X}_{v,i} : i \in \pi_{e,v}^{-1}(j)\}$ is independently

1. Sample a random hyperedge $e \sim E$. Let $\pi_v := \pi_{e,v}$ for all $v \in e$. Recall the predefined partition of the vertices of e as $e = e_X \cup e_Y$ where $|e_X| = |e_Y| = k$
2. Set $\mathbf{X}_v = \mathbf{Y}_v = \mathbf{0}$ for all $v \notin e$
3. Sample $a \in \{0, 1\}$ u.a.r. as the label of the point sampled below
4. For every label $j \in [m]$, do the following:
 - 4.1. Independently sample $S_j \stackrel{\text{u.a.r.}}{\sim} \binom{e_X}{t}$ and $S'_j \stackrel{\text{u.a.r.}}{\sim} \binom{e_Y}{t}$
 - 4.2. Sample a bit b_j to be 0 w.p. ζ and 1 w.p. $1 - \zeta$. If $b_j = 0$, for every $v \in e_X \setminus S_j$ and $i \in \pi_v^{-1}(j)$, sample $\mathbf{X}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q$. If $b_j = 1$, for every $v \in e_Y \setminus S'_j$ and $i \in \pi_v^{-1}(j)$, sample $\mathbf{Y}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q$
 - 4.3 If $a = 1$ do Step 4.4 otherwise do Step 4.5
 - 4.4 **Sampling a 1-Point** Do the following:
 1. Independently sample $u_{X,j} \sim S_j$ and $u_{Y,j} \sim S'_j$ u.a.r.
 2. For every $i \in \pi_{u_{X,j}}^{-1}(j)$, set $\mathbf{X}_{u_{X,j},i} \stackrel{\text{u.a.r.}}{\sim} \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$
 3. For every $i \in \pi_{u_{Y,j}}^{-1}(j)$, set $\mathbf{Y}_{u_{Y,j},i} \stackrel{\text{u.a.r.}}{\sim} \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$
 - 4.5 **Sampling a 0-Point** With probability $1 - 1/(\zeta(1 - \zeta)t)$, for every $u \in S_j$ and $v \in S'_j$, set $\mathbf{X}_{u,i}$ and $\mathbf{Y}_{v,i'}$ to $\mathbf{0}$ for each $i \in \pi_u^{-1}(j)$ and $i' \in \pi_v^{-1}(j)$. With the rest of the probability $1/(\zeta(1 - \zeta)t)$, do the following:
 1. If $b_j = 0$, choose $T_j \subseteq S_j$ by independently sampling vertices of S_j w.p. $(1 - \zeta)$. For every $v \in T_j$ and $i \in \pi_v^{-1}(j)$, sample $\mathbf{X}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$
 2. If $b_j = 1$, choose $T'_j \subseteq S'_j$ by independently sampling vertices of S'_j w.p. ζ . For every $v \in T'_j$ and $i \in \pi_v^{-1}(j)$, sample $\mathbf{Y}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$
5. For those $\mathbf{X}_{v,i}, \mathbf{Y}_{v,i}$ not assigned values yet, set them to $\mathbf{0}$. Output the point (\mathbf{X}, \mathbf{Y}) with label a .

Figure 3: Distribution $\mathcal{D}_{\text{global}}$

sampled u.a.r. from $\{0, 1\}^Q$, and if $b_j = 0$ they are all set to $\mathbf{0}$. If $v \in S_j$, with probability exactly $1/t$ each of $\{\mathbf{X}_{v,i} : i \in \pi_{e,v}^{-1}(j)\}$ is independently sampled u.a.r from $\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_Q\}$, and otherwise set to $\mathbf{0}$. The same analogously holds true for $\{\mathbf{Y}_{v,i} : i \in \pi_{e,v}^{-1}(j)\}$ for $v \in e_Y$, with \bar{b}_j replacing b_j , S'_j replacing S_j . Based on this we have the following observation. Let $\mathbf{X}_e^{(j)} := \cup_{v \in e_X} \{\mathbf{X}_{v,i} : i \in \pi_{e,v}^{-1}(j)\}$ and $\mathbf{Y}_e^{(j)} := \cup_{v \in e_Y} \{\mathbf{Y}_{v,i} : i \in \pi_{e,v}^{-1}(j)\}$ be the set of variables for e corresponding to $j \in [m]$.

Observation 3.2 *Conditioned on any fixation of $\{S_j, S'_j\}_{j \in [m]}$ the following holds.*

- Distributions \mathcal{D}_0^e and \mathcal{D}_1^e induce the same distribution on any variable $\mathbf{X}_{v,i}$ or $\mathbf{Y}_{v,i}$.
- Under both \mathcal{D}_0^e and \mathcal{D}_1^e , the variables $(\mathbf{X}_e^{(j)}, \mathbf{Y}_e^{(j)}, b_j)$, $j = 1, \dots, m$ are sampled independently (over the different $j \in [m]$).
- For each $j \in [m]$, after additionally fixing b_j , the variables

$$\{\mathbf{X}_{v,i} : v \in e_X, v \notin S_j, i \in \pi_{e,v}^{-1}(j)\} \cup \{\mathbf{Y}_{v,i} : v \in e_Y, v \notin S'_j, i \in \pi_{e,v}^{-1}(j)\}$$

have the same distribution in \mathcal{D}_0^e , and \mathcal{D}_1^e , and are independent of the variables corresponding to $v \in S_j$ and $v \in S'_j$.

We denote a halfspace over the variables \mathbf{X} and \mathbf{Y} as $\text{pos}(h(\mathbf{X}, \mathbf{Y}))$ where $h(\mathbf{X}, \mathbf{Y}) = \langle \mathbf{c}_X, \mathbf{X} \rangle +$

$\langle \mathbf{c}_Y, \mathbf{Y} \rangle + \theta$, where $\mathbf{c}_X, \mathbf{c}_Y \in \mathbb{R}^{V_{\mathcal{L}} \times [M] \times [q]}$ and θ is a constant. For any vertex $v \in V_{\mathcal{L}}$, $i \in [M]$ and $q \in [Q]$ we denote by (i) $\mathbf{c}_{X,v,i} := (c_{X,v,i,q})_{q=1}^Q$ the vector of coefficients corresponding to $\mathbf{X}_{v,i}$, and (ii) $\mathbf{c}_{X,v}$ the vector of coefficients corresponding to \mathbf{X}_v . Similarly for $\mathbf{c}_{Y,v,i}$ and $\mathbf{c}_{Y,v}$.

4 Completeness

Suppose \mathcal{L} is a YES instance. Then there exists labeling $\sigma : V \mapsto [M]$ such that for every hyperedge e we have $j_e \in [m]$ s.t. $\pi_{e,v}(\sigma(v)) = j_e$ for every $v \in e$. Consider the following two OR formulas:

$$\mathcal{C}_1 = \left(\bigvee_{v \in V} \bigvee_{q \in [Q]} X_{v,\sigma(v),q} \right), \quad \text{and} \quad \mathcal{C}_2 = \left(\bigvee_{v \in V} \bigvee_{q \in [Q]} Y_{v,\sigma(v),q} \right)$$

Note that \mathcal{C}_1 depends only on \mathbf{X} and \mathcal{C}_2 only on \mathbf{Y} . Fix a hyperedge $e = e_X \cup e_Y$.

Suppose $a = 1$. By construction of the distribution, there is $u = u_{X,j_e}$ and $v = u_{Y,j_e}$ such that for some $q, q' \in [Q]$ we have $X_{u,\sigma(u),q} = 1$ and $Y_{v,\sigma(v),q'} = 1$. This implies that $\mathcal{C}_1 = \mathcal{C}_2 = 1$, and $\mathcal{C}_1 \wedge \mathcal{C}_2$ evaluates to 1 when $a = 1$.

On the other hand, suppose $a = 0$. Note that all the variables corresponding to vertices not in e are set to zero. If $b_{j_e} = 0$, then $\mathbf{Y}_{v,\sigma(v)} = \mathbf{0}$ for every $v \in e_Y$ implying that $\mathcal{C}_2 = 0$. Otherwise if $b_{j_e} = 1$ then $\mathbf{X}_{v,\sigma(v)} = \mathbf{0}$ for every $v \in e_X$ so that $\mathcal{C}_1 = 0$. Therefore, we have $\mathcal{C}_1 \wedge \mathcal{C}_2 = 0$ whenever $a = 0$. Furthermore, since $b_{j_e} = 1$ w.p. $1 - \zeta$, $\mathcal{C}_1 = 0$ w.p. $1 - \zeta$ on points labeled 0.

5 Soundness Analysis

Suppose there exists a function f of ℓ -halfspaces $\text{pos}(h_1(\mathbf{X}, \mathbf{Y})), \dots, \text{pos}(h_\ell(\mathbf{X}, \mathbf{Y}))$ such that:

$$\Pr_{((\mathbf{X}, \mathbf{Y}), a) \sim \mathcal{D}_{\text{global}}} [f(\mathbf{X}, \mathbf{Y}) = a] \geq \frac{1}{2} + \nu. \quad (9)$$

Observe that by averaging, for at least $\nu/2$ -fraction of the hyperedges e , f is consistent with \mathcal{D}_e with probability at least $\frac{1}{2} + \frac{\nu}{2}$. Call such hyperedges *fine* and for each such hyperedge e we have

$$\Pr_{((\mathbf{X}, \mathbf{Y}), a) \sim \mathcal{D}^e} [f(\mathbf{X}, \mathbf{Y}) = a] \geq \frac{1}{2} + \frac{\nu}{2}. \quad (10)$$

Applying Label Cover Smoothness. Recall that $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell$ are the ℓ coefficient vectors for the variables (\mathbf{X}, \mathbf{Y}) . For a vertex v , we have the subsets $C_\tau(\mathbf{c}_{X,v}^{(s)})$, $C_\tau(\mathbf{c}_{Y,v}^{(s)})$, $C_\tau^{\leq K}(\mathbf{c}_{X,v}^{(s)})$, and $C_\tau^{\leq K}(\mathbf{c}_{Y,v}^{(s)})$, for $s \in [\ell]$ as defined in Section 2.1. Additionally, define the following subset of $[M]$:

$$L_v \left(\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell \right) := \bigcup_{s=1}^\ell I_v(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)}), \quad (11)$$

where

$$I_v(\mathbf{c}_X, \mathbf{c}_Y) = C_\tau^{\leq K}(\mathbf{c}_{X,v}) \cup C_\tau^{\leq K}(\mathbf{c}_{Y,v}) \cup \left\{ i \in [M] \setminus C_\tau(\mathbf{c}_{X,v}) : \|\mathbf{c}_{X,v,i}\|^2 > \frac{1}{d^8} \sum_{i' \notin C_\tau(\mathbf{c}_{X,v})} \|\mathbf{c}_{X,v,i'}\|_2^2 \right\}$$

$$\bigcup \left\{ i \in [M] \setminus C_\tau(\mathbf{c}_{Y,v}) : \|\mathbf{c}_{Y,v,i}\|^2 > \frac{1}{d^8} \sum_{i' \notin C_\tau(\mathbf{c}_{Y,v})} \|\mathbf{c}_{Y,v,i'}\|_2^2 \right\}. \quad (12)$$

The following is a consequence of the smoothness property of SMOOTH-2k-LABEL COVER and our setting of the parameters.

Lemma 5.1 *Given $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell$, at least $(1 - \nu/4)$ fraction of the hyperedges $e \in E_\mathcal{L}$ are nice w.r.t. $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell$ i.e., the following condition holds for every vertex $v \in e$: for any pair of labels $i_1, i_2 \in L_v \left(\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell \right)$, $\pi_{e,v}(i_1) \neq \pi_{e,v}(i_2)$.*

Proof. It is easy to see that the size of $I_v(\mathbf{c}_X, \mathbf{c}_Y)$ for any vertex v is at most $2(K + d^8)$, and thus the size of $L_v = L_v \left(\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell \right)$ is at most $2\ell(K + d^8)$. By the regularity of $G_\mathcal{L} = (V_\mathcal{L}, E_\mathcal{L})$, the smoothness of \mathcal{L} and union bound, we have,

$$\begin{aligned} & \Pr_{e \sim E_\mathcal{L}} [\exists v \in e, \exists i_1, i_2 \in L_v : \pi_{e,v}(i_1) = \pi_{e,v}(i_2)] \\ & \leq 2k \cdot \mathbb{E}_{v \in V_\mathcal{L}} \left[\Pr_{e \sim v} [\exists i_1, i_2 \in L_v : \pi_{e,v}(i_1) = \pi_{e,v}(i_2)] \right] \\ & \leq 2k \cdot \mathbb{E}_{v \in V_\mathcal{L}} [|L_v|^2 / J] \\ & \leq (8k\ell^2 \cdot (K + d^8)^2) / J \leq \nu/4, \end{aligned} \quad (13)$$

where the last inequality is obtained by our setting of J in (7). \square

For convenience we shall abuse notation to say that e is nice w.r.t. $(\mathbf{c}_X, \mathbf{c}_Y)$ if for each $v \in e$, for any pair of labels $i_1, i_2 \in I_v(\mathbf{c}_X, \mathbf{c}_Y)$, $\pi_{e,v}(i_1) \neq \pi_{e,v}(i_2)$. We have the following observation.

Observation 5.2 *If e is nice w.r.t. $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell$ then it is nice w.r.t. each $(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})$, $s \in [\ell]$.*

The Main Structural Lemma. Applying Lemma 5.1 to the coefficient vectors $(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})$ for h_s for $s \in [\ell]$ we obtain that there are at least $(\nu/2 - \nu/4) = (\nu/4)$ -fraction of hyperedges which are *fine* and *nice* i.e., they satisfy (10) as well as the niceness condition of Lemma 5.1 w.r.t. $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^\ell$. Let E^* represent the set such hyperedges. We now state the main structural lemma the proof of which is provided later in this section.

Lemma 5.3 *For each $e \in E^*$ there exist distinct vertices $u, v \in e$ such that at least one of the following is satisfied,*

I. *There exist $r, p \in [\ell]$ (r may equal p) s.t.*

$$\pi_{e,u} \left(C_\tau^{\leq K}(\mathbf{c}_{X,u}^{(r)}) \cup C_\tau^{\leq K}(\mathbf{c}_{Y,u}^{(r)}) \right) \cap \pi_{e,v} \left(C_\tau^{\leq K}(\mathbf{c}_{X,v}^{(p)}) \cup C_\tau^{\leq K}(\mathbf{c}_{Y,v}^{(p)}) \right) \neq \emptyset.$$

II. *For some $r \in [\ell]$, there exists $j \in \pi_{e,u} \left(C_\tau^{\leq K}(\mathbf{c}_{X,u}^{(r)}) \cup C_\tau^{\leq K}(\mathbf{c}_{Y,u}^{(r)}) \right)$ such that at least one of (14) or (15) given below is satisfied*

$$j \notin \pi_{e,v} \left(C_\tau^{\leq K}(\mathbf{c}_{X,v}^{(r)}) \right), \quad \text{and} \quad \sum_{i \in \pi_{e,v}^{-1}(j) \setminus C_\tau(\mathbf{c}_{X,v}^{(r)})} \|\mathbf{c}_{X,v,i}^{(r)}\|^2 > \tau^4 \sum_{i \in [M] \setminus C_\tau(\mathbf{c}_{X,v}^{(r)})} \|\mathbf{c}_{X,v,i}^{(r)}\|^2 \quad (14)$$

$$j \notin \pi_{e,v} \left(C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}^{(r)}) \right), \quad \text{and} \quad \sum_{i \in \pi_{e,v}^{-1}(j) \setminus C_{\tau}(\mathbf{c}_{Y,v}^{(r)})} \|\mathbf{c}_{Y,v,i}^{(r)}\|^2 > \tau^4 \sum_{i \in [M] \setminus C_{\tau}(\mathbf{c}_{Y,v}^{(r)})} \|\mathbf{c}_{Y,v,i}^{(r)}\|^2 \quad (15)$$

Labeling the Vertices of \mathcal{L} using Lemma 5.3. Consider the following randomized labeling for each $v \in V_{\mathcal{L}}$: with probability $1/2$ each do Step 2 or 3.

1. Choose $s \in [\ell]$ u.a.r.
2. Assign $\sigma(v) \stackrel{\text{unif}}{\sim} \left(C_{\tau}^{\leq K}(\mathbf{c}_{X,v}^{(s)}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}^{(s)}) \right)$.
3. W.p. $1/2$ each do (a) or (b):
 - (a) Assign $\sigma(v)$ a label $i \notin C_{\tau}(\mathbf{c}_{X,v}^{(s)})$ with probability

$$\|\mathbf{c}_{X,v,i}^{(s)}\|^2 / \left(\sum_{i \in [M] \setminus C_{\tau}(\mathbf{c}_{X,v}^{(s)})} \|\mathbf{c}_{X,v,i}^{(s)}\|^2 \right).$$

- (b) Assign $\sigma(v)$ a label $i \notin C_{\tau}(\mathbf{c}_{Y,v}^{(s)})$ with probability

$$\|\mathbf{c}_{Y,v,i}^{(s)}\|^2 / \left(\sum_{i \in [M] \setminus C_{\tau}(\mathbf{c}_{Y,v}^{(s)})} \|\mathbf{c}_{Y,v,i}^{(s)}\|^2 \right).$$

We now analyze the probability of the above labeling *weakly* satisfying a fixed hyperedge $e \in E^*$. Suppose that e satisfies Case I of Lemma 5.3 for $u, v \in e$. Then, with probability $\geq 1/(4\ell^2)$ both u and v are labeled using Step 1 of the labeling procedure with u choosing r and v choosing p . Therefore, in this case, $\pi_{e,u}(\sigma(u)) = \pi_{e,v}(\sigma(v))$ with probability at least

$$\frac{1}{4\ell^2} \cdot \frac{1}{\left| C_{\tau}^{\leq K}(\mathbf{c}_{X,u}^{(r)}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,u}^{(r)}) \right|} \cdot \frac{1}{\left| C_{\tau}^{\leq K}(\mathbf{c}_{X,v}^{(p)}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}^{(p)}) \right|} \geq \frac{1}{16K^2\ell^2}. \quad (16)$$

On the other hand, suppose e satisfies Case II of Lemma 5.3 such that (14) holds, for $u, v \in e$. Then, with probability $1/(8\ell^2)$ u is labeled according to Step 1 using $s = r$ and v according to Step 2.a of the labeling procedure using $s = r$. Thus, in this case $\pi_{e,u}(\sigma(u)) = \pi_{e,v}(\sigma(v))$ with probability at least

$$\frac{1}{8\ell^2} \cdot \frac{1}{\left| C_{\tau}^{\leq K}(\mathbf{c}_{X,u}^{(r)}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,u}^{(r)}) \right|} \cdot \tau^4 \geq \frac{\tau^4}{16K\ell^2}, \quad (17)$$

which also analogously holds when e satisfies Case II of Lemma 5.3 such that (15) is true.

Combining the above with the lower bound on the size of E^* , we obtain that the expected fraction of hyperedges satisfied by the labeling is at least,

$$\frac{1}{16\ell^2} \cdot \frac{\nu}{4} \cdot \min \left\{ \frac{1}{K^2}, \frac{\tau^4}{K} \right\} \quad (18)$$

Plugging in the parameters from (8) and $d = 4^z$ (as given in Theorem 2.6), the RHS of the above equation is at least $\nu(10k \log(16dk))^{-10} \geq \nu(10kz \log(64k))^{-10}$ (from the value of d in Theorem 2.6). On the other hand, from the guarantee of the NO case of Theorem 2.6, we know that any labeling can weakly satisfy at most $2k^2 2^{-\gamma_0 z}$ fraction of hyperedges in \mathcal{L} . By choosing z to be a large enough constant, the expected number of weakly satisfied hyperedges exceeds the soundness parameter, which completes the proof of Theorem 3.1.

5.1 Proof of Lemma 5.3

Fix an edge $e \in E^*$. Assume for a contradiction that e does not satisfy Lemma 5.3. For convenience, define

$$B_{s,v} := C_{\tau}^{\leq K}(\mathbf{c}_{X,v}^{(s)}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}^{(s)}) \quad (19)$$

for $s \in [\ell]$ and $v \in e$. Let us say that two triples (r, u, i_1) and (p, v, i_2) from $[\ell] \times e \times [M]$ are distinct if they differ in at least one coordinate. From the niceness of e w.r.t. $\{(\mathbf{c}_X^{(s)}, \mathbf{c}_Y^{(s)})\}_{s=1}^{\ell}$ given by Lemma 5.1 and the negation of Condition I of Lemma 5.3 we have the following observation.

Observation 5.4 *For any two distinct triples (r, u, i_1) and (p, v, i_2) s.t. $i_1 \in B_{r,u}$ and $i_2 \in B_{p,v}$, $\pi_{e,u}(i_1) \neq \pi_{e,v}(i_2)$.*

Combining the above with Observation 3.2 we obtain that after any fixation of $\{S_j, S'_j\}_{j=1}^m$ the variables

$$\{\mathbf{X}_{v,i} : v \in e_X, s \in [\ell], i \in C_{\tau}^{\leq K}(\mathbf{c}_{X,v}^{(s)})\} \cup \{\mathbf{Y}_{v,i} : v \in e_Y, s \in [\ell], i \in C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}^{(s)})\}, \quad (20)$$

are independently distributed under both \mathcal{D}_0^e and \mathcal{D}_1^e . Further, since their marginal distributions match under \mathcal{D}_0^e and \mathcal{D}_1^e (again by Observation 3.2), their joint distributions also match. Using this we describe in Figure 4 a distribution $\widehat{\mathcal{D}}_e$ on $((\mathbf{X}^0, \mathbf{Y}^0), (\mathbf{X}^1, \mathbf{Y}^1))$ with marginal distributions \mathcal{D}_0^e and \mathcal{D}_1^e respectively.

1. Sample $\{S_j, S'_j\}_{j=1}^m$ as according to \mathcal{D}^e .
2. Conditioned on $\{S_j, S'_j\}_{j=1}^m$, sample the variables in (20) by sampling b_j for $j \in \cup_{v \in e} \cup_{s \in [\ell]} \pi_{e,v}(B_{s,v})$.
3. Conditioned on the fixing till now sample all the rest of the b_j and then the rest of the variables in

$$\{\mathbf{X}_{v,i} : v \in e_X, v \notin S_j, i \in \pi_{e,v}^{-1}(j)\} \cup \{\mathbf{Y}_{v,i} : v \in e_Y, v \notin S'_j, i \in \pi_{e,v}^{-1}(j)\},$$

according to the distribution common to \mathcal{D}_0^e and \mathcal{D}_1^e (see Observation 3.2).

4. Call the above fixings Γ_e . Conditioned on Γ_e sample $(\mathbf{X}^0, \mathbf{Y}^0) \leftarrow \mathcal{D}_0^e$ and $(\mathbf{X}^1, \mathbf{Y}^1) \leftarrow \mathcal{D}_1^e$.
5. Output $((\mathbf{X}^0, \mathbf{Y}^0), (\mathbf{X}^1, \mathbf{Y}^1))$.

Figure 4: Distribution $\widehat{\mathcal{D}}_e$

From (10) we obtain that,

$$\mathbb{E}_{((\mathbf{X}^0, \mathbf{Y}^0), (\mathbf{X}^1, \mathbf{Y}^1)) \leftarrow \widehat{\mathcal{D}}_e} \left[\left| f(\mathbf{X}^1, \mathbf{Y}^1) - f(\mathbf{X}^0, \mathbf{Y}^0) \right| \right] \geq \frac{\nu}{2}. \quad (21)$$

By averaging, there is one halfspace $\text{pos}(h)(\cdot)$ on which f depends, satisfying:

$$\mathbb{E}_{((\mathbf{X}^0, \mathbf{Y}^0), (\mathbf{X}^1, \mathbf{Y}^1)) \leftarrow \widehat{\mathcal{D}}_e} \left[\left| \text{pos}(h(\mathbf{X}^1, \mathbf{Y}^1)) - \text{pos}(h(\mathbf{X}^0, \mathbf{Y}^0)) \right| \right] \geq \frac{\nu}{2\ell}. \quad (22)$$

Let the linear form h be given by $h(\mathbf{X}, \mathbf{Y}) = \langle \mathbf{c}_X, \mathbf{X} \rangle + \langle \mathbf{c}_Y, \mathbf{Y} \rangle + \theta$.

Applying Lemma 6.1 iteratively for each $v \in e$ on h , we obtain a *truncated* linear form $\tilde{h}(\mathbf{X}, \mathbf{Y}) = \langle \tilde{\mathbf{c}}_X, \mathbf{X} \rangle + \langle \tilde{\mathbf{c}}_Y, \mathbf{Y} \rangle + \theta$ satisfying

- (A) $\tilde{\mathbf{c}}_{X,u} = \mathbf{c}_{X,u}$ and $\tilde{\mathbf{c}}_{Y,u} = \mathbf{c}_{Y,u}$ for all $u \in V_{\mathcal{L}}, u \notin e$.
- (B) For $v \in e_X$: $\tilde{\mathbf{c}}_{X,v,i} = \mathbf{0}$ for all $i \in C_{\tau}(\mathbf{c}_{X,v}) \setminus C_{\tau}^{\leq K}(\mathbf{c}_{X,v})$ and $\tilde{\mathbf{c}}_{X,v,i} = \mathbf{c}_{X,v,i}$ otherwise,
- (C) For $v \in e_Y$: $\tilde{\mathbf{c}}_{Y,v,i} = \mathbf{0}$ for all $i \in C_{\tau}(\mathbf{c}_{Y,v}) \setminus C_{\tau}^{\leq K}(\mathbf{c}_{Y,v})$ and $\tilde{\mathbf{c}}_{Y,v,i} = \mathbf{c}_{Y,v,i}$ otherwise,

and by union bound,

$$\mathbb{E}_{((\mathbf{X},\mathbf{Y}),a) \sim \mathcal{D}_e} [|\text{pos}(h(\mathbf{X},\mathbf{Y})) - \text{pos}(\tilde{h}(\mathbf{X},\mathbf{Y}))|] \leq 2k \cdot \frac{\tau^{1/4}}{4k} = \frac{\tau^{1/4}}{2}. \quad (23)$$

Using the above we obtain,

$$\begin{aligned} & \mathbb{E}_{((\mathbf{x}^0, \mathbf{y}^0), (\mathbf{x}^1, \mathbf{y}^1)) \leftarrow \widehat{\mathcal{D}}_e} \left[\left| \text{pos}(\tilde{h}(\mathbf{x}^1, \mathbf{y}^1)) - \text{pos}(\tilde{h}(\mathbf{x}^0, \mathbf{y}^0)) \right| \right] \\ \geq & \mathbb{E}_{((\mathbf{x}^0, \mathbf{y}^0), (\mathbf{x}^1, \mathbf{y}^1)) \leftarrow \widehat{\mathcal{D}}_e} \left[\left| \text{pos}(h(\mathbf{x}^1, \mathbf{y}^1)) - \text{pos}(h(\mathbf{x}^0, \mathbf{y}^0)) \right| \right] \\ & - \left(\mathbb{E}_{((\mathbf{x}^0, \mathbf{y}^0)) \sim \mathcal{D}_e^0} [|\text{pos}(h(\mathbf{x}^0, \mathbf{y}^0)) - \text{pos}(\tilde{h}(\mathbf{x}^0, \mathbf{y}^0))|] \right. \\ & \left. + \mathbb{E}_{((\mathbf{x}^1, \mathbf{y}^1)) \sim \mathcal{D}_e^1} [|\text{pos}(h(\mathbf{x}^1, \mathbf{y}^1)) - \text{pos}(\tilde{h}(\mathbf{x}^1, \mathbf{y}^1))|] \right) \\ = & \frac{\nu}{2\ell} - 2\mathbb{E}_{((\mathbf{X},\mathbf{Y}),a) \sim \mathcal{D}_e} [|\text{pos}(h(\mathbf{X},\mathbf{Y})) - \text{pos}(\tilde{h}(\mathbf{X},\mathbf{Y}))|] \end{aligned} \quad (24)$$

$$\geq \frac{\nu}{2\ell} - \tau^{1/4}, \quad (25)$$

where the equality (24) is due to a being 0 or 1 with equal probability under \mathcal{D}^e , and the final inequality uses (22).

From the structural properties (A), (B) and (C) of \tilde{h} listed above it is easy to see that $(\tilde{\mathbf{c}}_X, \tilde{\mathbf{c}}_Y)$ satisfies (35), and

$$C_{\tau}(\tilde{\mathbf{c}}_{X,v}) = C_{\tau}^{\leq K}(\tilde{\mathbf{c}}_{X,v}) = C_{\tau}^{\leq K}(\mathbf{c}_{X,v}), \text{ and } \sum_{i \notin C_{\tau}(\tilde{\mathbf{c}}_{X,v})} \|\tilde{\mathbf{c}}_{X,v,i}\|_2^2 = \sum_{i \notin C_{\tau}(\mathbf{c}_{X,v})} \|\mathbf{c}_{X,v,i}\|_2^2 \text{ for } v \in e_X, \quad (26)$$

$$C_{\tau}(\tilde{\mathbf{c}}_{Y,v}) = C_{\tau}^{\leq K}(\tilde{\mathbf{c}}_{Y,v}) = C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}), \text{ and } \sum_{i \notin C_{\tau}(\tilde{\mathbf{c}}_{Y,v})} \|\tilde{\mathbf{c}}_{Y,v,i}\|_2^2 = \sum_{i \notin C_{\tau}(\mathbf{c}_{Y,v})} \|\mathbf{c}_{Y,v,i}\|_2^2 \text{ for } v \in e_Y, \quad (27)$$

Together with (B) and (C), (26) and (27) imply that $I_v(\mathbf{c}_X, \mathbf{c}_Y) \supseteq I_v(\tilde{\mathbf{c}}_X, \tilde{\mathbf{c}}_Y)$ and thus e remains nice w.r.t. $(\tilde{\mathbf{c}}_X, \tilde{\mathbf{c}}_Y)$. From our assumption that e does not satisfy the conditions from Lemma 5.3, we claim that $(\tilde{\mathbf{c}}_X, \tilde{\mathbf{c}}_Y)$ satisfies both the conditions of Lemma 7.1. In particular, by negating Condition I of Lemma 5.3, for any pair of vertices $u, v \in e$ we have

$$\begin{aligned} & \pi_{e,u} \left(C_{\tau}^{\leq K}(\mathbf{c}_{X,u}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,u}) \right) \cap \pi_{e,v} \left(C_{\tau}^{\leq K}(\mathbf{c}_{X,v}) \cup C_{\tau}^{\leq K}(\mathbf{c}_{Y,v}) \right) = \emptyset \\ \Rightarrow & \pi_{e,u} \left(C_{\tau}(\tilde{\mathbf{c}}_{X,u}) \cup C_{\tau}(\tilde{\mathbf{c}}_{Y,u}) \right) \cap \pi_{e,v} \left(C_{\tau}(\tilde{\mathbf{c}}_{X,v}) \cup C_{\tau}(\tilde{\mathbf{c}}_{Y,v}) \right) = \emptyset \end{aligned}$$

which gives us Condition 1 of Lemma 7.1. Towards establishing the Condition 2 of Lemma 7.1, we observe that for P as used in Lemma 7.1, any $v \in e_X$ and $j \in P \setminus \pi_{e,v} \left(C_{\tau}^{\leq K}(\tilde{\mathbf{c}}_{X,v}) \right)$ we have

$$\begin{aligned} \sum_{i \in \pi_{e,v}^{-1}(j) \setminus C_{\tau}(\tilde{\mathbf{c}}_{X,v})} \|\tilde{\mathbf{c}}_{X,v,i}\|_2^2 &= \sum_{i \in \pi_{e,v}^{-1}(j) \setminus C_{\tau}(\mathbf{c}_{X,v})} \|\tilde{\mathbf{c}}_{X,v,i}\|_2^2 + \sum_{i \in \pi_{e,v}^{-1}(j) \cap (C_{\tau}(\mathbf{c}_{X,v}) \setminus C_{\tau}^{\leq K}(\mathbf{c}_{X,v}))} \|\tilde{\mathbf{c}}_{X,v,i}\|_2^2 \\ &= \sum_{i \in \pi_{e,v}^{-1}(j) \setminus C_{\tau}(\mathbf{c}_{X,v})} \|\mathbf{c}_{X,v,i}\|_2^2 \leq \tau^4 \sum_{i \notin C_{\tau}(\mathbf{c}_{X,v})} \|\mathbf{c}_{X,v,i}\|_2^2 = \tau^4 \sum_{i \notin C_{\tau}(\tilde{\mathbf{c}}_{X,v})} \|\tilde{\mathbf{c}}_{X,v,i}\|_2^2 \end{aligned}$$

where the first equality uses the first part of (26), the second and the last equalities follow property (B) of \tilde{h} , and the inequality uses our assumption that e does not satisfy Condition II from Lemma 5.3. Similar guarantees also hold for any $v \in e_Y$, which together give us Condition 2 of Lemma 7.1. Therefore, by applying Lemma 7.1 we get that

$$\mathbb{E}_{\tilde{\mathcal{D}}} \left[\left| \text{pos}(\tilde{h}(\mathbf{X}^0, \mathbf{Y}^0)) - \text{pos}(\tilde{h}(\mathbf{X}^1, \mathbf{Y}^1)) \right| \right] \leq O(\tau),$$

which is a contradiction to (25) for small enough setting of τ , which can be achieved through (8) by setting z in Theorem 2.6 large enough.

6 Truncating Long Critical Index Lists

Fix a vertex $v \in V_{\mathcal{L}}$. Consider a linear form h given by $h(\mathbf{X}, \mathbf{Y}) = \langle \mathbf{c}_X, \mathbf{X} \rangle + \langle \mathbf{c}_Y, \mathbf{Y} \rangle + \theta$. Let \tilde{h} be the linear form of *truncated* coefficient vectors given by $\tilde{h}(\mathbf{X}, \mathbf{Y}) = \langle \tilde{\mathbf{c}}_X, \mathbf{X} \rangle + \langle \tilde{\mathbf{c}}_Y, \mathbf{Y} \rangle + \theta$, where

- $\tilde{\mathbf{c}}_{X,u} = \mathbf{c}_{X,u}$ and $\tilde{\mathbf{c}}_{Y,u} = \mathbf{c}_{Y,u}$ for all $u \in V_{\mathcal{L}}, u \neq v$.
- For $v \in e_X$: $\tilde{\mathbf{c}}_{X,v,i} = \mathbf{0}$ for all $i \in C_\tau(\mathbf{c}_{X,v}) \setminus C_\tau^{\leq K}(\mathbf{c}_{X,v})$ and $\tilde{\mathbf{c}}_{X,v,i} = \mathbf{c}_{X,v,i}$ otherwise.
- For $v \in e_Y$: $\tilde{\mathbf{c}}_{Y,v,i} = \mathbf{0}$ for all $i \in C_\tau(\mathbf{c}_{Y,v}) \setminus C_\tau^{\leq K}(\mathbf{c}_{Y,v})$ and $\tilde{\mathbf{c}}_{Y,v,i} = \mathbf{c}_{Y,v,i}$ otherwise.

This section proves the following lemma.

Lemma 6.1 *Given the above setting, for any hyperedge e , such that e is nice w.r.t. $(\mathbf{c}_X, \mathbf{c}_Y)$ (as given in Lemma 5.1), and $v \in e$ fixed above, the following holds:*

$$\mathbb{E}_{(\mathbf{X}, \mathbf{Y}, a) \sim \mathcal{D}_e} \left[\left| \text{pos}(h(\mathbf{X}, \mathbf{Y})) - \text{pos}(\tilde{h}(\mathbf{X}, \mathbf{Y})) \right| \right] \leq \frac{\tau^{1/4}}{4k}. \quad (28)$$

Proof. Assume that $v \in e_X$ (we shall handle the $v \in e_Y$ case analogously). Given this, we may further assume that $C_\tau(\mathbf{c}_{X,v}) \setminus C_\tau^{\leq K}(\mathbf{c}_{X,v}) \neq \emptyset$ implying that $|C_\tau^{\leq K}(\mathbf{c}_{X,v})| = K$, otherwise $h = \tilde{h}$ under \mathcal{D}_e . For ease of notation we relabel the indices in $[M]$ so that $C_\tau^{\leq K}(\mathbf{c}_{X,v}) = [K]$, and denote $\pi = \pi_{e,v}$. First we bound the difference between h and \tilde{h} as follows:

$$\begin{aligned} |h(\mathbf{X}, \mathbf{Y}) - \tilde{h}(\mathbf{X}, \mathbf{Y})| &\leq \sum_{i \in C_\tau(\mathbf{c}_{X,v}) \setminus [K]} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle \\ &\leq \sum_{i \in C_\tau(\mathbf{c}_{X,v}) \setminus [K]} \|\mathbf{c}_{X,v,i}\|_1 \\ \text{(By Cauchy-Schwartz)} &\leq \sqrt{Q} \sum_{i \in C_\tau(\mathbf{c}_{X,v}) \setminus [K]} \|\mathbf{c}_{X,v,i}\|_2 \\ \text{(By Proposition 2.4)} &\leq \sqrt{Q/\tau} (1-\tau)^{K/4} \|\mathbf{c}_{X,v,K/2}\|_2 \sum_{i=1}^{|C_\tau(\mathbf{c}_{X,v})| - K} (1-\tau)^{i/2} \\ &\leq \sqrt{Q/\tau} (1-\tau)^{K/4} (2/\tau) \|\mathbf{c}_{X,v,K/2}\|_2 \\ \text{(By our setting of } K) &\leq \tau^2 \|\mathbf{c}_{X,v,K/2}\|_2 \end{aligned} \quad (29)$$

Given a choice of $\{S_j, S'_j, b_j\}_{j=1}^m$ we can define the following subset:

$$A := \{i \in [K/4] : \pi(i) = j, b_j = 0, v \notin S_j\} \quad (30)$$

Let \mathcal{D}' be the restriction of \mathcal{D}^e fixing everything else except for values of $\mathbf{X}_{v,i}$ for $i \in A \cup (C_\tau(\mathbf{c}_{X,v}) \setminus [K])$. We have the following two claims which we shall prove later in this section.

Claim 6.2

$$\Pr_{\mathcal{D}'} \left[\left| \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \Delta \right| \leq \|\mathbf{c}_{X,v,K/2}\|_2 \right] \leq |A|^{-1/4}, \quad (31)$$

for any constant Δ .

Claim 6.3

$$\Pr_{\mathcal{D}^e} [|A| \leq K\zeta/8] \leq \exp(-K\zeta/64).$$

From the construction of $\tilde{\mathbf{c}}_{X,v}$ defined earlier, we obtain that under \mathcal{D}' ,

$$\begin{aligned} \tilde{h}(\mathbf{X}, \mathbf{Y}) &= \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \theta', \\ \Rightarrow h(\mathbf{X}, \mathbf{Y}) &= \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \theta' + h(\mathbf{X}, \mathbf{Y}) - \tilde{h}(\mathbf{X}, \mathbf{Y}) \end{aligned} \quad (32)$$

where θ' is some constant. The above implies that,

$$\begin{aligned} \Pr_{\mathcal{D}'} [\text{pos}(h(\mathbf{X}, \mathbf{Y})) \neq \text{pos}(\tilde{h}(\mathbf{X}, \mathbf{Y}))] &\leq \Pr_{\mathcal{D}'} \left[\left| \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \theta' \right| \leq |h(\mathbf{X}, \mathbf{Y}) - \tilde{h}(\mathbf{X}, \mathbf{Y})| \right] \\ &\stackrel{\text{(By (29))}}{\leq} \Pr_{\mathcal{D}'} \left[\left| \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \theta' \right| \leq \tau^2 \|\mathbf{c}_{X,v,K/2}\|_2 \right] \\ &\stackrel{\text{(By Claim 6.2)}}{\leq} |A|^{-1/4}. \end{aligned} \quad (33)$$

Using the above along with Claim 6.3 we can upper bound the LHS of (28) by

$$(K\zeta/8)^{-1/4} + \exp(-K\zeta/64), \quad (34)$$

which is at most $\tau^{1/4}/4k$ by our setting of parameters in (8) and large enough z in Theorem 2.6.

For the case when $v \in e_Y$, the proof is analogous to the above. The quantitative difference arises from defining A instead as $\{i \in [K/4] : \pi(i) = j, b_j = 1, v \notin S'_j\}$ and replacement of ζ by $(1 - \zeta)$ in the corresponding version of Claim 6.3 and in (34). By our setting of the parameters, the LHS of (28) remains bounded by $\tau^{1/4}/4k$. \square

Proof. (of Claim 6.2) Observe that by Proposition 2.4,

$$\|\mathbf{c}_{X,v,K/2}\|_2 \leq \sqrt{1/\tau}(1 - \tau)^{-K/4} \min_{i \in A} \|\mathbf{c}_{X,v,i}\|_2 \leq (4/K) \min_{i \in A} \|\mathbf{c}_{X,v,i}\|_2 \leq (1/|A|) \min_{i \in A} \|\mathbf{c}_{X,v,i}\|_2,$$

where the penultimate inequality follows from the setting of K which is large enough. Using this the LHS of (31) can be upper bounded by,

$$\Pr_{\mathcal{D}'} \left[\left| \sum_{i \in A} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle + \Delta \right| \leq \frac{\min_{i \in A} \|\mathbf{c}_{X,v,i}\|_2}{|A|} \right],$$

which is at most $|A|^{-1/4}$ by an application of Lemma A.1. \square

Proof. (of Claim 6.3) Since e is nice w.r.t. $(\mathbf{c}_X, \mathbf{c}_Y)$, we have that $|\pi([K/4])| = K/4$. Thus, each $i \in [K/4]$ is independently chosen to be in A w.p. $\zeta(1 - t/k) \geq \zeta/2$. An application of Chernoff bound completes the proof. \square

7 Main Structural Lemma

For the rest of this section we shall consider a linear form given by $h(\mathbf{X}, \mathbf{Y}) = \langle \mathbf{c}_X, \mathbf{X} \rangle + \langle \mathbf{c}_Y, \mathbf{Y} \rangle + \theta$, and an edge e which is nice w.r.t $(\mathbf{c}_X, \mathbf{c}_Y)$. Let $\pi_v := \pi_{e,v}$ for $v \in e$. Further, we assume that $(\mathbf{c}_X, \mathbf{c}_Y)$ satisfies

$$C_\tau(\mathbf{c}_{X,u}) = C_\tau^{\leq K}(\mathbf{c}_{X,u}), \forall u \in e_X, \text{ and } C_\tau(\mathbf{c}_{Y,v}) = C_\tau^{\leq K}(\mathbf{c}_{Y,v}), \forall v \in e_Y. \quad (35)$$

For convenience we define the following notation for each $v \in e$:

$$\begin{aligned} B_v &:= \begin{cases} C_\tau(\mathbf{c}_{X,v}) & \text{if } v \in e_X, \\ C_\tau(\mathbf{c}_{Y,v}) & \text{if } v \in e_Y. \end{cases} & \mathbf{c}_v &:= \begin{cases} \mathbf{c}_{X,v} & \text{if } v \in e_X, \\ \mathbf{c}_{Y,v} & \text{if } v \in e_Y. \end{cases} \\ \beta_v &:= \begin{cases} 0 & \text{if } v \in e_X, \\ 1 & \text{if } v \in e_Y. \end{cases} & \hat{S}_{v,j} &:= \begin{cases} S_j & \text{if } v \in e_X, \\ S'_j & \text{if } v \in e_Y. \end{cases} \quad \forall j \in [m]. \end{aligned} \quad (36)$$

Further, let $P \subseteq [m]$ denote $\cup_{v \in e} \pi_v(B_v)$ and,

$$\mathbf{c}_v^{(j)} := (\mathbf{c}_{v,i})_{i \in \pi_v^{-1}(j)}, \quad \mathbf{c}_v^{\text{reg}} := (\mathbf{c}_{v,i})_{i \in [m] \setminus (\pi_v^{-1}(P) \cup B_v)}, \quad \mathbf{c}^{\text{reg}} = (\mathbf{c}_v^{\text{reg}})_{v \in e}. \quad (37)$$

We have the following lemma.

Lemma 7.1 *Given the above setting, if the following two conditions are satisfied,*

1. **No Weak Intersections:** For every $u, v \in e$, $\pi_u(B_u) \cap \pi_v(B_v) = \emptyset$,
2. **No Large Regular Top Blocks:** For every $v \in e$, $j \in P \setminus \pi_v(B_v)$, $\|\mathbf{c}_v^{(j)}\|^2 \leq \tau^4 \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2$,

then the following holds,

$$E_{\hat{D}} \left[\left| \text{pos}(h(\mathbf{X}^0, \mathbf{Y}^0)) - \text{pos}(h(\mathbf{X}^1, \mathbf{Y}^1)) \right| \right] \leq O(\tau). \quad (38)$$

The rest of this section is devoted to proving the above lemma. To achieve a contradiction, we assume its two conditions. First we make a couple of easily verifiable observations from condition 1 of the lemma.

Observation 7.2 *For every $j \in P$, there exists a unique vertex v_j such that $B_{v_j} \cap \pi_{v_j}^{-1}(j) \neq \emptyset$.*

Observation 7.3 *The set P satisfies $|P| \leq \sum_{v \in e} |B_v| \leq (2k)K$.*

Furthermore, condition 2 of Lemma 7.1 yields the following.

Claim 7.4 *For every vertex $v \in e$, $\|\mathbf{c}_v^{\text{reg}}\|^2 \geq (1 - \tau) \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2$.*

Proof. For any vertex v , we have

$$\sum_{j \in P} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \|\mathbf{c}_{v,i}\|^2 = \sum_{j \in P \setminus \pi_v(B_v)} \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2 + \sum_{j \in \pi_v(B_v)} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \|\mathbf{c}_{v,i}\|^2. \quad (39)$$

The first term in the RHS of the above can be bounded by,

$$\leq \sum_{j \in P} \tau^4 \left(\sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2 \right) \leq \tau^4 |P| \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2 \leq (\tau/2) \sum_{i \in B_v} \|\mathbf{c}_{v,i}\|^2 \quad (40)$$

For the second term, observe that by the niceness of e w.r.t. $(\mathbf{c}_X, \mathbf{c}_Y)$, for any $j \in \pi_v(B_v)$, we have $|\pi_v^{-1}(j) \cap C_{\tau}^{\leq K}(\mathbf{c})| = 1$, and therefore for any other $i \in \pi_v^{-1}(j) \setminus B_v$ we must have that the value of $\|\mathbf{c}_{v,j}\|^2$ is at most $(1/d^8) \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2$. Since the number of such values in the summation is at most $d|B_v| \leq dK$, by setting z in Theorem 2.6 large enough, this summation can be bounded by $(\tau/2) \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2$. Therefore,

$$\|\mathbf{c}_v^{\text{reg}}\|^2 = \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2 - \sum_{j \in P} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \|\mathbf{c}_{v,j}\|^2 \geq (1 - \tau) \sum_{j \notin B_v} \|\mathbf{c}_{v,j}\|^2.$$

□

Claim 7.5 For every vertex $v \in e$, the coefficient vector $\mathbf{c}_v^{\text{reg}}$ is τ' -regular where $\tau' = \tau/(1 - \tau)$.

Proof. This follows from the fact that $\{\mathbf{c}_{v,i}\}_{i \notin B_v}$ is τ -regular, contains all of the $\mathbf{c}_{v,i}$ constituting $\mathbf{c}_v^{\text{reg}}$, and from Claim 7.4. □

The following lemma provides a useful concentration for some $v \in e$ the sum of squared coefficients $\|\mathbf{c}_{v,i}\|^2$ corresponding to those indices i constituting $\mathbf{c}_v^{\text{reg}}$ s.t. the corresponding variables $(\mathbf{X}_{v,i}$ or $\mathbf{Y}_{v,i}$ depending on whether v is in e_X or e_Y) are sampled u.a.r. from $\{0, 1\}^Q$.

Lemma 7.6 For any vertex $v \in e$, over the choice of $\{b_j, S_j, S'_j\}_{j=1}^m$,

$$\Pr \left[\sum_{\substack{j \in [m] \setminus P: (b_j = \beta_v) \\ \wedge (v \notin \hat{S}_{v,j})}} \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2 \leq \frac{\zeta}{8} \|\mathbf{c}_v^{\text{reg}}\|^2 \right] \leq \exp\left(-\frac{\zeta^2}{64\tau}\right), \quad (41)$$

where

Proof. Define the random variable $\phi_j := \mathbb{1}\{b_j = \beta_v, v \notin \hat{S}_{v,j}\} \cdot \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2$ for each $j \in [m] \setminus P$. Note that $\{\phi_j\}_{j \in [m] \setminus P}$ are independent non-negative random variables and the summation inside the probability expression on the LHS of (41) is precisely the random variable $\sum_{j \in [m] \setminus P} \phi_j$. Since $\{b_j = \beta_v, v \notin \hat{S}_{v,j}\}$ occurs with probability $(1 - \zeta)(1 - t/k)$ if $v \in e_Y$ and $\zeta(1 - t/k)$ if $v \in e_X$, letting $\alpha_v := (1 - t/k)(\beta_v \zeta + (1 - \beta_v)(1 - \zeta))$ we obtain,

$$\mathbb{E} \left[\sum_{j \in [m] \setminus P} \phi_j \right] = \alpha_v \|\mathbf{c}_v^{\text{reg}}\|^2. \quad (42)$$

Further, for each $j \in [m] \setminus P$.

$$(\max \phi_j)^2 = \left(\sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2 \right)^2 = \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^4 + \sum_{\substack{i, i' \in \pi_v^{-1}(j) \\ i \neq i'}} \|\mathbf{c}_{v,i}\|^2 \|\mathbf{c}_{v,i'}\|^2 \quad (43)$$

The first term on the RHS of (43) can be upper bounded using the τ' -regularity of $\mathbf{c}_v^{\text{reg}}$ (Claim 7.5) as follows:

$$\sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^4 \leq \tau' \|\mathbf{c}_v^{\text{reg}}\|^2 \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2. \quad (44)$$

On the other hand using the niceness of e and Claim 7.4, we obtain that all $i \in \pi_v^{-1}(j)$ except for at most one satisfy $\|\mathbf{c}_{v,i}\|^2 < (1/d^8) \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2 \leq (2/d^8) \|\mathbf{c}_v^{\text{reg}}\|^2$. Using this, the second term on the RHS of (43) is at most,

$$\begin{aligned} \sum_{\substack{i,i' \in \pi_v^{-1}(j) \\ i \neq i'}} (\max\{\|\mathbf{c}_{v,i}\|^2, \|\mathbf{c}_{v,i'}\|^2\} \cdot (2/d^8) \|\mathbf{c}_v^{\text{reg}}\|^2) &\leq \frac{4}{d^8} \sum_{i \in \pi_v^{-1}(j)} \left[\|\mathbf{c}_{v,i}\|^2 \sum_{\substack{i' \in \pi_v^{-1}(j) \\ i' \neq i}} \|\mathbf{c}_v^{\text{reg}}\|^2 \right] \\ &\leq \left(\frac{4}{d^7} \right) \|\mathbf{c}_v^{\text{reg}}\|^2 \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2. \end{aligned} \quad (45)$$

Combining the above with (44) and (43) we obtain,

$$\sum_{j \in [m] \setminus P} (\max \phi_j)^2 \leq (\tau' + 4/d^7) \|\mathbf{c}_v^{\text{reg}}\|^2 \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2 \leq 3\tau \|\mathbf{c}_v^{\text{reg}}\|^4. \quad (46)$$

Using the above along with (42), we apply the Chernoff-Hoeffding inequality (Theorem 2.3) to the sum $\sum_{j \in [m] \setminus P} \phi_j$ as follows:

$$\Pr \left[\sum_{j \in [m] \setminus P} \phi_j \leq \frac{\alpha_v}{2} \|\mathbf{c}_v^{\text{reg}}\|^2 \right] \leq 2 \exp \left(\frac{-\alpha_v^2 \|\mathbf{c}_v^{\text{reg}}\|^4}{6\tau \|\mathbf{c}_v^{\text{reg}}\|^4} \right) \leq \exp \left(-\frac{\alpha_v^2}{12\tau} \right), \quad (47)$$

which using the fact that $\alpha_v > \zeta/2$ (since $\zeta < 1/2$ and $t < k/2$) proves the lemma. \square

A repeated application of Lemma 7.6 for each vertex $v \in e$ along with a union bound yields,

$$\Pr \left[\sum_{v \in e} \sum_{\substack{j \in [m] \setminus P: (b_j = \beta_v) \\ \wedge (v \notin \hat{S}_{v,j})}} \sum_{i \in \pi_v^{-1}(j)} \|\mathbf{c}_{v,i}\|^2 \leq \frac{\zeta}{8} \|\mathbf{c}_v^{\text{reg}}\|^2 \right] \leq 2k \cdot \exp \left(-\frac{\zeta^2}{64\tau} \right) \quad (48)$$

Let us define,

$$\begin{aligned} \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j : j \in [m]\} \right) &:= \sum_{v \in e_X} \sum_{j \in [m] \setminus P: v \notin S_j} \sum_{i \in \pi_v^{-1}(j)} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i} \rangle \\ &+ \sum_{v \in e_Y} \sum_{j \in [m] \setminus P: v \notin S'_j} \sum_{i \in \pi_v^{-1}(j)} \langle \mathbf{c}_{Y,v,i}, \mathbf{Y}_{v,i} \rangle \end{aligned} \quad (49)$$

The above definition of \mathcal{S} captures the contribution due to the $\mathbf{c}_{v,i}$ constituting $\mathbf{c}_v^{\text{reg}}$ such that the variables ($\mathbf{X}_{v,i}$ or $\mathbf{Y}_{v,i}$ depending on whether v is in e_X or e_Y) are sampled u.a.r. from $\{0,1\}^Q$. We prove the following anti-concentration of \mathcal{S} .

Lemma 7.7

$$\Pr_{\mathcal{D}^e} \left[\left| \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j : j \in [m]\} \right) + \theta' \right| \leq \varepsilon_0 \right] \leq O(\tau) + \frac{64\varepsilon_0}{\|\mathbf{c}_v^{\text{reg}}\|_2 \sqrt{\zeta}} + 2k \cdot \exp \left(-\frac{\zeta^2}{64\tau} \right), \quad (50)$$

where $\varepsilon_0 \geq 0$ is a constant, and after fixing $\{b_j, S_j, S'_j\}_{j \in [m]}$ θ' does not depend on the variables $\mathbf{X}_{v,i}$ and $\mathbf{Y}_{v,i}$ in (49).

Proof. Define $\mathbf{Z}_v := \mathbf{X}_v - \frac{1}{2}\mathbf{1}$ if $v \in e_X$ and $\mathbf{Y}_v - \frac{1}{2}\mathbf{1}$ otherwise. Note that for any $v, j \in [m] \setminus P$ s.t. $b_j = \beta_v$ and $v \notin \hat{S}_{v,j}$, and $i \in \pi_v^{-1}(j)$, the vector $\mathbf{Z}_{v,i}$ is uniformly sampled from $\{-1/2, 1/2\}^Q$. Further we have,

$$\begin{aligned} \mathcal{S}(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j\}_{j \in [m]}) + \theta' &= \sum_{v \in e} \sum_{j \in [m] \setminus P: (b_j = \beta_v) \wedge (v \notin \hat{S}_{v,j})} \sum_{i \in \pi_v^{-1}(j)} \langle \mathbf{c}_{v,i}, \mathbf{Z}_{v,i} \rangle + \Theta \quad (51) \\ &= \sum_{v \in e} \sum_{j \in [m] \setminus P: (b_j = \beta_v) \wedge (v \notin \hat{S}_{v,j})} \sum_{i \in \pi_v^{-1}(j)} \sum_{q \in [Q]} c_{v,i,q} Z_{v,i,q} + \Theta, \quad (52) \end{aligned}$$

where Θ is a fixed constant after fixing $\{b_j, S_j\}_{j \in [m]}$ along with any value of θ' , and randomizing only on the values of the $\mathbf{Z}_{v,i}$ appearing on the RHS of (51). Any product $c_{v,i,q} Z_{v,i,q}$ on the RHS of (52) is a mean zero random variable with variance $c_{v,i,q}^2/4$ and third moment $|c_{v,i,q}^3|/8$. The maximum over all the random variables $c_{v,i,q} Z_{v,i,q}$ of the ratio of its third moment and second moments is the maximum value of $|c_{v,i,q}|/2$ over (v, i, q) appearing in (52). By Claim 7.5 this is at most $\tau' \|\mathbf{c}^{\text{reg}}\|$. Further, the sum of variances γ is (by (48)) at least $(\frac{\zeta}{32}) \|\mathbf{c}^{\text{reg}}\|^2$ except with probability $2k \cdot \exp(-\zeta^2/(64\tau))$ over the choice of $\{b_j, S_j, S'_j\}_{j \in [m]}$. Assuming this and applying the Berry-Esseen theorem (Theorem 2.1) we obtain that the CDF F of the RHS of (51) excluding Θ , satisfies $|F(x) - \Phi_\gamma(x)| \leq O(\tau') = O(\tau)$ for any $x \in (-\infty, \infty)$, where Φ_γ is the CDF of the normal distribution $N(0, \gamma)$. In this case, the LHS of (51) lies in a fixed interval of length $2\varepsilon_0$ with probability at most $O(\tau) + 16\varepsilon_0/(\sqrt{\zeta} \|\mathbf{c}^{\text{reg}}\|)$. Losing additional probability of $2k \cdot \exp(-\zeta^2/(64\tau))$ for our assumption by (48) completes the proof of Lemma 7.7. \square

7.1 Proof of Lemma 7.1

Our first goal is to bound the variance of, $|h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0)|$ under the distribution $\hat{\mathcal{D}}$. Let the difference for each $j \in [m]$ be represented by Δ_j so that $\mathbb{E}[h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0)] = \mathbb{E}[\sum_j \Delta_j]$. Observe that $\mathbb{E}[\Delta_j] = 0$ due to matching expectations of every variable under $\hat{\mathcal{D}}$. Explicitly,

$$\begin{aligned} \Delta_j := & \sum_{v \in e_X} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \mathbb{1}\{v = u_{X,j}\} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i}^1 \rangle + \sum_{v \in e_Y} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \mathbb{1}\{v = u_{Y,j}\} \langle \mathbf{c}_{Y,v,i}, \mathbf{Y}_{v,i}^1 \rangle \\ & - \sum_{v \in e_X} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \mathbb{1}\{v \in T_j\} \langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i}^0 \rangle + \sum_{v \in e_Y} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \mathbb{1}\{v \in T'_j\} \langle \mathbf{c}_{Y,v,i}, \mathbf{Y}_{v,i}^0 \rangle. \quad (53) \end{aligned}$$

Note that the $\mathbf{X}_{v,i}^1, \mathbf{X}_{v,i}^0, \mathbf{Y}_{v,i}^1$ or $\mathbf{Y}_{v,i}^0$ appearing in the above expression are independent random variables sampled u.a.r. from $\{\mathbf{e}_1, \dots, \mathbf{e}_Q\}$. Therefore,

$$\mathbb{E} \left[\left\langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i}^1 \right\rangle^2 \right] = \|\mathbf{c}_{X,v,i}\|^2 / Q, \quad (54)$$

and similarly for $\langle \mathbf{c}_{X,v,i}, \mathbf{X}_{v,i}^0 \rangle, \langle \mathbf{c}_{Y,v,i}, \mathbf{Y}_{v,i}^1 \rangle$, and $\langle \mathbf{c}_{Y,v,i}, \mathbf{Y}_{v,i}^0 \rangle$. There are at most $4kd$ inner products appearing in the above expression for Δ_j , with at most two corresponding to each (v, i) s.t. $v \in e, i \in \pi_v^{-1}(j) \setminus B_v$. Since $\mathbb{E}[\Delta_j] = 0$, we have $\text{Var}[\Delta_j] = \mathbb{E}[\Delta_j^2]$ which, by Cauchy-Schwartz and (54) is at most,

$$\left(\sqrt{4kd}\right) \cdot 2 \sum_{v \in e} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \left(\frac{\|\mathbf{c}_{v,i}\|^2}{Q}\right) = \frac{1}{\sqrt{Q}} \sum_{v \in e} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \|\mathbf{c}_{v,i}\|^2,$$

where the inequality is implied by the setting of Q in (8). Summing up over all $j \in [m]$ and using independence of $\{\Delta_j, j \in [m]\}$, we obtain that,

$$\begin{aligned} \text{Var} \left[\left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| \right] &\leq \frac{1}{\sqrt{Q}} \sum_{j \in [m]} \sum_{v \in \mathcal{E}} \sum_{i \in \pi_v^{-1}(j) \setminus B_v} \|\mathbf{c}_{v,i}\|^2 \\ &\leq \frac{1}{\sqrt{Q}} \sum_{v \in \mathcal{E}} \sum_{i \notin B_v} \|\mathbf{c}_{v,i}\|^2 \\ &\leq \frac{1}{\sqrt{Q}} (1 - \tau)^{-1} \sum_{v \in \mathcal{E}} \|\mathbf{c}_v^{\text{reg}}\|^2 \leq \frac{2\|\mathbf{c}^{\text{reg}}\|^2}{\sqrt{Q}} \end{aligned} \quad (55)$$

where the second last inequality follows from applying Claim 7.4 on each inner summation. Applying Chebyshev's inequality we obtain for any $\varepsilon_0 > 0$,

$$\Pr_{\widehat{\mathcal{D}}} \left[\left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| > \varepsilon_0 \right] < \frac{2\|\mathbf{c}^{\text{reg}}\|^2}{\varepsilon_0^2 \sqrt{Q}}. \quad (56)$$

To complete the proof of Lemma 7.1, observe that under $\widehat{\mathcal{D}}$,

$$\begin{aligned} \mathbf{X}_{v,i}^1 &= \mathbf{X}_{v,i}^0 =: \mathbf{X}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q, \quad \forall v \in e_X, j \in [m] \setminus P : b_j = 0, v \notin S_j, i \in \pi_v^{-1}(j), \\ \mathbf{Y}_{v,i}^1 &= \mathbf{Y}_{v,i}^0 =: \mathbf{Y}_{v,i} \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q, \quad \forall v \in e_Y, j \in [m] \setminus P : b_j = 1, v \notin S'_j, i \in \pi_v^{-1}(j). \end{aligned}$$

Thus, under $\widehat{\mathcal{D}}$

$$h(\mathbf{X}^0, \mathbf{Y}^0) = \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j : j \in [m]\} \right) + \theta' \quad (57)$$

$$\Rightarrow h(\mathbf{X}^1, \mathbf{Y}^1) = \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j : j \in [m]\} \right) + \theta' + h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0), \quad (58)$$

where θ' does not depend on the variables $\mathbf{X}_{v,i}$ and $\mathbf{Y}_{v,i}$ appearing in the expression (49) for \mathcal{S} after fixing $\{b_j, S_j, S'_j\}_{j \in [m]}$. The equations (57) and (58) also imply that (for some constant ε_0 that we shall choose shortly),

$$\begin{aligned} &\Pr_{\widehat{\mathcal{D}}} \left[\text{pos}(h(\mathbf{X}^1, \mathbf{Y}^1)) \neq \text{pos}(h(\mathbf{X}^0, \mathbf{Y}^0)) \right] \\ &\leq \Pr_{\widehat{\mathcal{D}}} \left[\left| \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j\}_{j \in [m]} \right) + \theta' \right| \leq \left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| \right] \\ &\leq \Pr_{\widehat{\mathcal{D}}} \left[\neg \left(\left(\left| \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j\}_{j \in [m]} \right) + \theta' \right| \geq \varepsilon_0 \right) \wedge \left(\left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| \leq \varepsilon_0 \right) \right) \right] \\ &\leq \Pr_{\widehat{\mathcal{D}}} \left[\left(\left| \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j\}_{j \in [m]} \right) + \theta' \right| \leq \varepsilon_0 \right) \vee \left(\left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| \geq \varepsilon_0 \right) \right] \\ &\leq \Pr_{\widehat{\mathcal{D}}} \left[\left| \mathcal{S} \left(\mathbf{c}_X, \mathbf{c}_Y, \mathbf{X}, \mathbf{Y}, \{b_j, S_j, S'_j\}_{j \in [m]} \right) + \theta' \right| \leq \varepsilon_0 \right] + \Pr_{\widehat{\mathcal{D}}} \left[\left| h(\mathbf{X}^1, \mathbf{Y}^1) - h(\mathbf{X}^0, \mathbf{Y}^0) \right| \geq \varepsilon_0 \right] \end{aligned} \quad (59)$$

Choosing ε_0 to be $\tau \sqrt{\zeta} \|\mathbf{c}^{\text{reg}}\| / 64$, and applying Lemma 7.7 and (56) we obtain that the LHS of (38) is bounded by,

$$O(\tau) + 2k \cdot \exp \left(-\frac{\zeta^2}{64\tau} \right) + O \left(\left(\tau^2 \zeta \sqrt{Q} \right)^{-1} \right),$$

which by our setting of the parameters in (8) is $O(\tau)$ for a large enough z in Theorem 2.6, completing the proof.

References

- [ABF⁺08] Misha Alekhnovich, Mark Braverman, Vitaly Feldman, Adam R Klivans, and Toniann Pitassi. The complexity of properly learning simple concept classes. *Journal of Computer and System Sciences*, 74(1):16–34, 2008.
- [ABX08] Benny Applebaum, Boaz Barak, and David Xiao. On basing lower-bounds for learning on worst-case assumptions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 211–220. IEEE, 2008.
- [Ale05] Michael Alekhnovich. Lower bounds for k-dnf resolution on random 3-cnfs. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 251–256. ACM, 2005.
- [AMO15] Albert Atserias, Moritz Müller, and Sergi Oliva. Lower bounds for dnf-refutations of a relativized weak pigeonhole principle. *The Journal of Symbolic Logic*, 80(2):450–476, 2015.
- [Baz09] Louay MJ Bazzi. Polylogarithmic independence can fool dnf formulas. *SIAM Journal on Computing*, 38(6):2220–2272, 2009.
- [BFJ⁺94] Avrim Blum, Merrick Furst, Jeffrey Jackson, Michael Kearns, Yishay Mansour, and Steven Rudich. Weakly learning dnf and characterizing statistical query learning using fourier analysis. In *STOC*, volume 94, pages 253–262, 1994.
- [BGS18] Arnab Bhattacharyya, Suprovat Ghoshal, and Rishi Saket. Hardness of learning noisy halfspaces using polynomial thresholds. In *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018.*, pages 876–917, 2018.
- [BMOS03] Nader Bshouty, Elchanan Mossel, Ryan O’Donnell, and Rocco A Servedio. Learning dnf from random walks. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 189–198. IEEE, 2003.
- [Bsh96] Nader H Bshouty. A subexponential exact learning algorithm for dnf using equivalence queries. *Information Processing Letters*, 59(1):37–39, 1996.
- [DHK⁺10] Ilias Diakonikolas, Prahladh Harsha, Adam Klivans, Raghu Meka, Prasad Raghavendra, Rocco A Servedio, and Li-Yang Tan. Bounding the average sensitivity and noise sensitivity of polynomial threshold functions. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 533–542. ACM, 2010.
- [DOSW11] Ilias Diakonikolas, Ryan O’Donnell, Rocco A. Servedio, and Yi Wu. Hardness results for agnostically learning low-degree polynomial threshold functions. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1590–1606, 2011.
- [DSS16] Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning dnf’s. In *Conference on Learning Theory*, pages 815–830, 2016.
- [Erd45] Paul Erdős. On a lemma of littlewood and offord. *Bulletin of the American Mathematical Society*, 51(12):898–902, 1945.
- [Fel06] Vitaly Feldman. Optimal hardness results for maximizing agreements with monomials. In *IEEE CCC 2006*, pages 226–236, 2006.

- [Fel09] Vitaly Feldman. Hardness of approximate two-level logic minimization and PAC learning with membership queries. *J. Comput. Syst. Sci.*, 75(1):13–26, 2009.
- [Fel12] Vitaly Feldman. Learning dnf expressions from fourier spectrum. In *Conference on Learning Theory*, pages 17–1, 2012.
- [FGKP09] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami. On agnostic learning of parities, monomials, and halfspaces. *SIAM J. Comput.*, 39(2):606–645, 2009.
- [FGRW12] Vitaly Feldman, Venkatesan Guruswami, Prasad Raghavendra, and Yi Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing*, 41(6):1558–1590, 2012.
- [GKKN18] Lee-Ad Gottlieb, Eran Kaufman, Aryeh Kontorovich, and Gabriel Nivasch. Learning convex polytopes with margin. In *Advances in Neural Information Processing Systems*, pages 5706–5716, 2018.
- [GMR13] Parikshit Gopalan, Raghu Meka, and Omer Reingold. Dnf sparsification and a faster deterministic counting algorithm. *Computational Complexity*, 22(2):275–310, 2013.
- [GR09] Venkatesan Guruswami and Prasad Raghavendra. Hardness of learning halfspaces with noise. *SIAM J. Comput.*, 39(2):742–765, 2009.
- [Jac97] Jeffrey C Jackson. An efficient membership-query algorithm for learning dnf with respect to the uniform distribution. *Journal of Computer and System Sciences*, 55(3):414–440, 1997.
- [KLW10] Adam R Klivans, Homin K Lee, and Andrew Wan. Mansour’s conjecture is true for random dnf formulas. In *COLT*, pages 368–380. Citeseer, 2010.
- [KOS04] Adam R Klivans, Ryan O’Donnell, and Rocco A Servedio. Learning intersections and thresholds of halfspaces. *Journal of Computer and System Sciences*, 68(4):808–840, 2004.
- [KS04] Adam R. Klivans and Rocco A. Servedio. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. *J. Comput. Syst. Sci.*, 68(2):303–318, 2004.
- [KS08a] Subhash Khot and Rishi Saket. Hardness of minimizing and learning dnf expressions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 231–240. IEEE, 2008.
- [KS08b] Subhash Khot and Rishi Saket. On hardness of learning intersection of two halfspaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 345–354. ACM, 2008.
- [KS09] Adam R Klivans and Alexander A Sherstov. Cryptographic hardness for learning intersections of halfspaces. *Journal of Computer and System Sciences*, 75(1):2–12, 2009.
- [KST09] Adam Tauman Kalai, Alex Samorodnitsky, and Shang-Hua Teng. Learning and smoothed analysis. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 395–404. IEEE, 2009.
- [Lee10] Homin K. Lee. Learning talagrand DNF formulas. In *COLT 2010 - The 23rd Conference on Learning Theory, Haifa, Israel, June 27-29, 2010*, pages 310–311, 2010.

- [LLZ18] Xin Li, Shachar Lovett, and Jiapeng Zhang. Sunflowers and quasi-sunflowers from randomness extractors. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [LZ19] Shachar Lovett and Jiapeng Zhang. Dnf sparsification beyond sunflowers. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 454–460. ACM, 2019.
- [NJS98] Richard Nock, Pascal Jappy, and Jean Sallantin. Generalized graph colorability and compressibility of boolean formulae. In *International Symposium on Algorithms and Computation*, pages 237–246. Springer, 1998.
- [O’D14] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [PV88] Leonard Pitt and Leslie G Valiant. Computational limitations on learning from examples. *Journal of the ACM (JACM)*, 35(4):965–984, 1988.
- [Ros58] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [SBI04] Nathan Segerlind, Sam Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k-dnf resolution. *SIAM Journal on Computing*, 33(5):1171–1200, 2004.
- [Ser04] Rocco A Servedio. On learning monotone dnf under product distributions. *Information and Computation*, 193(1):57–74, 2004.
- [Ser06] Rocco A Servedio. Every linear threshold function has a low-weight approximator. In *21st Annual IEEE Conference on Computational Complexity (CCC’06)*, pages 18–32. IEEE, 2006.
- [V⁺98] Vladimir Vapnik et al. *Statistical learning theory*, 1998.
- [Val84] Leslie G Valiant. A theory of the learnable. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 436–445. ACM, 1984.
- [Ver90] Karsten Verbeurgt. Learning dnf under the uniform distribution in quasi-polynomial time. In *Proceedings of the third annual workshop on Computational learning theory*, pages 314–326. Morgan Kaufmann Publishers Inc., 1990.

A Blockwise Small Ball Probability

Here we prove the following extension of the Littlewood-Offord-Erdős lemma.

Lemma A.1 *Let $\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_T \in \mathbb{R}^Q$ be such that $\|\mathbf{c}_1\| \geq \|\mathbf{c}_2\| \geq \dots \geq \|\mathbf{c}_T\|$. Furthermore, let $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_T$ be Q -dimensional Bernoulli vector random variables i.e., $\mathbf{X}_i \stackrel{\text{u.a.r.}}{\sim} \{0, 1\}^Q$ for every $i \in [T]$. Then*

$$\sup_{\theta \in \mathbb{R}} \Pr \left[\left| \sum_{i \in [T]} \langle \mathbf{c}_i, \mathbf{X}_i \rangle + \theta \right| \leq \frac{\|\mathbf{c}_T\|}{T^{1/2}} \right] \leq O(T^{-1/2})$$

Proof. Let $\eta > 0$ be a quantity which is fixed later. For any $i \in [T]$, we write $\mathbf{c}_i = (c_{i,q})_{q \in [Q]}$. We consider the following two cases.

Case (i)

$$\min_{i \in [T]} \max_{q \in [Q]} |c_{i,q}| \leq \eta \|\mathbf{c}_T\|$$

Then there exists $i^* \in [T]$ such that for every choice of $q \in [Q]$, we have $|c_{i^*,q}| \leq \eta \|\mathbf{c}_K\| \leq \eta \|\mathbf{c}_{i^*}\|$. For every $q \in [Q]$, define random variable $Z_q = c_{i^*,q} (X_{i^*,q} - \frac{1}{2})$, and let $\sigma_q^2 := \mathbb{E} Z_q^2$. Furthermore, define $\sigma^2 := \sum_{q \in [Q]} \sigma_q^2$. It is easy to verify that $\sigma_r^2 = c_{i^*,q}^2/2$ and $\sigma = \|\mathbf{c}_{i^*}\|/\sqrt{2}$.

Again for every choice of $q \in [Q]$, define $\tilde{Z}_q = Z_q/\sigma$. Then by construction we have (i) $\mathbb{E}[Z_q] = 0$ for every $q \in [Q]$ (ii) $\mathbb{E} \sum_{q \in [Q]} \tilde{Z}_q^2 = 1$ and (iii) $\mathbb{E} \sum_{q \in [Q]} |\tilde{Z}_q|^3 \leq \eta$. Therefore, using the Berry Esseen Theorem (Theorem 2.1), for any interval $I \subset \mathbb{R}$ we have

$$\Pr \left[\left| \sum_{q \in [Q]} \tilde{Z}_q \right| \in I \right] \leq O(\eta) + \Pr_{g \sim N(0,1)} [g \in I] \leq O(\eta) + |I|$$

Rolling back the sequence of transformations, for any fixing of $(\mathbf{X}_j)_{j \neq i^*}$ and any choice of $\theta \in \mathbb{R}$, we get that

$$\Pr_{\mathbf{X}_{i^*}} \left[\left| \sum_{i \neq i^*} \langle \mathbf{c}_i, \mathbf{X}_i \rangle + \langle \mathbf{c}_{i^*}, \mathbf{X}_{i^*} \rangle + \theta \right| \leq \eta \|\mathbf{c}_T\| \right] \leq O(\eta) + \frac{\eta \|\mathbf{c}_T\|}{\|\mathbf{c}_{i^*}\|} \leq O(\eta)$$

which gives us the bound for this case.

Case (ii)

$$\min_{i \in [T]} \max_{q \in [Q]} |c_{i,r}| > \eta \|\mathbf{c}_K\|$$

Here, for every choice of $i \in [T]$, there exists $q_i \in [Q]$ such that $|c_{i,q_i}| \geq \eta \|\mathbf{c}_T\|$. Then using Littlewood-Offord-Erdős Lemma (Lemma 2.2), for any $\theta \in \mathbb{R}$ we get that

$$\begin{aligned} \Pr_{(X_{i,q_i})_{i=1}^T} \left[\left| \sum_{i \in [T]} c_{i,q_i} X_{i,q_i} + \theta \right| \leq \eta \|\mathbf{c}_T\| \right] &= \Pr_{(X_{i,q_i})_{i=1}^T} \left[\left| \sum_{i \in [T]} \left(\frac{c_{i,q_i}}{\eta \|\mathbf{c}_T\|} \right) X_{i,q_i} + \frac{\theta}{\eta \|\mathbf{c}_T\|} \right| \leq 1 \right] \\ &\leq O(T^{-1/2}) \end{aligned}$$

Since the above bound holds independent of the realization of $(X_{i,\neq r(i)})_{i=1}^T$ we have

$$\Pr_{(\mathbf{X}_i)_{i=1}^T} \left[\left| \sum_{i \in [T]} \langle \mathbf{c}_i, \mathbf{X}_i \rangle + \theta \right| \leq \eta \|\mathbf{c}_T\| \right] \leq O(T^{-1/2})$$

for any fixed choice of $\theta \in \mathbb{R}$. Combining the two cases and setting $\eta = 1/\sqrt{T}$ completes the proof of the lemma. \square