SQ Lower Bounds for Learning Mixtures of Linear Classifiers

Ilias Diakonikolas[*]
University of Wisconsin-Madison
ilias@cs.wisc.edu

Daniel M. Kane[†]
University of California, San Diego
dakane@cs.ucsd.edu

Yuxin Sun[‡]
University of Wisconsin-Madison
yxsun@cs.wisc.edu

October 19, 2023

Abstract

We study the problem of learning mixtures of linear classifiers under Gaussian covariates. Given sample access to a mixture of r distributions on $\mathbb{R}^n$ of the form $(\mathbf{x}, y_\ell)$, $\ell \in [r]$, where $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_n)$ and $y_\ell = \text{sign}(\langle \mathbf{v}_\ell, \mathbf{x} \rangle)$ for an unknown unit vector $\mathbf{v}_\ell$, the goal is to learn the underlying distribution in total variation distance. Our main result is a Statistical Query (SQ) lower bound suggesting that known algorithms for this problem are essentially best possible, even for the special case of uniform mixtures. In particular, we show that the complexity of any SQ algorithm for the problem is $n^{\text{poly}(1/\Delta)\log(r)}$, where Δ is a lower bound on the pairwise ℓ_2-separation between the $\mathbf{v}_\ell$'s. The key technical ingredient underlying our result is a new construction of spherical designs that may be of independent interest.

1   Introduction

The motivation behind this work is to understand the computational complexity of learning high-dimensional latent variable (aka mixture) models. The task of learning various mixture models has a long history in statistics with an early focus on sample efficiency, starting with the pioneering work of Karl Pearson [Pea94] on learning Gaussian mixtures. During the past decades, an extensive line of work in machine learning and theoretical computer science has made significant progress on the computational aspects of this general question for a range of mixture models, including mixtures of high-dimensional Gaussians [Das99, AK01, VW02, AM05, KSV08, BV08, MV10, RV17, HL18, DKS18, DHKK20, LM20, BDJ+20, DKK+21], mixtures of linear regressions [SJA16, LL18, CLS20, DK20], and more generally mixtures of experts [JJ94, XA09, ME14, MVKO19, HYJ22].

In this paper, we focus on mixtures of *linear classifiers*, a classical supervised probabilistic model that has been intensely investigated from both statistical and algorithmic standpoints [BC92, SIM14, LLYH17, GMP20, CDV22]. A linear classifier (or halfspace) is any Boolean function $h : \mathbb{R}^n \to \{\pm 1\}$ of the form $h(\mathbf{x}) = \text{sign}(\langle \mathbf{v}, \mathbf{x} \rangle)$, where $\mathbf{v} \in \mathbb{R}^n$ is known as the weight vector and the univariate function sign is defined by $\text{sign}(t) = 1$ for $t \geq 0$ and $\text{sign}(t) = -1$ otherwise. For an integer $r \geq 2$, we can now formally describe an r-mixture of linear classifiers. The parameters of the model contain r unknown positive weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^r w_\ell = 1$, and r unknown unit vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{R}^n$. A random sample is drawn from the underlying distribution $D(\mathbf{x}, y)$ as follows: the sample oracle selects the index $\ell \in [r]$ with probability w_ℓ, and we then receive a random point $(\mathbf{x}, y) \in \mathbb{R}^n \times \{\pm 1\}$ where $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_n)$ and $y_\ell = \text{sign}(\langle \mathbf{v}_\ell, \mathbf{x} \rangle)$. The goal is to approximately estimate the model either by learning the underlying distribution $D(\mathbf{x}, y)$ in total variation distance (density estimation), or by approximately recovering the hidden parameters, i.e., w_ℓ and $\mathbf{v}_\ell$, $\ell \in [r]$ (parameter estimation). An algorithm for parameter estimation can be used for density estimation (because closeness in parameters can be shown to imply closeness in total variation distance). In that sense, parameter estimation is a harder problem.

Before we proceed to describe prior algorithmic work, we provide basic background on the sample complexity of the problem. We start by noting that density estimation for r-mixtures of linear classifiers on $\mathbb{R}^n$ is information-theoretically solvable using $\text{poly}(n, r)$ samples (with the optimal bound being $\tilde{\Theta}(nr/\epsilon^2)$ for total variation error ϵ) without any assumptions on the components. In contrast, for parameter estimation to be information-theoretically solvable with polynomial sample complexity, some further assumptions are needed. The typical assumption involves some kind of pairwise separation between the component vectors and a lower bound on the mixing weights. Let $\Delta \stackrel{\text{def}}{=} \min_{i \neq j} \{\|\mathbf{v}_i - \mathbf{v}_j\|_2, \|\mathbf{v}_i + \mathbf{v}_j\|_2\} > 0$ be the pairwise separation between the components and $w_{\min}$ be the minimum mixing weight. Under these assumptions, the parameter estimation problem is solvable using $\text{poly}(n, r, 1/\Delta, 1/w_{\min}, 1/\epsilon)$ samples to achieve parameter error of ϵ. In both cases, the term "information-theoretically solvable" is used to mean that a *sample-efficient* algorithm exists, without any constraints on computational efficiency. The main question addressed in this paper is to what extent a *computationally efficient* learning algorithm exists.

On the algorithmic front, [CDV22] gave parameter estimation algorithms with provable guarantees under a Δ-separation assumption. Specifically, [CDV22] provided two algorithms with different complexity guarantees. Their first algorithm has sample and computational complexity $\text{poly}(n^{O(\log(r)/\Delta^2)}, 1/w_{\min}, 1/\epsilon)$, while their second algorithm has complexity $\text{poly}((n/\Delta)^r, 1/w_{\min}, 1/\epsilon)$. Here we focus on settings where the number of components r is large and cannot be viewed as a constant. For the sake of intuition, it is instructive to simplify these upper bounds for the regime of uniform mixtures (corresponding to the special case that $w_\ell = 1/r$ for all $\ell \in [r]$) and ϵ is not too small. For this regime, the complexity upper bound achieved in [CDV22] is $\min\{n^{O(\log(r)/\Delta^2)}, (n/\Delta)^{O(r)}\}$. Concretely, if the separation Δ is $\Omega(1)$ or even $1/\text{poly}\log(r)$, the first term yields a quasi-polynomial

upper bound. On the other hand, for $\Delta = O(1/r^c)$, for a constant $c > 0$, the resulting upper bound is $n^{\text{poly}(r)}$. In both regimes, we observe a *super-polynomial* gap between the information-theoretic sample complexity — which is $\text{poly}(n, r, \Delta)$ — and the sample complexity of the [CDV22] algorithms. It is thus natural to ask if this gap is inherent.

> *What is the complexity of learning mixtures of linear classifiers?*
> *Is there an algorithm with significantly better sample-time tradeoff?*

We study these questions in a well-studied restricted model of computation, known as the Statistical Query (SQ) model [Kea98] (and, via [BBH+20], also for low-degree polynomial tests). *Our main result is that in both of these models the complexity of the above-mentioned algorithms is essentially best possible.* Along the way, we establish new results on the existence of spherical designs that may be of independent interest.

1.1 Our Results

In order to formally state our lower bounds, we require basic background on SQ algorithms.

Basics on SQ Model SQ algorithms are a class of algorithms that, instead of having direct access to samples, are allowed to query expectations of bounded functions of the distribution. Formally, an SQ algorithm has access to the following standard oracle.

Definition 1.1 (STAT Oracle). Let D be a distribution on $\mathbb{R}^n$. A *Statistical Query* is a bounded function $f : \mathbb{R}^n \to [-1, 1]$. For $\tau > 0$, the $\text{STAT}(\tau)$ oracle responds to the query f with a value v such that $|v - \mathbf{E}_{\mathbf{x} \sim D}[f(\mathbf{x})]| \leq \tau$. A *Statistical Query (SQ) algorithm* is an algorithm whose objective is to learn an unknown distribution D by making adaptive calls to the $\text{STAT}(\tau)$ oracle.

The SQ model was introduced in [Kea98]. Subsequently, the model has been extensively studied in a range of contexts [Fel16]). The class of SQ algorithms is broad and captures a range of known supervised learning algorithms. More broadly, several known algorithmic techniques in machine learning are known to be implementable using SQs [FGR+17, FGV17].

Our main result is a near-optimal SQ lower bound for mixtures of linear classifiers that applies even for the uniform case. Specifically, we establish the following:

Theorem 1.2 (Main Result: SQ Lower Bound for Uniform Mixtures). *Let $\epsilon \leq c\Delta/r$ for some universal constant $c > 0$ sufficiently small. Then, any SQ algorithm that learns a uniform mixture of linear classifiers with directions $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{S}^{n-1}$ satisfying $\min_{i \neq j}\{\|\mathbf{v}_i - \mathbf{v}_j\|_2, \|\mathbf{v}_i + \mathbf{v}_j\|_2\} \geq \Omega(\Delta)$ for some $r^{-1/10} \leq \Delta < 1$, within error ϵ in total variation distance must either use queries of tolerance $n^{-\text{poly}(1/\Delta)\log r}$, or make at least $2^{n^{\Omega(1)}}$ queries.*

Informally speaking, Theorem 1.2 shows that no SQ algorithm can perform density estimation for uniform mixtures of linear classifiers to small accuracy with a sub-exponential in $n^{\Omega(1)}$ many queries, unless using queries of very small tolerance – that would require at least $n^{\text{poly}(1/\Delta)\log r}$ samples to simulate. This result can be viewed as a near-optimal information-computation tradeoff for the problem, within the class of SQ algorithms. In more detail, for $\Delta = \Omega(1)$, we obtain a quasi-polynomial SQ lower bound of $n^{\Omega(\log r)}$; while for $\Delta = 1/r^c$, for some constant $0 < c < 1/2$, we obtain an SQ lower bound of $n^{\text{poly}(r)}$. In both cases, our SQ lower bounds qualitatively match the previously known algorithmic guarantees [CDV22] (that are efficiently implementable in the SQ model).

2

A conceptual implication of Theorem 1.2 is that the uniform (i.e., equal mixing weights) case is essentially as hard as the general case for density estimation of these mixtures. In contrast, for related mixture models, specifically for mixtures of Gaussians, there is recent evidence that restricting the weights may make the problem computationally easier [BS21].

Remark 1.3. We note that the condition $\Delta \geq r^{-c}$, for some constant $0 < c < 1$, is necessary in the statement of Theorem 1.2 for the following reason: the algorithmic result of [CDV22] has sample and computational complexity $\min\{n^{O(\log(r)/\Delta^2)}, (n/\Delta)^{O(r)}\}$, which will be $(n/\Delta)^{O(r)} \ll n^{\text{poly}(1/\Delta)\log r}$ if Δ is a sufficiently small inverse polynomial in r.

Remark 1.4. Our SQ lower bound result has immediate implications to another well-studied restricted computational model — that of low-degree polynomial tests [HS17, HKP+17, Hop18]. [BBH+20] established that (under certain assumptions) an SQ lower bound also implies a qualitatively similar lower bound in the low-degree model. This connection can be used as a black-box to show a similar lower bound for low-degree polynomials.

The key technical ingredient required for our SQ lower bound is a theorem establishing the existence of spherical designs with appropriate properties. The definition of a spherical design follows.

Definition 1.5 (Spherical Design). Let t be an odd integer. A set of points $\mathbf{x}_1, \ldots, \mathbf{x}_r \in \mathbb{S}^{n-1}$ is called a spherical t-design if $\mathbf{E}[p(\mathbf{x})] = \frac{1}{r}\sum_{i=1}^{r} p(\mathbf{x}_i)$ holds for every homogeneous n-variate polynomial p of degree t, where the expectation is taken over the uniform distribution on the unit sphere $\mathbb{S}^{n-1}$.

We note that this definition differs slightly from the usual definition of spherical design, which requires that the equation in Definition 1.5 holds for all polynomials p of degree *at most* t. However, by multiplying by powers of $\|\mathbf{x}\|_2^2$, we note that our definition implies that the equation holds for every odd polynomial of degree at most t, and it is sufficient for the requirement of establishing our SQ lower bound.

Spherical designs have been extensively studied in combinatorial design theory and a number of efficient constructions are known, see, e.g., [Ban79, DGS91, GP11, BRV13, Kan15, Wom18]. However, none of the known constructions seem to be compatible with our separation assumptions.

We establish the following result that may be of independent interest in this branch of mathematics.

Theorem 1.6 (Efficient Spherical Design). *Let t be an odd integer and $r \geq \binom{n+2t-1}{n-1}^5$. Let $\mathbf{y}_1, \ldots, \mathbf{y}_r$ be uniform random vectors over $\mathbb{S}^{n-1}$. Then, with probability at least $99/100$, there exist unit vectors $\mathbf{z}_i \in \mathbb{S}^{n-1}$ very close to $\mathbf{y}_i$ such that $(\mathbf{z}_1, \ldots, \mathbf{z}_r)$ form a spherical t-design.*

We mention here that the optimal sample complexity for the existence of spherical t-design (for even t) is $r = \Theta(\binom{n+t}{n})$ and our result matches this within a polynomial factor. Theorem 1.6 is essential for our construction of moment-matching. Roughly speaking, the resulting mixture matches moments in the sense we require if and only if the weight vectors $(\mathbf{v}_1, \ldots, \mathbf{v}_r)$ form a *spherical design.* since this leads to the separation of the hidden weight vectors in the mixture of linear classifiers. In order to guarantee the pairwise separation of the weight vectors $\mathbf{v}_i$'s, we note that with high probability the $\mathbf{y}_i$'s are pairwise separated and therefore if $\mathbf{z}_i$ is sufficiently close to $\mathbf{y}_i$ they will be too.

1.2 Technical Overview

Our starting point is the SQ lower bound technique of [DKS17] and its generalization in [DKPZ21]. In particular, our overall strategy is to find a mixture of homogeneous halfspaces in some "small"

number m of dimensions that matches its first k moments with the distribution on $(\mathbf{x}, y)$, where y is independent of $\mathbf{x}$. By appropriately embedding this low-dimensional construction into $\mathbb{R}^n$ via a random projection of $\mathbb{R}^n \to \mathbb{R}^m$, the SQ lower bound machinery of [DKPZ21] can be shown to imply that learning the projection requires either $2^{n^{\Omega(1)}}$ queries or some query of accuracy $n^{-\Omega(k)}$. This generic step reduces our problem to finding an appropriate m-dimensional construction. Such a construction is the main technical contribution of this work.

Note that specifying a mixture of halfspaces is equivalent to specifying a probability distribution with support consisting of at most r unit vectors and specifying the orthogonal vectors of the halfspaces. It is not hard to see that the resulting mixture matches moments in the sense we require if and only if these vectors form a *spherical k-design* [DGS91]: namely, that for any odd polynomial p of degree less than k that the average value of p over our distribution is the same as the average value of p over the unit sphere (namely, equal to zero). The bulk of our problem now reduces to finding constructions of weighted designs, where: (1) The support size of the design is relatively small. (2) The points in the design are pairwise separated. (3) For equally weighted mixtures, we require that the weight of each vector in the design is uniform.

In $m = 2$ two dimensions, there is a relatively simple explicit construction that can be given. In particular, if we take k evenly spaced points over the unit circle for some odd k, this matches the first $k-1$ moments and has separation $\Omega(1/k)$. (This is similar to an SQ construction in [DKKZ20] for a different setting.) Unfortunately, we cannot obtain better separation in two dimensions, since any k unit vectors in two dimensions will necessarily have some pair separated by at most $O(1/k)$. Therefore, if we want constructions with greater separation, we will need to pick larger values of m. Indeed, we show (Proposition 3.2) that a *random* collection of approximately $m^{O(k)}$ points can form the support of an appropriate design. This can be proved by applying linear programming duality. Unfortunately, this argument does not allow us to control the mixing weights. Specifically, it may be the case that the minimum weight is exponentially small which seems unnatural in practice.

The case of equal weights requires a significantly more sophisticated argument. In this case, we need to find a spherical design with uniform weight. Indeed, merely selecting a random support will no longer work. Although there is an extensive literature on finding efficient spherical designs [Ban79, DGS91, GP11, BRV13, Kan15, Wom18], none of the known constructions seem to be compatible with our separation assumptions. In particular, [BRV13] proves that for each $r \geq c_m k^m$, there exists a spherical k-design, where c_m is a constant depending only on the dimension m. Although it is plausible that their construction can be adapted to satisfy our separation requirement, their sample complexity has a potentially bad dependence on the dimension m. On the other hand, the work of [Kan15] achieves the optimal sample complexity r up to polynomials for all m and k. Unfortunately, the construction in [Kan15] definitely does not satisfy our separation requirement. We note that the sample complexity of $r = \Theta\left(\binom{m+k}{k}\right)$ should be optimal, and our results are polynomial in this bound.

In conclusion, to establish our SQ lower bound, we need to prove a new technical result. Our starting point to that end will be the work of [BRV13]. The basic plan of our approach will be to select random unit vectors $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_r \in \mathbb{R}^n$, and then — applying topological techniques — to show that there is some small perturbation set $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_r \in \mathbb{R}^n$ that is a spherical design. In particular, we will find a continuous function F mapping degree-k odd polynomials to sets of r unit vectors such that, for all unit-norm polynomials p, if $F(p) = (\mathbf{z}_1, \ldots, \mathbf{z}_r)$ then $\sum_{i=1}^r p(\mathbf{z}_i) > 0$. Given this statement, a standard fixed point theorem [CC06] will imply that our design can be found as $F(q)$ for some q with norm less than one. To construct the mapping F, we start with the points $\mathbf{y}_1, \ldots, \mathbf{y}_r$ and perturb each $\mathbf{y}_i$ in the direction of $\nabla_o p(\mathbf{y}_i)$ in order to try to increase the average value of p, where $\nabla_o p(\mathbf{y}_i)$ is the component of $\nabla p(\mathbf{y}_i)$ orthogonal to the direction $\mathbf{y}_i$.

Intuitively, this construction should work because with high probability the average value of $p(\mathbf{y}_i)$ is already small (since the empirical average should approximate the true average), the average gradient of $p(\mathbf{y}_i)$ is not too small, and the contributions to $p(\mathbf{z}_i)$ coming from higher-order terms will not be large as long as $\mathbf{z}_i$ is sufficiently close to $\mathbf{y}_i$. These facts can all be made precise with some careful analysis of spherical harmonics (Lemma 4.8).

Finally, we need to show that the hardness of learning the appropriate projection implies hardness of learning the mixture. By standard facts, it suffices to show that mixtures of linear classifiers are far from the distribution $(\mathbf{x}, y)$, where the uniform label y is independent of $\mathbf{x}$, in total variation distance. To show this, we prove that the total variation distance between any projection and the distribution where y is independent of $\mathbf{x}$ is at least $\Omega(\Delta/r)$ (Lemma 3.8).

2 Preliminaries

Notation For $n \in \mathbb{Z}_+$, we denote $[n] \overset{\text{def}}{=} \{1, \ldots, n\}$. For two distributions p, q over a probability space Ω, let $d_{\text{TV}}(p, q) = \sup_{S \subseteq \Omega} |p(S) - q(S)|$ denote the total variation distance between p and q. In this article, we typically use small letters to denote random variables and vectors. For a real random variable x, we use $\mathbf{E}[x]$ to denote the expectation. We use $\mathbf{Pr}[\mathcal{E}]$ and $\mathbb{I}[\mathcal{E}]$ for the probability and the indicator of event $\mathcal{E}$. Let $\mathcal{N}_n$ denote the standard n-dimensional Gaussian distribution and $\mathcal{N}$ denote the standard univariate Gaussian distribution. Let $\mathbb{S}^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 = 1\}$ denote the n-dimensional unit sphere. For a subset $S \subseteq \mathbb{R}^n$, we will use $\mathcal{U}(S)$ to denote the uniform distribution over S. We will use small boldface letters for vectors and capital boldface letters for matrices. Let $\|\mathbf{x}\|_2$ be the ℓ^2-norm of the vector $\mathbf{x} \in \mathbb{R}^n$. For vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, we use $\langle \mathbf{u}, \mathbf{v} \rangle$ to denote their inner product. We denote by $\mathcal{L}^2(\mathbb{R}^n, \mathcal{N}_n)$ the function space of all functions $f : \mathbb{R}^n \to \mathbb{R}$ such that $\mathbf{E}_{\mathbf{z} \in \mathcal{N}_n}[f^2(\mathbf{z})] < \infty$. The usual inner product for this space is $\mathbf{E}_{\mathbf{z} \in \mathcal{N}_n}[f(\mathbf{z})g(\mathbf{z})]$.

For a matrix $\mathbf{P} \in \mathbb{R}^{m \times n}$, we denote $\|\mathbf{P}\|_2, \|\mathbf{P}\|_F$ to be its spectral norm and Frobenius norm respectively. For a tensor $A \in (\mathbb{R}^n)^{\otimes k}$, let $\|A\|_2$ denote its spectral norm and $\|A\|_F$ denote its Frobenius norm. We will use $A_{i_1, \ldots, i_k}$ to denote the coordinate of the k-tensor A indexed by the k-tuple $(i_1, \ldots, i_k)$. The inner product between k-tensors is defined by thinking of tensors as vectors with n^k coordinates.

We use the framework of Statistical Query (SQ) algorithms for problems over distributions [FGR$^+$17] and require the following standard definition.

Definition 2.1 (Decision/Testing Problem over Distributions). Let D be a distribution and $\mathcal{D}$ be a family of distributions over $\mathbb{R}^n$. We denote by $\mathcal{B}(\mathcal{D}, D)$ the decision (or hypothesis testing) problem in which the input distribution D' is promised to satisfy either (a) $D' = D$ or (b) $D' \in \mathcal{D}$, and the goal of the algorithm is to distinguish between these two cases.

Probabilistic Facts We recall the definition of VC dimension for a set system.

Definition 2.2 (VC-Dimension). For a class $\mathcal{C}$ of boolean functions on a set $\mathcal{X}$, the *VC-dimension* of $\mathcal{C}$ is the largest integer d such that there exist d points $x_1, \ldots, x_d \in \mathcal{X}$ such that for any boolean function $g : \{x_1, \ldots, x_d\} \to \{\pm 1\}$, there exists an $f \in \mathcal{C}$ satisfying $f(x_i) = g(x_i), 1 \leq i \leq d$.

We will use the following probabilistic inequality.

Lemma 2.3 (VC inequality, see, e.g., [Ver18]). *Let $\mathcal{C}$ be a class of boolean functions on $\mathcal{X}$ with VC-dimension d, and let X be a distribution on $\mathcal{X}$. Let $\epsilon > 0$ and let N be an integer at least a*

5

sufficiently large constant multiple of d/ϵ^2. Then, if $X_1, X_2, \ldots, X_N$ are i.i.d. samples from X, we have that:

$$\mathbf{Pr}\left[\sup_{f \in \mathcal{C}} \left| \frac{\sum_{j=1}^{N} f(X_j)}{N} - \mathbf{E}[f(X)] \right| > \epsilon\right] = \exp(-\Omega(N\epsilon^2)).$$

To apply the VC inequality in our context, we additionally need the following fact which gives the VC-dimension of the class of bounded-degree polynomial threshold functions (PTFs).

Fact 2.4. *Let $\mathcal{C}_k$ denote the class of degree-k polynomial threshold functions (PTFs) on $\mathbb{R}^m$, namely the collection of functions of the form $f(\mathbf{x}) = \mathrm{sign}(p(\mathbf{x}))$ for some degree at most k real polynomial p. Then, the VC-dimension of $\mathcal{C}_k$ is $\mathrm{VC}(\mathcal{C}_k) = \binom{m+k}{k}$.*

3 Warmup: SQ Lower Bounds for General Weight Mixtures

SQ Lower Bound Machinery We start by defining the family of distributions that we will use to prove our SQ hardness result.

Definition 3.1 ([DKPZ21])**.** *Given a function $g : \mathbb{R}^m \to [-1, +1]$, we define $\mathcal{D}_g$ to be the class of distributions over $\mathbb{R}^n \times \{\pm 1\}$ of the form $(\mathbf{x}, y)$ such that $\mathbf{x} \sim \mathcal{N}_n$ and $\mathbf{E}[y \mid \mathbf{x} = \mathbf{z}] = g(\mathbf{U}\mathbf{z})$, where $\mathbf{U} \in \mathbb{R}^{m \times n}$ with $\mathbf{U}\mathbf{U}^\intercal = \mathbf{I}_m$.*

The following proposition states that if g has zero low-degree moments, then distinguishing $\mathcal{D}_g$ from the distribution $(\mathbf{x}, y)$ with $\mathbf{x} \sim \mathcal{N}_n, y \sim \mathcal{U}(\{\pm 1\})$ is hard in the SQ model.

Proposition 3.2 ([DKPZ21])**.** *Let $g : \mathbb{R}^m \to [-1, 1]$ be such that $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_m}[g(\mathbf{x})p(\mathbf{x})] = 0$, for every polynomial $p : \mathbb{R}^m \to \mathbb{R}$ of degree less than k, and $\mathcal{D}_g$ be the class of distributions from Definition 3.1. Then, if $m \leq n^a$, for some constant $a < 1/2$, any SQ algorithm that solves the decision problem $\mathcal{B}(\mathcal{D}_g, \mathcal{N}_n \times \mathcal{U}(\{\pm 1\}))$ must either use queries of tolerance $n^{-\Omega(k)}$, or make at least $2^{n^{\Omega(1)}}$ queries.*

We will apply Proposition 3.2 to establish our SQ lower bound for learning mixtures of linear classifiers. The main technical contribution required to achieve this is the construction of a class of distributions $\mathcal{D}_g$, where each element in $\mathcal{D}_g$ represents a distribution of mixture of linear classifiers. In particular, we will carefully choose some appropriate unit vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{R}^m$ and non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^{r} w_\ell = 1$, such that $\mathbf{v}_\ell$ are pairwise separated by $\Delta > 0$.

Let $g(\mathbf{z}) = \sum_{\ell=1}^{r} w_\ell \mathrm{sign}(\mathbf{v}_\ell^\intercal \mathbf{z}), \mathbf{z} \in \mathbb{R}^m$. For an arbitrary matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$ with $\mathbf{U}\mathbf{U}^\intercal = \mathbf{I}_m$, we denote by $D_\mathbf{U}$ the instance of mixture of linear classifiers with weight vectors $\mathbf{U}^\intercal \mathbf{v}_1, \ldots, \mathbf{U}^\intercal \mathbf{v}_r$ and weights $w_1, \ldots, w_r$. In this way, we have that

$$\mathbf{E}_{(\mathbf{x}, y) \sim D_\mathbf{U}}[y \mid \mathbf{x} = \mathbf{z}] = \sum_{\ell=1}^{r} w_\ell \mathrm{sign}((\mathbf{U}^\intercal \mathbf{v}_\ell)^\intercal \mathbf{z}) = \sum_{\ell=1}^{r} w_\ell \mathrm{sign}(\mathbf{v}_\ell^\intercal \mathbf{U}\mathbf{z}) = g(\mathbf{U}\mathbf{z}), \mathbf{z} \in \mathbb{R}^n.$$

3.1 Low-degree Moment Matching

The following proposition shows that there exist unit vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r$ and non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^{r} w_\ell = 1$ such that $\mathbf{v}_\ell$ are pairwise separated by some parameter $\Delta > 0$ and the low-degree moments of g vanish. Note that since g is odd (as the sign function is odd), we only require $\mathbf{E}_{\mathbf{z} \in \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ for every odd polynomial p of degree less than k.

Proposition 3.3. *Let* $m = \frac{c \log r}{\log(1/\Delta)}$, *where* $r^{-1/10} \leq \Delta < 1$ *and* $1.99 \leq c \leq 2$ *is a universal constant. Let* k *be a positive integer such that* $r \geq C\binom{m+k}{k}$ *for some constant* $C > 0$ *sufficiently large. There exist vectors* $\mathbf{v}_1, \ldots, \mathbf{v}_r$ *over the unit sphere* $\mathbb{S}^{m-1}$ *and non-negative weights* $w_1, \ldots, w_r$ *with* $\sum_{\ell=1}^r w_\ell = 1$, *such that*

- $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ *holds for every odd polynomial* p *of degree less than* k, *where* $g(\mathbf{z}) = \sum_{\ell=1}^r w_\ell \text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})$.

- $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2 \geq \Omega(\Delta), \forall 1 \leq i < j \leq r$.

The proof proceeds as follows: We uniformly sample unit vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{S}^{m-1}$. We will prove by the probabilistic method that both statements in Proposition 3.3 hold with high probability. We begin by showing that $\mathbf{Pr}[\min_{1 \leq i < j \leq r}\{\|\mathbf{v}_i - \mathbf{v}_j\|_2, \|\mathbf{v}_i + \mathbf{v}_j\|_2\} < O(\Delta)]$ is small. By symmetry, it suffices to bound $\mathbf{Pr}[\min_{1 \leq i < j \leq r} \|\mathbf{v}_i - \mathbf{v}_j\|_2 < O(\Delta)]$. Let $\Theta = \min_{1 \leq i < j \leq r} \arccos\langle \mathbf{v}_i, \mathbf{v}_j \rangle$. We require the following facts:

Fact 3.4 (Proposition 3.5 in [BRS+18]). $\mathbf{Pr}\left[\Theta \geq \gamma r^{-\frac{2}{m-1}}\right] \geq 1 - \frac{\kappa_{m-1}}{2}\gamma^{m-1}$, *where* $\kappa_m = \frac{\Gamma((m+1)/2)}{m\sqrt{\pi}\Gamma(m/2)} \in \left[\frac{1}{m}\sqrt{\frac{m-1}{2\pi}}, \frac{1}{m}\sqrt{\frac{m+1}{2\pi}}\right]$.

Applying Fact 3.4 by taking $\gamma = \left(\frac{1}{50\kappa_{m-1}}\right)^{\frac{1}{m-1}}$ and $r = \frac{1}{\sqrt{50\kappa_{m-1}\Delta^{m-1}}}$ yields that $\mathbf{Pr}[\Theta \geq \Delta] \geq 99/100$. Given some fixed random vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{S}^{m-1}$, we will prove that with high probability there exist non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ holds for every odd polynomial p of degree less than k, where $g(\mathbf{z}) = \sum_{\ell=1}^r w_\ell \text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})$. Noting that $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = \sum_{\ell=1}^r w_\ell \mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[p(\mathbf{z})\text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})]$, it suffices to bound the probability that there exist non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\sum_{\ell=1}^r w_\ell \mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[p(\mathbf{z})\text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})] = 0$ holds for every odd polynomial p of degree less than k. Our main technical lemma is the following:

Lemma 3.5. *Let* $p : \mathbb{R}^m \to \mathbb{R}$ *be a polynomial of degree less than* k *and* $f \in \mathcal{L}^2(\mathbb{R}, \mathcal{N})$. *Let* $\mathbf{v} \in \mathbb{S}^{m-1}$ *be a unit vector. We have that* $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[p(\mathbf{z})f(\mathbf{v}^\intercal \mathbf{z})]$ *is a polynomial in* $\mathbf{v}$ *of degree less than* k.

The proof proceeds by analyzing p and f as Hermite expansions. For completeness, we defer the proof of Lemma 3.5 to Appendix B. From Lemma 3.5, we can write $\sum_{\ell=1}^r w_\ell \mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[p(\mathbf{z})\text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})] = \sum_{\ell=1}^r w_\ell q(\mathbf{v}_\ell)$ for some odd polynomial q of degree less than k (since sign is odd). In this way, it suffices to show that with high probability there exist non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\sum_{\ell=1}^r w_\ell q(\mathbf{v}_\ell) = 0$ for all odd polynomials q of degree less than k. To prove this, we leverage the following lemma:

Lemma 3.6. *The following two statements are equivalent.*

1. *There exist non-negative weights* $w_1, \ldots, w_r$ *with* $\sum_{\ell=1}^r w_\ell = 1$ *such that* $\sum_{\ell=1}^r w_\ell q(\mathbf{v}_\ell) = 0$ *for all odd polynomials* q *of degree less than* k.

2. *There does not exist any odd polynomial* q *of degree less than* k *such that* $q(\mathbf{v}_\ell) > 0, 1 \leq \ell \leq r$.

The proof of Lemma 3.6 follows via a careful application of LP duality; see Appendix B.

Proof of Proposition 3.3. Let $\Theta = \min_{1 \leq i < j \leq r} \arccos\langle \mathbf{v}_i, \mathbf{v}_j \rangle$. Applying Fact 3.4 by taking $\gamma = \left(\frac{1}{50\kappa_{m-1}}\right)^{\frac{1}{m-1}}$ yields $\mathbf{Pr}[\Theta \geq \Delta] \geq \mathbf{Pr}\left[\Theta \geq \gamma r^{-\frac{2}{m-1}}\right] \geq 99/100$ (since our choice of m will imply

7

that $r = 1/\sqrt{50\kappa_{m-1}\Delta^{m-1}}$). This will imply that

$$\mathbf{Pr}\Big[\min_{1\leq i<j\leq r}\{\|\mathbf{v}_i-\mathbf{v}_j\|_2, \|\mathbf{v}_i+\mathbf{v}_j\|_2\}\geq\Omega(\Delta)\Big]\geq 99/100\,.$$

We then show that there exist non-negative weights $w_1,\ldots,w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\mathbf{E}_{\mathbf{z}\sim\mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ holds for every odd polynomial p with degree less than k, where $g(\mathbf{z}) = \sum_{\ell=1}^r w_\ell\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{z})$. By Lemma 3.5, it suffices to show that there exist non-negative weights $w_1,\ldots,w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\sum_{\ell=1}^r w_\ell q(\mathbf{v}_\ell) = 0$ for all odd polynomials q of degree less than k. By Fact 2.4, we have that $r \geq C\binom{m+k}{k} = C\mathrm{VC}(\mathcal{C}_k)$ for some sufficiently large constant $C > 0$. Note that for any odd polynomial q in m variables of degree less than k, we have that $\mathbf{E}[\mathrm{sign}(q(\mathbf{v}))] = 0$. Therefore, by the VC-inequality (Lemma 2.3), we have that

$$\mathbf{Pr}[\exists \text{ odd polynomial } q \text{ of degree less than } k \text{ such that } q(\mathbf{v}_\ell) > 0, 1\leq\ell\leq r]$$

$$\leq\mathbf{Pr}\left[\sup_{f\in\mathcal{C}_{k-1}}\frac{\sum_{\ell=1}^r f(\mathbf{v}_\ell)}{r} = 1\right]\leq\mathbf{Pr}\left[\sup_{f\in\mathcal{C}_k}\left|\frac{\sum_{\ell=1}^r f(\mathbf{v}_\ell)}{r} - \mathbf{E}[f(\mathbf{v})]\right|\geq 1\right]$$

$$\leq\exp(-\Omega(r))\leq 1/100.$$

Finally, applying Lemma 3.6 completes our proof. $\qquad\square$

3.2 Proof of SQ Lower Bound for General Weights

Let k be a positive integer such that $r \geq C\binom{m+k}{k}$ for some constant $C > 0$ sufficiently large. Let $m = \frac{c\log r}{\log(1/\Delta)}$, where $r^{-1/10}\leq\Delta < 1$ and $1.99\leq c\leq 2$. By Proposition 3.3, there exist unit vectors $\mathbf{v}_1,\ldots,\mathbf{v}_r\in\mathbb{R}^m$ such that

- There exist non-negative weights $w_1,\ldots,w_r$ with $\sum_{\ell=1}^r w_\ell = 1$ such that $\mathbf{E}_{\mathbf{z}\sim\mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ holds for every odd polynomial p with degree less than k, where $g(\mathbf{z}) = \sum_{\ell=1}^r w_\ell\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{z})$.

- $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2\geq\Omega(\Delta), \forall 1\leq i < j\leq r$ for some $\Delta > 0$.

For an arbitrary matrix $\mathbf{U}\in\mathbb{R}^{m\times n}$ with $\mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m$, denote by $D_\mathbf{U}$ the instance of mixture of linear classifiers with weight vectors $\mathbf{U}^\mathsf{T}\mathbf{v}_1,\ldots,\mathbf{U}^\mathsf{T}\mathbf{v}_r$ and weights $w_1,\ldots,w_r$. Let $\mathcal{D}_g = \{\mathcal{D}_\mathbf{U}\mid\mathbf{U}\in\mathbb{R}^{m\times n}, \mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m\}$. We have that $\|\mathbf{U}^\mathsf{T}\mathbf{v}_i\pm\mathbf{U}^\mathsf{T}\mathbf{v}_j\|_2 = \|\mathbf{v}_i\pm\mathbf{v}_j\|_2\geq\Omega(\Delta), 1\leq i < j\leq r$, and $\mathbf{E}_{(\mathbf{x},y)\sim D_\mathbf{U}}[y\mid\mathbf{x} = \mathbf{z}] = \sum_{\ell=1}^r w_\ell\mathrm{sign}((\mathbf{U}^\mathsf{T}\mathbf{v}_\ell)^\mathsf{T}\mathbf{z}) = \sum_{\ell=1}^r w_\ell\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{z}) = g(\mathbf{U}\mathbf{z}), \mathbf{z}\in\mathbb{R}^n$. By Proposition 3.2, any SQ algorithm for the decision problem $\mathcal{B}(\mathcal{D}_g,\mathcal{N}_n\times\mathcal{U}(\{\pm 1\}))$ must either use queries of tolerance $n^{-\Omega(k)}$, or make at least $2^{n^{\Omega(1)}}$ queries. The last step is to reduce the decision problem $\mathcal{B}(\mathcal{D}_g,\mathcal{N}_n\times\mathcal{U}(\{\pm 1\}))$ to the problem of learning mixture of linear classifiers.

Claim 3.7 (see, e.g., Lemma 8.5 in [DK23]). *Suppose there exists an SQ algorithm to learn an unknown distribution in a family $\mathcal{D}$ to total variation distance ϵ using at most N statistical queries of tolerance τ. Suppose furthermore that for each $D'\in\mathcal{D}$ we have that $d_{\mathrm{TV}}(D, D') > 2(\tau+\epsilon)$. Then there exists an SQ algorithm that solves the testing problem $\mathcal{B}(\mathcal{D}, D)$ using at most $n + 1$ queries of tolerance τ.*

To apply Claim 3.7, we need to show that the distribution $D_\mathbf{U}$ in the class $\mathcal{D}_g$ is sufficiently far from the null hypothesis $D_0 = \mathcal{N}_n\times\mathcal{U}(\{\pm 1\})$ in total variation distance.

Lemma 3.8. *Let $\mathbf{U}\in\mathbb{R}^{m\times n}$ with $\mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m$. We have that $d_{\mathrm{TV}}(D_\mathbf{U}, D_0)\geq\Omega(\Delta/r)$.*

We briefly sketch the proof idea and defer the proof details to Appendix B. We consider $w_\ell \text{sign}(\mathbf{v}_\ell^\intercal \mathbf{x})$ a halfspace of heaviest weight in $D_\mathbf{U}$, and pick points $\mathbf{z}$ and $\mathbf{z}'$ close to the defining hyperplane that are mirrors of each other over it. We note that under D_0 the expectations of y conditioned on $\mathbf{x}$ being $\mathbf{z}$ or $\mathbf{z}'$ are both 0, whereas under $D_\mathbf{U}$ they likely (in particular, unless they are also on opposite sides of another halfspace) differ by at least w_ℓ.

4   SQ Lower Bound for Uniform Mixtures via Spherical Designs

In this section, we prove our SQ lower bound for mixture of linear classifiers with uniform weights, thereby establishing Theorem 1.2. Our basic lower bound technique is essentially the same as in the previous section, but we need to construct a spherical design with *uniform weight*.

Proposition 4.1. *Let $m = \frac{c \log r}{\log(1/\Delta)}$, where $r^{-1/10} \leq \Delta < 1$ and $1.99 \leq c \leq 2$ is a constant. Let k be an odd integer such that $r \geq \binom{m+2k-1}{2k}^5$. There exist vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r$ on $\mathbb{S}^{m-1}$ such that:*

- $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ *holds for every odd polynomial p with degree less than k, where $g(\mathbf{z}) = \frac{1}{r}\sum_{\ell=1}^r \text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})$.*

- $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2 \geq \Omega(\Delta) - O\big(1/\binom{m+2k-1}{2k}\big), \forall 1 \leq i < j \leq r$.

By Lemma 3.5, we can write $\sum_{\ell=1}^r \mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}\big[p(\mathbf{z})\text{sign}(\mathbf{v}_\ell^\intercal \mathbf{z})\big] = \sum_{\ell=1}^r q(\mathbf{v}_\ell)$ for some odd polynomial q of degree less than k. Therefore, it suffices to show that with high probability there exist unit vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r$ such that $\sum_{\ell=1}^r q(\mathbf{v}_\ell) = 0$ holds for every odd polynomial q of degree less than k. To achieve this, we will leverage some techniques from [BRV13].

4.1   Spherical Design Construction

Notation   We start by introducing some additional notation we will use throughout this section. Let $\mathcal{P}_t^d$ denote the set of homogeneous polynomials in d variables of some odd degree t. For any $p, q \in \mathcal{P}_t^d$, we consider the inner product $\langle p, q \rangle = \mathbf{E}_{\mathbf{x} \sim \mathcal{U}(\mathbb{S}^{d-1})}[p(\mathbf{x})q(\mathbf{x})]$. For any $p \in \mathcal{P}_t^d$, we define $\|p\|_2^2 = \langle p, p \rangle$. We denote by $N_{t,d} = \binom{t+d-1}{d-1}$ to be the dimension of $\mathcal{P}_t^d$ and $\Omega_t^d = \{p \in \mathcal{P}_t^d \mid \|p\|_2 \leq 1\}$. Let $\partial\Omega_t^d$ denote the boundary of Ω_t^d, i.e., $\partial\Omega_t^d = \{p \in \mathcal{P}_t^d \mid \|p\|_2 = 1\}$. In the remaining part of this section, we will assume that the underlying distribution (over the expectation) is $\mathcal{U}(\mathbb{S}^{d-1})$.

We will require the following facts.

Fact 4.2 (see, e.g., Lemma 28 in [Kan15])**.** *For any $p \in \Omega_t^d$, we have that*

$$\sup_{\|\mathbf{x}\|_2=1} |p(\mathbf{x})| \leq \sqrt{N_{t,d}}\sqrt{\mathbf{E}[p(\mathbf{x})^2]} = \sqrt{N_{t,d}}\|p\|_2.$$

Fact 4.3. *Let $t \geq 2$ and $p, q \in \mathcal{P}_t^d$. Then, we have that*

$$t \int_{\|\mathbf{x}\|_2=1} p(\mathbf{x})q(\mathbf{x})d\mathbf{x} = \frac{1}{d+2t-2}\int_{\|\mathbf{x}\|_2=1} \langle \nabla p(\mathbf{x}), \nabla q(\mathbf{x}) \rangle d\mathbf{x} + \frac{1}{d+2t-2}\int_{\|\mathbf{x}\|_2=1} p(\mathbf{x})\nabla^2 q(\mathbf{x})d\mathbf{x}.$$

Proof. Applying the Gaussian Divergence theorem for the function $p(\mathbf{x})\nabla p(\mathbf{x})$ over the unit ball, we have that

$$
\begin{aligned}
t\int_{\|\mathbf{x}\|_2=1} p(\mathbf{x})q(\mathbf{x})d\mathbf{x} &= \int_{\|\mathbf{x}\|_2=1}\langle p(\mathbf{x})\nabla q(\mathbf{x}),\mathbf{x}\rangle d\mathbf{x} = \int_{\|\mathbf{x}\|_2\leq1}\nabla\cdot(p(\mathbf{x})\nabla q(\mathbf{x}))d\mathbf{x}\\
&= \int_{\|\mathbf{x}\|_2\leq1}\langle\nabla p(\mathbf{x}),\nabla q(\mathbf{x})\rangle d\mathbf{x} + \int_{\|\mathbf{x}\|_2\leq1}p(\mathbf{x})\nabla^2 q(\mathbf{x})d\mathbf{x}\\
&= \int_0^1 r^{d-1}dr\int_{\|\mathbf{x}\|_2=1}\langle\nabla p(r\mathbf{x}),\nabla q(r\mathbf{x})\rangle d\mathbf{x} + \int_0^1 r^{d-1}dr\int_{\|\mathbf{x}\|_2=1}p(r\mathbf{x})\nabla^2 q(r\mathbf{x})d\mathbf{x}\\
&= \int_0^1 r^{2t+d-3}dr\int_{\|\mathbf{x}\|_2=1}\langle\nabla p(\mathbf{x}),\nabla q(\mathbf{x})\rangle d\mathbf{x} + \int_0^1 r^{2t+d-3}dr\int_{\|\mathbf{x}\|_2=1}p(\mathbf{x})\nabla^2 q(\mathbf{x})d\mathbf{x}\\
&= \frac{1}{d+2t-2}\int_{\|\mathbf{x}\|_2=1}\langle\nabla p(\mathbf{x}),\nabla q(\mathbf{x})\rangle d\mathbf{x} + \frac{1}{d+2t-2}\int_{\|\mathbf{x}\|_2=1}p(\mathbf{x})\nabla^2 q(\mathbf{x})d\mathbf{x}.
\end{aligned}
$$

This completes the proof. $\qquad\square$

The following lemma provides upper and lower bounds for the expectation of the L^2-norm square of the gradient of any homogeneous polynomial $p\in\Omega_t^d$ over the unit sphere $\mathbb{S}^{d-1}$.

Lemma 4.4. *Let t be an odd positive integer. For any $p\in\mathcal{P}_t^d$, we have that $\mathbf{E}[\|\nabla_o p(\mathbf{x})\|_2^2]\geq (d-1)\|p\|_2^2$ and $\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2]\leq t(d+2t-2)\|p\|_2^2$.*

Proof. By Fact 4.3, we have that

$$
t(d+2t-2)\|p\|_2^2 = \mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2] + \mathbf{E}[p(\mathbf{x})\nabla^2 p(\mathbf{x})].
$$

We bound $\mathbf{E}[p(\mathbf{x})\nabla^2 p(\mathbf{x})]$ as follows. We consider the linear transformations $\mathcal{A}_t:\mathcal{P}_t^d\to\mathcal{P}_{t+2}^d, \mathcal{B}_t:\mathcal{P}_t^d\to\mathcal{P}_{t-2}^d$ as follows: $\mathcal{A}_t(p) = \mathbf{x}^\mathsf{T}\mathbf{x}p(\mathbf{x}), \mathcal{B}_t(p) = \nabla^2 p(\mathbf{x}), p\in\mathcal{P}_t^d$. We first show that for any $t\geq2$, both $\mathcal{A}_{t-2}\mathcal{B}_t$ and $\mathcal{B}_{t+2}\mathcal{A}_t$ are symmetric. For any $p,q\in\mathcal{P}_t^d$, applying Fact 4.3 yields

$$
\begin{aligned}
\langle\mathcal{A}_{t-2}\mathcal{B}_t p, q\rangle = \langle\mathcal{B}_{t+2}\mathcal{A}_t p, q\rangle &= \mathbf{E}[\nabla^2 p(\mathbf{x})q(\mathbf{x})]\\
&= t(d+2t-2)\mathbf{E}[p(\mathbf{x})q(\mathbf{x})] - \mathbf{E}[\langle\nabla p(\mathbf{x}),\nabla q(\mathbf{x})\rangle]\\
&= \mathbf{E}[\nabla^2 q(\mathbf{x})p(\mathbf{x})] = \langle\mathcal{A}_{t-2}\mathcal{B}_t q, p\rangle = \langle\mathcal{B}_{t+2}\mathcal{A}_t q, p\rangle.
\end{aligned}
$$

Therefore, by the eigendecomposition of symmetric linear transformations, we have that $\lambda_1\|p\|_2^2\leq \langle\mathcal{A}_{t-2}\mathcal{B}_t p, p\rangle = \mathbf{E}[p(\mathbf{x})\nabla^2 p(\mathbf{x})]\leq\lambda_t\|p\|_2^2, \forall p\in\Omega_t^d$, where $\lambda_1\leq\cdots\leq\lambda_t$ denote the eigenvalues of $\mathcal{A}_{t-2}\mathcal{B}_t$. In addition, by elementary calculation, for any $p\in\mathcal{P}_t^d$,

$$
\begin{aligned}
\mathcal{B}_{t+2}\mathcal{A}_t p = \nabla^2\mathbf{x}^\mathsf{T}\mathbf{x}p(\mathbf{x}) = \nabla\cdot(2p(\mathbf{x})\mathbf{x}+\mathbf{x}^\mathsf{T}\mathbf{x}\nabla p(\mathbf{x})) &= \sum_{i=1}^d\frac{\partial(2p(\mathbf{x})x_i+\mathbf{x}^\mathsf{T}\mathbf{x}(\nabla p(\mathbf{x}))_i)}{\partial x_i}\\
&= 2dp(\mathbf{x})+4\langle\mathbf{x},\nabla p(\mathbf{x})\rangle+\mathbf{x}^\mathsf{T}\mathbf{x}\nabla^2 p(\mathbf{x}) = (\mathcal{A}_{t-2}\mathcal{B}_t+2d+4t)p.
\end{aligned}
$$

If $\mathcal{A}_{t-2}\mathcal{B}_t$ has an eigenvector p^* corresponding to some eigenvalue λ^*, then $(\mathcal{A}_t\mathcal{B}_{t+2})(\mathcal{A}_t p^*) = \mathcal{A}_t\mathcal{A}_{t-2}\mathcal{B}_t p^* + (2d+4t)\mathcal{A}_t p^* = (\lambda^*+2d+4t)\mathcal{A}_t p^*$, which implies that $\mathcal{A}_t p^*$ is an eigenvector of $\mathcal{A}_t\mathcal{B}_{t+2}$ corresponding to the eigenvalue $\lambda^*+2d+4t$. Note that since $\mathcal{B}_{t+2}$ maps $\mathcal{P}_{t+2}^d$ to $\mathcal{P}_t^d$, we have that $\ker(\mathcal{B}_{t+2})\geq N_{t+2,d}-N_{t,d}$, which implies that $\mathcal{A}_t\mathcal{B}_{t+2}$ has eigenvalue 0 with multiplicity at least $N_{t+2,d}-N_{t,d}$. Therefore, the eigenvalues of $\mathcal{A}_t\mathcal{B}_{t+2}$ are $0<\lambda_1+2d+4t\leq\cdots\leq\lambda_t+2d+4t$, where the multiplicity of eigenvalue 0 is $N_{t+2,d}-N_{t,d}$ and the multiplicity of eigenvalue $\lambda_i+2d+4t$

10

is the same as the multiplicity of eigenvalue λ_i of $\mathcal{A}_{t-2}\mathcal{B}_t$. Therefore, we have that $\lambda_1 = 0$ and $\lambda_t = (t-1)(d+t-1)$, which implies that

$$\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2] = t(d+2t-2)\|p\|_2^2 - \mathbf{E}[p(\mathbf{x})\nabla^2 p(\mathbf{x})] \in [(t^2+d-1)\|p\|_2^2, t(d+2t-2)\|p\|_2^2].$$

Therefore, we have that $\mathbf{E}[\|\nabla_o p(\mathbf{x})\|_2^2] = \mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2 - \langle \mathbf{x}, \nabla p(\mathbf{x}) \rangle^2] = \mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2] - t^2\|p\|_2^2 \geq (d-1)\|p\|_2^2$, completing the proof. $\square$

4.1.1 Existence of Spherical Designs

The following theorem — a more detailed version of Theorem 1.6 — establishes the existence of a spherical t-design of size $\text{poly}(N_{2t,d})$ with the desired pairwise separation properties.

Theorem 4.5 (Spherical Design Construction). *Let t be an odd integer and $r \geq N_{2t,d}^5$. Let $\mathbf{y}_1, \ldots, \mathbf{y}_r$ be uniform random vectors over $\mathbb{S}^{d-1}$. Then, with probability at least $99/100$, there exist unit vectors $\mathbf{z}_1, \ldots, \mathbf{z}_r \in \mathbb{S}^{d-1}$ such that $\|\mathbf{z}_i - \mathbf{y}_i\|_2 \leq O(1/N_{2t,d}), i \in [r]$, and $(\mathbf{z}_1, \ldots, \mathbf{z}_r)$ form a spherical t-design.*

To prove Theorem 4.5, we will start from the following result.

Theorem 4.6 ([BRV13]). *If there exists a continuous mapping $F : \mathcal{P}_t^d \to (\mathbb{S}^{d-1})^r$ such that for all $p \in \partial\Omega_t^d$, $\sum_{i=1}^r p(\mathbf{x}_i(p)) > 0$, where $F(p) = (\mathbf{x}_1(p), \ldots, \mathbf{x}_r(p))$, then there exists a polynomial $p^* \in \Omega_t^d$ such that $\mathbf{E}[p(\mathbf{x})] = \frac{1}{r}\sum_{i=1}^r p(\mathbf{x}_i(p^*))$ holds for every polynomial $p \in \mathcal{P}_t^d$.*

To apply Theorem 4.6, we need to find a continuous function F mapping $\mathcal{P}_t^d$ to $(\mathbb{S}^{d-1})^r$ such that for any $p \in \partial\Omega_t^d$, $\sum_{i=1}^r p(\mathbf{z}_i) > 0$, where $F(p) = (\mathbf{z}_1, \ldots, \mathbf{z}_r)$. We will construct the mapping F as follows: we sample $\mathbf{y}_1, \ldots, \mathbf{y}_r$ uniformly over the unit sphere $\mathbb{S}^{d-1}$, and then try to make the value of $p(\mathbf{y}_i)$ larger by moving each point $\mathbf{y}_i$ in the direction of the gradient. In particular, we let $\mathbf{z}_i = \frac{\mathbf{y}_i + \delta\nabla_o p(\mathbf{y}_i)}{\|\mathbf{y}_i + \delta\nabla_o p(\mathbf{y}_i)\|_2}$ for some $\delta > 0$ sufficiently small, where $\nabla_o p(\mathbf{y})$ is the component of $\nabla p(\mathbf{y})$ orthogonal to the direction $\mathbf{y}$. We will prove that with high probability, for any $p \in \partial\Omega_t^d$, $\sum_{i=1}^r p(\mathbf{z}_i) > 0$. Intuitively, this works because of two facts:

1. With high probability over the choice of $\mathbf{y}_i$, for all $p \in \Omega_t^d$, the average value of $p(\mathbf{y})$ is already close to zero.

2. Moving in the direction of $\nabla_o p(\mathbf{y}_i)$ increases $p(\mathbf{y}_i)$ by a notable amount.

Lemma 4.7. *Let Ω be a subspace of polynomials in d variables with mean zero. Let N be the dimension of Ω. Let $\mathbf{x}_1, \ldots, \mathbf{x}_r$ be i.i.d. random vectors over $\mathbb{S}^{d-1}$. Then, with probability at least $1 - \frac{N}{r\eta^2}$, we have that for any $p \in \Omega$, $\left|\frac{1}{r}\sum_{i=1}^r p(\mathbf{x}_i)\right| \leq \eta\|p\|_2$.*

To prove Lemma 4.7, we take an orthonormal basis $p_1, \ldots, p_N \in \Omega$ and define $\mathbf{p}(\mathbf{x}) \overset{\text{def}}{=} [p_1(\mathbf{x}), \ldots, p_N(\mathbf{x})]^\mathsf{T}$. Noting that $\mathbf{p}(\mathbf{x})$ is a random vector with mean zero and covariance identity, applying Markov's inequality will yield the result. We defer the proof details to Appendix C.

Lemma 4.8. *Let $\mathbf{y} \in \mathbb{S}^{d-1}$ and $0 < \delta \leq 1/N_{2t,d}^2$. For any $p \in \partial\Omega_t^d$, let $\mathbf{z} = \frac{\mathbf{y}+\delta\nabla_o p(\mathbf{y})}{\|\mathbf{y}+\delta\nabla_o p(\mathbf{y})\|_2}$. We have that $p(\mathbf{z}) - p(\mathbf{y}) \geq C\delta\|\nabla_o p(\mathbf{y})\|_2^2$ for some universal constant $0 < C < 1$.*

The proof of Lemma 4.8 follows by Taylor expansion, where the contributions to $p(\mathbf{z})$ coming from higher-order terms will not be large as long as $\mathbf{z}$ is sufficiently close to $\mathbf{y}$. We defer the proof details to Appendix C. By applying the above two lemmas, we establish the following:

11

Theorem 4.9. *Let* $\mathbf{y}_1, \ldots, \mathbf{y}_r$ *be i.i.d. random vectors over* $\mathbb{S}^{d-1}$. *Let* $\delta = 1/N_{2t,d}^2$ *and* $r \geq N_{2t,d}^5$. *We consider the mapping* $F : \mathcal{P}_t^d \to (\mathbb{S}^{d-1})^r$ *as follows: for any* $p \in \mathcal{P}_t^d$, *let*

$$\mathbf{z}_i = \frac{\mathbf{y}_i + \delta \nabla_o p(\mathbf{y}_i)}{\|\mathbf{y}_i + \delta \nabla_o p(\mathbf{y}_i)\|_2} \ ,$$

where $\nabla_o p(\mathbf{y})$ *is the component of* $\nabla p(\mathbf{y})$ *orthogonal to the direction* $\mathbf{y}$. *Let* $F(p) := (\mathbf{z}_1, \ldots, \mathbf{z}_r)$. *Then, with probability at least* $99/100$, *we have that for any* $p \in \partial \Omega_t^d$, $\sum_{i=1}^r p(\mathbf{z}_i) > 0$.

To prove Theorem 4.9, we consider $\widetilde{p}(\mathbf{y}) = p(\mathbf{y}) + C\delta(\|\nabla_o p(\mathbf{y})\|_2^2 - \mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2])$, where C, δ comes from Lemma 4.8. Noting that $\mathbf{E}[\widetilde{p}(\mathbf{y})] = 0$ and $\widetilde{p}(\mathbf{y})$ is a polynomial containing only monomials of degree $2t, 2t - 2, t, 0$, we are able to apply Lemma 4.7 to obtain the desired result. See Appendix C for the proof.

4.1.2 Proof of Theorem 4.5

Let $\delta = 1/N_{2t,d}^2$. We consider the mapping $F : \mathcal{P}_t^d \to (\mathbb{S}^{d-1})^r$ as follows: for any $p \in \mathcal{P}_t^d$, let $\mathbf{z}_i = \frac{\mathbf{y}_i + \delta \nabla_o p(\mathbf{y}_i)}{\|\mathbf{y}_i + \delta \nabla_o p(\mathbf{y}_i)\|_2}$, where $\nabla_o p(\mathbf{y})$ is the component of $\nabla p(\mathbf{y})$ orthogonal to the direction $\mathbf{y}$. By Theorem 4.9, with probability at least $99/100$, we have that for any $p \in \partial \Omega_t^d$, $\sum_{i=1}^r p(\mathbf{z}_i) > 0$. Applying Theorem 4.6 yields that there exists some $p^* \in \Omega_t^d$ such that $F(p^*) = (\mathbf{z}_1^*, \ldots, \mathbf{z}_r^*)$ form a spherical t-design. By elementary calculation, we have that

$$\|\mathbf{z}_i^* - \mathbf{y}_i\|_2 = \left\| \frac{\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i)}{\|\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i)\|_2} - \mathbf{y}_i \right\|_2 = \frac{\|\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i) - \|\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i)\|_2 \mathbf{y}_i\|_2}{\|\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i)\|_2}$$

$$\leq \frac{|1 - \|\mathbf{y}_i + \delta \nabla_o p^*(\mathbf{y}_i)\|_2| + \delta \|\nabla p^*(\mathbf{y}_i)\|_2}{1 - \delta \|\nabla p^*(\mathbf{y}_i)\|_2} \leq \frac{2\delta \|\nabla p^*(\mathbf{y}_i)\|_2}{1 - \delta \|\nabla p^*(\mathbf{y}_i)\|_2} \leq O(1/N_{2t,d}) \ ,$$

where the last inequality follows from $\forall \mathbf{y} \in \mathbb{S}^{d-1}, \|\nabla p^*(\mathbf{y})\|_2 \leq \sqrt{t(d + 2t - 2)N_{2(t-1),d}\|p^*\|_2^2} \leq N_{2t,d}$ by Lemma 4.4.

4.2 Proof of Theorem 1.2

We first prove Proposition 4.1 based on our construction of spherical t-design. To achieve this, it suffices to show that with high probability there exist $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{S}^{m-1}$ such that

- $\mathbf{v}_1, \ldots, \mathbf{v}_r$ is a spherical k-design.

- $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2 \geq \Omega(\Delta), \forall 1 \leq i < j \leq r$ for some $\Delta > 0$.

Proof of Proposition 4.1. Let $\Theta = \min_{1 \leq i < j \leq r} \arccos \langle \mathbf{y}_i, \mathbf{y}_j \rangle$. Applying Fact 3.4 by taking $\gamma = \left(\frac{1}{50\kappa_{m-1}} \right)^{\frac{1}{m-1}}$ yields $\mathbf{Pr}[\Theta \geq \Delta] \geq \mathbf{Pr}\left[\Theta \geq \gamma r^{-\frac{2}{m-1}} \right] \geq 99/100$. This will give

$$\mathbf{Pr}\left[\min_{1 \leq i < j \leq r} \{ \|\mathbf{y}_i - \mathbf{y}_j\|_2, \|\mathbf{y}_i + \mathbf{y}_j\|_2 \} \geq \Omega(\Delta) \right] \geq 99/100 \ ,$$

since our choice of m will imply that $r = 1/\sqrt{50\kappa_{m-1}\Delta^{m-1}}$. By Theorem 4.5, with probability at least $99/100$, there exist unit vectors $\mathbf{z}_1^*, \ldots, \mathbf{z}_r^* \in \mathbb{S}^{m-1}$ such that $(\mathbf{z}_1^*, \ldots, \mathbf{z}_r^*)$ form a spherical k-design and $\|\mathbf{z}_i^* - \mathbf{y}_i\|_2 \leq O(1/N_{2k,m}), i \in [r]$. Therefore, for any odd homogeneous polynomial p in m variables of odd degree $t < k$, we have that

$$\frac{1}{r} \sum_{i=1}^r p(\mathbf{z}_i^*) = \frac{1}{r} \sum_{i=1}^r (\|\mathbf{z}_i^*\|_2^2)^{\frac{k-t}{2}} p(\mathbf{z}_i^*) = \mathbf{E}\left[(\|\mathbf{z}\|_2^2)^{\frac{k-t}{2}} p(\mathbf{z}) \right] = 0 \ ,$$

12

which implies that $\sum_{i=1}^{r} q(\mathbf{z}_i^*) = 0$ holds for any polynomial q in m variables of degree less than k. In addition, for every $1 \leq i < j \leq r$, we have that

$$
\begin{aligned}
\|\mathbf{z}_i^* \pm \mathbf{z}_j^*\|_2 &= \|(\mathbf{y}_i \pm \mathbf{y}_j) + (\mathbf{z}_i^* - \mathbf{y}_i) \pm (\mathbf{z}_j^* - \mathbf{y}_j)\|_2 \geq \|\mathbf{y}_i \pm \mathbf{y}_j\| - \|\mathbf{z}_i^* - \mathbf{y}_i\|_2 - \|\mathbf{z}_j^* - \mathbf{y}_j\|_2 \\
&\geq \|\mathbf{y}_i \pm \mathbf{y}_j\|_2 - O(1/N_{2k,m}) \,.
\end{aligned}
$$

This completes the proof. $\qquad\square$

Proof of Theorem 1.2. Let $m = \frac{c \log r}{\log(1/\Delta)}$, where $r^{-1/10} \leq \Delta < 1$ and $1.99 \leq c \leq 2$ is a constant. Let $c' = \frac{(1/e)(1/\Delta)^{1/(5c)} - 1}{2}$ and $k = c'm = \frac{((1/e)(1/\Delta)^{1/(5c)} - 1)c \log r}{2 \log(1/\Delta)}$. In this way, we will have that

$$
\begin{aligned}
N_{2k,m}^5 &\leq \binom{m+2k}{2k}^5 = \binom{(1+2c')m}{m}^5 \leq 2^{5(1+2c')mH\left(\frac{1}{1+2c'}\right)} = 2^{5(1+2c')m\left(\frac{\log(1+2c')}{1+2c'} + \frac{2c' \log(1+1/2c')}{1+2c'}\right)} \\
&= 2^{5m(\log(1+2c')+2c' \log(1+1/2c'))} \leq 2^{\frac{5c \log r(\log(1+2c')+2c' \log(1+1/2c'))}{\log(1/\Delta)}} \leq r,
\end{aligned}
$$

where $H(p) = -p \log p - (1-p) \log(1-p)$, $p \in [0,1]$, is the standard binary entropy function. Therefore, by Proposition 4.1, there exist vectors $\mathbf{v}_1, \ldots, \mathbf{v}_r \in \mathbb{S}^{m-1}$ such that

- $\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[g(\mathbf{z})p(\mathbf{z})] = 0$ holds for every odd polynomial p of degree less than k, where $g(\mathbf{z}) = \frac{1}{r} \sum_{\ell=1}^{r} \text{sign}(\mathbf{v}_\ell^\mathsf{T} \mathbf{z})$.

- $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2 \geq \Omega(\Delta) - O(1/N_{2k,m}), \forall 1 \leq i < j \leq r$.

For an arbitrary matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$ with $\mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m$, denote by $D_{\mathbf{U}}$ the instance of mixture of linear classifiers with weight vectors $\mathbf{U}^\mathsf{T}\mathbf{v}_1, \ldots, \mathbf{U}^\mathsf{T}\mathbf{v}_r$. Let $\mathcal{D}_g = \{\mathcal{D}_{\mathbf{U}} \mid \mathbf{U} \in \mathbb{R}^{m \times n}, \mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m\}$. By definition, we have that $\mathbf{E}_{(\mathbf{x},y) \sim D_{\mathbf{U}}}[y \mid \mathbf{x} = \mathbf{z}] = \frac{1}{r} \sum_{\ell=1}^{r} \text{sign}((\mathbf{U}^\mathsf{T}\mathbf{v}_\ell)^\mathsf{T}\mathbf{z}) = \frac{1}{r} \sum_{\ell=1}^{r} \text{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{z}) = g(\mathbf{U}\mathbf{z}), \mathbf{z} \in \mathbb{R}^n$, and $\|\mathbf{U}^\mathsf{T}\mathbf{v}_i \pm \mathbf{U}^\mathsf{T}\mathbf{v}_j\|_2 = \|\mathbf{v}_i \pm \mathbf{v}_j\|_2 \geq \Omega(\Delta), 1 \leq i < j \leq r$, since

$$
\begin{aligned}
N_{2k,m} &= \binom{m+2k-1}{m-1} = \binom{(1+2c')m-1}{m-1} \geq \left(\frac{(1+2c')m-1}{m-1}\right)^{m-1} \geq (1+2c')^{m-1} \\
&\geq (1+2c')^{\frac{1.99 \log r}{\log(1/\Delta)}-1} = \left((1/e)(1/\Delta)^{1/5c}\right)^{\frac{1.89 \log r}{\log(1/\Delta)}} \geq \Omega((1/\Delta)^{1.89}) \,.
\end{aligned}
$$

Therefore by Proposition 3.2, any SQ algorithm that solves the decision problem $\mathcal{B}(\mathcal{D}_g, \mathcal{N}_n \times \mathcal{U}(\{\pm 1\}))$ must either use queries of tolerance $n^{-\Omega(k)}$, or make at least $2^{n^{\Omega(1)}}$ queries. By Lemma 3.8, for any $\mathbf{U} \in \mathbb{R}^{m \times n}$ with $\mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m$, we have that $d_{\text{TV}}(D_{\mathbf{U}}, D_0) \geq \Omega(\Delta/r) \geq 2(n^{-\Omega(k)} + \epsilon)$, where $D_0 = \mathcal{N}_n \times \mathcal{U}(\{\pm 1\})$. Therefore, by Claim 3.7, any SQ algorithm that learns a distribution in $\mathcal{D}_g$ within error ϵ in total variation distance must either use queries of tolerance $n^{-\Omega(\log r/(\Delta^{1/(5c)} \log(1/\Delta)))}$, or make at least $2^{n^{\Omega(1)}}$ queries. This completes the proof. $\qquad\square$

5 Conclusion

This work establishes a near-optimal Statistical Query (SQ) lower bound for learning uniform mixtures of linear classifiers under the Gaussian distribution. Our lower bound nearly matches prior algorithmic work on the problem. Our result applies for the simplest (and well-studied) distributional setting where the covariates are drawn from the standard Gaussian distribution. This directly implies similar information-computation tradeoffs for the setting that the covariates are drawn from a more general distribution family (e.g., an unknown subgaussian or a log-concave distribution) that includes the standard normal.

From a technical perspective, we believe that our new efficient construction of spherical designs is a mathematical contribution of independent interest that could be used to establish SQ lower bounds for other related latent variable models (e.g., various mixtures of experts). A natural direction is to establish information-computation tradeoffs for a fixed non-Gaussian distribution on covariates (e.g., the uniform distribution over the Boolean hypercube), for which a different hardness construction is needed.

References

[AK01]   S. Arora and R. Kannan. Learning mixtures of arbitrary Gaussians. In *Proceedings of the 33rd Symposium on Theory of Computing*, pages 247–257, 2001.

[AM05]   D. Achlioptas and F. McSherry. On spectral learning of mixtures of distributions. In *Proceedings of the Eighteenth Annual Conference on Learning Theory (COLT)*, pages 458–469, 2005.

[Ban79]   E. Bannai. On tight spherical designs. *Journal of Combinatorial Theory, Series A*, 26(1):38–47, 1979.

[BBH⁺20] M. Brennan, G. Bresler, S. B. Hopkins, J. Li, and T. Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint arXiv:2009.06107*, 2020. In *34th Annual Conference on Learning Theory (COLT 2021)*.

[BC92]   A. Blum and P. Chalasani. Learning switching concepts. In David Haussler, editor, *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory, COLT 1992, Pittsburgh, PA, USA, July 27-29, 1992*, pages 231–242. ACM, 1992.

[BDJ⁺20] A. Bakshi, I. Diakonikolas, H. Jia, D. M. Kane, P. K. Kothari, and S. S. Vempala. Robustly learning mixtures of k arbitrary gaussians. *CoRR*, abs/2012.02119, 2020. Conference version in STOC'22.

[BRS⁺18] J. S. Brauchart, A. B. Reznikov, E. B. Saff, I. H. Sloan, Y. G. Wang, and R. S. Womersley. Random point sets on the sphere—hole radii, covering, and separation. *Experimental Mathematics*, 27(1):62–81, 2018.

[BRV13]   A. Bondarenko, D. Radchenko, and M. Viazovska. Optimal asymptotic bounds for spherical designs. *Annals of mathematics*, pages 443–452, 2013.

[BS21]   R. D. Buhai and D. Steurer. Beyond parallel pancakes: Quasi-polynomial time guarantees for non-spherical gaussian mixtures. *CoRR*, abs/2112.05445, 2021.

[BV08]   S. C. Brubaker and S. Vempala. Isotropic PCA and Affine-Invariant Clustering. In *Proc. 49th IEEE Symposium on Foundations of Computer Science*, pages 551–560, 2008.

[CC06]   Y. J. Cho and Y. Q. Chen. *Topological degree theory and applications*. CRC Press, 2006.

[CDV22]   A. Chen, A. De, and A. Vijayaraghavan. Algorithms for learning a mixture of linear classifiers. In *International Conference on Algorithmic Learning Theory*, pages 205–226. PMLR, 2022.

[CLS20] S. Chen, J. Li, and Z. Song. Learning mixtures of linear regressions in subexponential time via fourier moments. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 587–600, 2020.

[Das99] S. Dasgupta. Learning mixtures of Gaussians. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, pages 634–644, 1999.

[DGS91] P. Delsarte, J. Goethals, and J. J. Seidel. Spherical codes and designs. In *Geometry and Combinatorics*, pages 68–93. Elsevier, 1991.

[DHKK20] I. Diakonikolas, S. B. Hopkins, D. Kane, and S. Karmalkar. Robustly learning any clusterable mixture of gaussians. *CoRR*, abs/2005.06417, 2020. Conference version in FOCS'20.

[DK20] I. Diakonikolas and D. M. Kane. Small covers for near-zero sets of polynomials and learning latent variable models. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2020)*, 2020.

[DK23] I. Diakonikolas and D. M. Kane. *Algorithmic High-Dimensional Robust Statistics*. Cambridge university press, 2023.

[DKK+21] I. Diakonikolas, D. M. Kane, D. Kongsgaard, J. Li, and K. Tian. Clustering mixture models in almost-linear time via list-decodable mean estimation. *CoRR*, abs/2106.08537, 2021. Conference version in STOC'22.

[DKKZ20] I. Diakonikolas, D. M. Kane, V. Kontonis, and N. Zarifis. Algorithms and SQ lower bounds for PAC learning one-hidden-layer relu networks. In *Conference on Learning Theory, COLT 2020*, volume 125 of *Proceedings of Machine Learning Research*, pages 1514–1539. PMLR, 2020.

[DKPZ21] I. Diakonikolas, D. M. Kane, T. Pittas, and N. Zarifis. The optimality of polynomial regression for agnostic learning under gaussian marginals in the SQ model. In *Conference on Learning Theory, COLT 2021*, volume 134 of *Proceedings of Machine Learning Research*, pages 1552–1584. PMLR, 2021.

[DKS17] I. Diakonikolas, D. M. Kane, and A. Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 73–84, 2017. Full version at http://arxiv.org/abs/1611.03473.

[DKS18] I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 1047–1060, 2018. Full version available at https://arxiv.org/abs/1711.07211.

[Fel16] V. Feldman. Statistical query learning. In *Encyclopedia of Algorithms*, pages 2090–2095. Springer New York, 2016.

[FGR+17] V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM*, 64(2):8:1–8:37, 2017.

[FGV17]   V. Feldman, C. Guzman, and S. S. Vempala. Statistical query algorithms for mean vector estimation and stochastic convex optimization. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 1265–1277. SIAM, 2017.

[GMP20]   V. Gandikota, A. Mazumdar, and S. Pal. Recovery of sparse linear classifiers from mixture of responses. *Advances in Neural Information Processing Systems*, 33:14688–14698, 2020.

[GP11]   M. Gräf and D. Potts. On the computation of spherical designs by a new optimization approach based on fast spherical fourier transforms. *Numerische Mathematik*, 119(4):699–724, 2011.

[HKP$^+$17]   S. B. Hopkins, P. K. Kothari, A. Potechin, P. Raghavendra, T. Schramm, and D. Steurer. The power of sum-of-squares for detecting hidden structures. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 720–731. IEEE Computer Society, 2017.

[HL18]   S. B. Hopkins and J. Li. Mixture models, robustness, and sum of squares proofs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 1021–1034, 2018.

[Hop18]   S. B. Hopkins. *Statistical inference and the sum of squares method*. PhD thesis, Cornell University, 2018.

[HS17]   S. B. Hopkins and D. Steurer. Efficient bayesian estimation from few samples: Community detection and related problems. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 379–390. IEEE Computer Society, 2017.

[HYJ22]   N. Ho, C. Yang, and M. I. Jordan. Convergence rates for gaussian mixtures of experts. *The Journal of Machine Learning Research*, 23(1):14523–14603, 2022.

[JJ94]   M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.

[Kan15]   D. M. Kane. Small designs for path-connected spaces and path-connected homogeneous spaces. *Transactions of the American Mathematical Society*, 367(9):6387–6414, 2015.

[Kea98]   M. J. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6):983–1006, 1998.

[KSV08]   R. Kannan, H. Salmasian, and S. Vempala. The spectral method for general mixture models. *SIAM J. Comput.*, 38(3):1141–1156, 2008.

[LL18]   Y. Li and Y. Liang. Learning mixtures of linear regressions with nearly optimal complexity. In *Conference On Learning Theory*, pages 1125–1144. PMLR, 2018.

[LLYH17]   W. Li, P. Liang, X. Yuan, and J. Hu. Non-local information for a mixture of multiple linear classifiers. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 3741–3746. IEEE, 2017.

[LM20]   A. Liu and A. Moitra. Settling the robust learnability of mixtures of gaussians. *coRR*, 2020. Conference version in STOC'21.

[ME14]   S. Masoudnia and R. Ebrahimpour. Mixture of experts: a literature survey. *The Artificial Intelligence Review*, 42(2):275, 2014.

[MV10]   A. Moitra and G. Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *FOCS*, pages 93–102, 2010.

[MVKO19] A. Makkuva, P. Viswanath, S. Kannan, and S. Oh. Breaking the gridlock in mixture-of-experts: Consistent and efficient algorithms. In *International Conference on Machine Learning*, pages 4304–4313. PMLR, 2019.

[Pea94]   K. Pearson. Contribution to the mathematical theory of evolution. *Phil. Trans. Roy. Soc. A*, 185:71–110, 1894.

[RV17]   O. Regev and A. Vijayaraghavan. On learning mixtures of well-separated gaussians. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 85–96, 2017.

[SIM14]   Y. Sun, S. Ioannidis, and A. Montanari. Learning mixtures of linear classifiers. In *International Conference on Machine Learning*, pages 721–729. PMLR, 2014.

[SJA16]   H. Sedghi, M. Janzamin, and A. Anandkumar. Provable tensor methods for learning mixtures of generalized linear models. In *Artificial Intelligence and Statistics*, pages 1223–1231. PMLR, 2016.

[Ver18]   R. Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

[VW02]   S. Vempala and G. Wang. A spectral algorithm for learning mixtures of distributions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science*, pages 113–122, 2002.

[Wom18]   R. S. Womersley. Efficient spherical designs with good geometric properties. *Contemporary computational mathematics-A celebration of the 80th birthday of Ian Sloan*, pages 1243–1285, 2018.

[XA09]   L. Xu and S. Amari. Combining classifiers and learning mixture-of-experts. In *Encyclopedia of artificial intelligence*, pages 318–326. IGI Global, 2009.

Appendix

A Background on Hermite Polynomials

Recall the definition of the probabilist's Hermite polynomials:

$$He_n(x) = (-1)^n e^{x^2/2} \cdot \frac{d^2}{dx^2} e^{-x^2/2}.$$

Under this definition, the first four Hermite polynomials are

$$He_0(x) = 1, He_1(x) = x, He_2(x) = x^2 - 1, He_3(x) = x^3 - 3x.$$

In our work, we will consider the *normalized* Hermite polynomial of degree n to be $h_n(x) = \frac{He_n(x)}{\sqrt{n!}}$. These normalized Hermite polynomials form a complete orthogonal basis for inner product space $\mathcal{L}^2(\mathbb{R}, \mathcal{N})$. To obtain an orthogonal basis for $\mathcal{L}^2(\mathbb{R}^d, \mathcal{N}_d)$, we will use a multi-index $J = (j_1, \ldots, j_d) \in \mathbb{N}^d$ to define the d-variate normalized Hermite polynomial as $H_J(\mathbf{x}) = \prod_{i=1}^d H_{j_i}(x_i)$. Let the total degree of H_J be $|J| = \sum_{i=1}^d j_i$. Given a function $f \in \mathcal{L}^2(\mathbb{R}^d, \mathcal{N}_d)$, we can express it uniquely as $f(\mathbf{x}) = \sum_{J \in \mathbb{N}^d} \widehat{f}(J) H_J(\mathbf{x})$, where $\widehat{f}(J) = \mathbf{E}_{\mathbf{x} \in \mathcal{N}_d}[f(\mathbf{x}) H_J(\mathbf{x})]$. We denote by $f^{[k]}(\mathbf{x})$ the degree k part of the Hermite expansion of f, i.e., $f^{[k]}(\mathbf{x}) = \sum_{|J|=k} \widehat{f}(J) H_J(\mathbf{x})$.

Definition A.1. We say that a polynomial q in d variables is harmonic of degree k if it is a linear combination of degree k Hermite polynomials. That is, q is harmonic if it can be written as

$$q(\mathbf{x}) = q^{[k]}(\mathbf{x}) = \sum_{J:|J|=k} c_J H_J(\mathbf{x}).$$

Notice that, since for a single-dimensional Hermite polynomial it holds $h'_m(x) = \sqrt{m} h_{m-1}(x)$, we have that $\nabla H_M^{(i)}(\mathbf{x}) = \sqrt{m_i} H_{M-E_i}(\mathbf{x})$, where $M = (m_1, \ldots, m_d)$. From this fact and the orthogonality of Hermite polynomials, we obtain

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_d}[\langle \nabla H_M(\mathbf{x}), \nabla H_L(\mathbf{x}) \rangle] = |M| \, \mathbb{I}[M = L] \, .$$

We will also require the following standard facts:

Fact A.2. *Let p be a polynomial of degree k in d variables. Then p is harmonic of degree k if and only if for all $\mathbf{x} \in \mathbb{R}^d$ it holds that $kp(\mathbf{x}) = \langle \mathbf{x}, \nabla p(\mathbf{x}) \rangle - \nabla^2 p(\mathbf{x})$.*

Fact A.3 (see, e.g., [DKPZ21]). *Let p, q be harmonic polynomials of degree k. Then,*

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_d} \left[\langle \nabla^\ell p(\mathbf{x}), \nabla^\ell q(\mathbf{x}) \rangle \right] = k(k-1) \ldots (k - \ell + 1) \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_d}[p(\mathbf{x}) q(\mathbf{x})].$$

In particular,

$$\langle \nabla^k p(\mathbf{x}), \nabla^k q(\mathbf{x}) \rangle = k! \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_d}[p(\mathbf{x}) q(\mathbf{x})].$$

B Omitted Proofs from Section 3

B.1 Proof of Lemma 3.5

We start with the following claim:

Claim B.1. *Let $p : \mathbb{R}^{n_1} \to \mathbb{R}$ and $q : \mathbb{R}^{n_2} \to \mathbb{R}$, where p is a polynomial of degree at most k and $q \in \mathcal{L}^2(\mathbb{R}^{n_2}, \mathcal{N}_{n_2})$. Let $\mathbf{U} \in \mathbb{R}^{n_1 \times n}, \mathbf{V} \in \mathbb{R}^{n_2 \times n}$ such that $\mathbf{U}\mathbf{U}^{\mathsf{T}} = \mathbf{I}_{n_1}, \mathbf{V}\mathbf{V}^{\mathsf{T}} = \mathbf{I}_{n_2}$. Then, we have that $\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n}[p(\mathbf{U}\mathbf{x})q(\mathbf{V}\mathbf{x})] = \sum_{m=0}^{k} \frac{1}{m!}\langle (\mathbf{U}^{\mathsf{T}})^{\otimes m}\mathbf{R}_1^m, (\mathbf{V}^{\mathsf{T}})^{\otimes m}\mathbf{R}_2^m \rangle$, where $\mathbf{R}_1^m = \nabla^m p^{[m]}(\mathbf{x}), \mathbf{R}_2^m = \nabla^m q^{[m]}(\mathbf{x})$.*

We require the following lemma:

Lemma B.2. *Let p be a harmonic polynomial of degree k. Let $\mathbf{V} \in \mathbb{R}^{m \times n}$ with $\mathbf{V}\mathbf{V}^{\mathsf{T}} = \mathbf{I}_m$. Then the polynomial $p(\mathbf{V}\mathbf{x})$ is harmonic of degree k.*

Proof. Let $f(\mathbf{x}) = p(\mathbf{V}\mathbf{x})$. By Fact A.2, it suffices to show that for all $\mathbf{x} \in \mathbb{R}^n$ it holds that $kf(\mathbf{x}) = \langle \mathbf{x}, \nabla f(\mathbf{x}) \rangle - \nabla^2 f(\mathbf{x})$. Since $\mathbf{V}\mathbf{V}^{\mathsf{T}} = \mathbf{I}_m$, applying Fact A.2 yields

$$\langle \mathbf{x}, \nabla f(\mathbf{x}) \rangle - \nabla^2 f(\mathbf{x}) = \langle \mathbf{V}\mathbf{x}, \nabla p(\mathbf{V}\mathbf{x}) \rangle - \nabla^2 p(\mathbf{V}\mathbf{x}) = kp(\mathbf{V}\mathbf{x}) = kf(\mathbf{x}) .$$

This completes the proof. $\qquad\square$

Proof of Claim B.1. For $m \in \mathbb{N}$, let $f^{(m)}(\mathbf{x}) = p^{[m]}(\mathbf{U}\mathbf{x})$ and $g^{(m)}(\mathbf{x}) = q^{[m]}(\mathbf{V}\mathbf{x})$. We can write $p(\mathbf{U}\mathbf{x}) \sim \sum_{m=0}^{k} f^{(m)}(\mathbf{x})$ and $q(\mathbf{V}\mathbf{x}) \sim \sum_{m=0}^{\infty} g^{(m)}(\mathbf{x})$. Then applying Fact A.3 and Lemma B.2 yields

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n}[p(\mathbf{U}\mathbf{x})q(\mathbf{V}\mathbf{x})] = \sum_{m_1=0}^{k} \sum_{m_2=0}^{\infty} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n}[f^{(m_1)}(\mathbf{x})g^{(m_2)}(\mathbf{x})] = \sum_{m=0}^{k} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n}[f^{(m)}(\mathbf{x})g^{(m)}(\mathbf{x})]$$

$$= \sum_{m=0}^{k} \frac{1}{m!} \left\langle \nabla^m f^{(m)}(\mathbf{x}), \nabla^m g^{(m)}(\mathbf{x}) \right\rangle = \sum_{m=0}^{k} \frac{1}{m!} \left\langle \nabla^m p^{[m]}(\mathbf{U}\mathbf{x}), \nabla^m q^{[m]}(\mathbf{V}\mathbf{x}) \right\rangle .$$

Denote by $\mathcal{U} \subseteq \mathbb{R}^n$ the image of the linear map $\mathbf{U}^{\mathsf{T}}$. Applying the chain rule, for any function $h(\mathbf{U}\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$, it holds $\nabla h(\mathbf{U}\mathbf{x}) = \partial_i h(\mathbf{U}\mathbf{x})U_{ij} \in \mathcal{U}$, where we applied Einstein's summation notation for repeated indices. Applying the above rule m times, we have that

$$\nabla^m h(\mathbf{U}\mathbf{x}) = \partial_{i_m} \ldots \partial_{i_1} h(\mathbf{U}\mathbf{x})U_{i_1,j_1} \ldots U_{i_m,j_m} \in \mathcal{U}^{\otimes m}.$$

Moreover, denote $\mathbf{S}_m = \nabla^m p^{[m]}(\mathbf{U}\mathbf{x}) = (\mathbf{U}^{\mathsf{T}})^{\otimes m}\mathbf{R}_1^m \in \mathcal{U}^{\otimes m}$, and $\mathbf{T}_m = \nabla^m q^{[m]}(\mathbf{V}\mathbf{x}) = (\mathbf{V}^{\mathsf{T}})^{\otimes m}\mathbf{R}_2^m \in \mathcal{V}^{\otimes m}$. We have that

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n}[f(\mathbf{x})g(\mathbf{x})] = \sum_{m=0}^{k} \frac{1}{m!} \left\langle \nabla^m p^{[m]}(\mathbf{U}\mathbf{x}), \nabla^m q^{[m]}(\mathbf{V}\mathbf{x}) \right\rangle = \sum_{m=0}^{k} \frac{1}{m!}\langle \mathbf{S}_m, \mathbf{T}_m \rangle$$

$$= \sum_{m=0}^{k} \frac{1}{m!}\langle (\mathbf{U}^{\mathsf{T}})^{\otimes m}\mathbf{R}_1^m, (\mathbf{V}^{\mathsf{T}})^{\otimes m}\mathbf{R}_2^m \rangle.$$

This proves the claim. $\qquad\square$

Proof of Lemma 3.5. Applying Claim B.1 by taking $\mathbf{U} = \mathbf{I}_m$ and $\mathbf{V} = \mathbf{v}^{\mathsf{T}}$, we have that

$$\mathbf{E}_{\mathbf{z} \sim \mathcal{N}_m}[p(\mathbf{z})f(\mathbf{v}^{\mathsf{T}}\mathbf{z})] = \sum_{d=0}^{k-1} \frac{1}{d!}\langle \mathbf{R}_1^d, \mathbf{v}^{\otimes d}\mathbf{R}_2^d \rangle,$$

which is a polynomial in $\mathbf{v}$ of degree less than k, since $\mathbf{R}_1^d = \nabla^d p^{[d]}(\mathbf{x})$ and $\mathbf{R}_2^d = \nabla^d f^{[d]}(\mathbf{x})$ are constants only depending on p and f. This completes the proof of Lemma 3.5. $\qquad\square$

B.2 Proof of Lemma 3.6

We start by proving that *"there exist non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^{r} w_\ell = 1$ such that $\sum_{\ell=1}^{r} w_\ell q(\mathbf{v}_\ell) = 0$ for all odd polynomials q of degree less than k"* implies *"there does not exist any odd polynomial q of degree less than k such that $q(\mathbf{v}_\ell) > 0, 1 \leq \ell \leq r$."* Suppose for contradiction that there exists an odd polynomial q^* of degree less than k such that $q^*(\mathbf{v}_\ell) > 0, 1 \leq \ell \leq r$. For arbitrary non-negative weights $w_1, \ldots, w_r$ with $\sum_{\ell=1}^{r} w_\ell = 1$, we have that $\sum_{\ell=1}^{r} w_\ell q^*(\mathbf{v}_\ell) \geq \min\{q^*(\mathbf{v}_1), \ldots, q^*(\mathbf{v}_r)\} > 0$, which contradicts to the first statement.

We then prove the opposite direction. We will use the following version of Farkas' lemma.

Fact B.3 (Farkas' lemma). *Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. Then exactly one of the following two assertions is true:*

- *There exists an $\mathbf{x} \in \mathbb{R}^n$ such that $A\mathbf{x} = b$ and $\mathbf{x} \geq 0$.*

- *There exists a $\mathbf{y} \in \mathbb{R}^m$ such that $\mathbf{y}^\intercal A \geq 0$ and $\mathbf{y}^\intercal b < 0$.*

Suppose for contradiction that there does not exist $w_1, \ldots, w_r$ with $\sum_{\ell=1}^{r} w_\ell = 1$ such that $\sum_{\ell=1}^{r} w_\ell q(\mathbf{v}_\ell) = 0$ holds for every odd polynomial q of degree less than k. Let $s_{k,m}$ denote the total number of m-variate odd monomials of degree less than k, and $\{q_j^{k,m}\}_{1 \leq j \leq s_{k,m}}$ denote such monomials. We consider the following LP with variables $\mathbf{w} = (w_1, \ldots, w_r)^\intercal$: $\sum_{\ell=1}^{r} w_\ell q_j^{k,m}(\mathbf{v}_\ell) = 0, 1 \leq j \leq s_{k,m}, \sum_{\ell=1}^{r} w_\ell = 1, w_\ell \geq 0, 1 \leq \ell \leq r$. By our assumption, the LP is infeasible. In order to applying the Farkas Lemma (Fact B.3), we write the linear system as $A\mathbf{w} = b$, where

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ q_1^{k,m}(\mathbf{v}_1) & q_1^{k,m}(\mathbf{v}_2) & \cdots & q_1^{k,m}(\mathbf{v}_r) \\ \vdots & \vdots & \ddots & \vdots \\ q_{s_{k,m}}^{k,m}(\mathbf{v}_1) & q_{s_{k,m}}^{k,m}(\mathbf{v}_2) & \cdots & q_{s_{k,m}}^{k,m}(\mathbf{v}_r) \end{bmatrix}, \mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_r \end{bmatrix}, \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

By Fact B.3, the original linear system is infeasible if and only if there exists a vector $\mathbf{u} = [u_0, u_1, \ldots, u_{s_{k,d}}]^\intercal$, $\mathbf{u}^\intercal \mathbf{A} \geq 0$ and $\mathbf{u}^\intercal \mathbf{b} < 0$, which is equivalent to $u_0 + \sum_{j=1}^{s_{k,m}} u_j q_j^{k,m}(\mathbf{v}_\ell) \geq 0, \forall 1 \leq \ell \leq r$ and $u_0 < 0$. Let $q^*(\mathbf{v}) = \sum_{j=1}^{s_{k,m}} u_j q_j^{k,m}(\mathbf{v}), \mathbf{v} \in \mathbb{R}^m$, which is an odd polynomial of degree less than k. By our definition of q^*, we have that $q^*(\mathbf{v}_\ell) = \sum_{j=1}^{s_{k,m}} u_j q_j^{k,m}(\mathbf{v}_\ell) \geq -u_0 > 0, \forall 1 \leq \ell \leq r$, which contradicts to our assumption that *there does not exist any odd polynomial q of degree less than k such that $q(\mathbf{v}_\ell) > 0, \forall 1 \leq \ell \leq r$*. This completes the proof.

B.3 Proof of Lemma 3.8

We denote by $G(\mathbf{x})$ to be the standard Gaussian density. By definition, we have that

$$d_{\mathrm{TV}}(D_{\mathbf{U}}, D_0) = (1/2) \int_{\mathbf{x} \in \mathbb{R}^n} \sum_{y \in \{\pm 1\}} |D_{\mathbf{U}}(\mathbf{x}, y) - D_0(\mathbf{x}, y)| d\mathbf{x}$$

$$= (1/2) \int_{\mathbf{x} \in \mathbb{R}^n} G(\mathbf{x}) \sum_{y \in \{\pm 1\}} \left| \sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\intercal \mathbf{U} \mathbf{x}) = y] - (1/2) \right| d\mathbf{x}$$

$$= (1/2) \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n} \left[\sum_{y \in \{\pm 1\}} \left| \sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\intercal \mathbf{U} \mathbf{x}) = y] - (1/2) \right| \right]$$

$$= (1/2) \sum_{y \in \{\pm 1\}} \mathbf{E}_{\mathbf{x} \sim \mathcal{N}_n} \left[\left| \sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\intercal \mathbf{U} \mathbf{x}) = y] - (1/2) \right| \right].$$

Therefore, it suffices to show that

$$\mathbf{E}_{\mathbf{x}\sim\mathcal{N}_n}\left[\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{x}) = y] - (1/2)\right|\right] \geq \Omega(\Delta/r), \quad \forall y \in \{\pm 1\}.$$

We assume that $w_{\ell_0} \geq 1/r$ for some $\ell_0 \in [r]$. Let $\mathbf{v}^*$ be an arbitrary vector satisfying $\mathbf{v}_{\ell_0}^\mathsf{T}\mathbf{v}^* = 0$. We denote by

$$\mathcal{X}_1 = \{\mathbf{x} \in \mathbb{R}^m \mid \mathrm{sign}(\mathbf{v}_{\ell_0}^\mathsf{T}\mathbf{x}) > 0, \mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}) = \mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{v}^*), \ell \in [r] \setminus \{\ell_0\}\},$$
$$\mathcal{X}_2 = \{\mathbf{x} \in \mathbb{R}^m \mid \mathrm{sign}(\mathbf{v}_{\ell_0}^\mathsf{T}\mathbf{x}) < 0, \mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}) = \mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{v}^*), \ell \in [r] \setminus \{\ell_0\}\}.$$

Roughly speaking, $\mathcal{X}_1$ and $\mathcal{X}_2$ denote the subsets of vectors which are very close to the boundary of the halfspace with direction $\mathbf{v}_{\ell_0}$ and maintain the same label with the boundary for the other halfspaces. By definition, for any $\mathbf{x}_1 \in \mathcal{X}_1, \mathbf{x}_2 \in \mathcal{X}_2$, we have that

$$\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}_1) = y] - \sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}_2) = y]\right| = w_{\ell_0} \geq 1/r.$$

Therefore, we have either

$$\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}_1) = y] - (1/2)\right| \geq 1/2r, \quad \forall \mathbf{x}_1 \in \mathcal{X}_1,$$

or

$$\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{x}_2) = y] - (1/2)\right| \geq 1/2r, \quad \forall \mathbf{x}_2 \in \mathcal{X}_2.$$

Since $\mathbf{U}\mathbf{x}$ is a standard Gaussian for any $\mathbf{U}\mathbf{U}^\mathsf{T} = \mathbf{I}_m$ and $\|\mathbf{v}_i + \mathbf{v}_j\|_2, \|\mathbf{v}_i - \mathbf{v}_j\|_2 \geq \Omega(\Delta), 1 \leq i < j \leq r$, we have that for $y \in \{\pm 1\}$,

$$\mathbf{E}_{\mathbf{x}\sim\mathcal{N}_n}\left[\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{x}) = y] - (1/2)\right|\right]$$

$$\geq \mathbf{Pr}_{\mathbf{x}\sim\mathcal{N}_n}[\mathbf{U}\mathbf{x} \in \mathcal{X}_1] \cdot \mathbf{E}_{\mathbf{x}\sim\mathcal{N}_n}\left[\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{x}) = y] - (1/2)\right| \mid \mathbf{U}\mathbf{x} \in \mathcal{X}_1\right]$$

$$+ \mathbf{Pr}_{\mathbf{x}\sim\mathcal{N}_n}[\mathbf{U}\mathbf{x} \in \mathcal{X}_2] \cdot \mathbf{E}_{\mathbf{x}\sim\mathcal{N}_n}\left[\left|\sum_{\ell=1}^{r} w_\ell \mathbb{I}[\mathrm{sign}(\mathbf{v}_\ell^\mathsf{T}\mathbf{U}\mathbf{x}) = y] - (1/2)\right| \mid \mathbf{U}\mathbf{x} \in \mathcal{X}_2\right]$$

$$\geq \Omega(\Delta/r).$$

C   Omitted Proofs from Section 4

We start by introducing the following technical result which provides a universal upper bound for the L_2^2-norm of the high-order gradient of any homogeneous polynomial $p \in \Omega_t^d$.

Lemma C.1. *For any $p \in \Omega_t^d$ and any $1 \leq j \leq t$, we have that*

$$\sup_{\|\mathbf{x}\|_2=1} \left\|\frac{\partial^j p(\mathbf{y})}{\partial \mathbf{y}^j}\right\|_2^2 \leq t^j (d + 2t - 2)^j N_{2(t-j),d} \|p\|_2^2.$$

Proof. Note that $\|\nabla p(\mathbf{x})\|_2^2 \in \Omega_{2(t-1)}^d$, by Fact 4.2, we have that

$$\sup_{\|\mathbf{x}\|_2=1} \|\nabla p(\mathbf{x})\|_2^2 \le \sqrt{N_{2(t-1),d}}\sqrt{\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^4]} \le \sqrt{N_{2(t-1),d}}\sqrt{\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2]}\sqrt{\sup_{\|\mathbf{x}\|_2=1}\|\nabla p(\mathbf{x})\|_2^2},$$

which implies that $\sup_{\|\mathbf{x}\|_2=1}\|\nabla p(\mathbf{x})\|_2^2 \le N_{2(t-1),d}\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2] \le t(d+2t-2)N_{2(t-1),d}\|p\|_2^2$.

Since $\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_2^2 \le \left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2$, it suffices to obtain an upper bound for $\sup_{\|\mathbf{x}\|_2=1}\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2$.
Noting that $\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2 \in \Omega_{2(t-j)}^d$, by Fact 4.2, we have that

$$\sup_{\|\mathbf{x}\|_2=1}\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2 \le \sqrt{N_{2(t-j),d}}\sqrt{\mathbf{E}\left[\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^4\right]}$$

$$\le \sqrt{N_{2(t-j),d}}\sqrt{\mathbf{E}\left[\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2\right]}\sqrt{\sup_{\|\mathbf{x}\|_2=1}\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2},$$

which implies that $\sup_{\mathbf{x}\in\mathbb{S}^{d-1}}\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2 \le N_{2(t-j),d}\mathbf{E}\left[\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2\right]$. Noting that $\frac{\partial p(\mathbf{x})}{\partial x_i} \in \Omega_{t-1}^d$,
by Lemma 4.4, we have that

$$\mathbf{E}\left[\left\|\frac{\partial^2 p(\mathbf{x})}{\partial \mathbf{x}^2}\right\|_F^2\right] = \mathbf{E}\left[\sum_{i_1,i_2\in[d]}\left(\frac{\partial^2 p(\mathbf{x})}{\partial x_{i_1}\partial x_{i_2}}\right)^2\right] = \sum_{i_1=1}^d \mathbf{E}\left[\sum_{i_2=1}^d\left(\frac{\partial}{\partial x_{i_2}}\left(\frac{\partial p(\mathbf{x})}{\partial x_{i_1}}\right)\right)^2\right]$$

$$\le (t-1)(d+2t-4)\sum_{i_1=1}^d \mathbf{E}\left[\left(\frac{\partial p(\mathbf{x})}{\partial x_{i_1}}\right)^2\right] \le t(d+2t-2)\mathbf{E}[\|\nabla p(\mathbf{x})\|_2^2] \le t^2(d+2t-2)^2\|p\|_2^2.$$

In general, noting that $\frac{\partial^{j-1}p(\mathbf{x})}{\partial x_{i_1}\cdots\partial x_{i_{j-1}}} \in \Omega_{t-j+1}^d$, by Lemma 4.4, we have that

$$\mathbf{E}\left[\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2\right] = \mathbf{E}\left[\sum_{i_1,\ldots,i_j\in[d]}\left(\frac{\partial^2 p(\mathbf{x})}{\partial x_{i_1}\ldots\partial x_{i_j}}\right)^2\right]$$

$$= \sum_{i_1,\ldots,i_{j-1}\in[d]}\mathbf{E}\left[\sum_{i_j=1}^d\left(\frac{\partial}{\partial x_{i_j}}\left(\frac{\partial^{j-1}p(\mathbf{x})}{\partial x_{i_1}\ldots x_{i_{j-1}}}\right)\right)^2\right]$$

$$\le (t-j+1)(d+2(t-j))\sum_{i_1,\ldots,i_{j-1}\in[d]}\mathbf{E}\left[\left(\frac{\partial^{j-1}p(\mathbf{x})}{\partial x_{i_1}\ldots x_{i_{j-1}}}\right)^2\right]$$

$$\le t(d+2t-2)\mathbf{E}\left[\left\|\frac{\partial^{j-1}p(\mathbf{x})}{\partial \mathbf{x}^{j-1}}\right\|_F^2\right] \le t^j(d+2t-2)^j\|p\|_2^2.$$

Therefore, we have that

$$\sup_{\|\mathbf{x}\|_2=1}\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2 \le N_{2(t-j),d}\mathbf{E}\left[\left\|\frac{\partial^j p(\mathbf{x})}{\partial \mathbf{x}^j}\right\|_F^2\right] \le t^j(d+2t-2)^j N_{2(t-j),d}\|p\|_2^2.$$

This completes the proof. $\qquad\square$

C.1   Proof of Lemma 4.7

Let $p_1, \ldots, p_N \in \Omega$ be an orthonormal basis, i.e., $\mathbf{E}[p_i(\mathbf{x})p_j(\mathbf{x})] = \mathbb{I}[i = j]$. Let vector $\mathbf{p}(\mathbf{x}) \overset{\text{def}}{=} [p_1(\mathbf{x}), \ldots, p_N(\mathbf{x})]$. We have that $\mathbf{E}[\mathbf{p}(\mathbf{x})] = 0$ and $\mathbf{Cov}[\mathbf{p}(\mathbf{x})] = \mathbf{I}_N$.

$$\mathbf{Pr}\left[\left\|\frac{1}{r}\sum_{i=1}^{r}\mathbf{p}(\mathbf{x}_i)\right\|_2 \geq \eta\right] = \mathbf{Pr}\left[\frac{1}{r^2}\left\|\sum_{i=1}^{r}\mathbf{p}(\mathbf{x}_i)\right\|_2^2 \geq \eta^2\right] = \mathbf{Pr}\left[\frac{1}{r^2}\sum_{j=1}^{N}\left(\sum_{i=1}^{r}p_j(\mathbf{x}_i)\right)^2 \geq \eta^2\right]$$

$$\leq \frac{1}{\eta^2 r^2}\sum_{j=1}^{N}\mathbf{E}\left[\left(\sum_{i=1}^{r}p_j(\mathbf{x}_i)\right)^2\right] = \frac{N}{r\eta^2}.$$

We now assume that $\frac{1}{r}\|\sum_{i=1}^{r}\mathbf{p}(\mathbf{x}_i)\|_2 \leq \eta$. Let $p \in \Omega$ be an arbitrary polynomial. We can write $p(\mathbf{x}) = \sum_{j=1}^{N}\alpha_j p_j(\mathbf{x})$, where $\|p\|_2^2 = \sum_{j=1}^{N}\alpha_j^2$. We have that

$$\frac{1}{r}\left|\sum_{i=1}^{r}p(\mathbf{x}_i)\right| = \frac{1}{r}\left|\sum_{i=1}^{r}\sum_{j=1}^{N}\alpha_j p_j(\mathbf{x}_i)\right| \leq \frac{1}{r}\sum_{j=1}^{N}|\alpha_j|\left|\sum_{i=1}^{r}p_j(\mathbf{x}_i)\right|$$

$$\leq \frac{1}{r}\sqrt{\sum_{j=1}^{N}\alpha_j^2}\sqrt{\sum_{j=1}^{N}\left(\sum_{i=1}^{r}p_j(\mathbf{x}_i)\right)^2} = \frac{\|p\|_2}{r}\left\|\sum_{i=1}^{r}\mathbf{p}(\mathbf{x}_i)\right\|_2 \leq \eta\|p\|_2,$$

where the second inequality follows from Cauchy-Schwarz. This completes the proof.

C.2   Proof of Lemma 4.8

By definition of $\nabla_o p(\mathbf{y})$, we have that

$$p(\mathbf{z}) - p(\mathbf{y}) = \frac{p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y}))}{\|\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})\|_2^t} - p(\mathbf{y})$$

$$= \frac{p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y})}{(1 + \delta^2\|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}} - \left(1 - \frac{1}{(1 + \delta^2\|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}}\right)p(\mathbf{y})$$

$$\geq \frac{p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y})}{(1 + \delta^2\|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}} - \left(1 - \exp(-t\delta^2\|\nabla_o p(\mathbf{y})\|_2^2/2)\right)|p(\mathbf{y})|$$

$$\geq \frac{p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y})}{(1 + \delta^2\|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}} - t\delta^2\|\nabla_o p(\mathbf{y})\|_2^2|p(\mathbf{y})|/2 \;.$$

We bound $p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y})$ as follows: Let $f(s) = p(\mathbf{y} + s\mathbf{v})$ for some unit vector $v \in \mathbb{R}^d$. Noting that p is a degree-t homogeneous polynomial, by Taylor expansion, we have that $f(s) = f(0) + \sum_{j=1}^{t}\frac{f^{(j)}(0)s^j}{j!}$. By elementary calculation, we have that $f'(0) = \mathbf{v}^\intercal\nabla p(\mathbf{y}), f''(0) = \mathbf{v}^\intercal\frac{\partial^2 p(\mathbf{y})}{\partial\mathbf{y}^2}\mathbf{v}, \ldots, f^{(t)}(0) = \left\langle\mathbf{v}^{\otimes t}, \frac{\partial^t p(\mathbf{y})}{\partial\mathbf{y}^t}\right\rangle$. By taking $\mathbf{v}$ to be the direction of $\nabla_o p(\mathbf{y})$, i.e., $\mathbf{v} = \frac{\nabla_o p(\mathbf{y})}{\|\nabla_o p(\mathbf{y})\|_2}$, we have that

$$p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y}) = f(\delta\|\nabla_o p(\mathbf{y})\|_2) - f(0) = \sum_{j=1}^{t}\frac{\left\langle\nabla_o p(\mathbf{y})^{\otimes j}, \frac{\partial^j p(\mathbf{y})}{\partial\mathbf{y}^j}\right\rangle\delta^j}{j!} \;.$$

Noting that the first order term is $\delta\|\nabla_o p(\mathbf{y})\|_2^2$, it suffices to show that the absolute value of $\sum_{j=2}^t \frac{\left\langle \nabla_o p(\mathbf{y})^{\otimes j}, \frac{\partial^j p(\mathbf{y})}{\partial \mathbf{y}^j} \right\rangle \delta^j}{j!}$ is sufficiently small. Applying Lemma C.1 yields

$$\left| \sum_{j=2}^t \frac{\left\langle \nabla_o p(\mathbf{y})^{\otimes j}, \nabla^j p(\mathbf{y}) \right\rangle \delta^j}{j!} \right| \leq \sum_{j=2}^t \frac{\delta^j \|\nabla_o p(\mathbf{y})\|_2^j \left\| \frac{\partial^j p(\mathbf{y})}{\partial \mathbf{y}^j} \right\|_2}{j!}$$

$$= \delta \|\nabla_o p(\mathbf{y})\|_2^2 \sum_{j=2}^t \frac{\delta^{j-1} \|\nabla_o p(\mathbf{y})\|_2^{j-2} \left\| \frac{\partial^j p(\mathbf{y})}{\partial \mathbf{y}^j} \right\|_2}{j!}$$

$$\leq \delta \|\nabla_o p(\mathbf{y})\|_2^2 \left(\sum_{j=2}^t \frac{\delta^{j-1} \|\nabla p(\mathbf{y})\|_2^{2j-4}}{2j!} + \sum_{j=2}^t \frac{\delta^{j-1} \left\| \frac{\partial^j p(\mathbf{y})}{\partial \mathbf{y}^j} \right\|_2^2}{2j!} \right)$$

$$\leq \delta \|\nabla_o p(\mathbf{y})\|_2^2 \left(\sum_{j=2}^t \frac{\delta^{j-1} (t(d + 2t - 2) N_{2(t-1),d})^{j-2}}{2j!} + \sum_{j=2}^t \frac{\delta^{j-1} t^j (d + 2t - 2)^j N_{2(t-j),d}}{2j!} \right) .$$

Therefore, we will have that $p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y}) \geq C' \delta \|\nabla_o p(\mathbf{y})\|_2^2$ for some universal constant $0 < C' < 1$, as long as $\delta \leq 1/N_{2t,d}^2$. Thus, by Lemma 4.4, we have that

$$p(\mathbf{z}) - p(\mathbf{y}) \geq \frac{p(\mathbf{y} + \delta \cdot \nabla_o p(\mathbf{y})) - p(\mathbf{y})}{(1 + \delta^2 \|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}} - t\delta^2 \|\nabla_o p(\mathbf{y})\|_2^2 |p(\mathbf{y})|/2$$

$$= \frac{C'\delta \|\nabla_o p(\mathbf{y})\|_2^2}{(1 + \delta^2 \|\nabla_o p(\mathbf{y})\|_2^2)^{t/2}} - t\delta^2 \|\nabla_o p(\mathbf{y})\|_2^2 |p(\mathbf{y})|/2$$

$$= C'\delta \|\nabla_o p(\mathbf{y})\|_2^2 \exp(-t\delta^2 \|\nabla_o p(\mathbf{y})\|_2^2/2) - t\delta^2 \|\nabla_o p(\mathbf{y})\|_2^2 |p(\mathbf{y})|/2$$

$$\geq C'\delta \|\nabla_o p(\mathbf{y})\|_2^2 \left(1 - t\delta^2 \|\nabla p(\mathbf{y})\|_2^2/2 - t\delta |p(\mathbf{y})|/2C' \right)$$

$$\geq \delta \|\nabla_o p(\mathbf{y})\|_2^2 \left(C'(1 - t^2 \delta^2 (d + 2t - 2) N_{2(t-1),d}/2) - t\delta \sqrt{N_{t,d}}/2 \right)$$

$$\geq C\delta \|\nabla_o p(\mathbf{y})\|_2^2 ,$$

for some universal constant $0 < C < 1$, as long as $\delta \leq 1/N_{2t,d}^2$. This completes the proof.

C.3 Proof of Theorem 4.9

Let $p \in \partial \Omega_t^d$. Since

$$\|\nabla p(\mathbf{y})\|_2^2 = \|\nabla_o p(\mathbf{y})\|_2^2 + \langle \mathbf{y}, \nabla p(\mathbf{y}) \rangle^2 = \|\nabla_o p(\mathbf{y})\|_2^2 + t^2 p(\mathbf{y})^2 ,$$

by Lemma 4.8, we have that $p(\mathbf{z}) \geq p(\mathbf{y}) + C\delta(\|\nabla p(\mathbf{y})\|_2^2 - t^2 p(\mathbf{y})^2)$. Let

$$q(\mathbf{y}) = p(\mathbf{y}) + C\delta \|\nabla_o p(\mathbf{y})\|_2^2 = p(\mathbf{y}) + C\delta(\|\nabla p(\mathbf{y})\|_2^2 - t^2 p(\mathbf{y})^2) .$$

By definition, we have that $q(\mathbf{y}) - \mathbf{E}[q(\mathbf{y})] = p(\mathbf{y}) + C\delta(\|\nabla p(\mathbf{y})\|_2^2 - t^2 p(\mathbf{y})^2) - C\delta \mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]$, which is a polynomial of degree at most $2t$ and contains only monomials of degree $2t, 2t-2, t, 0$. Let Ω be the subspace of polynomials in d-variables containing all monomials of degree $2t, 2t - 2, t, 0$. In this way, the dimension of Ω is

$$N = \binom{d + 2t - 1}{d - 1} + \binom{d + 2t - 3}{d - 1} + \binom{d + t - 1}{d - 1} + 1 \leq 3N_{2t,d} .$$

24

Applying [Lemma 4.7](#) yields that with probability at least $1-\frac{N}{r\eta^2}$, we have that $\left|\frac{1}{r}\sum_{i=1}^r q(\mathbf{y}_i) - \mathbf{E}[q(\mathbf{y})]\right| \le \eta\|q(\mathbf{y}) - \mathbf{E}[q(\mathbf{y})]\|_2, \forall q \in \Omega$. Therefore, we have that

$$\frac{1}{r}\sum_{i=1}^r p(\mathbf{z}_i) \ge \frac{1}{r}\sum_{i=1}^r q(\mathbf{y}_i) \ge \mathbf{E}[q(\mathbf{y})] - \eta\|q(\mathbf{y}) - \mathbf{E}[q(\mathbf{y})]\|_2 = \mathbf{E}[q(\mathbf{y})] - \eta\sqrt{\mathbf{E}[q(\mathbf{y})^2] - \mathbf{E}[q(\mathbf{y})]^2}.$$

By elementary calculation, we have that

$$\mathbf{E}[q(\mathbf{y})^2] - \mathbf{E}[q(\mathbf{y})]^2 = \mathbf{E}[(p(\mathbf{y}) + C\delta\|\nabla_o p(\mathbf{y})\|_2^2)^2] - C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]^2$$
$$= \mathbf{E}[p(\mathbf{y})^2] + 2C\delta\mathbf{E}[p(\mathbf{y})\|\nabla_o p(\mathbf{y})\|_2^2] + C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^4] - C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]^2$$
$$= \mathbf{E}[p(\mathbf{y})^2] + C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^4] - C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]^2$$
$$= 1 + C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^4] - C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]^2,$$

where the second equality is due to $p(\mathbf{y})$ being odd and

$$\|\nabla_o p(-\mathbf{y})\|_2^2 = \|\nabla p(-\mathbf{y})\|_2^2 - t^2 p(-\mathbf{y})^2 = \|\nabla(-p(\mathbf{y}))\|_2^2 - t^2(-p(\mathbf{y}))^2$$
$$= \|\nabla p(\mathbf{y})\|_2^2 - t^2 p(\mathbf{y})^2 = \|\nabla_o p(\mathbf{y})\|_2^2.$$

By [Lemma 4.4](#) and [Lemma C.1](#), we have that $\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2] \ge d-1$ and

$$\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^4] \le \mathbf{E}[\|\nabla p(\mathbf{y})\|_2^2] \sup_{\|\mathbf{y}\|_2=1}\|\nabla p(\mathbf{y})\|_2^2 \le t^2(d+2t-2)^2 N_{2(t-1),d}.$$

Therefore, we have that

$$\frac{1}{r}\sum_{i=1}^r p(\mathbf{z}_i) \ge \mathbf{E}[q(\mathbf{y})] - \eta\sqrt{\mathbf{E}[q(\mathbf{y})^2] - \mathbf{E}[q(\mathbf{y})]^2}$$

$$\ge C\delta\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2] - \eta\sqrt{1 + C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^4] - C^2\delta^2\mathbf{E}[\|\nabla_o p(\mathbf{y})\|_2^2]^2}$$

$$\ge C\delta(d-1) - \eta\sqrt{1 + C^2\delta^2(t^2(d+2t-2)^2 N_{2(t-1),d} - (d-1)^2)}$$

$$= C\delta\left(d-1-\eta\sqrt{\frac{1}{C^2\delta^2} + t^2(d+2t-2)^2 N_{2(t-1),d} - (d-1)^2}\right).$$

Taking $\delta = 1/N_{2t,d}^2$ and $\eta = \frac{Cd}{3N_{2t,d}^2}$ yields that with probability at least

$$1 - \frac{N}{r\eta^2} \ge 1 - \frac{27}{C^2 d^2} \ge 99/100\,,$$

$$\frac{1}{r}\sum_{i=1}^r p(\mathbf{z}_i) \ge C\delta\left(d-1-\eta\sqrt{N_{2t,d}^4/C^2 + t^2(d+2t-2)^2 N_{2(t-1),d} - (d-1)^2}\right)$$
$$> C\delta\left(d/2 - \eta\sqrt{2N_{2t,d}^4/C^2}\right) \ge 0\,.$$