

Approximation Schemes for ReLU Regression

Ilias Diakonikolas

University of Wisconsin, Madison

ILIAS@CS.WISC.EDU

Surbhi Goel

Sushrut Karmalkar

Adam R. Klivans

University of Texas at Austin

SURBHI@CS.UTEXAS.EDU

SUSRUTK@CS.UTEXAS.EDU

KLIVANS@CS.UTEXAS.EDU

Mahdi Soltanolkotabi

University of Southern California

SOLTANOL@USC.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We consider the fundamental problem of ReLU regression, where the goal is to output the best fitting ReLU with respect to square loss given access to draws from some unknown distribution. We give the first efficient, constant-factor approximation algorithm for this problem assuming the underlying distribution satisfies some weak concentration and anti-concentration conditions (and includes, for example, all log-concave distributions). This solves the main open problem of Goel et al., who proved hardness results for any exact algorithm for ReLU regression (up to an additive ϵ). Using more sophisticated techniques, we can improve our results and obtain a polynomial-time approximation scheme for any subgaussian distribution. Given the aforementioned hardness results, these guarantees can not be substantially improved.

Our main insight is a new characterization of *surrogate losses* for nonconvex activations. While prior work had established the existence of convex surrogates for monotone activations, we show that properties of the underlying distribution actually induce strong convexity for the loss, allowing us to relate the global minimum to the activation’s *Chow parameters*.

1. Introduction

Finding the best-fitting ReLU with respect to square-loss – also called “ReLU Regression” – is a fundamental primitive in the theory of neural networks. Many authors have recently studied the problem both in terms of finding algorithms that succeed under various assumptions and proving hardness results (Manurangsi and Reichman, 2018; Soltanolkotabi, 2017; Goel et al., 2019; Yehudai and Shamir, 2019; Goel et al., 2017; Manurangsi and Reichman, 2018). In this work, we consider the agnostic model of learning where *no assumptions* are made on the noise.

Recall the ReLU function $\text{ReLU}_w : \mathbb{R}^d \rightarrow \mathbb{R}$ parameterized by w is defined as $\text{ReLU}_w(x) := \text{ReLU}(\langle w, x \rangle) = \max\{0, \langle w, x \rangle\}$ (for simplicity, let $\|w\|_2 \leq 1$). Given samples (x, y) drawn from a distribution \mathcal{D} over $\mathbb{R}^d \times \mathbb{R}$, the objective of the learner is to find a hypothesis $h : \mathbb{R}^d \rightarrow \mathbb{R}$ that has square loss at most $\text{opt} + \epsilon$, where $\text{opt} \ll 1$ is defined to be the loss of the best fitting ReLU, i.e.,

$$\text{opt} := \min_{w \in \mathbb{R}^d} \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle w, x \rangle) - y)^2] .$$

There are several hardness results known for this problem. A recent result shows that finding a hypothesis achieving a loss of $O(\text{opt}) + \epsilon$ is NP-hard when there are no distributional assumptions

on $\mathcal{D}_{\mathcal{X}}$, the marginal of \mathcal{D} on the examples (Manurangsi and Reichman, 2018). Recent work due to Goel et al. (2019) gives hardness results for achieving error $\text{opt} + \epsilon$, even if the underlying distribution is the standard Gaussian. This work also provides an algorithm that achieves error $O(\text{opt}^{2/3}) + \epsilon$ under the assumption that $\mathcal{D}_{\mathcal{X}}$ is log-concave. The main problem open problem posed by Goel et al. (2019) is the following:

Question 1 *For the problem of ReLU regression, is it possible to recover a hypothesis achieving error of $O(\text{opt}) + \epsilon$ in time $\text{poly}(d, 1/\epsilon)$?*

In this paper we answer this question in the affirmative. Specifically, we show that there is a fully polynomial time algorithm which can recover a vector w such that the loss of the corresponding ReLU function, ReLU_w , is at most $O(\text{opt}) + \epsilon$. More formally, we prove the following:

Theorem 1 *If $\mathcal{D}_{\mathcal{X}}$ is isotropic log-concave, there is an algorithm that takes $\tilde{O}(d/\epsilon^2)$ samples and runs in time $\tilde{O}(d^2/\epsilon^2)$ and returns a vector w such that ReLU_w has square loss $O(\text{opt}) + \epsilon$ with high probability.*

The sample complexity of our algorithm is nearly linear in the problem dimension and hence information-theoretically optimal up to logarithmic factors. To establish this near-optimal sample complexity, we leverage intricate tools involving uniform *one-sided* concentration of empirical processes of log-concave distributions.

Additionally, we show that under stronger distributional assumptions and if the algorithm is allowed to be improper, i.e., if the hypothesis need not be the ReLU of a linear function, then it is possible to return a hypothesis that achieves a loss of $(1 + \eta) \cdot \text{opt} + \epsilon$ in polynomial time for any constant $\eta > 0$ as long as $\text{opt} < 1$.

Theorem 2 *If $\mathcal{D}_{\mathcal{X}}$ is ν -subgaussian for $\nu \leq O(1)$, then for any constant $\eta > 0$, there is an algorithm with sample complexity and running time $O\left(\frac{1}{\epsilon^2} \cdot \left(\frac{d}{\eta^3 \nu^2}\right)^{1/\eta^3}\right)$ that outputs a hypothesis $h : \mathbb{R}^d \rightarrow \mathbb{R}$ whose square loss is at most $(1 + \eta) \cdot \text{opt} + \epsilon$ with high probability.*

Given the hardness results of Goel et al. (2019), the aforementioned accuracy guarantees are essentially best-possible.

1.1. Our Approach

A major barrier to minimizing the square loss for the ReLU regression problem is that it is nonconvex. In such settings, gradient descent-based algorithms can potentially fail due to the presence of poor local minima. In the case of ReLU regression, the number of these bad local minima for the square loss can be as large as exponential in the dimension (Auer et al., 1996).

Despite this fact, for well-structured noise models, it is possible to learn a ReLU with respect to square loss by applying results on isotonic regression (Kalai and Sastry, 2009; Kakade et al., 2011; Klivans and Meka, 2017). These results show that if the noise is bounded and has zero mean, it is possible to learn conditional mean functions of the form $\sigma_w : x \mapsto \sigma(\langle w, x \rangle)$ where σ is a monotone and Lipschitz activation. This is proven via an analysis similar to that of the perceptron algorithm. It is not clear, however, how to extend these results to harder noise models.

In retrospect, one way to interpret the algorithms from [Kalai and Sastry \(2009\)](#) and [Kakade et al. \(2011\)](#) is to view them as implicitly minimizing a *surrogate loss*¹. The intuition is as follows: although a monotone and Lipschitz function need not be convex, it is not difficult to see that its *integral* is convex. This motivates the following definition of a surrogate loss:

$$L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\int_0^{\langle \mathbf{w}, \mathbf{x} \rangle} (\sigma(a) - y) da \right].$$

Properties of this loss were explored early on in the work of [Auer et al. \(1996\)](#) who gave a formal proof that the loss is convex (a succinct write-up of properties of this loss can also be found in notes due to [Kanade \(2018\)](#)). Thus, we can efficiently minimize this loss using gradient descent. What is more subtle is the relationship of the minima of the surrogate loss to the minima of the original square-loss.

The main insight of the current work is that algorithms that directly minimize this surrogate loss have strong noise-tolerance properties if the underlying marginal distribution satisfies some mild conditions. As a consequence, we prove that the GLMtron algorithm of [Kakade et al. \(2011\)](#) (or equivalently projected gradient descent on the surrogate loss) achieves a constant-factor approximation for ReLU regression. The proof of this relies on three key structural observations:

- The first insight concerns the notion of the *Chow parameters* of a function. The Chow parameters $\chi_{\mathcal{D}}^f$ of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with respect to a distribution \mathcal{D} are defined to be the first moments of f with respect to $\mathcal{D}_{\mathcal{X}}$, i.e., $\chi_{\mathcal{D}}^f := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}} [f(\langle \mathbf{w}, \mathbf{x} \rangle) \mathbf{x}]$. We show that the Chow parameters of a strictly monotone and Lipschitz activation function σ *robustly characterize* the function, i.e., two functions with approximately the same Chow parameters have approximately the same loss. More precisely, any \mathbf{w} that satisfies $\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} = \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot \mathbf{x}]$ induces a concept $\sigma_{\mathbf{w}}$ with square loss $O(\text{opt})$.
- The second observation is that the gradient of the surrogate loss at \mathbf{w} is the difference between the Chow parameters of $\sigma_{\mathbf{w}}$ and the first moments of the labels, $\chi_{\mathcal{D}} := \mathbb{E}_{(x,y) \sim \mathcal{D}} [y \cdot \mathbf{x}]$, i.e.,

$$\nabla_{\mathbf{w}} L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} - \chi_{\mathcal{D}}.$$

- The third insight is that if the underlying distribution $\mathcal{D}_{\mathcal{X}}$ satisfies some concentration and anti-concentration properties (satisfied, for instance, by log-concave distributions), then the surrogate loss is *strongly* convex. In particular, this holds for any activation that is strictly monotone and 1-Lipschitz, including ReLUs.

Any strongly convex function achieves its minimum at a point where the gradient is zero. The first two observations now imply that the point where the surrogate loss has zero gradient corresponds to a weight vector achieving a loss of $O(\text{opt}) + \epsilon$.

A naive analysis for the concentration of empirical gradients results in a sample complexity of roughly $O(d^4)$. To achieve the near-linear sample complexity of $O(d \text{ polylog}(d))$ in [Theorem 1](#), we show that while the gradient is not uniformly concentrated in all directions, it does concentrate from below in the direction going from the current estimate to the minimizer of the loss.

1. The analysis of [Kalai and Sastry \(2009\)](#) and [Kakade et al. \(2011\)](#) works directly with square loss and does not use the existence of a surrogate loss for its analysis.

Theorem 1 achieves a constant factor approximation to the ReLU regression problem when the underlying distribution is log-concave. It is not clear how to show that minimizing the surrogate loss alone can go beyond a constant factor approximation. Still, it turns out that under a slightly stronger distributional assumption on $\mathcal{D}_{\mathcal{X}}$ (sub-gaussianity), we can give a polynomial-time approximation scheme (PTAS) for ReLU regression.

To achieve this, we build on the localization framework used to solve the problem of learning halfspaces under various noise models (Daniely, 2015; Awasthi et al., 2017). The problem of learning halfspaces, however, differs from the problem of ReLU regression. One crucial difference is that for the problem of learning halfspaces, the agnostic noise model is equivalent to the noise model where an opt fraction of the labels are corrupted. In the case of ReLU regression, every point’s label can potentially be corrupted.

Our approach broadly proceeds in two stages:

- First, we use our constant-factor approximation algorithm to recover a vector \mathbf{w} satisfying $\|\mathbf{w} - \mathbf{w}^*\|_2^2 \leq O(\text{opt})$, where \mathbf{w}^* is the vector achieving an error of opt. We use this to partition the space into three regions for a certain choice of a parameter t . Our three regions are $T = \{\mathbf{u} \in \mathbb{R}^d : |\langle \mathbf{w}, \mathbf{u} \rangle| \leq t\}$, $T_+ = \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle > t\}$, and $T_- = \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle < -t\}$.
- In each of these regions we find functions whose loss competes with that of the best fitting ReLU (i.e., $\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle)$).

Observe that $\text{ReLU}_{\mathbf{w}^*}(\mathbf{x})$ takes the value $\langle \mathbf{w}^*, \mathbf{x} \rangle$ for most of the region T_+ . Intuitively, the best-fitting linear function \mathbf{w}_+ must achieve a loss comparable to $\text{ReLU}_{\mathbf{w}^*}(\mathbf{x})$ for T_+ . Similar reasoning shows that for the region T_- , 0 is a good hypothesis. Using results from approximation theory, we show that the function $\text{ReLU}_{\mathbf{w}^*}(\mathbf{x})$ in the region T is closely approximated by a polynomial of degree $O\left(\frac{1}{\eta^3}\right)$. To find a function which achieves a comparable loss to the concept, we perform polynomial regression to find the best-fitting polynomial of appropriate degree in this region. Finally, our algorithm returns the following hypothesis h .

$$h(\mathbf{x}) = \begin{cases} \langle \mathbf{w}_+, \mathbf{x} \rangle, & \mathbf{x} \in T_+ \\ P(\mathbf{x}), & \mathbf{x} \in T \\ 0, & \mathbf{x} \in T_- \end{cases} .$$

The paper by Daniely (2015) shows this result only for the uniform distribution on the sphere, while our result works for sub-gaussian distributions. The analysis of this algorithm is nontrivial. In particular, in addition to using tools from approximation theory to derive the polynomial approximation, the choice of the parameter t to partition our space is delicate, and we need to calculate approximations with respect to complicated marginal distributions that do not have nice closed-form expressions.

1.2. Prior and Related Work

Here we provide an overview of the most relevant prior work. Goel et al. (2017) give an efficient algorithm for ReLU regression that succeeds with respect to any distribution supported on the unit sphere, but has sample complexity and running time exponential in $1/\epsilon$. Soltanolkotabi (2017) shows that SGD efficiently learns a ReLU in the realizable setting when the underlying distribution is assumed to be the standard Gaussian. Goel et al. (2018) gives a learning algorithm for one

convolutional layer of ReLUs for any symmetric distribution (including Gaussians). [Goel et al. \(2019\)](#) gives an efficient algorithm for ReLU regression with error guarantee of $O(\text{opt}^{2/3}) + \epsilon$.

[Yehudai and Shamir \(2019\)](#) shows that it is hard to learn a single ReLU activation via stochastic gradient descent, when the hypothesis used to learn the ReLU function is of the form $N(\mathbf{x}) := \sum_{i=1}^r u_i f_i(\mathbf{x})$ and the functions $f_i(\mathbf{x})$ are random feature maps drawn from a fixed distribution. In particular, they show that any $N(\mathbf{x})$ which approximates $\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle + b)$ (where $\|\mathbf{w}\|_2 = d^2$ and $b \in \mathbb{R}$) up to a small constant square loss, must have one of the $|u_i|$ being exponentially large in d for some i or have exponentially many random features in the sum (i.e., $r \geq \exp(\Omega(d))$). Their paper makes the point that regression using *random features* cannot learn the ReLU function in polynomial time. Our results use different techniques to learn the unknown ReLU function that are not captured by this model.

We note that Chow parameters have been previously used in the context of learning halfspaces under well-behaved distributions, see, e.g., [O’Donnell and Servedio \(2008\)](#); [De et al. \(2012\)](#); [Diakonikolas et al. \(2019\)](#) and references therein. The technique of localization has been used extensively in the context of learning halfspaces over various structured distributions. Specifically, [Awasthi et al. \(2017\)](#) use this technique to learn origin-centered halfspaces with respect to log-concave distributions in the presence of agnostic noise, obtaining an error guarantee of $O(\text{opt}) + \epsilon$. Subsequently, [Daniely \(2015\)](#) uses an adaptation of the localization technique in conjunction with the polynomial approximation technique from [Kalai et al. \(2005\)](#) to obtain a PTAS for the problem of agnostically learning origin-centered halfspaces under the uniform distribution over the sphere. More recently, [Diakonikolas et al. \(2018\)](#) obtain similar guarantees in the presence of nasty noise, where the halfspace need not be origin-centered.

While the problem of learning halfspaces is related to that of ReLU regression, we stress that for ReLU regression *every* label may be corrupted (possibly by arbitrarily large values), while in the context of learning halfspaces only an opt fraction of the labels are corrupted. This is because the loss for halfspace learning is 0/1 instead of the square-loss. Indeed, a black-box application of the results for halfspace learning in the context ReLU regression results in the suboptimal guarantee of $O(\text{opt}^{2/3})$ ([Goel et al., 2019](#)).

2. Preliminaries

Notation. For $n \in \mathbb{Z}_+$, we denote $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. We will use small boldface characters for vectors. For $\mathbf{x} \in \mathbb{R}^d$, and $i \in [d]$, x_i denotes the i -th coordinate of \mathbf{x} , and $\|\mathbf{x}\|_2 \stackrel{\text{def}}{=} (\sum_{i=1}^d x_i^2)^{1/2}$ denotes the ℓ_2 -norm of \mathbf{x} . We will use $\langle \mathbf{x}, \mathbf{y} \rangle$ for the inner product between $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. We will use $\mathbb{E}[X]$ for the expectation of random variable X and $\Pr[\mathcal{E}]$ for the probability of event \mathcal{E} . For two functions f, g let $f \lesssim g$ mean that there exists a $C > 0$ such that $f(x) \leq Cg(x)$ for all $x > C$ and $f \gtrsim g$ denote $g \lesssim f$. $B(d, W)$ denotes the d -dimensional Euclidean ball at the origin with radius W , that is, $B(d, W) := \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\|_2 \leq W\}$. We say $f = O(g)$ if $f \lesssim g$, also we use \tilde{O} to hide log factors of the input. We will use $\sigma'(\mathbf{x})$ to denote a subgradient of σ at the point \mathbf{x} .

Learning Models. We start by reviewing the PAC learning model [Vapnik \(1982\)](#); [Valiant \(1984\)](#). Let \mathcal{C} be the target (concept) class of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$, \mathcal{H} be a hypothesis class, and $\ell : \mathcal{H} \times \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a loss function. In the (distribution-specific) agnostic PAC model [Haussler \(1992\)](#); [Kearns et al. \(1994\)](#), we are given a multi-set of labeled examples $(\mathbf{x}^{(i)}, y^{(i)})$ that are i.i.d. samples drawn from a distribution $\mathcal{D} = (\mathcal{D}_{\mathcal{X}}, \mathcal{D}_{\mathcal{Y}})$ on $\mathcal{X} \times \mathcal{Y}$, where $\mathbf{x}^{(i)} \sim \mathcal{D}_{\mathbf{x}}$. The marginal distribution

$\mathcal{D}_{\mathbf{x}}$ is assumed to lie in a family of well-behaved distributions. The goal is to find a hypothesis $h \in \mathcal{H}$ that approximately minimizes the expected loss $L_{\mathcal{D}}(h) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\ell(h(\mathbf{x}), y)]$, compared to $\text{opt}_{\mathcal{D}}(\mathcal{C}) := \min_{f \in \mathcal{C}} L_{\mathcal{D}}(f)$. In this paper, we will have $\mathcal{X} = \mathbb{R}^d$, $\mathcal{Y} = \mathbb{R}$, and $\ell(h(\mathbf{x}), y) = (h(\mathbf{x}) - y)^2$. We will focus on constant factor approximation algorithms, that is, we will want a hypothesis which satisfies $L_{\mathcal{D}}(h) \leq C \cdot \text{opt}_{\mathcal{D}}(\mathcal{C}) + \epsilon$ for some universal constant $C > 1$ and $\epsilon \in (0, 1)$. If the hypothesis $h \in \mathcal{C}$ then the learner is *proper* else it is called *improper*.

Problem Setup. We consider the concept class of Generalized Linear Models (GLMs) $\mathcal{C}_{\sigma} := \{\mathbf{x} \rightarrow \sigma(\langle \mathbf{w}, \mathbf{x} \rangle)\}$ for activation functions $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ which are non-decreasing and 1-Lipschitz. Common activations such as ReLU and Sigmoid satisfy this assumption. We use the L_2 -error as our loss function, i.e., $L_{\mathcal{D}}(h) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[(h(\mathbf{x}) - y)^2]$. We overload the definition by setting $L_{\mathcal{D}}(f, g) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[(f(\mathbf{x}) - g(\mathbf{x}))^2]$. Our goal is to design a proper constant-approximation PAC learner for class \mathcal{C}_{σ} in time and sample complexity polynomial in the input parameters.

In this paper, we focus primarily on the ReLU activation, that is, $\text{ReLU}(a) = \max(0, a)$. We also restrict ourselves to isotropic distributions, that is, $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}}[\mathbf{x}] = 0$ and $\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}}[\mathbf{x}\mathbf{x}^T] = \mathbf{I}$. We also assume that the labels are bounded in absolute value by 1 for ease of presentation. For approximate learning guarantees, our results go through if we assume the distribution of labels is sub-exponential.

Definition 3 (Chow parameters) Given a distribution \mathcal{D} over $\mathbb{R}^d \times \mathbb{R}$, for any function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, define the (degree-1) Chow parameters of f w.r.t. \mathcal{D} as $\chi_{\mathcal{D}}^f := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_{\mathcal{X}}}[f(\mathbf{x})\mathbf{x}]$.

For a sample S drawn from \mathcal{D} , we also define the corresponding empirical Chow parameter with respect to S as $\widehat{\chi}_S^f := \frac{1}{|S|} \sum_{(\mathbf{x}, y) \in S} f(\mathbf{x})\mathbf{x}$.

We overload notation by defining the true Chow parameters as $\chi_{\mathcal{D}} = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[y\mathbf{x}]$ and its corresponding empirical true Chow parameter w.r.t. S as $\widehat{\chi}_S := \frac{1}{|S|} \sum_{(\mathbf{x}, y) \in S} y\mathbf{x}$.

Definition 4 (Chow distance) Given distribution \mathcal{D} over $\mathbb{R}^d \times \mathbb{R}$, for any functions $f, g : \mathbb{R}^d \rightarrow \mathbb{R}$, define the Chow distance between f and g w.r.t. \mathcal{D} as $\text{chow}_{\mathcal{D}}(f, g) = \|\chi_{\mathcal{D}}^f - \chi_{\mathcal{D}}^g\|_2$, that is, the Euclidean distance between the corresponding Chow parameters.

Lemma 5 (Chow distance to function distance) Let \mathcal{D} be such that the marginal on \mathcal{X} is isotropic. For any functions f and g , $\|\chi_{\mathcal{D}}^f - \chi_{\mathcal{D}}^g\|_2 \leq \sqrt{L_{\mathcal{D}}(f, g)}$.

Proof We have

$$\begin{aligned} \|\chi_{\mathcal{D}}^f - \chi_{\mathcal{D}}^g\|_2 &= \|\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[(f(\mathbf{x}) - g(\mathbf{x}))\mathbf{x}]\|_2 \\ &= \max_{\|\mathbf{u}\|_2=1} \mathbb{E}_{\mathcal{D}}[(f(\mathbf{x}) - g(\mathbf{x}))\langle \mathbf{u}, \mathbf{x} \rangle] \\ &\leq \sqrt{L_{\mathcal{D}}(f, g)} \max_{\|\mathbf{u}\|_2 \leq 1} \sqrt{\mathbb{E}_{\mathcal{D}}[\langle \mathbf{u}, \mathbf{x} \rangle^2]} = \sqrt{L_{\mathcal{D}}(f, g)}. \end{aligned}$$

Here the first equality follows from the variational form of the Euclidean norm and the last follows from applying Cauchy-Schwartz inequality and using isotropy of the underlying distribution on \mathcal{X} .

■

Corollary 6 (Chow-distance from true Chow vector) *Let \mathcal{D} be such that the marginal on \mathcal{X} is isotropic. For any activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ and vector $\mathbf{w} \in \mathbb{R}^d$, we have $\|\chi_{\mathcal{D}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}}\|_2 \leq \sqrt{L_{\mathcal{D}}(\sigma_{\mathbf{w}})}$.*

Proof Letting $f = \mathbb{E}[y|\mathbf{x}]$ and $g = \sigma_{\mathbf{w}}$ in Lemma 5 gives us,

$$\begin{aligned} \|\chi_{\mathcal{D}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}}\|_2^2 &\leq \mathbb{E}_{\mathcal{D}}[(\mathbb{E}[y|\mathbf{x}] - \sigma_{\mathbf{w}}(\mathbf{x}))^2] \\ &\leq \mathbb{E}_{\mathcal{D}}[(y - \sigma_{\mathbf{w}}(\mathbf{x}))^2] = L_{\mathcal{D}}(\sigma_{\mathbf{w}}). \end{aligned}$$

Here the last inequality follows from an application of Jensen's inequality. ■

Organization. In Section 3, we give an algorithm to find a weight vector that matches the true Chow parameters for the class of GLMs. In Section 4, we show that under certain assumptions on the activation function, the so obtained weight vector in fact gives us the approximate learning guarantee. In Section 5, we show that, for isotropic log-concave distributions, the ReLU satisfies our assumptions, and combining the previous techniques gives us the desired approximate learning result. Finally, in Section 6 we give an algorithm that improves the approximation factor to $1 + \eta$ for any constant $0 < \eta \leq 1$ at the cost of improper learning.

3. Matching Chow Parameters via Projected Gradient Descent

In this section, we show that projected gradient descent on the surrogate loss outputs a hypothesis $\sigma_{\mathbf{w}}$ whose Chow parameters nearly match the true Chow parameters, $\mathbb{E}[y\mathbf{x}]$. More formally, we redefine the surrogate loss as follows:

$$L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} \left[\int_0^{\langle \mathbf{w}, \mathbf{x} \rangle} (\sigma(a) - y) da \right] = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\tilde{\sigma}(\langle \mathbf{w}, \mathbf{x} \rangle) - y \langle \mathbf{w}, \chi_{\mathcal{D}} \rangle].$$

Here $\tilde{\sigma}$ is the anti-derivative of σ . For example, for the ReLU activation, we have that $\widetilde{\text{ReLU}}(a) = 0$ for all $a \leq 0$ and $\widetilde{\text{ReLU}}(a) = a^2/2$ otherwise. We correspondingly define the empirical version of the surrogate loss over sample set S as \hat{L}_S^{surr} .

We note that the gradient of $L_{\mathcal{D}}^{\text{surr}}$ is directly related to the Chow parameters as follows

$$\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \mathbb{E}[\sigma'(\langle \mathbf{w}, \mathbf{x} \rangle) \mathbf{x}] - \chi_{\mathcal{D}} = \chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} - \chi_{\mathcal{D}}.$$

Furthermore, the Hessian can be computed as

$$\nabla^2 L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \mathbb{E}[\sigma''(\langle \mathbf{w}, \mathbf{x} \rangle) \mathbf{x} \mathbf{x}^T] \succeq 0.$$

Where σ' is a subgradient. Here the last inequality follows from the non-decreasing property of σ . Thus, we have that $L_{\mathcal{D}}^{\text{surr}}$ is convex. Moreover, since σ is 1-Lipschitz, and our distribution is isotropic, we have that $1 \succeq \nabla^2 L_{\mathcal{D}}^{\text{surr}}(\mathbf{w})$ implying that $L_{\mathcal{D}}^{\text{surr}}$ is 1-smooth. Since minimizing the surrogate loss minimizes the gradient norm of the loss, loss minimization matches the Chow parameters of the GLM to the true Chow parameters.

By standard Projected Gradient Descent analysis with approximate gradients, we have the following theorem, the proof of which is in Section D of the appendix.

Algorithm 1 Projected Gradient Descent on Surrogate Loss

Input: Set $S = (\mathbf{x}^{(i)}, y^{(i)})_{i=1}^m$ i.i.d. samples drawn from \mathcal{D}

Parameter: Learning rate $\eta > 0$ and weight bound W

$\mathbf{w}^{(0)} := \mathbf{0}^d$

for $t = 0, \dots, T$ **do**

$\mathbf{v}^{(t+1)} := \mathbf{w}^{(t)} - \eta \nabla \widehat{L}_S^{\text{surr}}(\mathbf{w}^{(t)})$

$\mathbf{w}^{(t+1)} := \Pi_{B(d,W)}(\mathbf{v}^{(t+1)})$ (Projection step)

end

Theorem 7 Suppose m is sufficiently large so that for all $\mathbf{w} \in B(d, W)$ we have

$$\|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}_S^{\text{surr}}(\mathbf{w})\|_2 \leq \epsilon.$$

Also suppose that the minimizer of $L_{\mathcal{D}}^{\text{surr}}$ lies in $B(d, W)$. Then Algorithm 1 when run on m samples from \mathcal{D} with weight bound W and $\eta < 1/4$ for $T \geq \frac{32W^2}{4\epsilon W + \epsilon^2}$ iterations has an iteration $T' \leq T$ such that

$$\|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}^{(T')}}} - \chi_{\mathcal{D}}\|_2^2 \leq 8\epsilon W + 2\epsilon^2.$$

Subsequently, we can use a fresh batch of samples and choose the hypothesis with the smallest gradient. Assuming our distribution satisfies certain concentration properties, we can bound the number of samples needed by the above algorithm using the following lemma whose proof we defer to Section C of the appendix.

Lemma 8 If \mathcal{D} is a distribution such that for every \mathbf{v} , $\langle \mathbf{x}, \mathbf{v} \rangle$ has a density bounded above by $\exp(-\langle \mathbf{x}, \mathbf{v} \rangle^t)$ for some $t > 0$, then for $m \geq \Omega\left(\left(W \frac{d}{\epsilon} \log \frac{W}{\epsilon} \log \frac{1}{\delta}\right)^{2/t}\right)$, for all $\mathbf{w} \in B(d, W)$ we have that

$$\Pr_{S \sim \mathcal{D}^m} \left[\|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}_S^{\text{surr}}(\mathbf{w})\|_2 \leq \epsilon \right] \geq 1 - \delta.$$

Faster Rates under Strong Convexity If we assume that $L_{\mathcal{D}}^{\text{surr}}$ is strongly convex and restrict to a bounded fourth moment distribution, we can get much faster rates and improved sample complexity (in fact linear in the dimension d up to log factors).

Definition 9 (Strong-Convexity) We say that the activation σ satisfies μ -strong convexity w.r.t. distribution \mathcal{D} , if for all \mathbf{u}, \mathbf{v} there exists $\mu > 0$ such that

$$\langle \chi_{\mathcal{D}}^{\sigma_{\mathbf{u}}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{v}}}, \mathbf{u} - \mathbf{v} \rangle \geq \mu \|\mathbf{u} - \mathbf{v}\|_2^2.$$

Theorem 10 Let \mathcal{D} be such that $\mathcal{D}_{\mathcal{X}}$ is isotropic log-concave. Suppose that the minimizer of $L_{\mathcal{D}}^{\text{surr}}$ lies in $B(d, W)$. If σ satisfies μ -strong convexity w.r.t. \mathcal{D} then for Algorithm 1 (without the projection step) run with $\eta \leq 1/16$ and

$$m \geq \tilde{\Omega} \left(\frac{(\mu + 1)}{\mu^2 \epsilon^2} d \log^4 \left(\frac{d}{\delta} \right) (W + 1)^2 + \frac{d}{\mu^2} \log \left(\frac{W + 1}{\mu \delta} \right) \right) \quad \text{where } 0 \leq \epsilon \leq W$$

after $T \geq \frac{2 \log \left(\frac{9W}{\epsilon} \right)}{\log \left(1 - \frac{\mu \eta}{6} \right)}$ iterations, $\|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}^{(T)}}} - \chi_{\mathcal{D}}\|_2 \leq \epsilon$ holds with probability at least $1 - \delta$ as long as $\delta \geq e^{-O(\sqrt{d})}$.

The proof of Theorem 10 is deferred to Section B in the Appendix.

4. Matching Chow Parameters Suffices for Approximate Learning

In this section, we show that under certain assumptions on the activation function, matching Chow vectors implies small loss of the surrogate minimizer. We subsequently show that commonly used activation functions such as ReLU satisfy this assumption.

Definition 11 (Chow Learnability) *We say that an activation function satisfies β -Chow Learnability w.r.t. some distribution \mathcal{D} if for all $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$ and some fixed constant $\beta > 0$, we have that*

$$L_{\mathcal{D}}(\sigma_{\mathbf{u}}, \sigma_{\mathbf{v}}) \leq \beta \cdot \|\chi_{\mathcal{D}}^{\sigma_{\mathbf{u}}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{v}}}\|_2^2.$$

We will require the following lemma, proved in Section E.

Lemma 12 *If a 1-Lipschitz activation σ satisfies μ -strong convexity w.r.t. \mathcal{D} such that $\mathcal{D}_{\mathcal{X}}$ is isotropic, then the activation also satisfies μ -Chow Learnability.*

Remark 13 *Observe that Chow learnability may be a much weaker notion than strong convexity, since strong convexity requires parameter closeness. For activations with bounded ranges, such as sigmoid, it is possible for the loss to be small and Chow parameters to be close while the vectors themselves may be far.*

If the activation satisfies the Chow learnability condition, then we can show that a hypothesis nearly matching the Chow parameters attains small loss.

Theorem 14 *Let σ be such that it satisfies β -Chow Learnability w.r.t. \mathcal{D} with $\mathcal{D}_{\mathcal{X}}$ being isotropic. Suppose \mathbf{w} is such that $\|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} - \chi_{\mathcal{D}}\|_2^2 \leq \epsilon$. Then we have*

$$L_{\mathcal{D}}(\sigma_{\mathbf{w}}) \leq 2 \operatorname{opt}_{\mathcal{D}}(C_{\sigma}) (1 + 2\beta) + 4\beta\epsilon.$$

Proof Let $\sigma_{\mathbf{w}^*}$ be the function attaining the loss $\operatorname{opt}_{\mathcal{D}}(C_{\sigma})$. By assumption on σ , we have

$$\begin{aligned} L_{\mathcal{D}}(\sigma_{\mathbf{w}}, \sigma_{\mathbf{w}^*}) &\leq \beta \cdot \|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{w}^*}}\|_2^2 \\ &\leq 2\beta \left(\|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}}} - \chi_{\mathcal{D}}\|_2^2 + \|\chi_{\mathcal{D}}^{\sigma_{\mathbf{w}^*}} - \chi_{\mathcal{D}}\|_2^2 \right) \\ &\leq 2\beta (\epsilon + \operatorname{opt}_{\mathcal{D}}(C_{\sigma})). \end{aligned}$$

Here the last inequality follows by Corollary 6. Also using triangle inequality,

$$L_{\mathcal{D}}(\sigma_{\mathbf{w}}) \leq 2 \operatorname{opt}_{\mathcal{D}}(C_{\sigma}) + 2 L_{\mathcal{D}}(\sigma_{\mathbf{w}}, \sigma_{\mathbf{w}^*}).$$

Combining the above gives us the desired result. ■

Remark 15 *In the above guarantee, we can replace $\operatorname{opt}_{\mathcal{D}}(C_{\sigma})$ by $\min_{c \in C_{\sigma}} \mathbb{E}[(\mathbb{E}[y|\mathbf{x}] - c(\mathbf{x}))^2]$ (see proof of Lemma 6). In the p -concept setting, where $\mathbb{E}[y|\mathbf{x}] = c^*(\mathbf{x})$ this is potentially a tighter guarantee. This is because $\min_{c \in C_{\sigma}} \mathbb{E}[(\mathbb{E}[y|\mathbf{x}] - c(\mathbf{x}))^2]$ is in fact 0 whereas $\operatorname{opt}_{\mathcal{D}}(C_{\sigma})$ might be large. Since we are focused on the agnostic setting, we will stick to using $\operatorname{opt}_{\mathcal{D}}(C_{\sigma})$ in our results.*

5. Constant Factor Approximation for ReLU Regression

In this section, we present a constant factor approximation algorithm for ReLU regression over any isotropic log-concave distribution using the techniques developed in the previous sections.

Theorem 16 *Let \mathcal{D} be such that $\mathcal{D}_{\mathcal{X}}$ is isotropic log-concave and assume the labels are bounded. Let $\text{ReLU}_{\mathbf{w}^*}$ achieve loss $\text{opt}_{\mathcal{D}}(C_{\text{ReLU}})$ and assume that $\|\mathbf{w}^*\|_2 \leq W_{\text{opt}}$. Then Algorithm 1 outputs a vector \mathbf{w} such that*

$$L_{\mathcal{D}}(\text{ReLU}_{\mathbf{w}}) \leq O(\text{opt}_{\mathcal{D}}(C_{\text{ReLU}})) + \epsilon ,$$

with probability $1-\delta$ using $m \gtrsim \frac{d}{\epsilon^2} \log^4\left(\frac{d}{\delta}\right) (W_{\text{opt}} + 1)^2$ samples, for $0 \leq \epsilon \leq W_{\text{opt}}$, and $O\left(dm \log\left(\frac{W}{\epsilon}\right)\right)$ time.

Our main observation is that the ReLU activation satisfies the strong convexity condition w.r.t. any isotropic log-concave distribution.

Lemma 17 (Strong Convexity of ReLU) *Let \mathcal{D} be such that $\mathcal{D}_{\mathcal{X}}$ is isotropic log-concave. Then there exists some fixed constant $\mu > 0$ such that ReLU is μ -strongly convex w.r.t. \mathcal{D} .*

Proof Sketch. Since the ReLU is 1-Lipschitz and non-decreasing, we have

$$\begin{aligned} (\chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{v}}} - \chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{u}}})^T(\mathbf{v} - \mathbf{u}) &= \mathbb{E}[(\text{ReLU}(\langle \mathbf{v}, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{u}, \mathbf{x} \rangle))(\langle \mathbf{v} - \mathbf{u}, \mathbf{x} \rangle)] \\ &\geq \mathbb{E}[(\text{ReLU}(\langle \mathbf{v}, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{u}, \mathbf{x} \rangle))^2] . \end{aligned}$$

Now our goal is to bound from below the error between the two ReLUs by the distance between the corresponding vectors. Due to the anti-concentration properties of log-concave distributions, there is sufficient probability mass in a constant radius ball around the origin. This enables us to exploit the linear region of the corresponding ReLUs to establish the lower bound. We defer the full proof to Section F in the Appendix.

Proof of Theorem 16 By Lemma 17, the ReLU activation satisfies μ -strong convexity w.r.t. \mathcal{D} for some constant $\mu > 0$. This implies that $L_{\mathcal{D}}^{\text{surr}}$ is strongly-convex and therefore the minimizer of $L_{\mathcal{D}}^{\text{surr}}$ (say \mathbf{w}) satisfies $\chi_{\mathcal{D}} = \chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{w}}}$. Using Lemma 6 and the strong convexity of ReLU, we have that

$$\begin{aligned} \|\mathbf{w}^* - \mathbf{w}\|_2 &\lesssim \|\chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{w}^*}} - \chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{w}}}\|_2 \\ &= \|\chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{w}^*}} - \chi_{\mathcal{D}}\|_2 \\ &\leq \sqrt{L_{\mathcal{D}}(\text{ReLU}_{\mathbf{w}^*})} = \sqrt{\text{opt}_{\mathcal{D}}(C_{\text{ReLU}})} . \end{aligned}$$

Therefore, $\|\mathbf{w}\|_2 \leq W_{\text{opt}} + O\left(\sqrt{\text{opt}_{\mathcal{D}}(C_{\text{ReLU}})}\right)$. It is not hard to see that with bounded labels $\text{opt}_{\mathcal{D}}(C_{\text{ReLU}}) \leq O(W_{\text{opt}}^2 + 1)$. Therefore, we can now apply Theorem 10 to find a hypothesis with Chow distance at most ϵ . The result now follows directly from Theorem 14.

6. A PTAS for ReLU Regression

In this section, we show that if the activation is the ReLU function we can solve the problem of finding the best fitting ReLU up to a $(1 + \eta)$ -approximation, when the underlying marginal over the input is sub-gaussian. We assume that $\text{opt} := \text{opt}_{\mathcal{D}}(C_{\text{ReLU}}^1) \leq c \leq 1$, for some constant c .

We define sub-gaussian distributions here:

Definition 18 A distribution \mathcal{D} on \mathbb{R}^d is called ν -subgaussian, $\nu > 0$, if for any direction \mathbf{v} the probability density function of $\langle \mathbf{x}, \mathbf{v} \rangle$ where $\mathbf{x} \sim \mathcal{D}$, $p_{\mathbf{v}}(\mathbf{x})$ satisfies $p_{\mathbf{v}}(\mathbf{x}) = O\left(\frac{1}{\nu} \cdot \exp\left(-\frac{(\mathbf{v} \cdot \mathbf{x})^2}{2\nu^2}\right)\right)$.

Our algorithm (Algorithm 2) works by partitioning the domain into three parts T_-, T, T_+ , where

$$\begin{aligned} T &= \{\mathbf{u} \in \mathbb{R}^d : |\langle \mathbf{w}, \mathbf{u} \rangle| \leq \gamma \sqrt{\text{opt}}\} \\ T_+ &= \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle > \gamma \sqrt{\text{opt}}\} \\ T_- &= \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle < -\gamma \sqrt{\text{opt}}\}. \end{aligned}$$

The hypothesis $h(\mathbf{x})$ behaves as a different function in each of these parts. For $\mathbf{x} \in T_-$, the hypothesis is the 0 function. For $\mathbf{x} \in T_+$, the hypothesis takes the value of $\langle \mathbf{w}_+, \mathbf{x} \rangle$, which is the best fitting linear function over T_+ . Finally, over T the hypothesis outputs the value that the best fitting ℓ_1 -norm bounded polynomial of degree $1/\eta^3$. Our main theorem of this section is the following:

Theorem 19 Let $\mathcal{D}_{\mathbf{x}}$ be ν -subgaussian for $\nu \leq O(1)$, $\|\mathbf{w}^*\|_2 \leq 1$ and $y \in [0, 1]$ for every $(\mathbf{x}, y) \sim \mathcal{D}$, then there is an algorithm that takes $O\left(\frac{1}{\epsilon^2} \cdot \left(\frac{d}{\eta^3 \nu^2}\right)^{1/\eta^3}\right)$ samples and time, and returns a hypothesis h that with high probability satisfies

$$\mathbb{E}_{\mathcal{D}} [(h(\mathbf{x}) - y)^2] \leq (1 + \eta)\text{opt} + \epsilon.$$

Remark 20 We note that if the distribution is uniform over \mathbb{S}^{n-1} , then the sample complexity of our algorithm scales as $2^{1/\eta^3}$, instead of d^{1/η^3} , since the distribution is $(1/\sqrt{d})$ -subgaussian. That is, under the uniform distribution over the unit sphere, the sample complexity is independent of d .

The proof Theorem 19 follows from a direct application of the following properties of Algorithm 2 with the specified parameters.

Lemma 21 Let $\mathcal{D}_{\mathbf{x}}$ be ν -subgaussian for $\nu \leq O(1)$ and let S be a set of i.i.d. samples drawn from \mathcal{D} . If $m = |S| = \Omega\left(\frac{k^k \cdot d^k}{\nu^{2k}} \cdot \frac{1}{\epsilon^2}\right)$, where $k = \frac{1}{\eta^3}$, $\|\mathbf{w}^*\|_2 \leq 1$ and $y \in [0, 1]$, then for $\gamma = \Omega\left(\sqrt{\log\left(\frac{1}{\eta}\right)}\right)$, we have

1. $\mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}_+, \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] \leq \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_{T_+}(\mathbf{x})] + \frac{\eta}{3} \cdot \text{opt} + \epsilon \Pr_{\mathcal{D}}[T_+]$.
2. $\mathbb{E}_{\mathcal{D}} [(0 - y)^2 1_{T_-}(\mathbf{x})] \leq \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_{T_-}(\mathbf{x})] + \frac{\eta}{3} \cdot \text{opt}$.
3. $\mathbb{E}_{\mathcal{D}} [(P(\mathbf{x}) - y)^2 1_T(\mathbf{x})] \leq \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_T(\mathbf{x})] + \frac{\eta}{3} \cdot \text{opt} + \epsilon \Pr_{\mathcal{D}}[T]$.

The proof of Lemma 21 is in Appendix G.

Proof [Proof of Theorem 19] This follows by putting together the cases in Lemma 21. ■

Algorithm 2 PTAS for ReLU regression

Input: $0 < \text{opt} \leq 1$ and access to i.i.d. samples from \mathcal{D}

Parameter : $r \in \mathbb{N}, \epsilon, \gamma, \epsilon, W > 0$

- 1: Find \mathbf{w} using algorithm from previous section. This takes $O(\frac{d}{\epsilon^2} \log \frac{d}{\delta})$ samples and satisfies $\|\mathbf{w} - \mathbf{w}^*\|_2 \leq O(1/\nu) \cdot \sqrt{\text{opt}}$.
- 2: Let $T = \{\mathbf{u} \in \mathbb{R}^d : |\langle \mathbf{w}, \mathbf{u} \rangle| \leq \gamma \sqrt{\text{opt}}\}$, $T_+ = \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle > \gamma \sqrt{\text{opt}}\}$ and $T_- = \{\mathbf{u} \in \mathbb{R}^d : \langle \mathbf{w}, \mathbf{u} \rangle < -\gamma \sqrt{\text{opt}}\}$.
- 3: Find a degree $k = \frac{1}{\eta^3}$, d -variate polynomial, P , the ℓ_1 -norm of whose coefficients is at most $\nu \cdot O(4^k) = O(4^k)$, using L_2 -polynomial regression on $m_{pol} = O(\frac{k^k \cdot d^k}{\nu^{2k}} \cdot \frac{1}{\epsilon^2})$ samples such that

$$\mathbb{E}_{S|T}[(P(\mathbf{x}) - y)^2] \leq \min_{P' \in \text{POL}_{r,d}} \mathbb{E}_{\mathcal{D}|T}[(P'(\mathbf{x}) - y)^2] + \epsilon,$$

where $h_{\mathbf{w}^*}$ is the optimal ReLU classifier w.r.t. \mathcal{D} .

- 4: Find $\mathbf{w}_+ \in B(d, 1)$ using least squares with $m_{ls} = O(1/\epsilon^2)$ to get

$$\mathbb{E}_{S|T_+}[(\langle \mathbf{w}_+, \mathbf{x} \rangle - y)^2] \leq \min_{\mathbf{w}' \in B(d,W)} \mathbb{E}_{\mathcal{D}|T_+}[(\langle \mathbf{w}', \mathbf{x} \rangle - y)^2] + \epsilon.$$

- 5: Output the following classifier:

$$h(\mathbf{x}) = \begin{cases} \langle \mathbf{w}_+, \mathbf{x} \rangle, & \mathbf{x} \in T_+ \\ P(\mathbf{x}), & \mathbf{x} \in T \\ 0, & \mathbf{x} \in T_- \end{cases}$$

7. Conclusions

In this work, we gave the first constant approximation scheme for ReLU regression under the assumption of log-concavity. We proved that optimizing a convex surrogate loss suffices for obtaining approximate guarantees. We further proposed a PTAS for ReLU regression under the assumption of sub-gaussianity, which refines the so obtained solution using ideas from localization and polynomial approximation.

Our work here was focussed on the ReLU activation and we leave open the extensions to other activation functions. We believe that the Chow learnability condition is potentially satisfied under log-concavity for activations that approximate thresholds such as sigmoid.

The underlying surrogate loss approach seems powerful and exploring further applications is an interesting direction for future work. Further, designing approximation schemes for a linear combination of activations functions is an interesting open question.

Acknowledgements

This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing for the Summer 2019 program on the Foundations of Deep Learning. ID was supported

by NSF Award CCF-1652862 (CAREER), a Sloan Research Fellowship, and a DARPA Learning with Less Labels (LwLL) grant. SG was supported by the JP Morgan AI Phd Fellowship. SK was supported by NSF award CNS 1414082 and ID’s startup grant. AK was supported by NSF awards CCF 1909204 and CCF 1717896. MS was supported by the Packard Fellowship in Science and Engineering, a Sloan Research Fellowship in Mathematics, an NSF-CAREER under award #1846369, the Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) under award #FA 9550-18-1-0078, DARPA Learning with Less Labels (LwLL) and FastNICs programs, an NSF-CIF award #1813877, and a Google faculty research award.

References

- Radosław Adamczak, Alexander Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of the American Mathematical Society*, 23(2):535–561, 2010.
- Peter Auer, Mark Herbster, and Manfred K. K Warmuth. Exponentially many local minima for single neurons. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 316–322. MIT Press, 1996. URL <http://papers.nips.cc/paper/1028-exponentially-many-local-minima-for-single-neurons.pdf>.
- P. Awasthi, M. F. Balcan, and P. M. Long. The power of localization for efficiently learning linear separators with noise. *J. ACM*, 63(6):50:1–50:27, 2017.
- V Bentkus. An inequality for tail probabilities of martingales with differences bounded from one side. *Journal of Theoretical Probability*, 16(1):161–173, 2003.
- Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- A. Daniely. A PTAS for agnostically learning halfspaces. In *Proceedings of The 28th Conference on Learning Theory, COLT 2015*, pages 484–502, 2015.
- A. De, I. Diakonikolas, V. Feldman, and R. Servedio. Near-optimal solutions for the Chow Parameters Problem and low-weight approximation of halfspaces. In *Proc. 44th ACM Symposium on Theory of Computing (STOC)*, pages 729–746, 2012.
- I. Diakonikolas, D. M. Kane, and A. Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, pages 1061–1073, 2018.
- I. Diakonikolas, D. Kane, and P. Manurangsi. Nearly tight bounds for robust proper learning of halfspaces with a margin. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, pages 10473–10484, 2019.
- Surbhi Goel, Varun Kanade, Adam Klivans, and Justin Thaler. Reliably learning the relu in polynomial time. In *Conference on Learning Theory*, pages 1004–1042, 2017.

- Surbhi Goel, Adam R. Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. *International Conference on Machine Learning*, 2018.
- Surbhi Goel, Sushrut Karmalkar, and Adam Klivans. Time/accuracy tradeoffs for learning a relu with respect to gaussian marginals. In *Advances in Neural Information Processing Systems*, pages 8582–8591, 2019.
- D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.
- Sham M Kakade, Varun Kanade, Ohad Shamir, and Adam Kalai. Efficient learning of generalized linear and single index models with isotonic regression. In *Advances in Neural Information Processing Systems*, pages 927–935, 2011.
- A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. Agnostically learning halfspaces. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 11–20, 2005.
- Adam Tauman Kalai and Ravi Sastry. The isotron algorithm: High-dimensional isotonic regression. COLT, 2009.
- Varun Kanade. Lecture notes: Learning real-valued functions. 2018.
- M. Kearns, R. Schapire, and L. Sellie. Toward Efficient Agnostic Learning. *Machine Learning*, 17 (2/3):115–141, 1994.
- Adam R. Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 343–354. IEEE Computer Society, 2017. ISBN 978-1-5386-3464-6. URL <http://doi.ieeecomputersociety.org/10.1109/FOCS.2017.39>.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.
- Pasin Manurangsi and Daniel Reichman. The computational complexity of training relu (s). *arXiv preprint arXiv:1810.04207*, 2018.
- R. O’Donnell and R. Servedio. The Chow Parameters Problem. In *Proc. 40th STOC*, pages 517–526, 2008.
- Alexander A. Sherstov. Making polynomials robust to noise. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC 12, page 747758, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450312455. doi: 10.1145/2213977.2214044. URL <https://doi.org/10.1145/2213977.2214044>.
- Galen R Shorack and Jon A Wellner. *Empirical processes with applications to statistics*. SIAM, 2009.
- Mahdi Soltanolkotabi. Learning relus via gradient descent. In *Advances in neural information processing systems*, pages 2007–2017, 2017.

- L. G. Valiant. A theory of the learnable. In *Proc. 16th Annual ACM Symposium on Theory of Computing (STOC)*, pages 436–445. ACM Press, 1984.
- V. Vapnik. *Estimation of Dependences Based on Empirical Data: Springer Series in Statistics*. Springer-Verlag, Berlin, Heidelberg, 1982. ISBN 0387907335.
- Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *CoRR*, abs/1904.00687, 2019. URL <http://arxiv.org/abs/1904.00687>.

Appendix A. Useful Properties

We use the following fact about sub-gaussian distributions.

Fact 1 *If \mathcal{D} is ν -subgaussian then if P is an d -variate degree k polynomial then taking an expectation over m_0 samples yeilds*

$$\Pr [|\mathbb{E}_S[P(\mathbf{x})] - \mathbb{E}_{\mathcal{D}}[P(\mathbf{x})]| \geq \epsilon] \leq \exp\left(-\frac{m\epsilon^2}{\nu^2 \text{Var}[P(\mathbf{x})]}\right)^{1/k}.$$

Appendix B. Proof of Theorem 10

We begin by stating a few auxiliary lemmas that play a crucial role in our proof.

Lemma 22 *Consider the assumptions of Theorem 10. Also assume $(\mathbf{x}_i, y_i)_{i=1}^m$ are generated i.i.d. with \mathbf{x}_i having a log-concave marginal and y_i obeying $|y_i| \leq 1$. Furthermore, assume $\|\widehat{\mathbf{w}}\|_2 \leq W$ and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is an activation obeying $|\sigma(z)| \leq B|z|$. Then, as long as*

$$m \geq \tilde{\Omega}\left(\frac{d}{\xi^2} \log^4(d/\delta) (W+1)^2\right),$$

we have that

$$\left\| \frac{1}{m} \sum_{i=1}^m (\sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle) - y_i) \mathbf{x}_i \right\|_2 \leq \xi$$

holds with probability at least $1 - \delta$.

Next we show that the gradient of the surrogate loss obeys a certain correlation inequality with the proof deferred to end of the section.

Lemma 23 *As long as $\epsilon \leq W$, $\delta \geq e^{-O(\sqrt{d})}$ and*

$$m \geq \tilde{\Omega}\left(\frac{\gamma^2}{\mu^2} d \log\left(\frac{W+1}{\mu\delta}\right)\right),$$

we have

$$\langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w} - \widehat{\mathbf{w}} \rangle \geq \alpha \|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2 + \beta \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \quad (1)$$

holds for all $\mathbf{w} \in \mathbb{R}^d$ obeying $\frac{\epsilon}{3} \leq \|\mathbf{w} - \widehat{\mathbf{w}}\|_2 \leq 2W$ with $\alpha = \frac{\mu}{3}$ and $\beta = \frac{1}{8}$ with probability at least $1 - \delta$.

With these two key lemmas in place we are now ready to prove the main theorem. First note that since $\mathbf{w}_0 = 0$ we have $\|\mathbf{w}_0 - \widehat{\mathbf{w}}\|_2 \leq W \leq 2W$ and thus by Lemma 23 the correlation inequality (1) holds at \mathbf{w}_0 with high probability. Furthermore, as we show next when the correlation inequality (1) holds subsequent iterations also obey $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq 4W$ allowing us to apply the correlation inequality (1) in an inductive fashion.

Let us now consider the progress from one iteration to the next. We can write:

$$\begin{aligned}
 \|\mathbf{w}_{\tau+1} - \widehat{\mathbf{w}}\|_2^2 &= \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 - \eta \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle + \eta^2 \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau)\|_2^2 \\
 &= \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 - \eta \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle - \eta \langle \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle \\
 &\quad + \eta^2 \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau)\|_2^2 \\
 &\leq \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 - \eta \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle - \eta \langle \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle \\
 &\quad + 2\eta^2 \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 + 2\eta^2 \|\nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \\
 &\stackrel{(a)}{\leq} (1 - \eta\alpha) \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 - \eta(\beta - 2\eta) \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \\
 &\quad + 2\eta^2 \|\nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 - \eta \langle \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w}_\tau - \widehat{\mathbf{w}} \rangle \\
 &\stackrel{(b)}{\leq} \left(1 - \frac{\alpha\eta}{2}\right) \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 - \eta(\beta - 2\eta) \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}_\tau) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \\
 &\quad + \eta \left(2\eta + \frac{1}{2\alpha}\right) \|\nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \\
 &\stackrel{(c)}{\leq} \left(1 - \frac{\alpha\eta}{2}\right) \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 + \eta \left(2\eta + \frac{1}{2\alpha}\right) \|\nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 \\
 &\stackrel{(d)}{\leq} \left(1 - \frac{\alpha\eta}{2}\right) \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 + \eta \left(2\eta + \frac{1}{2\alpha}\right) \frac{\alpha^2}{24\alpha\beta + 12} \epsilon^2.
 \end{aligned}$$

Here, (a) follows from (1) and (b) from $\langle \mathbf{a}, \mathbf{b} \rangle \leq \frac{1}{2\alpha} \|\mathbf{a}\|_2^2 + \frac{\alpha}{2} \|\mathbf{b}\|_2^2$, (c) from $\eta \leq \frac{\beta}{2}$, and (d) from Lemma 22 with $\xi = \frac{\alpha\epsilon}{\sqrt{18\alpha\beta+9}} = \frac{\mu\epsilon}{3\sqrt{\frac{3\mu}{4}+1}}$. Thus, iterating the above in all subsequent iterations we have

$$\begin{aligned}
 \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 &\leq W^2 + \frac{\eta \left(2\eta + \frac{1}{2\alpha}\right)}{1 - \left(1 - \frac{\alpha\eta}{2}\right)} \frac{\alpha^2}{2\alpha\beta + 1} \epsilon^2 \\
 &= W^2 + \frac{1}{\alpha} \left(4\eta + \frac{1}{\alpha}\right) \frac{\alpha^2}{18\alpha\beta + 9} \epsilon^2 \\
 &\leq W^2 + \frac{1}{\alpha} \left(2\beta + \frac{1}{\alpha}\right) \frac{\alpha^2}{18\alpha\beta + 9} \epsilon^2 \\
 &\leq W^2 + \frac{\epsilon^2}{9} \\
 &< 4W^2,
 \end{aligned}$$

where in the last inequality we used the fact that $\epsilon \leq W$. Therefore, $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq 2W$ for all $\tau \geq 1$ and use of the correlation inequality is justified. Furthermore, iterating the above lemma we conclude that as long as $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \geq \frac{\epsilon}{3}$ it holds

$$\begin{aligned}
 \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 &\leq \left(1 - \frac{\alpha\eta}{2}\right)^\tau \|\mathbf{w}_0 - \widehat{\mathbf{w}}\|_2^2 + \frac{\eta \left(2\eta + \frac{1}{2\alpha}\right) \xi^2}{1 - \left(1 - \frac{\alpha\eta}{2}\right)} \\
 &= \left(1 - \frac{\alpha\eta}{2}\right)^\tau \|\mathbf{w}_0 - \widehat{\mathbf{w}}\|_2^2 + \frac{1}{\alpha} \left(4\eta + \frac{1}{\alpha}\right) \xi^2
 \end{aligned}$$

$$\begin{aligned}
 &\leq \left(1 - \frac{\alpha\eta}{2}\right)^\tau \|\mathbf{w}_0 - \widehat{\mathbf{w}}\|_2^2 + \frac{1}{\alpha} \left(2\beta + \frac{1}{\alpha}\right) \xi^2 \\
 &= \left(1 - \frac{\alpha\eta}{2}\right)^\tau \|\mathbf{w}_0 - \widehat{\mathbf{w}}\|_2^2 + \frac{\epsilon^2}{9}.
 \end{aligned} \tag{2}$$

Thus, after $\tau \geq T := \frac{2 \log(\frac{\epsilon}{3W})}{\log(1 - \frac{\mu\eta}{6})}$ we have

$$\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 \leq \frac{\epsilon^2}{9} + \frac{\epsilon^2}{9} = \frac{2\epsilon^2}{9} \quad \Rightarrow \quad \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \frac{2}{3}\epsilon.$$

Note that above was carried out under the assumption that for all $t = 1, 2, \dots, T$ we have $\|\mathbf{w}_t - \widehat{\mathbf{w}}\|_2 \geq \frac{\epsilon}{3}$. We note that if this assumption is violated at some iteration \tilde{t} we have $\|\mathbf{w}_{\tilde{t}} - \widehat{\mathbf{w}}\|_2 \leq \frac{\epsilon}{3}$. Now either $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \frac{\epsilon}{3}$ for all $\tau \geq \tilde{t}$ in which case after $\tau \geq T := \frac{2 \log(\frac{\epsilon}{3W})}{\log(1 - \frac{\mu\eta}{6})}$ we have $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \frac{1}{3}\epsilon$. If not at some iteration $t \geq \tilde{t}$ we have $\|\mathbf{w}_t - \widehat{\mathbf{w}}\|_2 \leq \frac{1}{3}\epsilon$ and $\|\mathbf{w}_{t+1} - \widehat{\mathbf{w}}\|_2 \geq \frac{1}{3}\epsilon$. Thus,

$$\begin{aligned}
 \|\mathbf{w}_{t+1} - \widehat{\mathbf{w}}\|_2 &= \|\mathbf{w}_t - \widehat{\mathbf{w}} - \eta \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\mathbf{w}_t)\|_2 \\
 &= \|\mathbf{w}_t - \widehat{\mathbf{w}} - \eta (\nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\mathbf{w}_t) - \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}})) - \eta \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}})\|_2 \\
 &\leq \|\mathbf{w}_t - \widehat{\mathbf{w}} - \eta (\nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\mathbf{w}_t) - \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}}))\|_2 + \eta \|\nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}})\|_2 \\
 &\stackrel{(a)}{\leq} \|\mathbf{w}_t - \widehat{\mathbf{w}}\|_2 + \eta \|\nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}})\|_2 \\
 &\stackrel{(b)}{\leq} \frac{2}{3}\epsilon \leq 2W.
 \end{aligned}$$

In the above, (a) follows from the fact that $\tilde{\sigma}$ is convex which implies that for any two scalars z, \tilde{z} we have $(\tilde{\sigma}'(z) - \tilde{\sigma}'(\tilde{z}))(z - \tilde{z}) \geq 0$ which implies $(\sigma(z) - \sigma(\tilde{z}))(z - \tilde{z}) \geq 0$. This in turn implies that

$$\langle \nabla \widehat{\mathcal{L}}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w} - \widehat{\mathbf{w}} \rangle = \frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle)) (\mathbf{x}_i^T (\mathbf{w} - \widehat{\mathbf{w}})) \geq 0,$$

so that $\|\mathbf{w}_t - \widehat{\mathbf{w}} - \eta (\nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\mathbf{w}_t) - \nabla \widehat{\mathcal{L}}^{\text{surr}}(\widehat{\mathbf{w}})(\widehat{\mathbf{w}}))\|_2 \leq \|\mathbf{w}_t - \widehat{\mathbf{w}}\|_2$. Also (b) follows from the fact that $\eta \leq \frac{\beta}{2} = \frac{1}{16}$ and Lemma 22 with $\xi \leq \frac{16}{3}\epsilon$. As a result, we are in a region where the correlation inequality applies. Furthermore, using an argument similar to (3) for all $\tau \geq t$, where $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \geq \frac{\epsilon}{3}$, we have

$$\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2^2 \leq \left(1 - \frac{\alpha\eta}{2}\right)^{\tau - (t+1)} \|\mathbf{w}_{t+1} - \widehat{\mathbf{w}}\|_2^2 + \frac{\epsilon^2}{9} \leq \|\mathbf{w}_{t+1} - \widehat{\mathbf{w}}\|_2^2 + \frac{\epsilon^2}{9} \leq \frac{5}{9}\epsilon^2 \quad \Rightarrow \quad \|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \epsilon. \tag{3}$$

Of course, if at some point we again have $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \frac{\epsilon}{3}$, we repeat the above arguments. In conclusion, in all cases after $\tau \geq T := \frac{2 \log(\frac{\epsilon}{3W})}{\log(1 - \frac{\mu\eta}{6})}$, we have $\|\mathbf{w}_\tau - \widehat{\mathbf{w}}\|_2 \leq \epsilon$ completing the proof.

B.1. Proof of Lemma 22

Define the random vector

$$\mathbf{z} = \frac{1}{m} \sum_{i=1}^m (\sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle) - y_i) \mathbf{x}_i.$$

Note that

$$\|\mathbf{z}\|_2^2 = \sum_{j=1}^m |z_j|^2,$$

so that it suffices to bound the square of the individual entries of the vector \mathbf{z} . To bound this quantity we bound individual entries of the vector. Note that any such entry can be written in the form

$$\frac{1}{m} \sum_{i=1}^m (\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i) X_i,$$

where $Z_i = \langle \frac{\widehat{\mathbf{w}}}{\|\widehat{\mathbf{w}}\|_2}, \mathbf{x}_i \rangle$ and $X_i = \langle \mathbf{x}_i, \mathbf{e}_j \rangle$ are sub-exponential random variables with constant $\|\cdot\|_{\psi_1}$ norm (where $\|\cdot\|_{\psi_p} := \inf\{k \in (0, \infty) \mid \mathbb{E}[\exp((|x|/k)^p) - 1] \leq 1\}$ – this characterizes the limiting behavior of the probability density function). Furthermore, $|y_i| \leq 1$ implying that y_i is a sub-Gaussian random variable. Therefore,

$$\|\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i\|_{\psi_1} \leq c(W+1) \quad \Rightarrow \quad \|(\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i) X_i\|_{\psi_{1/2}} \leq c(W+1).$$

Thus, using a well-known result of Talagrand (specifically combining (Ledoux and Talagrand, 2013, Theorem 6.21) and (Shorack and Wellner, 2009, Lemma 2.2.2)), we have

$$\left\| \frac{1}{m} \sum_{i=1}^m (\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i) X_i \right\|_{\psi_{1/2}} \leq c \frac{\log m}{\sqrt{m}} (W+1).$$

Therefore,

$$\mathbb{P} \left\{ \frac{1}{m} \sum_{i=1}^m (\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i) X_i \geq ct \frac{\log m}{\sqrt{m}} (W+1) \right\} \leq Ce^{-\sqrt{t}}.$$

Thus, using the union bound

$$\begin{aligned} \mathbb{P} \left\{ \|\mathbf{z}\|_2 \geq c \frac{\sqrt{d}}{\sqrt{m}} t \log m (W+1) \right\} &\leq d \mathbb{P} \left\{ \frac{1}{m} \sum_{i=1}^m (\sigma(\|\widehat{\mathbf{w}}\|_2 Z_i) - y_i) X_i \geq ct \frac{\log m}{\sqrt{m}} (W+1) \right\} \\ &\leq dCe^{-\sqrt{t}}. \end{aligned}$$

Setting $t = \log^2(dC/\delta)$ completes the proof.

B.2. Proof of Lemma 23

For any vector $\mathbf{w} \in \mathbb{R}^d$ and $\widehat{\mathbf{w}}$, we have

$$\begin{aligned} \frac{1}{\|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2} \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w} - \widehat{\mathbf{w}} \rangle &= \frac{1}{m} \sum_{i=1}^m \frac{(\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle)) (\mathbf{x}_i^T (\mathbf{w} - \widehat{\mathbf{w}}))}{\|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2} \\ &:= \frac{1}{m} \sum_{i=1}^m \mathcal{Y}_i(\mathbf{w}), \end{aligned}$$

where we define the random processes $\mathcal{Y}_i(\mathbf{w}) := \frac{(\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle)) (\mathbf{x}_i^T (\mathbf{w} - \widehat{\mathbf{w}}))}{\|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2}$.

Thus, for the random process $\mathcal{X}_i(\mathbf{w}) := \mathbb{E}[\mathcal{Y}_i(\mathbf{w})] - \mathcal{Y}_i(\mathbf{w})$ we have

$$\begin{aligned} \mathcal{X}_i(\mathbf{w}) &= \mathbb{E}[\mathcal{Y}_i(\mathbf{w})] - \mathcal{Y}_i(\mathbf{w}) \\ &\stackrel{(a)}{\leq} \mathbb{E}[\mathcal{Y}_i(\mathbf{w})] \\ &\stackrel{(b)}{\leq} \frac{\mathbb{E}[|\mathbf{x}_i^T(\mathbf{w} - \widehat{\mathbf{w}})|^2]}{\|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2} \\ &\stackrel{(c)}{=} 1. \end{aligned}$$

Here, (a) follows from the fact that $\tilde{\sigma}$ is convex which implies that for any two scalars z, \widehat{z} we have

$$(\tilde{\sigma}'(z) - \tilde{\sigma}'(\widehat{z}))(z - \widehat{z}) \geq 0 \quad \Rightarrow \quad (\sigma(z) - \sigma(\widehat{z}))(z - \widehat{z}) \geq 0.$$

Thus, we always have $\mathcal{Y}_i(\mathbf{w}) := (\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle)) (\mathbf{x}_i^T(\mathbf{w} - \widehat{\mathbf{w}})) \geq 0$, (b) from 1-Lipschitzness, and (c) from the isotropic assumption on \mathbf{x}_i .

Also we have

$$\mathbb{E}[\mathcal{X}_i^2(\mathbf{w})] = \mathbb{E}[\mathcal{Y}_i^2(\mathbf{w})] - (\mathbb{E}[\mathcal{Y}_i(\mathbf{w})])^2 \leq \mathbb{E}[\mathcal{Y}_i^2(\mathbf{w})] \leq \mathbb{E}[(\mathbf{x}_i^T(\mathbf{w} - \widehat{\mathbf{w}}))^4] \leq \gamma^2 \|\mathbf{w} - \widehat{\mathbf{w}}\|_2^4,$$

where in the penultimate inequality we used 1-Lipschitz property of ReLU and in the last inequality we used boundedness of fourth moments of the distribution. We will now apply Lemma 7.13 of [Candes et al. \(2015\)](#) (also see [Bentkus, 2003](#)) for a fixed \mathbf{w} with $v = \gamma$, $b = 1$, and $y = m\xi$ to conclude that

$$\Pr \left[\frac{1}{m} \sum_{i=1}^m \mathcal{X}_i(\mathbf{w}) \geq \xi\mu \right] \leq e^{-m \frac{\mu^2}{\gamma^2} \xi^2}. \quad (4)$$

Now note that by μ -strong convexity of the surrogate loss, we have

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \mathcal{Y}_i(\mathbf{w}) &= \frac{1}{m} \sum_{i=1}^m \mathbb{E}[\mathcal{Y}_i(\mathbf{w})] - \frac{1}{m} \sum_{i=1}^m \mathcal{X}_i(\mathbf{w}) \\ &\geq \mu - \frac{1}{m} \sum_{i=1}^m \mathcal{X}_i(\mathbf{w}). \end{aligned}$$

Using (4) in the latter, we conclude that

$$\frac{1}{m} \sum_{i=1}^m \mathcal{Y}_i(\mathbf{w}) \geq (1 - \xi)\mu,$$

holds with probability at least $1 - e^{-m \frac{\mu^2}{\gamma^2} \xi^2}$. To continue, define $\mathbf{h} = \frac{\mathbf{w} - \widehat{\mathbf{w}}}{\|\mathbf{w} - \widehat{\mathbf{w}}\|_2}$ and $s = \|\mathbf{w} - \widehat{\mathbf{w}}\|_2$, and note that \mathcal{Y}_i can be alternatively be written in the form of the stochastic process

$$\mathcal{Y}_i(\mathbf{h}; s) := \frac{(\sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle))}{s} (\mathbf{x}_i^T \mathbf{h}).$$

Thus, based on the argument above for a fixed $\mathbf{h} \in \mathbb{S}^{d-1}$ and a fixed $0 \leq s \leq CW$, we have that

$$\mathcal{Z}(\mathbf{h}; s) := \frac{1}{m} \sum_{i=1}^m \mathcal{Y}_i(\mathbf{h}; s) \geq \left(1 - \frac{\xi}{3}\right)\mu \quad (5)$$

holds with probability at least $1 - e^{-m \frac{\mu^2}{9\gamma^2} \xi^2}$. To continue, we prove the following simple lemma.

Lemma 24 For any $\mathbf{h}, \tilde{\mathbf{h}} \in \mathbb{S}^{d-1}$, we have

$$|\mathcal{Y}_i(\mathbf{h}; s) - \mathcal{Y}_i(\tilde{\mathbf{h}}; s)| \leq (|\mathbf{x}_i^T \mathbf{h}| + |\mathbf{x}_i^T \tilde{\mathbf{h}}| + 1) |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})| .$$

Proof Define

$$f(z) := \frac{(\sigma(x + sz) - \sigma(x))}{s} z$$

and note that for some $0 \leq t \leq 1$ we have

$$\begin{aligned} & |f(z) - f(\tilde{z})| \\ &= |f'(tz + (1-t)\tilde{z})(z - \tilde{z})| \\ &= \left| \frac{(\sigma(x + s(tz + (1-t)\tilde{z})) - \sigma(x))}{s} + \sigma'(x + s(tz + (1-t)\tilde{z}))(tz + (1-t)\tilde{z}) \right| |z - \tilde{z}| \\ &\leq \frac{|\sigma(x + s(tz + (1-t)\tilde{z})) - \sigma(x)|}{s} |z - \tilde{z}| + |\sigma'(x + s(tz + (1-t)\tilde{z}))(tz + (1-t)\tilde{z})| |z - \tilde{z}| \\ &\leq |tz + (1-t)\tilde{z}| |z - \tilde{z}| + L |z - \tilde{z}| \\ &\leq (|z| + |\tilde{z}|) |z - \tilde{z}| + |z - \tilde{z}| . \end{aligned}$$

The proof is complete by noting that $\mathcal{Y}_i(\mathbf{h}; s) = f(\mathbf{x}_i^T \mathbf{h})$. ■

Define a matrix $\mathbf{X} \in \mathbb{R}^{m \times d}$ with rows given by $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m \in \mathbb{R}^d$ and note that using the previous lemma allows us to handle the deviation of the process $\mathcal{Z}(\mathbf{h}; s)$. Specifically, using the triangle inequality we have

$$\begin{aligned} |\mathcal{Z}(\mathbf{h}; s) - \mathcal{Z}(\tilde{\mathbf{h}}; s)| &\leq \frac{1}{m} \sum_{i=1}^m (|\mathbf{x}_i^T \mathbf{h}| + |\mathbf{x}_i^T \tilde{\mathbf{h}}| + 1) |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})| \\ &= \frac{1}{m} \sum_{i=1}^m (|\mathbf{x}_i^T \mathbf{h}| + |\mathbf{x}_i^T \tilde{\mathbf{h}}|) |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})| + \frac{1}{m} \sum_{i=1}^m |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})| \\ &\leq \sqrt{\frac{1}{m} \sum_{i=1}^m (|\mathbf{x}_i^T \mathbf{h}| + |\mathbf{x}_i^T \tilde{\mathbf{h}}|)^2} \sqrt{\frac{1}{m} \sum_{i=1}^m |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})|^2} + \sqrt{\frac{1}{m} \sum_{i=1}^m |\mathbf{x}_i^T (\mathbf{h} - \tilde{\mathbf{h}})|^2} \\ &= \frac{1}{m} \|\mathbf{X}\mathbf{h} + \mathbf{X}\tilde{\mathbf{h}}\|_2 \|\mathbf{X}(\mathbf{h} - \tilde{\mathbf{h}})\|_2 + \frac{1}{\sqrt{m}} \|\mathbf{X}(\mathbf{h} - \tilde{\mathbf{h}})\|_2 \\ &\leq \frac{1}{m} (\|\mathbf{X}\mathbf{h}\|_2 + \|\mathbf{X}\tilde{\mathbf{h}}\|_2) \|\mathbf{X}(\mathbf{h} - \tilde{\mathbf{h}})\|_2 + \frac{1}{\sqrt{m}} \|\mathbf{X}(\mathbf{h} - \tilde{\mathbf{h}})\|_2 \\ &\leq \frac{2}{m} \|\mathbf{X}\|^2 \|\mathbf{h} - \tilde{\mathbf{h}}\|_2 + \frac{1}{\sqrt{m}} \|\mathbf{X}\| \|\mathbf{h} - \tilde{\mathbf{h}}\|_2 . \end{aligned}$$

To continue further, note that under the log-concave density, centered and isotropy assumption using (Adamczak et al., 2010), as long as $m \geq Cd$, we have that

$$\|\mathbf{X}\| \leq 2\sqrt{m}$$

holds with probability at least $1 - e^{-c\sqrt{d}}$. Thus, for any $\mathbf{h}, \tilde{\mathbf{h}} \in \mathbb{S}^{d-1}$ and any $s \leq CW$ we have

$$|\mathcal{Z}(\mathbf{h}; s) - \mathcal{Z}(\tilde{\mathbf{h}}; s)| \leq 10 \|\mathbf{h} - \tilde{\mathbf{h}}\|_2 \quad (6)$$

holds with high probability. Now let us consider an \mathcal{N}_η -cover of the unit sphere with $\eta := \frac{\xi\mu}{30}$. Using the union bound combined with (5) for any $\tilde{\mathbf{h}} \in \mathcal{N}_\eta$, we have that $\mathcal{Z}(\tilde{\mathbf{h}}; s) \geq \left(1 - \frac{\xi}{3}\right)\mu$ holds with probability at least

$$1 - \left(\frac{3}{\eta}\right)^d e^{-m \frac{\mu^2}{9\gamma^2} \xi^2} = 1 - \left(\frac{90}{\xi\mu}\right)^d e^{-m \frac{\mu^2}{9\gamma^2} \xi^2} = 1 - e^{d \log\left(\frac{90}{\xi\mu}\right) - m \frac{\mu^2}{9\gamma^2} \xi^2} \geq 1 - e^{-m \frac{\mu^2}{10\gamma^2} \xi^2},$$

as long as $m \geq 90 \frac{\gamma^2}{\mu^2} d \frac{\log\left(\frac{90}{\xi\mu}\right)}{\xi^2}$. Therefore, using (6) for any $\mathbf{h} \in \mathbb{S}^{d-1}$ there exists $\tilde{\mathbf{h}} \in \mathcal{N}_\eta$ with $\|\mathbf{h} - \tilde{\mathbf{h}}\|_2 \leq \eta := \frac{\xi}{30}$. Thus,

$$\mathcal{Z}(\mathbf{h}; s) \geq \mathcal{Z}(\tilde{\mathbf{h}}; s) - |\mathcal{Z}(\mathbf{h}; s) - \mathcal{Z}(\tilde{\mathbf{h}}; s)| \geq \left(1 - \frac{2}{3}\xi\right)\mu.$$

In conclusion, for all $\mathbf{h} \in \mathbb{S}^{d-1}$ and a fixed $0 \leq s \leq CW$ we have that

$$\mathcal{Z}(\mathbf{h}; s) \geq \left(1 - \frac{2}{3}\xi\right)\mu \quad (7)$$

holds with probability at least $1 - e^{-m \frac{\mu^2}{10\gamma^2} \xi^2} - e^{-c\sqrt{d}}$. We now turn our attention to making the result also hold uniformly for all $\frac{\epsilon}{3} \leq s \leq CW$. To this aim, we state the following lemma.

Lemma 25 *Let $s \geq \epsilon/3$, then*

$$|\mathcal{Y}_i(\mathbf{h}; s) - \mathcal{Y}_i(\mathbf{h}; \tilde{s})| \leq \frac{6}{\epsilon} |\mathbf{x}_i^T \mathbf{h}|^2 |s - \tilde{s}|.$$

Proof Define

$$f(s) := \frac{(\sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle) - \sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle))}{s} (\mathbf{x}_i^T \mathbf{h})$$

and note that for $s \geq \frac{\epsilon}{3}$

$$\begin{aligned} |f'(s)| &= \left| \frac{s\langle \mathbf{h}, \mathbf{x}_i \rangle \sigma'(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle) - (\sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle) - \sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle))}{s^2} (\mathbf{x}_i^T \mathbf{h}) \right| \\ &\leq \frac{1}{s} |\sigma'(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle)| |\mathbf{x}_i^T \mathbf{h}|^2 + \frac{|\sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle + s\langle \mathbf{h}, \mathbf{x}_i \rangle) - \sigma(\langle \hat{\mathbf{w}}, \mathbf{x}_i \rangle)|}{s^2} |\mathbf{x}_i^T \mathbf{h}| \\ &\leq \frac{2}{s} |\mathbf{x}_i^T \mathbf{h}|^2 \\ &\leq \frac{6}{\epsilon} |\mathbf{x}_i^T \mathbf{h}|^2. \end{aligned}$$

Thus, by the mean value theorem we have

$$\begin{aligned} |f(s) - f(\tilde{s})| &= |f'(ts + (1-t)s)(s - \tilde{s})| \\ &\leq \frac{6}{\epsilon} |\mathbf{x}_i^T \mathbf{h}|^2 |s - \tilde{s}|. \end{aligned}$$

■

Applying the above lemma, we have

$$\begin{aligned}
 |\mathcal{Z}(\mathbf{h}; s) - \mathcal{Z}(\mathbf{h}; \tilde{s})| &= \frac{6}{\epsilon} |s - \tilde{s}| \left(\frac{1}{m} \sum_{i=1}^m |\mathbf{x}_i^T \mathbf{h}|^2 \right) \\
 &= \frac{6}{\epsilon} |s - \tilde{s}| \frac{\|\mathbf{X}\mathbf{h}\|_2^2}{m} \\
 &\leq \frac{24}{\epsilon} |s - \tilde{s}|. \tag{8}
 \end{aligned}$$

Now let us consider an \mathcal{N}_η -cover of the $\frac{\epsilon}{3} \leq s \leq 4W$ with $\eta := \frac{\epsilon\mu}{72}$. Using the union bound combined with (5) for any $\tilde{s} \in \mathcal{N}_\eta$ we have $\mathcal{Z}(\mathbf{h}; \tilde{s}) \geq \left(1 - \frac{2\xi}{3}\right)\mu$ holds for all $\mathbf{h} \in \mathbb{S}^{d-1}$ with probability at least

$$1 - \frac{4W}{\eta} e^{-m \frac{\mu^2}{10\gamma^2} \xi^2} = 1 - \left(\frac{288W}{\mu\xi} \right) e^{-m \frac{\mu^2}{10\gamma^2} \xi^2} = 1 - e^{\log\left(\frac{288W}{\mu\xi}\right) - m \frac{\mu^2}{10\gamma^2} \xi^2} \geq 1 - e^{-m \frac{\mu^2}{11\gamma^2} \xi^2},$$

as long as $m \geq 110 \frac{\gamma^2 \log\left(\frac{288W}{\mu\xi}\right)}{\mu^2 \xi^2}$. Therefore, using (8) for all $\mathbf{h} \in \mathbb{S}^{d-1}$ and all $\frac{\epsilon}{3} \leq s \leq 4W$ there exists $\tilde{s} \in \mathcal{N}_\eta$ with $|s - \tilde{s}| \leq \eta := \frac{\epsilon\mu}{72}$. Thus,

$$\mathcal{Z}(\mathbf{h}; s) \geq \mathcal{Z}(\mathbf{h}; \tilde{s}) - |\mathcal{Z}(\mathbf{h}; s) - \mathcal{Z}(\mathbf{h}; \tilde{s})| \geq (1 - \xi)\mu.$$

In conclusion, for all $\mathbf{h} \in \mathbb{S}^{d-1}$ and all $\frac{\epsilon}{3} \leq s \leq CW$ we have that

$$\mathcal{Z}(\mathbf{h}; s) \geq (1 - \xi)\mu$$

holds with probability at least $1 - e^{-m \frac{\mu^2}{11\gamma^2} \xi^2} - e^{-c\sqrt{d}}$, as long as $m \geq 110 \frac{\gamma^2 \log\left(\frac{288(W+1)}{\mu\xi}\right)}{\mu^2 \xi^2}$; which in turn with $\xi = \frac{1}{3}$ implies that

$$\langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w} - \widehat{\mathbf{w}} \rangle \geq \frac{2}{3} \mu \|\mathbf{w} - \widehat{\mathbf{w}}\|_2^2 \tag{9}$$

holds for all \mathbf{w} obeying $\frac{\epsilon}{3} \leq \|\mathbf{w} - \widehat{\mathbf{w}}\|_2 \leq 4W$ with probability at least $1 - e^{-m \frac{\mu^2}{99\gamma^2}} - e^{-c\sqrt{d}} \geq 1 - \delta$ as long as $m \gtrsim \frac{\gamma^2}{\mu^2} d \log\left(\frac{(W+1)}{\mu\delta}\right)$.

Now note that for any vector $\mathbf{w} \in \mathbb{R}^d$ and $\widehat{\mathbf{w}}$ we have

$$\begin{aligned}
 \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}}), \mathbf{w} - \widehat{\mathbf{w}} \rangle &= \frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle)) (\mathbf{x}_i^T (\mathbf{w} - \widehat{\mathbf{w}})) \\
 &\geq \frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) - \sigma(\langle \widehat{\mathbf{w}}, \mathbf{x}_i \rangle))^2 \\
 &= \frac{1}{m} \|\sigma(\mathbf{X}\mathbf{w}) - \sigma(\mathbf{X}\widehat{\mathbf{w}})\|_2^2. \tag{10}
 \end{aligned}$$

Also note that

$$\|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\widehat{\mathbf{w}})\|_2^2 = \frac{1}{m^2} \left\| \sum_{i=1}^m (\sigma(\mathbf{w}^T \mathbf{x}_i) - \sigma(\widehat{\mathbf{w}}^T \mathbf{x}_i)) \mathbf{x}_i \right\|_2^2$$

$$\begin{aligned}
 &= \frac{1}{m^2} \|\mathbf{X}^T (\sigma(\mathbf{X}\mathbf{w}) - \sigma(\mathbf{X}\hat{\mathbf{w}}))\|_2^2 \\
 &\leq \frac{\|\mathbf{X}\|_2^2}{m^2} \|\sigma(\mathbf{X}\mathbf{w}) - \sigma(\mathbf{X}\hat{\mathbf{w}})\|_2^2 \\
 &\stackrel{(a)}{\leq} \frac{4}{m} \|\sigma(\mathbf{X}\mathbf{w}) - \sigma(\mathbf{X}\hat{\mathbf{w}})\|_2^2 \\
 &\stackrel{(b)}{\leq} 4 \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}}), \mathbf{w} - \hat{\mathbf{w}} \rangle.
 \end{aligned}$$

Here, (a) follows from the fact that under the log-concave density, centered and isotropy assumptions using [Adamczak et al. \(2010\)](#) as long as $m \geq Cd$ we have that

$$\|\mathbf{X}\| \leq 2\sqrt{m}$$

holds with probability at least $1 - e^{-c\sqrt{d}}$ and (b) follows from (10). Therefore,

$$\langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}}), \mathbf{w} - \hat{\mathbf{w}} \rangle \geq \frac{1}{4} \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}})\|_2^2. \quad (11)$$

Combining (9) and (11), we have that

$$\begin{aligned}
 &\langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}}), \mathbf{w} - \hat{\mathbf{w}} \rangle \\
 &= \frac{1}{2} \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}}), \mathbf{w} - \hat{\mathbf{w}} \rangle + \frac{1}{2} \langle \nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}}), \mathbf{w} - \hat{\mathbf{w}} \rangle \\
 &\geq \frac{\mu}{3} \|\mathbf{w} - \hat{\mathbf{w}}\|_2^2 + \frac{1}{8} \|\nabla \widehat{L}^{\text{surr}}(\mathbf{w}) - \nabla \widehat{L}^{\text{surr}}(\hat{\mathbf{w}})\|_2^2,
 \end{aligned}$$

which gives us the desired result.

Appendix C. Proof of Lemma 8

Recall that $\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}) = \mathbb{E}[\sigma(\langle \mathbf{w}, \mathbf{x} \rangle) \mathbf{x}] - \chi_{\mathcal{D}} = \chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}$. For any fixed \mathbf{w} observe that if y is at most 1, then for any 1-Lipschitz and monotone function σ , $(\sigma(\langle \mathbf{w}, \mathbf{x} \rangle) - y) \cdot \left\langle \mathbf{x}, \frac{\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}}{\|\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}\|_2} \right\rangle$ is the product of random variables with tails bounded by $\exp(-\Omega(x^t))$. This implies

$$\Pr_{S \sim \mathcal{D}^m} \left[\left(\frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \left\langle \mathbf{x}^{(i)}, \frac{\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}}{\|\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}\|_2} \right\rangle \right) - \|\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}\|_2 \geq \epsilon \|\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}\|_2 \right] \leq \exp(-(\sqrt{m}\epsilon)^t).$$

Using the variational form of the norm (i.e., $\|\mathbf{v}\|_2 := \max_{\mathbf{u}, \|\mathbf{u}\|_2=1} \langle \mathbf{u}, \mathbf{v} \rangle$) on

$$\left(\frac{1}{m} \sum_{i=1}^m (\text{ReLU}(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \mathbf{x}^{(i)} \right) - (\chi_{\mathbf{w}} - \chi),$$

we see that for any fixed $\mathbf{w} \in B(d, W)$

$$\Pr_{S \sim \mathcal{D}^m} \left[\left\| \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \mathbf{x}^{(i)} \right) - (\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}) \right\|_2 \geq \epsilon W \right] \leq \exp(-(\sqrt{m}\epsilon)^t).$$

Taking a union bound over a γ -net N_γ for $B(d, W)$ gives us

$$\Pr_{\mathcal{D}} \left[\forall \mathbf{w} \in N_\gamma \mid \left\| \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \mathbf{x}^{(i)}) - (\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}) \right\|_2 \geq \epsilon W \right] \leq \exp(-(\sqrt{m}\epsilon)^t) \cdot \left(\frac{3W}{\gamma} \right)^d,$$

i.e., for some constant C depending on the distribution, we get

$$\begin{aligned} \Pr_{\mathcal{D}} \left[\forall \mathbf{w} \in B(d, W) \mid \left\| \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \mathbf{x}^{(i)}) - (\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}) \right\|_2 \geq \epsilon W + C\gamma \right] \\ \leq \exp(-(\sqrt{m}\epsilon)^t) \cdot \left(\frac{3W}{\gamma} \right)^d. \end{aligned}$$

Hence, rescaling γ we see that when $m \geq \Omega\left(\left(\frac{d}{\epsilon} \log \frac{W}{\gamma} \log \frac{1}{\delta}\right)^{2/t}\right)$ we have with probability $1 - \delta$

$$\left\| \left(\frac{1}{m} \sum_{i=1}^m (\sigma(\langle \mathbf{w}, \mathbf{x}^{(i)} \rangle) - y^{(i)}) \cdot \mathbf{x}^{(i)}) - (\chi_{\mathcal{D}}^{\sigma \mathbf{w}} - \chi_{\mathcal{D}}) \right\|_2 \leq \epsilon W + \gamma.$$

Substituting $\gamma = \epsilon/2$ and rescaling ϵ we get the lemma.

Appendix D. Proof of Theorem 7

Let \mathbf{w}^{surr} be the minimizer of $L_{\mathcal{D}}^{\text{surr}}$, then we have, for all t

$$\begin{aligned} & \|\mathbf{w}^{(t+1)} - \mathbf{w}^{\text{surr}}\|_2^2 \\ & \leq \|\mathbf{w}^{(t+1)} - \mathbf{w}^{\text{surr}}\|_2^2 \\ & = \|\mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}}\|_2^2 - \eta \langle \nabla \hat{L}_S^{\text{surr}}(\mathbf{w}^{(t)}), \mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}} \rangle + \eta^2 \|\nabla \hat{L}_S^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 \\ & \leq \|\mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}}\|_2^2 - \eta \langle \nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)}), \mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}} \rangle - \eta \langle \nabla \hat{L}_S^{\text{surr}}(\mathbf{w}^{(t)}) - \nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)}), \mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}} \rangle \\ & \quad + 2\eta^2 \|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 + 2\eta^2 \|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)}) - \nabla \hat{L}_S^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 \\ & \leq \|\mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}}\|_2^2 - \eta(1 - 2\eta) \|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 + 2\eta\epsilon W + 2\eta^2\epsilon^2. \end{aligned}$$

For the final inequality, we use the smoothness and strong convexity of $L_{\mathcal{D}}^{\text{surr}}$ to see

$$L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{\text{surr}}) \geq L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^t) - \langle \nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^t), \mathbf{w}^t - \mathbf{w}^{\text{surr}} \rangle + \frac{\mu}{2} \|\mathbf{w} - \mathbf{w}'\|_2^2$$

and

$$-\|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{\text{surr}})\|_2^2 \leq L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{\text{surr}}) - L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^t).$$

Additionally, for $\eta < 1/4$, either $\|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 \leq 4\frac{\epsilon W + \eta\epsilon^2}{1 - 2\eta}$ or $\|\mathbf{w}^{(t+1)} - \mathbf{w}^{\text{surr}}\|_2^2 \leq \|\mathbf{w}^{(t)} - \mathbf{w}^{\text{surr}}\|_2^2 - 2\eta\epsilon W - 2\eta^2\epsilon^2$. Therefore, after $T \geq \frac{2W^2}{\eta\epsilon W + \eta^2\epsilon^2}$ iterations, there must exist some $t \leq T$ such that $\|\nabla L_{\mathcal{D}}^{\text{surr}}(\mathbf{w}^{(t)})\|_2^2 \leq 4\frac{\epsilon W + \eta\epsilon^2}{1 - 2\eta} \leq 8\epsilon W + 2\epsilon^2$. Scaling ϵ appropriately gives us the result.

Appendix E. Proof of Lemma 12

By Definition 9, we have that for all \mathbf{u}, \mathbf{v} ,

$$\langle \chi_{\mathcal{D}}^{\sigma_{\mathbf{u}}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{v}}}, \mathbf{u} - \mathbf{v} \rangle \geq \mu \|\mathbf{u} - \mathbf{v}\|_2^2 \implies \|\chi_{\mathcal{D}}^{\sigma_{\mathbf{u}}} - \chi_{\mathcal{D}}^{\sigma_{\mathbf{v}}}\|_2^2 \geq \mu^2 \|\mathbf{u} - \mathbf{v}\|_2^2.$$

Also by 1-Lipschitzness of σ and isotropy of $\mathcal{D}_{\mathcal{X}}$, we have

$$L_{\mathcal{D}}(\sigma_{\mathbf{u}}, \sigma_{\mathbf{v}}) \leq \mathbb{E}[\langle \mathbf{u} - \mathbf{v}, \mathbf{x} \rangle^2] = \|\mathbf{v} - \mathbf{u}\|_2^2.$$

Combining the above gives us the desired result.

Appendix F. Proof of Lemma 17

If $\mathbf{u} = \mathbf{v}$, then the claim follows directly. Let $\mathbf{u} \neq \mathbf{v}$. In the calculation below, we use monotonicity as well as 1-Lipschitzness.

$$\begin{aligned} (\chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{v}}} - \chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{u}}})^T(\mathbf{v} - \mathbf{u}) &= \mathbb{E}[(\text{ReLU}(\langle \mathbf{v}, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{u}, \mathbf{x} \rangle)) \langle \mathbf{v} - \mathbf{u}, \mathbf{x} \rangle] \\ &\geq \mathbb{E}[(\text{ReLU}(\langle \mathbf{v}, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{u}, \mathbf{x} \rangle))^2]. \end{aligned}$$

The above term depends only on $\langle \mathbf{u}, \mathbf{x} \rangle$ and $\langle \mathbf{v}, \mathbf{x} \rangle$. To bound this, it is sufficient to work with $d = 2$. Let $f(x_1, x_2)$ be the density function for the log-concave distribution. The above term can be bounded based on the following two cases,

- **Case 1** ($\theta(\mathbf{u}, \mathbf{v}) \leq \pi/2$) We have,

$$\begin{aligned} &(\chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{v}}} - \chi_{\mathcal{D}}^{\text{ReLU}_{\mathbf{u}}})^T(\mathbf{v} - \mathbf{u}) \\ &\geq \mathbb{E}[\langle (\mathbf{v} - \mathbf{u}) \cdot \mathbf{x} \rangle^2 \mathbb{1}[\langle \mathbf{u}, \mathbf{x} \rangle \geq 0, \langle \mathbf{v}, \mathbf{x} \rangle \geq 0]] \\ &\geq \|\mathbf{v} - \mathbf{u}\|_2^2 \mathbb{E}[\langle \overline{\mathbf{v} - \mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0]] \\ &= \|\mathbf{v} - \mathbf{u}\|_2^2 \int_{\mathbb{R}^2} \langle \overline{\mathbf{v} - \mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0] f(x_1, x_2) dx_1 dx_2 \\ &\geq c \|\mathbf{v} - \mathbf{u}\|_2^2 \int_{\|\mathbf{x}\|_2 \leq 1/9} \langle \overline{\mathbf{v} - \mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0] dx_1 dx_2. \end{aligned}$$

Here the last inequality follows from the anti-concentration of the log-concave distribution. To prove strong convexity of the surrogate loss, it is sufficient to bound from below the above integral by a constant. Since the angle between \mathbf{u} and \mathbf{v} is less than $\frac{\pi}{2}$, we see that with respect to the uniform measure the set $\{\mathbf{x} \mid \langle \mathbf{x}, \bar{\mathbf{u}} \rangle > 0 \text{ and } \langle \mathbf{x}, \bar{\mathbf{v}} \rangle > 0\}$ has mass $> \frac{\text{vol}(B(1/9))}{4}$. Thus,

$$\int_{\|\mathbf{x}\|_2 \leq 1/9} \langle \overline{\mathbf{v} - \mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0] dx_1 dx_2 \geq \min_{\substack{\mathbf{w}: \|\mathbf{w}\|_2=1 \\ S \subseteq B(1/9): \text{vol}(S) = \frac{\text{vol}(B(1/9))}{4}}} \int_{\mathbf{x} \in S} \langle \mathbf{w}, \mathbf{x} \rangle^2 dx_1 dx_2.$$

To bound from below the integral above, let $\mathbf{x} = \alpha \mathbf{w} + \beta \mathbf{w}^{\perp}$ – here we abuse notation slightly to and use \mathbf{w} to denote the minimizer of the integral above. Since we are in 2 dimensions, the set that minimizes the integral is given by the region that minimizes $\langle \mathbf{w}, \mathbf{x} \rangle$, subject to the volume constraint. Using the fact that in a ball of radius $1/9$ the distribution is lower bounded

by a log-concave distribution, we see that it is sufficient to lower bound the integral for the above set $\{\alpha\mathbf{w} + \beta\mathbf{w}^\perp \mid |\alpha| < \gamma, \alpha^2 + \beta^2 \leq \frac{1}{9}\}$, for some constant γ .

$$\begin{aligned}
 & \int_{\|\mathbf{x}\|_2 \leq 1/9} (\langle \bar{\mathbf{v}} - \bar{\mathbf{u}}, \mathbf{x} \rangle)^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0] dx_1 dx_2 \\
 & \geq \min_{S \mid \Pr[S] = \Pr[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \geq 0]} \int_{\alpha^2 + \beta^2 < 1/9} \alpha^2 \mathbb{1}[S] d\alpha d\beta \\
 & \geq \int_{\alpha^2 + \beta^2 < 1/9, |\alpha| < \gamma} \alpha^2 d\alpha d\beta \\
 & \geq \int_{\alpha^2 + \beta^2 < 1/9, \frac{\gamma}{2} < \alpha < \gamma} \alpha^2 d\alpha d\beta \\
 & \geq \int_{\alpha^2 + \beta^2 < 1/9, \frac{\gamma}{2} < \alpha < \gamma} \frac{\gamma^2}{4} d\alpha d\beta \\
 & \geq \frac{\gamma^2}{4} \cdot c'.
 \end{aligned}$$

Here c' is a constant satisfying $c' = \text{vol}(\{(\alpha, \beta) \mid \alpha^2 + \beta^2 < 1/9, \alpha \in [\gamma/2, \gamma]\})$.

- **Case 2** ($\theta(\mathbf{u}, \mathbf{v}) > \pi/2$): We assume w.l.o.g. that $\|\mathbf{u}\|_2 \geq \|\mathbf{v}\|_2$. Similar to the previous case, we have that

$$\begin{aligned}
 (\chi_{\mathcal{D}}^{\text{ReLU}_v} - \chi_{\mathcal{D}}^{\text{ReLU}_u})^T(\mathbf{v} - \mathbf{u}) & \geq \mathbb{E}[\langle \mathbf{u}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \mathbf{u}, \mathbf{x} \rangle \geq 0, \langle \mathbf{v}, \mathbf{x} \rangle \leq 0]] \\
 & \geq \|\mathbf{u}\|_2^2 \mathbb{E}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \leq 0]] \\
 & \geq c \|\mathbf{u}\|_2^2 \int_{\|\mathbf{x}\|_2 \leq 1/9} \langle \bar{\mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \leq 0] dx_1 dx_2 \\
 & \geq \frac{c}{2} \|\mathbf{u} - \mathbf{v}\|_2^2 \int_{\|\mathbf{x}\|_2 \leq 1/9} \langle \bar{\mathbf{u}}, \mathbf{x} \rangle^2 \mathbb{1}[\langle \bar{\mathbf{u}}, \mathbf{x} \rangle \geq 0, \langle \bar{\mathbf{v}}, \mathbf{x} \rangle \leq 0] dx_1 dx_2.
 \end{aligned}$$

Since the angle between \mathbf{u} and \mathbf{v} is more than $\frac{\pi}{2}$ we see that with respect to the uniform measure the set $\{\mathbf{x} \mid \langle \mathbf{x}, \bar{\mathbf{u}} \rangle > 0 \text{ and } \langle \mathbf{x}, \bar{\mathbf{v}} \rangle < 0\}$ has mass $> \frac{\text{vol}(B(1/9))}{4}$. The final integral above can again be lower bounded as in Case 1.

Appendix G. Proof of Lemma 21

Let \mathbf{w}, \mathbf{w}^* be as defined in Algorithm 2. We first project \mathcal{D}_x down to two dimensions. Let $S = \{\mathbf{x} \mid \langle \mathbf{w}^*, \mathbf{x} \rangle > 0\}$. Let V be the 2-dimensional space spanned by \mathbf{w}, \mathbf{w}^* and let P_V be the orthogonal projection onto V . If $\mathbf{x} \in \bar{S} \wedge T_+$, then $\langle \mathbf{w}, \mathbf{x} \rangle \geq \gamma\sqrt{\text{opt}}$ and $\langle \mathbf{w}^*, \mathbf{x} \rangle \leq 0$. Since \mathbf{w} is a constant factor approximation for a ν -subgaussian distribution, with probability $1 - \delta$ we have $\|\mathbf{w} - \mathbf{w}^*\|_2 \leq \frac{c_\chi}{\nu} \cdot \sqrt{\text{opt}}$, for some constant c_χ . This is easy to check via the Cauchy-Schwartz inequality and using the structural lemmas from previous subsections. Additionally

$$-\langle \mathbf{w}^*, \mathbf{x} \rangle = -\langle \mathbf{w}^*, P_V(\mathbf{x}) \rangle = -\langle \mathbf{w}^* - \mathbf{w}, P_V(\mathbf{x}) \rangle - \langle \mathbf{w}, P_V(\mathbf{x}) \rangle \quad (12)$$

$$\leq -\langle \mathbf{w}^* - \mathbf{w}, P_V(\mathbf{x}) \rangle \leq \|\mathbf{w}^* - \mathbf{w}\|_2 \|P_V(\mathbf{x})\|_2 \leq \frac{c_\chi}{\nu} \sqrt{\text{opt}} \|P_V(\mathbf{x})\|_2. \quad (13)$$

A similar calculation for $\mathbf{x} \in S \wedge T_-$ implies $\langle \mathbf{w}^*, \mathbf{x} \rangle \leq \frac{c_\chi}{\nu} \sqrt{\text{opt}} \|P_V(\mathbf{x})\|_2$. We now bound above the error in the region T_+ . Since $\langle \mathbf{w}_+, \mathbf{x} \rangle$ is the best fitting linear function over T_+ , the loss of $\langle \mathbf{w}^*, \mathbf{x} \rangle$




Figure 1: We partition the space into three regions depending on $\langle \mathbf{w}, \mathbf{x} \rangle$. Our hypothesis returns 0, $\langle \mathbf{w}, \mathbf{x} \rangle$ or the value of $P(\mathbf{x})$ depending on the region. We set $t = \gamma\sqrt{\text{opt}}$.

is necessarily larger than that of $\langle \mathbf{w}_+, \mathbf{x} \rangle$. An application of Lemma 31 in the first step implies

$$\begin{aligned} \mathbb{E}_S [(\langle \mathbf{w}_+, \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] &= \min_{\mathbf{w}' \in B(d, W)} \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}', \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] + \epsilon \Pr[T_+] \\ &\leq \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] + \epsilon \Pr[T_+]. \end{aligned}$$

Observe that for $\mathbf{x} \in S$, $\langle \mathbf{w}^*, \mathbf{x} \rangle = \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle)$. Since $1_{T_+}(\mathbf{x}) = 1_{T_+ \wedge S}(\mathbf{x}) + 1_{T_+ \wedge \bar{S}}(\mathbf{x})$, we get

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] &= \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle - y)^2 1_{T_+ \wedge S}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle - y)^2 1_{T_+ \wedge \bar{S}}(\mathbf{x})] \\ &= \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_{T_+ \wedge S}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}} [y^2 1_{T_+ \wedge \bar{S}}(\mathbf{x})] \\ &\quad + \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle)(\langle \mathbf{w}^*, \mathbf{x} \rangle - 2y) 1_{T_+ \wedge \bar{S}}(\mathbf{x})]. \end{aligned}$$

It remains to show that the terms corresponding to $1_{T_+ \wedge \bar{S}}$ contribute a small error overall. Note that $\mathbb{E}_{\mathcal{D}} [y^2 1_{T_+ \wedge \bar{S}}(\mathbf{x})] = \mathbb{E}_{\mathcal{D}} [(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_+ \wedge \bar{S}}(\mathbf{x})]$. This implies

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle - y)^2 1_{T_+}(\mathbf{x})] \\ &= \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_{T_+}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}} [(\langle \mathbf{w}^*, \mathbf{x} \rangle)(\langle \mathbf{w}^*, \mathbf{x} \rangle - 2y) 1_{T_+ \wedge \bar{S}}(\mathbf{x})] \\ &\leq \mathbb{E}_{\mathcal{D}} [(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_{T_+}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}} [|\langle \mathbf{w}^*, \mathbf{x} \rangle| (|\langle \mathbf{w}^*, \mathbf{x} \rangle| + 2) 1_{T_+ \wedge \bar{S}}(\mathbf{x})]. \end{aligned}$$

To bound above the second term, we use bounds from Equation (12), and an application of Lemma 27.

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}} [|\langle \mathbf{w}^*, \mathbf{x} \rangle| (|\langle \mathbf{w}^*, \mathbf{x} \rangle| + 2) 1_{T_+ \wedge \bar{S}}(\mathbf{x})] \\ &\leq \mathbb{E}_{\mathcal{D}} \left[\frac{c_X}{\nu} \sqrt{\text{opt}} \|P_V(\mathbf{x})\|_2 \left(\frac{c_X}{\nu} \sqrt{\text{opt}} \|P_V(\mathbf{x})\|_2 + 2 \right) 1_{T_+ \wedge \bar{S}}(\mathbf{x}) \right] \\ &= \frac{c_X}{\nu} \cdot \sqrt{\text{opt}} \cdot \mathbb{E}_{\mathcal{D}} \left[\left(\frac{c_X}{\nu} \sqrt{\text{opt}} \|P_V(\mathbf{x})\|_2^2 + 2 \|P_V(\mathbf{x})\|_2 \right) 1_{T_+ \wedge \bar{S}}(\mathbf{x}) \right] \\ &\leq \frac{\eta}{6} \cdot \text{opt}. \end{aligned}$$

Overall, the three equation blocks above imply that the first condition of the lemma is true. The analysis of the error of h in T_- is done similarly. We can write:

$$\mathbb{E}_{\mathcal{D}} [y^2 1_{T_-}(\mathbf{x})]$$

$$\begin{aligned}
 &= \mathbb{E}_{\mathcal{D}}[y^2 1_{T_- \wedge S}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}}[y^2 1_{T_- \wedge \bar{S}}(\mathbf{x})] \\
 &= \mathbb{E}_{\mathcal{D}}[(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) + \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_- \wedge S}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}}[(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_- \wedge \bar{S}}(\mathbf{x})] \\
 &= \mathbb{E}_{\mathcal{D}}[(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_-}(\mathbf{x})] + \mathbb{E}_{\mathcal{D}}[\langle \mathbf{w}^*, \mathbf{x} \rangle (2y - \langle \mathbf{w}^*, \mathbf{x} \rangle) 1_{T_- \wedge S}(\mathbf{x})] \\
 &\leq \mathbb{E}_{\mathcal{D}}[(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_-}(\mathbf{x})] + 2\mathbb{E}_{\mathcal{D}}[|\langle \mathbf{w}^*, \mathbf{x} \rangle| (2 - |\langle \mathbf{w}^*, \mathbf{x} \rangle|) 1_{T_- \wedge S}(\mathbf{x})] \\
 &\leq \mathbb{E}_{\mathcal{D}}[(y - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_{T_-}(\mathbf{x})] + \frac{\eta}{6} \cdot \text{opt} .
 \end{aligned}$$

Finally, we analyze the error of our hypothesis in the region T .

$$\begin{aligned}
 &\mathbb{E}_{\mathcal{D}}[(P(\mathbf{x}) - y)^2 1_T(\mathbf{x})] \\
 &= \mathbb{E}_{\mathcal{D}}[(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_T(\mathbf{x})] + 2\mathbb{E}_{\mathcal{D}}[(P(\mathbf{x}) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle)) (\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y) 1_T(\mathbf{x})] \\
 &\quad + \mathbb{E}_{\mathcal{D}}[(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - P(\mathbf{x}))^2 1_T(\mathbf{x})] .
 \end{aligned}$$

The final error term is bounded above via applications of Lemma 31 and Lemma 29 by $\frac{\eta^2}{10} \cdot \text{opt}$. To bound from above the cross term, we use the fact that $1_T(\mathbf{x})^2 = 1_T(\mathbf{x})$

$$\begin{aligned}
 &2\mathbb{E}_{\mathcal{D}}[(P(\mathbf{x}) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle)) 1_T(\mathbf{x}) (\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y) 1_T(\mathbf{x})] \\
 &\leq \sqrt{\mathbb{E}_{\mathcal{D}}[(P(\mathbf{x}) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 1_T(\mathbf{x})]} \sqrt{\mathbb{E}_{\mathcal{D}}[(\text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle) - y)^2 1_T(\mathbf{x})]} \\
 &\leq \frac{\eta}{10} \cdot \text{opt} .
 \end{aligned}$$

Putting these together gives us the desired result.

Appendix H. Lemmas used in Appendix G

Lemma 26 Suppose $\|\mathbf{w} - \mathbf{w}^*\|_2 \leq O\left(\frac{\sqrt{\text{opt}}}{\nu}\right)$ and let $\theta(\mathbf{w}, \mathbf{w}^*)$ represent the angle between \mathbf{w} and \mathbf{w}^* with $\|\mathbf{w}^*\|_2 = 1$. If $\theta(\mathbf{w}, \mathbf{w}^*) \leq \pi/2$ then

$$\frac{\theta(\mathbf{w}, \mathbf{w}^*)}{2} \leq \sin(\theta(\mathbf{w}, \mathbf{w}^*)) \leq O\left(\frac{\sqrt{\text{opt}}}{\nu}\right) .$$

Proof Since $\|\mathbf{w} - \mathbf{w}^*\|_2 \leq O\left(\frac{\sqrt{\text{opt}}}{\nu}\right)$,

$$\begin{aligned}
 O(\text{opt}/\nu^2) &\geq \|\mathbf{w} - \mathbf{w}^*\|_2^2 \\
 &= \|\mathbf{w}\|_2^2 + \|\mathbf{w}^*\|_2^2 - 2\|\mathbf{w}\|_2 \|\mathbf{w}^*\|_2 \cos(\theta(\mathbf{w}, \mathbf{w}^*)) \\
 &= \|\mathbf{w}\|_2^2 + 1 - 2\|\mathbf{w}\|_2 \cos(\theta(\mathbf{w}, \mathbf{w}^*)) .
 \end{aligned}$$

This implies

$$\cos(\theta(\mathbf{w}, \mathbf{w}^*)) \geq \frac{\|\mathbf{w}\|_2^2 + 1 - O(\text{opt}/\nu^2)}{2\|\mathbf{w}\|_2} \geq \sqrt{1 - O(\text{opt}/\nu^2)} .$$

i.e., $\sin(\theta(\mathbf{w}, \mathbf{w}^*)) \leq O\left(\frac{\sqrt{\text{opt}}}{\nu}\right)$. Since $\theta(\mathbf{w}, \mathbf{w}^*) \in [0, \pi/2)$ we have $\frac{\theta(\mathbf{w}, \mathbf{w}^*)}{2} \leq \sin(\theta(\mathbf{w}, \mathbf{w}^*))$. ■

Lemma 27 *If \mathbf{z} is drawn from a ν -subgaussian distribution, $\mathbf{z} \in A := \{\mathbf{x} \mid \langle \mathbf{x}, \mathbf{w} \rangle \geq \gamma\sqrt{\text{opt}}$ and $\langle \mathbf{w}^*, \mathbf{x} \rangle \leq 0\}$ and $\|\mathbf{w} - \mathbf{w}^*\|_2 \leq O(\sqrt{\text{opt}/\nu})$, then*

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\left(\frac{\Theta(1)}{\nu} \sqrt{\text{opt}} \|\mathbf{z}\|_2^2 + 2\|\mathbf{z}\|_2 \right) 1_A(\mathbf{z}) \right] \leq \nu \cdot \sqrt{\text{opt}} \cdot \frac{\eta}{10}.$$

Proof Since $\mathbf{z} \in A$,

$$\gamma\sqrt{\text{opt}} \leq (\mathbf{w} - \mathbf{w}^*) \cdot \mathbf{z} \leq \|\mathbf{w} - \mathbf{w}^*\|_2 \|\mathbf{z}\|_2 \lesssim \frac{1}{\nu} \sqrt{\text{opt}} \|\mathbf{z}\|_2.$$

Hence, $\|\mathbf{z}\|_2 \geq \Omega(\nu\gamma)$. To bound from above the expectation in question, we integrate in polar coordinates. Specifically, we can write

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}} \left[\left(\frac{\Theta(1)}{\nu} \sqrt{\text{opt}} \|\mathbf{z}\|_2^2 + 2\|\mathbf{z}\|_2 \right) 1_{T_+ \wedge \bar{S}}(\mathbf{z}) \right] \\ & \lesssim \frac{1}{\nu} \int_0^{2\pi} \int_0^\infty \left(\frac{\Theta(1)}{\nu} \sqrt{\text{opt}} r^3 + 2r^2 \right) 1_{T_+ \wedge \bar{S}}(r, \theta) \exp(-r^2/2\nu^2) dr d\theta \\ & \lesssim \frac{\theta(\mathbf{w}, \mathbf{w}^*)}{\nu} \int_{\Omega(\nu\gamma)}^\infty \left(\frac{\Theta(1)}{\nu} \sqrt{\text{opt}} r^3 + 2r^2 \right) \exp(-r^2/2\nu^2) dr \\ & \lesssim \theta(\mathbf{w}, \mathbf{w}^*) \int_{\Omega(\gamma)}^\infty \left(\frac{1}{\nu} \sqrt{\text{opt}} (\nu s)^3 + (\nu s)^2 \right) \exp(-s^2/2) ds \\ & = \nu^2 \theta(\mathbf{w}, \mathbf{w}^*) \int_{\Omega(\gamma)}^\infty (\sqrt{\text{opt}} s^3 + s^2) \exp(-s^2/2) ds, \end{aligned}$$

where the final inequality is a consequence of $r = s\nu$ and $\gamma \geq 0$. We now use the following facts about the gaussian integrals of a x^2 and x^3 to bound from above the previous integral:

$$\int_t^\infty x^2 \exp\left(-\frac{x^2}{2}\right) dx \leq O\left(\text{erfc}\left(\frac{t}{\sqrt{2}}\right) + 2t \exp(-t^2/2)\right) \leq O(1+t) \cdot \exp(-t^2/2) \quad (14)$$

$$\int_t^\infty x^3 \exp\left(-\frac{x^2}{2}\right) dx = (2+t^2) \cdot \exp(-t^2/2) \quad (15)$$

The final inequality in Equation (14) follows from $\text{erfc}(x) \leq 2 \exp(-x^2/2)$. Making these substitutions and using Lemma 26 to get the bound $\theta(\mathbf{w}, \mathbf{w}^*) \leq O\left(\frac{\sqrt{\text{opt}}}{\nu}\right)$, we see

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} [((c/\nu)\sqrt{\text{opt}} \|\mathbf{z}\|_2^2 + 2\|\mathbf{z}\|_2) 1_{T_+ \wedge \bar{S}}(\mathbf{z})] & \leq O(\theta(\mathbf{w}, \mathbf{w}^*) \sqrt{\text{opt}} \cdot \nu^2 \cdot (2+\gamma^2) \cdot \exp(-\gamma^2/2)) \\ & \quad + O(\theta(\mathbf{w}, \mathbf{w}^*) \cdot \nu^2 \cdot (1+\gamma) \cdot \exp(-\gamma^2/2)) \\ & \lesssim \nu \sqrt{\text{opt}} \cdot \max\{(1+\gamma^2), (1+\gamma)\} \exp(-\gamma^2/2). \end{aligned}$$

The choice $\gamma = \Omega\left(\sqrt{\log\left(\frac{1}{\eta}\right)}\right)$ ensures that this is at most $\nu \cdot \sqrt{\text{opt}} \cdot \frac{\eta}{10}$. ■

Lemma 28 (Probability of being in the band) *If \mathcal{D} is ν -subgaussian, and \mathbf{w} is the minimizer of the surrogate loss*

$$\Pr_{\mathcal{D}} [|\langle \mathbf{w}, \mathbf{x} \rangle| \leq \gamma\sqrt{\text{opt}}] \leq \frac{c'\gamma\sqrt{\text{opt}}}{\nu}.$$

Proof By standard properties of $\mathcal{N}(0, \nu I)$, we get the result for $\|\mathbf{w}\|_2 = 1$. However, we know that $\|\mathbf{w}\|_2 \geq \|\mathbf{w}^*\|_2 - \sqrt{\eta} = 1 - \sqrt{\eta}$ and we know that $1 - \sqrt{\eta}$ is larger than a constant which gives us the desired result. \blacksquare

H.1. Polynomial Approximation in the Band

We show that there exists a low-degree polynomial approximation of the ReLU in squared error over the band.

Lemma 29 *If \mathcal{D} is ν -subgaussian, then there exists a degree $O(\frac{1}{\eta^3})$ polynomial P satisfying*

$$\mathbb{E}_{\mathcal{D}|T_{d,\gamma}} \left[(P(\langle \mathbf{w}^*, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 \right] \lesssim \eta^2 \cdot \nu \cdot \sqrt{\text{opt}} \quad (16)$$

Here, $T_{d,\gamma}(\mathbf{w}) = \{\mathbf{u} \in \mathbb{R}^d \mid |\langle \mathbf{w}, \mathbf{u} \rangle| \leq \gamma \sqrt{\text{opt}}\}$, where $\gamma = \sqrt{\log(1/\eta)}$.

Proof An application of Jackson's theorem implies there is a degree $\frac{6s}{\tau}$ polynomial such that

$$\|P(t) - \text{ReLU}(t)\|_{[-s,s],\infty} \leq \tau.$$

We apply the following Theorem 30 from Sherstov (2012) to see that $P(t)$ satisfies

$$|P(t) - \text{ReLU}(t)| < 2(4t/s)^{\frac{6s}{\tau}} \text{ for } t \in \mathbb{R} \setminus [-s, s].$$

Lemma 30 (Sherstov) *Let $p(t) := \sum_{i=1}^d a_i t^i$ be a given polynomial, then*

$$\sum_i |a_i| \leq 4^d \max_{i=0,\dots,d} \left| p\left(\frac{d-2j}{d}\right) \right|.$$

To estimate $\mathbb{E}_{\mathcal{D}|T} \left[(P(\langle \mathbf{w}^*, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 \right]$, we bound from above the density function of $\langle \mathbf{w}^*, \mathbf{x} \rangle$ where $\mathbf{x} \sim \mathcal{D}_T$ (denoted by $\rho_{\gamma,\delta}$).

Since we only need to look at the distribution $\langle \mathbf{w}^*, \mathbf{x} \rangle$, it is sufficient to project the distribution \mathcal{D} to $\text{span}(\mathbf{w}, \mathbf{w}^*)$. Henceforth, we will abuse notation and let \mathcal{D} refer to this projected distribution.

Let $\mathbf{x} \sim \mathcal{D}_T$ and $\bar{\mathbf{w}} = \mathbf{w}/\|\mathbf{w}\|_2$. Express \mathbf{x} in the basis $\bar{\mathbf{w}}, \bar{\mathbf{w}}^\perp$ to get $\mathbf{x} = \alpha_x \bar{\mathbf{w}} + \beta_x \bar{\mathbf{w}}^\perp$. We now study the random variable $\langle \mathbf{w}^*, \mathbf{x} \rangle$. Let $\theta = \theta(\mathbf{w}, \mathbf{w}^*)$ and note that $\|\mathbf{w}^*\|_2 = 1$. For any $\mathbf{x} \in T_{\gamma,d}(\mathbf{w})$. Hence,

$$\begin{aligned} \langle \mathbf{w}^*, \mathbf{x} \rangle &= \alpha_x \langle \mathbf{w}^*, \bar{\mathbf{w}} \rangle + \beta_x \langle \mathbf{w}^*, \bar{\mathbf{w}}^\perp \rangle \\ &= \alpha_x \cos(\theta) + \beta_x \sin(\theta), \end{aligned}$$

where $\mathbf{w}^{*\perp} = \mathbf{w}^* - \bar{\mathbf{w}}(\bar{\mathbf{w}} \cdot \mathbf{w}^*)$ and $\|\mathbf{w}^{*\perp}\|_2 = \sin(\theta)$. We now individually upper bound the pdf of $\alpha_x \cos(\theta)$ and $\beta_x \langle \mathbf{w}^{*\perp}, \bar{\mathbf{w}}^\perp \rangle$ respectively, when $\mathbf{x} \sim \mathcal{D}_T$. Recall that \mathcal{D} is ν -subgaussian, hence the pdf of $\alpha_x, \rho_{\alpha_x}$ is upper bounded by $\frac{1}{\nu} \exp(-x^2/2\nu^2)$. Similarly, for β_x . If $f_{a,b}$ is the pdf of the joint distribution over $(a, b) \in \mathbb{R}^2$, then the pdf of $a + b$ is given by $\rho_{a+b}(z) = \int_{-\infty}^{\infty} f_{a,b}(t, z - t) dt$. For $\mathbf{x} \sim \mathcal{D}_T$, the pdf of \mathbf{x} in the basis above (i.e., the pdf of (α_x, β_x)) is upper bounded

by $\rho_{\mathcal{D}}(\mathbf{x}) \leq \frac{1}{\Pr_{\mathcal{D}}[T]} \frac{100}{\nu^2} \cdot \exp(-(\alpha_{\mathbf{x}}^2 + \beta_{\mathbf{x}}^2)/2\nu^2)$, when $\alpha_{\mathbf{x}} \leq \gamma\sqrt{\text{opt}}\|\mathbf{w}\|_2$, and 0 otherwise. Since $\langle \mathbf{w}^*, \mathbf{x} \rangle = \alpha_{\mathbf{x}} \cos(\theta) + \beta_{\mathbf{x}} \sin(\theta)$, the pdf of $\langle \mathbf{w}^*, \mathbf{x} \rangle$, $\rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}$ is bounded by

$$\rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(z) \lesssim \frac{1}{\Pr[T]} \int_{-b}^b \frac{1}{\nu^2 \sin(\theta) \cos(\theta)} \cdot \exp(-(t^2/\cos^2(\theta) + (z-t)^2/\sin^2(\theta))/2\nu^2) dt.$$

We bound the above expression in two ways. First, since

$$\exp(-(z-t)^2/2\sin^2(\theta)\nu^2) \leq \exp(-t^2/2\sin^2(\theta)\nu^2) \leq 1,$$

we can write

$$\begin{aligned} \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(z) &\lesssim \frac{1}{\Pr[T]} \int_{-b}^b \frac{1}{\nu^2 \sin(\theta) \cos(\theta)} \cdot \exp(-(t^2/\cos^2(\theta) + (z-t)^2/\sin^2(\theta))/2\nu^2) dt \\ &\lesssim \frac{1}{\Pr[T]} \int_{-b}^b \frac{1}{\nu^2 \sin(\theta) \cos(\theta)} \cdot \exp(-(t^2/\cos^2(\theta) + t^2/\sin^2(\theta))/2\nu^2) dt \\ &\lesssim \frac{1}{\Pr[T]\nu \sin(\theta)} \int_{-b}^b \frac{1}{\nu \cos(\theta)} \cdot \exp(-(t^2/2\nu^2 \cos^2(\theta))) dt \\ &\lesssim \frac{1}{\nu \sin(\theta)}. \end{aligned}$$

Also, if for some l it holds $|z| \geq l$, we have that

$$\begin{aligned} \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(z) &\lesssim \frac{1}{\Pr[T]} \int_{-b}^b \frac{1}{\nu^2 \sin(\theta) \cos(\theta)} \cdot \exp(-(t^2/\cos^2(\theta) + (z-t)^2/\sin^2(\theta))/2\nu^2) dt \\ &\lesssim \frac{1}{\Pr[T]} \int_{-b}^b \frac{1}{\nu^2 \sin(\theta) \cos(\theta)} \cdot \exp(-(t^2/\cos^2(\theta) + (z-l)^2/\sin^2(\theta))/2\nu^2) dt \\ &\lesssim \exp(-(z-l)^2/2\nu^2 \sin^2(\theta)) \cdot (1/\nu \sin(\theta)). \end{aligned}$$

We will use the above bound with $l = \gamma\sqrt{\text{opt}}\|\mathbf{w}\|_2$. We now use the above upper bounds to bound the total error in the band

$$\begin{aligned} &\mathbb{E}_{\mathcal{D}|T} [(P(\langle \mathbf{w}^*, \mathbf{x} \rangle) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2] \\ &\leq \int_{-s}^s (P(t) - \text{ReLU}(t))^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(t) dt + 2 \int_s^\infty (P(t) - \text{ReLU}(t))^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(t) dt. \end{aligned}$$

Recall that P is at most τ -away from ReLU in the range $[-s, s]$. If $\tau = \nu^{1/2} \text{opt}^{1/4} \cdot (\eta/10)$, the first integral above can be upper bounded as follows:

$$\int_{-s}^s (P(t) - \text{ReLU}(t))^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(t) dt \leq \int_{-s}^s \tau^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(t) dt \lesssim \nu \cdot \eta^2 \cdot \text{opt}^{1/2}.$$

Additionally, if $t \geq 2l$, $(t-l)^2 \leq t^2/4$, and so under the condition $s \geq 2l$, the final integral can be bounded above by

$$\begin{aligned} &2 \int_s^\infty (P(t) - \text{ReLU}(t))^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle}(t) dt \\ &\lesssim \frac{1}{\nu \cdot \sin(\theta)} \int_s^\infty \max\{(4t/s)^{6s/\tau}, (4t)^{s/\tau}\} \exp\left(-\frac{(t-l)^2}{2 \cdot \nu^2 \cdot \sin^2(\theta)}\right) dt \end{aligned}$$

$$\lesssim \frac{1}{\nu \cdot \sin(\theta)} \int_s^\infty (4t/s)^{6s/\tau} \exp\left(-\frac{t^2}{8 \cdot \nu^2 \cdot \sin^2(\theta)}\right) dt .$$

Setting $t = 2\nu \sin(\theta)p$ and $r = s/(\nu \sin(\theta))$, we see

$$\begin{aligned} & 2 \int_s^\infty (P(t) - \text{ReLU}(t))^2 \rho_{\langle \mathbf{w}^*, \mathbf{x} \rangle_a}(t) dt \\ & \lesssim \int_{s/(\nu \sin(\theta))}^\infty (4p \cdot (\nu \sin(\theta)/s))^{6s/\tau} \exp\left(-\frac{p^2}{8}\right) dp \\ & \lesssim \int_r^\infty (4p \cdot (1/r))^{6s/\tau} \exp\left(-\frac{p^2}{8}\right) dp \\ & \lesssim \mathbb{E}_{p \sim N(0,1)}[1_{p \geq r} (4p/r)^{3s/\tau}] \\ & \lesssim (4/r)^{3s/\tau} \cdot \Pr[p \geq r]^{1/2} \cdot \mathbb{E}_{p \sim N(0,1)}[p^{6s/\tau}]^{1/2} \\ & \lesssim (4/r)^{3s/\tau} \cdot \exp(-r^2/4) \cdot (6s/\tau)^{3s/\tau} \\ & \lesssim s/(r^2 \tau)^{3s/\tau} \cdot \exp(-r^2/4) . \end{aligned}$$

Taking logs, we see that the following inequality needs to be satisfied

$$\left(\frac{s}{\tau} + \frac{s}{\tau} \log\left(\frac{s}{\tau}\right) + \frac{s}{\tau} \log\left(\frac{1}{r}\right) - \log(\nu^{1/2} \eta \text{opt}^{1/4}) \right) \lesssim r^2 . \quad (17)$$

Let $s = \nu(\nu \sin(\theta))^{1/2}/(\nu^{1/2} \eta^2 \text{opt}^{1/4})$, then $r = \nu^{1/2} \eta^{-2} (\nu \sin(\theta))^{-1/2} \text{opt}^{-1/4}$. Then, $s/\tau \leq f(\eta)$ for $f(x) = 1/x^3$. Substituting for s/τ we see that it is sufficient to check

$$\begin{aligned} & \left(f(\eta) + f(\eta) \log(f(\eta)) + f(\eta) \log\left(\nu^{-1/2} \eta^2 (\nu \sin(\theta))^{1/2} \text{opt}^{1/4}\right) - \log(\nu^{1/2} \eta \text{opt}^{1/4}) \right) \\ & \lesssim r^2 . \end{aligned}$$

Since $f(\eta) \geq 1$ and $\nu \sin(\theta) \leq \text{opt}^{1/2}$, it is sufficient to check that

$$\begin{aligned} & \left(f(\eta) + f(\eta) \log(\eta^2 f(\eta)) + (f(\eta) - 1) \log\left((\nu \sin(\theta))^{1/2} \text{opt}^{1/4}\right) - \log(\eta) + f(\eta) \log(\nu^{-1/2}) - \log(\nu^{1/2}) \right) \\ & \lesssim r^2 . \end{aligned}$$

Substituting for $f(\eta)$ and multiplying both sides by η^{-3} , we get

$$\left(\log(1/\eta) + \log\left(\nu^{-1/2} (\nu \sin(\theta))^{1/2} \text{opt}^{1/4}\right) \right) \lesssim \nu \eta^{-1} \cdot (\nu \sin(\theta))^{-1} \cdot \text{opt}^{-1/2} , \quad (18)$$

i.e., rescaling, we see

$$\log(r^{-2} \eta^{-3}) \lesssim r^2 \eta^3 .$$

Since $r \geq \nu^{1/2} \eta^{-2} \sin(\theta)^{-1} \nu^{-1} \geq \eta^{-2} \nu^{-1/2} > 1$ for opt, ν less than some constant, we see that this is true for small enough η .

Substituting this back we see that the overall error is bounded by

$$\mathbb{E}_{\mathcal{D}|T} \left[(P(\mathbf{x}) - \text{ReLU}(\langle \mathbf{w}^*, \mathbf{x} \rangle))^2 \right] \lesssim \nu \cdot \eta \cdot \sqrt{\text{opt}} . \quad (19)$$

This completes the proof. ■

Lemma 31 *Let*

$$A := \left\{ p \mid p \text{ is a degree } k, d\text{-variate polynomial with coefficients } a_i \text{ such that } \sum_i |a_i| \leq O(4^k) \right\}.$$

If $S \sim \mathcal{D}^m$ is a set of m iid samples and $m \gtrsim k^k \cdot \frac{1}{\epsilon^2} \cdot \left(\frac{4d}{\nu^2}\right)^{\Omega(k)}$, then

$$\mathbb{E}_{S|T}[(P(\mathbf{x}) - y)^2] \leq \min_{P' \in A} \mathbb{E}_{\mathcal{D}|T}[(P'(\mathbf{x}) - y)^2] + \epsilon.$$

Proof Let $p(\mathbf{x}) = \sum a_i m_i$, where $\{m_i\}$ are monomials that correspond to a_i . Then

$$\begin{aligned} (\mathbb{E}_{S|T}[P(\mathbf{x})] - \mathbb{E}_{\mathcal{D}|T}[P(\mathbf{x})]) &\leq \left(\sum_i |a_i| \right) \left(\sum_i |\mathbb{E}_{S|T}[m_i(\mathbf{x})] - \mathbb{E}_{\mathcal{D}|T}[m_i(\mathbf{x})]| \right) \\ &\leq (4)^k \left(\sum_i |\mathbb{E}_{S|T}[m_i(\mathbf{x})] - \mathbb{E}_{\mathcal{D}|T}[m_i(\mathbf{x})]| \right). \end{aligned}$$

Since there are at most d^k different monomials, it is sufficient to approximate each monomial m_i up to an accuracy of $\epsilon \cdot \frac{1}{4^k} \cdot \frac{1}{d^k}$.

Let $M_{d,k}$ denote the set of d -variate monomials of degree k . Applying Fact 1 on these monomials and noting that the variance of any monomial is at most $\nu^{2k} \cdot O(k)^k$ with respect to the ν -subgaussian distribution where $\nu = O(1)$, we get

$$\Pr \left[m \in M_{d,k} \mid |\mathbb{E}_{S|T}[m(\mathbf{x})] - \mathbb{E}_{\mathcal{D}|T}[m(\mathbf{x})]| \geq \epsilon \right] \leq d^k \cdot \exp \left(- \frac{m \epsilon^2 \cdot \frac{1}{4^k} \cdot \frac{1}{d^k}}{\nu^{2k} k^k} \right)^{1/k}.$$

Setting $m \gtrsim O(k)^k \cdot \frac{1}{\epsilon^2} \cdot \left(\frac{4d}{\nu^2}\right)^{O(k)}$ we obtain our desired bounds. ■