

Efficiently Learning Any One Hidden Layer ReLU Network From Queries

Sitan Chen*
 sitanc@mit.edu
 MIT

Adam R. Klivans†
 klivans@cs.utexas.edu
 UT Austin

Raghuram Meka‡
 raghum@cs.ucla.edu
 UCLA

November 9, 2021

Abstract

Model extraction attacks have renewed interest in the classic problem of learning neural networks from queries. This work gives the first polynomial-time algorithm for learning one hidden layer neural networks provided black-box access to the network. Formally, we show that if F is an arbitrary one hidden layer neural network with ReLU activations, there is an algorithm with query complexity and running time that is polynomial in all parameters that outputs a network F' achieving low square loss relative to F with respect to the Gaussian measure. While a number of works in the security literature have proposed and empirically demonstrated the effectiveness of certain algorithms for this problem, ours is the first with fully polynomial-time guarantees of efficiency for worst-case networks (in particular our algorithm succeeds in the overparameterized setting).

1 Introduction

The problem of learning neural networks given random labeled examples continues to be a fundamental algorithmic challenge in machine learning theory. There has been a resurgence of interest in developing algorithms for learning one hidden layer neural networks, namely functions of the form $F(x) = \sum_{i=1}^k s_i \sigma(\langle w_i, x \rangle - b_i)$ where $w_i \in \mathbb{R}^d$, $s_i \in \{\pm 1\}$, $b_i \in \mathbb{R}$ are the unknown parameters. Here we consider the ubiquitous setting where σ is the ReLU activation, the input distribution on x is Gaussian, and the goal is to output a ReLU network F' with $\mathbb{E}[(F(x) - F'(x))^2] \leq \epsilon$.

Despite much recent work on this problem [JSA15, ZSJ⁺17, GKM18, GLM18, GKLW18, ZYWG19, BJW19, DKKZ20, CKM20, LMZ20], obtaining a polynomial-time algorithm remains open. All known works take additional assumptions on the unknown weights and coefficients or do not run in polynomial-time in all the important parameters. In fact, there are a number of lower bounds, both for restricted models of computation like correlational statistical queries [GGJ⁺20, DKKZ20] as well as under various average-case assumptions [DV20, DV21], that suggest that a truly polynomial-time algorithm may be impossible to achieve.

In this work, we strengthen the learner by allowing it query access to the unknown network. That is, the learner may select an input x of its choosing and receive the output value of the

*This work was supported in part by an NSF CAREER Award CCF-1453261 and NSF Large CCF-1565235.

†Supported by NSF awards AF-1909204, AF-1717896, and the NSF AI Institute for Foundations of Machine Learning (IFML).

‡Supported by NSF CAREER Award CCF-1553605.

unknown network $F(x)$. Our main result is the first polynomial-time algorithm for learning one hidden layer ReLU networks with respect to Gaussian inputs when the learner is given query access:

Theorem 1.1. *Let $F(x) = \sum_{i=1}^k s_i \sigma(\langle w_i, x \rangle - b_i)$ with $\|w_i\|_2 \leq R$, $s_i \in \{\pm 1\}$, and $b_i \leq B$. Given black-box access to F , there exists a (randomized) algorithm that will output with probability at least $1 - \delta$, a one hidden layer network F' such that $\mathbb{E}_{x \sim \mathcal{N}(0, \text{Id})}[(F(x) - F'(x))^2] \leq \epsilon$. The algorithm has query complexity and running time that is polynomial in $d, B, R, k, 1/\epsilon$, and $\log(1/\delta)$.*

In light of the aforementioned lower bounds [GGJ⁺20, DKKZ20, DV20, DV21], Theorem 1.1 suggests there is a separation between learning from random samples and learning from queries in the context of neural networks.

In addition to being a naturally motivated challenge in computational learning theory, the question of learning neural networks from queries has also received strong interest from the security and privacy communities in light of so-called *model extraction attacks* [TJ⁺16, MSDH19, PMG⁺17, JCB⁺20, RK20, JWZ20]. These attacks attempt to reverse-engineer neural networks found in publicly deployed real-world systems. Since the target network is publicly available, an attacker therefore obtains black-box query access to the network – exactly the learning model we are working in. Of particular interest is the work of Carlini et al. [CJM20] (see also [JCB⁺20]), who gave several heuristics for learning deep neural networks given black-box access (they also empirically verified their results on networks trained on MNIST). Another relevant work is that of [MSDH19] which gave theoretical guarantees for the problem we study under quite strong separation/linear independence assumptions on the weight vectors; in fact they are able to *exactly* recover the network, but in the absence of such strong assumptions this is impossible (see Section 1.2). A very recent follow-up work [DG21] improved upon the results of [MSDH19] by giving a similar guarantee under a somewhat milder assumption and notably also giving an algorithm for learning ReLU networks with *two* hidden layers under such assumptions. Since these papers bear strong parallels with our techniques, in Section 1.2 we provide a more detailed description of why these approaches break down in our setting and highlight the subtleties that we address to achieve the guarantee of Theorem 1.1.

1.1 Our Approach

Here we describe our approach at a high level. Fix an unknown one hidden-layer network $F(x) = \sum_{i=1}^k s_i \sigma(\langle w_i, x \rangle - b_i)$ with weight vectors $w_i \in \mathbb{R}^d$, signs $s_i \in \{\pm 1\}$, and bias terms $b_i \in \mathbb{R}$. Consider a line $L = \{x_0 + t \cdot v\}_{t \in \mathbb{R}}$ for $x_0, v \in \mathbb{R}^d$ and the *restriction* $F|_L(t) \triangleq F(x_0 + t \cdot v)$ given by

$$F|_L(t) = \sum_{i=1}^k s_i \sigma(\langle w_i, x_0 \rangle - b_i + t \langle w_i, v \rangle).$$

Note that as a univariate function, $F|_L(t)$ is a piecewise-linear function. We call a point $t \in \mathbb{R}$ a *critical point* of $F|_L$ if the slope of the piecewise-linear function changes at t . Our starting point is that if we can identify the critical points of $F|_L$, then we can use estimates of the gradient of F at those critical points to estimate the weights of the hidden units in F . More concretely, suppose we have identified a critical point t_i of $F|_L$ for a line L . Now, back in the d -dimensional space, if we look at a sufficiently small neighborhood of the point $x' = x_0 + t_i \cdot v$, then the set of activated units around x' changes by exactly one (in fact, by this reasoning, for a random L , $F|_L$ will have one unique critical point for each neuron in F , provided no two neurons are exact multiples of each other). We can exploit this to retrieve the w_i 's and b_i 's by finite differencing: for each such x' , compute $(F(x' + \delta) - F(x'))/\|\delta\|$, for several sufficiently small perturbations $\delta \in \mathbb{R}^d$, and this

will recover $s_i w_i$ (see Algorithm 2 and 4.9 below for the details). We can recover $s_i b_i$ in a similar fashion (see Algorithm 1). We remark that the approach we have described appears to be quite similar to that of [CJM20, JCB⁺20, MSDH19].

However, the above description leaves out the following key interrelated challenges for getting provable guarantees for arbitrary networks:

1. How do we identify the critical points of $F|_L$ for a given line L ?
2. If we make no assumptions on how well-separated the weight vectors are, it is actually impossible to recover all of the weight vectors and biases in a bounded number of queries. Using the neurons that we do approximately recover, is it possible to piece them together to approximate F ?

A rigorous solution to these issues turns out to be quite tricky and technically involved. For instance, what if some critical points are arbitrarily close to each other (imagine the piecewise linear function being a tiny peak of exponentially small in d width on the real line) on the line $F|_L$. In this case, identifying them would not be possible efficiently. We develop several structural results about cancellations in ReLU networks (culminating in Lemma 3.21) that show that such a situation only arises in degenerate cases where some sub-network of F contains a cluster of many similar neurons (we formalize the notion of similarity in Definition 3.7). We show in Section 3 that much smaller networks can approximate these sub-networks. On the other hand, for pairs of neurons which are sufficiently far from each other, the critical points on a random line L will be well-separated.

With the structural results in hand, roughly speaking it suffices to identify one “representative” neuron for every cluster of similar neurons in F . To do this, we discretize L into intervals of equal length and look at differences between gradients/biases of F at the midpoints of these intervals. The key result leveraging our structural results is Theorem 4.8 which shows that the set of all such differences comprises a rich enough collection of neurons that we can assemble some linear combination of them that approximates F . To find a suitable linear combination, we run linear regression on the features computed by these neurons (see Section 4.4).

1.2 Comparison to Previous Approaches

Our general approach of looking for critical lines along random restrictions of $F|_L$ is also the approach taken in the empirical works of [CJM20, JCB⁺20] and in the theoretical work of [MSDH19, DG21], but we emphasize that there are a variety of subtleties that arise because we are interested in learning *worst-case* networks from queries. In contrast, the empirical works consider trained networks arising in practice, while [MSDH19] makes a strong assumption that the weight vectors w_i are *linearly independent* and that they are angularly separated. Note that such assumptions cannot hold in the practically relevant setting where F is overparameterized. The recent work of [DG21] likewise makes non-degeneracy assumptions that allow them to avoid dealing with closely spaced critical points, arguably the central technical challenge in the present work. Indeed, these sorts of assumptions are powerful enough that [MSDH19, DG21] are able to *exactly recover* the parameters of the network in a finite number of samples, whereas at the level of generality of the present work, this is clearly impossible.

To give a sense for the issues that these existing techniques run up against when it comes to worst-case networks, imagine that the one-dimensional piecewise linear function given by the restriction $F|_L$ were simply a “bump,” that is, it is zero over most of the line except in a small interval $[a, a + \delta]$. For instance, this would arise in the following example:

Example 1.2. *Consider the one-dimensional one hidden layer network $F : \mathbb{R} \rightarrow \mathbb{R}$ given by $F(x) = \sigma(x - a) + \sigma(x - a - \delta) - \sigma(2x - 2a - \delta)$. This function looks like the zero function except over the interval $[a, a + \delta]$, where it looks like a small “bump.”*

The proposal in the works mentioned above is to run a binary search to find a critical point. Namely, they initialize to some large enough interval $[-\tau, \tau]$ in L and check whether the gradient at the left or right endpoint differs from the gradient at the midpoint $t_{\text{mid}} = 0$. If the former, then they restrict the search to $[-\tau, t_{\text{mid}}]$ and recurse until they end up with a sufficiently small interval, at which point they return the midpoint of that interval as an approximation to a critical point. The intuition is that n steps of binary search suffice to identify a critical point up to n bits of precision.¹

It is clear, however, from the “bump” example that such a binary search procedure is doomed to fail: if the bump does not occur in the middle of the starting interval $[-\tau, \tau]$, then at the outset, we don’t know which half to recurse on because the gradients at the endpoints and the gradient at the midpoint are all zero! Indeed, to locate the bump, it is clear that we need to query a number of points which is at least inverse in the width of the bump to even find an input to the network with nonzero output.

We remark that this is in stark contrast to related active learning settings where the ability to query the function can sometimes yield *exponential savings* on the dependence on $1/\epsilon$ (see e.g. [Han09]). Indeed, if we took the width of the bump to be of size $\text{poly}(\epsilon)$ to make it ϵ -far in square distance from the zero function, we would still need $\text{poly}(1/\epsilon)$, rather than $\log(1/\epsilon)$, queries to locate the bump and distinguish the function from the zero function.

The core reason binary search fails for arbitrary networks is that there can be many disjoint linear pieces of $F|_L$ which all have the same slope. Indeed, because the gradient of F at a given point is some linear combination of the weight vectors, if F is overparameterized so that there are many linear dependencies among the weight vectors, it is not hard to design examples like Example 1.2 where there may be many repeated slopes on any given random restriction of F to a line.

Apart from these technical issues that arise in the worst-case setting we consider and not in previous empirical or theoretical works on model extraction, we also emphasize that our results are the first theoretical guarantees for learning *general* one hidden-layer ReLU networks with bias terms in any learning model, including PAC learning from samples. As will quickly become evident in the sequel, biases introduce many technical hurdles in their own right. To our knowledge, the only other theoretical guarantees for learning networks with biases are the work of [JSA15] which considers certain activations with nonzero second/third-order derivatives, precluding the ReLU, and the recent work of [ATV21] which considers ReLU networks whose weight vectors are well-conditioned.

Remark 1.3. *As discussed above, the “bump” construction in Example 1.2 shows that unlike in related active learning contexts, $\text{poly}(1/\epsilon)$ dependence is necessary in our setting. In fact this example also tells us another somewhat counterintuitive fact about our setting. Naively, one might hope to upgrade Theorem 1.1 by replacing the dependence on the scaling parameters R, B with one solely on the L_2 norm of the function. To see why this is impossible, consider scaling the function in Example 1.2 by a factor of $\delta^{-3/2}$ to have unit L_2 norm. To learn F , we have to figure out the location of the bump, but this requires $\Omega(1/\delta)$ queries, and δ can be taken to be arbitrarily small.*

Limitations and Societal Impact The line of work on learning neural networks from black-box access does pose a risk, for instance, that proprietary models offered via publicly-hosted APIs may be stolen by attackers who are only able to interact with the models through queries. The attackers might then use the extracted parameters of the model to learn sensitive information about

¹We remark that the refinement of binary search given in [CJM20, JCB⁺20] can speed this up to $O(1)$ steps for networks that arise in practice, but the issue that we describe poses a challenge for their technique as well.

the data the model was trained on, or perhaps to construct adversarial examples. That being said, understanding why this query learning problem can be easy for theorists is the first step towards building provable safeguards to ensure that it is hard for actual attackers. For instance, it is conceivable that one can prove information-theoretic or computational hardness for such problems if appropriate levels and kinds of noise are injected into the responses to the queries. Furthermore, query complexity lower bounds can inform how many accesses an API should allow any given user.

2 Preliminaries

Notation We let \mathbb{S}^{d-1} denote the unit sphere in \mathbb{R}^d . Let e_j denote the j -th standard basis vector in \mathbb{R}^d . Given vectors u, v , let $\angle(u, v) \triangleq \arccos\left(\frac{\langle u, v \rangle}{\|u\| \|v\|}\right)$. Given a matrix M , let $\|M\|_{\text{op}}$ and $\|M\|_F$ denote the operator and Frobenius norms respectively. Given a function h which is square-integrable with respect to the Gaussian measure, we will use $\|h\|$ to denote $\mathbb{E}_{x \sim \mathcal{N}(0, \text{Id})}[h(x)^2]^{1/2}$. Given a collection of indices $S \subseteq \mathbb{Z}$, we say that $i, j \in S$ are *neighboring* if there does not exist $i < \ell < j$ for which $\ell \in S$.

The following elementary fact will be useful:

Fact 2.1. $|\sin(x + y)| = |\sin(x) \cos(y) + \sin(y) \cos(x)| \leq |\sin(x)| + |\sin(y)|$ for any $x, y \in \mathbb{R}$

2.1 Neural Networks, Restrictions, and Critical Points

Definition 2.2. A neuron is a pair (v, b) where $v \in \mathbb{R}^d$ and $b \in \mathbb{R}$; it corresponds to the function $x \mapsto \sigma(\langle v, x \rangle - b)$, which we sometimes denote by $\sigma(\langle v, \cdot \rangle - b)$.

As mentioned in the overview, we will be taking random restrictions of the underlying network F , for which we use the following notation:

Definition 2.3. Given a line $L \subset \mathbb{R}^d$ parametrized by $L = \{x_0 + t \cdot v\}_{t \in \mathbb{R}}$, and a function $F : \mathbb{R}^d \rightarrow \mathbb{R}$, define the restriction of F to L by $F|_L(t) \triangleq F(x_0 + t \cdot v)$.

Definition 2.4. Given a line $L \subset \mathbb{R}^d$ and a restriction $F|_L$ of a piecewise linear function $F : \mathbb{R}^d \rightarrow \mathbb{R}$ to that line, the critical points of $F|_L$ are the points $t \in \mathbb{R}$ at which the slope of $F|_L$ changes.

2.2 Concentration and Anti-Concentration

We will need the following standard tail bounds and anti-concentration bounds:

Fact 2.5 (Concentration of norm of Gaussian vector). Given Gaussian vector $h \sim \mathcal{N}(0, \Sigma)$, $\Pr\left[\|h\| \geq O(\|\Sigma^{1/2}\|_{\text{op}}(\sqrt{r} + \sqrt{\log(1/\delta)}))\right] \leq \delta$, where r is the rank of Σ .

Fact 2.6 (Uniform bound on entries of Gaussian vector). For covariance matrix $\Sigma \in \mathbb{R}^{m \times m}$, given $h \sim \mathcal{N}(0, \Sigma)$ we have that $|h_i| \leq O\left(\sqrt{\Sigma_{i,i}}\sqrt{\log(m/\delta)}\right)$ for all $i \in [m]$ with probability at least $1 - \delta$.

Proof. For every $i \in [m]$, $h_i \sim \mathcal{N}(0, \Sigma_{i,i})$, so $|h_i| \leq O(\Sigma_{i,i}^{1/2} \sqrt{\log(m/\delta)})$ with probability at least $1 - \delta/m$, from which the claim follows by a union bound and the fact that the largest diagonal entry of a psd matrix is the largest entry of that matrix. \square

Fact 2.7 (Carbery-Wright [CW01]). There is an absolute constant $C > 0$ such that for any $\nu > 0$ and quadratic polynomial $p : \mathbb{R}^d \rightarrow \mathbb{R}$, $\Pr_{g \sim \mathcal{N}(0, \text{Id})}[|p(g)| \leq \nu \cdot \mathbb{V}[p(g)]^{1/2}] \leq C\sqrt{\nu}$.

Lemma 2.8 (Anti-concentration of norm of Gaussian vector). *There is an absolute constant $C > 0$ such that given any Gaussian vector $h \sim \mathcal{N}(\mu, \Sigma)$, $\Pr[\|h\| \geq \sqrt{\nu} \|\Sigma\|_F^{1/2}] \geq 1 - C\sqrt{\nu}$.*

Proof. Define the polynomial $p(g) \triangleq (g + \mu)^\top \Sigma (g + \mu)$. Note that for $g \sim \mathcal{N}(0, \text{Id})$, $p(g)$ is distributed as $\|h\|^2$ for $h \sim \mathcal{N}(\mu, \Sigma)$. We have $\mathbb{E}_{g \sim \mathcal{N}(0, \text{Id})}[p(g)] = \text{Tr}(\Sigma) + \mu^\top \Sigma \mu$, so

$$\begin{aligned}\mathbb{V}[p(g)] &= \mathbb{E}[(g^\top \Sigma g + 2g^\top \Sigma \mu - \text{Tr}(\Sigma))^2] \\ &= \mathbb{E}[(g^\top \Sigma g)^2] + \mathbb{E}[(2g^\top \Sigma \mu - \text{Tr}(\Sigma))^2] + 2\mathbb{E}[(g^\top \Sigma g)(2g^\top \Sigma \mu - \text{Tr}(\Sigma))] \\ &= (2\text{Tr}(\Sigma^2) + \text{Tr}(\Sigma)^2) + (4\text{Tr}(\Sigma \mu \mu^\top \Sigma) + \text{Tr}(\Sigma)^2) - 2\text{Tr}(\Sigma)^2 \\ &= 2\langle \Sigma^2, \text{Id} + 2\mu \mu^\top \rangle \geq \|\Sigma\|_F^2,\end{aligned}$$

so by Fact 2.7 we conclude that $\Pr[p(g) \leq \nu \|\Sigma\|_F] \leq C\sqrt{\nu}$. \square

Lemma 2.9 (Anti-concentration for random unit vectors). *For random $v \in \mathbb{S}^{d-1}$, $\Pr\left[|v_1| < \frac{\delta}{2\sqrt{d} + O(\sqrt{\log(1/\delta)})}\right] \leq \delta$.*

Proof. For $g \sim \mathcal{N}(0, \text{Id})$, $g/\|g\|$ is identical in distribution to v . $\|g\| \leq \sqrt{d} + O(\sqrt{\log(1/\delta)})$ with probability at least $1 - \delta/2$ for absolute constant $c > 0$, and furthermore $\Pr_{\gamma \sim \mathcal{N}(0, 1)}[|g| > t] \geq 1 - t$ for any $t > 0$, from which the claim follows by a union bound. \square

3 ReLU Networks with Cancellations

In the following section we prove several general results about approximating one hidden-layer networks with many ‘‘similar’’ neurons by much smaller networks.

3.1 Stability Bounds for ReLUs

The main result of this subsection will be the following stability bound for (non-homogeneous) ReLUs with the same bias.

Lemma 3.1. *Fix any $\Delta < 1$. For orthogonal $v, v' \in \mathbb{R}^d$ for which $\|v - v'\| \leq \Delta \|v\|$, and $b \in \mathbb{R}$, we have*

$$\mathbb{E}[(\sigma(\langle v, x \rangle) - b) - \sigma(\langle v', x \rangle - b))^2] \leq O\left(\Delta^{2/5} \|v\|^2\right)$$

To prove this, we will need to collect some standard facts about stability of homogeneous ReLUs and affine threshold functions, given in Fact 3.2, Lemma 3.3, Lemma 3.4, and Lemma 3.5.

The following formula is standard [CS09]:

Fact 3.2. $\mathbb{E}[\sigma(\langle v, x \rangle) \sigma(\langle v', x \rangle)] = \frac{1}{2\pi} \|v\| \|v'\| (\sin \angle(v, v') + (\pi - \angle(v, v')) \cos \angle(v, v')).$ For $\langle v, v' \rangle \geq 0$, note that this is at least $\frac{1}{6} \|v\| \|v'\| + \frac{1}{3} \langle v, v' \rangle$.

As a consequence, we obtain the following stability result for homogeneous ReLUs:

Lemma 3.3. *For any $v, v' \in \mathbb{R}^d$ for which $\langle v, v' \rangle \geq 0$, we have*

$$\mathbb{E}[(\sigma(\langle v, x \rangle) - \sigma(\langle v', x \rangle))^2] \leq \frac{1}{2} \|v - v'\|^2 + \frac{2}{3} \|v\| \|v'\| (1 - \cos \angle(v, v'))$$

Proof. We can expand the expectation and apply Fact 3.2 to get

$$\begin{aligned}
\mathbb{E}[(\sigma(\langle v, x \rangle) - \sigma(\langle v', x \rangle))^2] &= \mathbb{E}[\sigma(\langle v, x \rangle)^2] + \mathbb{E}[\sigma(\langle v', x \rangle)^2] - 2\mathbb{E}[\sigma(\langle v, x \rangle)\sigma(\langle v', x \rangle)] \\
&\leq \frac{1}{2}\|v\|^2 + \frac{1}{2}\|v'\|^2 - 2\left(\frac{1}{6}\|v\|\|v'\| + \frac{1}{3}\langle v, v' \rangle\right) \\
&= \frac{1}{2}\|v - v'\|^2 + \frac{2}{3}(\|v\|\|v'\| - \langle v, v' \rangle) \\
&= \frac{1}{2}\|v - v'\|^2 + \frac{2}{3}\|v\|\|v'\|(1 - \cos \angle(v, v'))
\end{aligned}$$

as claimed. \square

We will also need the following stability result for affine linear thresholds.

Lemma 3.4 (Lemma 5.7 in [CM20]). *Given $v, v' \in \mathbb{R}^d$ and $b \in \mathbb{R}$,*

$$\Pr[\langle v, x \rangle > b \wedge \langle v', x \rangle \leq b] \leq O(\|v - v'\|/b).$$

Lemma 3.5. *For any $v \in \mathbb{R}^d$ and $b \leq b'$,*

$$\mathbb{E}[(\sigma(\langle v, x \rangle - b) - \sigma(\langle v, x \rangle - b'))^2] \leq (b' - b)^2$$

Proof. Note that $\langle v, x \rangle \sim \mathcal{N}(0, \|v\|^2)$, so it suffices to show that for the univariate function $f(z) \triangleq \sigma(z - b) - \sigma(z - b')$, $\mathbb{E}_{z \sim \mathcal{N}(0, \|v\|^2)}[f(z)^2] \leq (b' - b)^2$. Observe that $f(z) = b' - b$ for $z > b'$, $f(z) = 0$ for $z < b$, and $f(z) = z - b$ for $z \in [b, b']$. In particular, $|f(z)| \leq b' - b$, from which the claim follows. \square

The following basic lemma giving L_2 bounds for Lipschitz functions which are bounded with high probability will be useful throughout.

Lemma 3.6. *Let $\epsilon(x) : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be any square-integrable function with respect to the Gaussian measure. If $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is an L -Lipschitz continuous piecewise linear function and satisfies $\Pr_{x \sim \mathcal{N}(0, Id)}[|f(x)| \leq \epsilon(x)] \geq 1 - \zeta$ and $|f(0)| \leq M$, then $\mathbb{E}_{x \sim \mathcal{N}(0, Id)}[f(x)^2] \leq 2\zeta M^2 + L^2\zeta^{1/2}(d^2 + 2d) + \mathbb{E}[\epsilon(x)^4]^{1/2}$.*

Proof. Because f is L -Lipschitz, $f(x)^2 \leq (M + L\|x\|)^2 \leq 2M^2 + L^2\|x\|^2$. Then

$$\begin{aligned}
\mathbb{E}[f(x)^2] &\leq \mathbb{E}[f(x)^2 \mathbf{1}[f(x) > \epsilon(x)]] + \mathbb{E}[\epsilon(x)^2 \mathbf{1}[f(x) \leq \epsilon(x)]] \\
&\leq 2\zeta M^2 + L^2 \mathbb{E}[\|x\|^2 \mathbf{1}[f(x) > \rho\|x\|]] + \mathbb{E}[\epsilon(x)^4]^{1/2}(1 - \zeta)^{1/2} \\
&\leq 2\zeta M^2 + L^2 \zeta^{1/2} \mathbb{E}[\|x\|^4] + \mathbb{E}[\epsilon(x)^4]^{1/2}(1 - \zeta)^{1/2} \\
&= 2\zeta M^2 + 3L^2 \zeta^{1/2} d^2 + \mathbb{E}[\epsilon(x)^4]^{1/2},
\end{aligned}$$

as claimed. \square

Putting all of these ingredients together, we can now complete the proof of the main Lemma 3.1 of this subsection.

Proof. Suppose $b \geq \Delta^{1/5}\|v\|$. By Lemma 3.4, $\text{sgn}(\langle v, x \rangle - b) \neq \text{sgn}(\langle v', x \rangle - b)$ with probability at most $O(\Delta\|v\|/b)$. So with probability at least $1 - O(\Delta\|v\|/b)$, the function $(\sigma(\langle v, x \rangle - b) - \sigma(\langle v', x \rangle - b))$ is at most $\langle v - v', x \rangle \leq \Delta\|v\|\|x\|$. Furthermore, this function is L -Lipschitz for

$L = \|v\| + \|v'\| \leq O(\|v\|)$. By Lemma 3.6 applied to the projection of f to the two-dimensional subspace spanned by v, v' ,

$$\mathbb{E}[(\sigma(\langle v, x \rangle) - b) - \sigma(\langle v', x \rangle - b))^2] \lesssim \|v\|^2 \left(\sqrt{\Delta \|v\|/b} + \Delta^2 \right) \lesssim \Delta^{2/5} \|v\|^2.$$

Now suppose $b < \Delta^{1/5} \|v\|$. Then $\|\sigma(\langle v, \cdot \rangle - b) - \sigma(\langle v, \cdot \rangle)\|^2 \leq \Delta^{2/5} \|v\|^2$ and $\|\sigma(\langle v', \cdot \rangle - b) - \sigma(\langle v', \cdot \rangle)\|^2 \leq \Delta^{2/5} \|v'\|^2$. By triangle inequality, it suffices to bound $\|\sigma(\langle v, \cdot \rangle) - \sigma(\langle v', \cdot \rangle)\|^2$. By Lemma 3.3, we have

$$\|\sigma(\langle v, \cdot \rangle) - \sigma(\langle v', \cdot \rangle)\|^2 \lesssim \Delta^2 \|v\|^2 + \|v\|^2 \cdot (1 - \cos \angle(v, v')) \lesssim \Delta^2 \|v\|^2,$$

where the last step follows by the fact that $\|v - v'\| \leq \Delta \|v\|$ implies that $\cos \angle(v, v') \geq \sqrt{1 - \Delta^2} \geq 1 - \Delta^2$. \square

3.2 (Δ, α) -Closeness of Neurons

We now formalize a notion of geodesic closeness between two neurons and record some useful properties. This notion is motivated by Lemma 4.4 in Section 4.1 where we study the critical points of random restrictions of one hidden-layer networks.

Definition 3.7. *Given $v, v' \in \mathbb{R}^d$ and $b, b' \in \mathbb{R}$, we say that (v, b) and (v', b') are (Δ, α) -close if the following two conditions are satisfied:*

1. $|\sin \angle(v, v')| \leq \Delta$
2. $\|bv' - b'v\| \leq \alpha \|v\| \|v'\|$.

Note that this is a measure of angular closeness between $(v, b), (v', b') \in \mathbb{R}^{d+1}$. For instance, if $(v, b) = (\lambda v^*, \lambda b^*)$ and $(v', b') = (\lambda' v^*, \lambda' b^*)$ for some (v^*, b^*) , then (v, b) and (v', b') are $(0, 0)$ -close.

We first collect some elementary consequences of closeness. The following intuitively says that if we scale two (Δ, α) -neurons to have similar norm, then their biases will be close.

Lemma 3.8. *If (v, b) and (v', b') are (Δ, α) -close, and $v = \gamma v' + v^\perp$ for v^\perp orthogonal to v' , then $|\gamma b' - b| \leq \alpha \|v\|$.*

Proof. We know that $\|bv' - b'v\| \leq \alpha \|v\| \|v'\|$. The left-hand side of this is $\|(b - \gamma b')v' - b'v^\perp\| \geq |b - \gamma b'| \|v'\|$, where the inequality follows from orthogonality of v, v' . Therefore, $|\gamma b' - b| \leq \alpha \|v\|$ as claimed. \square

Note that when two neurons are (Δ, α) -close, their weight vectors are either extremely correlated or extremely anti-correlated. In fact, given a collection of neurons that are all pairwise close, they will exhibit the following ‘‘polarization’’ effect.

Lemma 3.9. *Suppose $\Delta < \sqrt{2}/2$. If $(v_1, b_1), \dots, (v_k, b_k)$ are all pairwise (Δ, α) -close for some $\alpha > 0$, then there is a partition $[k] = S_1 \sqcup S_2$ for which $\langle v_i, v_j \rangle \geq 0$ for any $i \in S_1, j \in S_1$ or $i \in S_2, j \in S_2$, and for which $\langle v_i, v_j \rangle < 0$ for any $i \in S_1, j \in S_2$ or $i \in S_2, j \in S_1$.*

Proof. Let S_1 be the set of $i \in [k]$ for which $\langle v_i, v_1 \rangle \geq 0$, and let S_2 be the remaining indices. First consider any $i, j \in S_1$ and note that $\angle(v_i, v_j) \leq \angle(v_i, v_1) + \angle(v_j, v_1) \leq 2 \arcsin \Delta$, and because $\langle v_i, v_1 \rangle, \langle v_j, v_1 \rangle \geq 0$, this is less than $\pi/4$ for $\Delta < \sqrt{2}/2$. By the same reasoning, we can show that for any $i, j \in S_2$, $\angle(v_i, v_j) < \pi/2$ if $\Delta < \sqrt{2}/2$. Finally, consider $i \in S_1$ and $j \in S_2$. We have $\angle(v_i, v_j) \geq \angle(v_j, v_1) - \angle(v_i, v_1)$. If $\Delta < \sqrt{2}/2$, then $\angle(v_j, v_1) > 3\pi/4$ while $\angle(v_i, v_1) < \pi/4$, concluding the proof. \square

In the rest of the paper we will take Δ to be small, so Lemma 3.9 will always apply. As such, it will be useful to define the following terminology:

Definition 3.10. *Given $(v_1, b_1), \dots, (v_k, b_k)$ which are all pairwise-close, we will call the partition $S_1 \sqcup S_2$ given in Lemma 3.9 the orientation induced by $\{(v_i, b_i)\}$.*

We note that (Δ, α) -closeness satisfies triangle inequality.

Lemma 3.11. *If (v_1, b_1) and (v_2, b_2) are (Δ, α) -close, and (v_2, b_2) and (v_3, b_3) are (Δ', α') -close, then (v_1, b_1) and (v_3, b_3) are $(\Delta + \Delta', 2\alpha + 2\alpha')$ -close.*

Proof. As $\angle(v_1, v_3) \leq \angle(v_1, v_2) + \angle(v_2, v_3)$, it is clear from Fact 2.1 that $|\sin \angle(v_1, v_3)| \leq \Delta + \Delta'$.

Now write the orthogonal decompositions $v_1 = \gamma_1 v_2 + v_1^\perp$ and $v_3 = \gamma_3 v_2 + v_3^\perp$, noting that $\gamma_1 \|v_2\| \leq \|v_1\|, \gamma_3 \|v_2\| \leq \|v_3\|$. We can write

$$b_1 v_3 - b_3 v_1 = (b_1 \gamma_3 - b_3 \gamma_1) v_2 + (b_1 v_3^\perp - b_3 v_1^\perp). \quad (1)$$

We will handle these two terms separately. First note that (Δ, α) -closeness of $(v_1, b_1), (v_2, b_2)$ and Lemma 3.8 imply $|b_2 \gamma_1 - b_1| \leq \alpha \|v_1\|$, so in particular $|b_2 \gamma_1 \gamma_3 - b_1 \gamma_3| \leq \alpha \gamma_3 \|v_1\|$. Similarly, $|b_2 \gamma_1 \gamma_3 - b_3 \gamma_1| \leq \alpha' \gamma_1 \|v_3\|$. This allows us to conclude by triangle inequality that

$$|b_1 \gamma_3 - b_3 \gamma_1| \cdot \|v_2\| \leq \alpha \gamma_1 \|v_3\| + \alpha' \gamma_3 \|v_1\| \|v_2\| \leq (\alpha + \alpha') \|v_1\| \|v_3\|. \quad (2)$$

It remains to handle the second term on the right-hand side of (1). Note that Lemma 3.8 also tells us that

$$\|b_1 v_3^\perp - b_3 v_1^\perp\| \leq \|b_2 \gamma_1 v_3^\perp - b_1 v_3^\perp\| + \|b_2 \gamma_3 v_1^\perp - b_3 v_1^\perp\| \leq \alpha \|v_1\| \|v_3^\perp\| + \alpha' \|v_3\| \|v_1^\perp\| \leq (\alpha + \alpha') \|v_1\| \|v_3\|, \quad (3)$$

so by (1), (2), and (3), $\|b_1 v_3 - b_3 v_1\| \leq 2(\alpha + \alpha') \|v_1\| \|v_3\|$. \square

3.3 Merging Neurons

In this section we begin to apply the tools we have developed in the preceding sections to show our main results about approximating neural networks with many close neurons by smaller networks. The goal of this subsection is to prove that a one hidden-layer network where all neurons are (Δ, α) -close to some neuron can be approximated by at most two neurons:

Lemma 3.12. *Given $F(x) = \sum_{i=1}^k s_i \sigma(\langle w_i, x \rangle - b_i)$ for $s_i \in \{\pm 1\}$ and $(v^*, b^*) \in \mathbb{R}^d \times \mathbb{R}$ for which (w_i, b_i) is (Δ, α) -close to (v^*, b^*) for all $i \in [k]$, there exist coefficients $a^+, a^- \in \mathbb{R}$ for which*

$$\mathbb{E}_{x \sim \mathcal{N}(0, Id)} \left[(F(x) - a^+ \sigma(\langle v^*, x \rangle - b^*) - a^- \sigma(\langle -v^*, x \rangle + b^*))^2 \right] \leq O(k^2(\Delta^{2/5} + \alpha^2)) \|v^*\|^2. \quad (4)$$

Furthermore, we have that

$$|a^+| \|v^*\|, |a^-| \|v^*\| \leq \sum_i \|w_i\| \quad \text{and} \quad |a^+ b^*|, |a^- b^*| \leq \alpha \sum_i \|w_i\| + \sum_i |b_i|. \quad (5)$$

Our starting point for showing this is the following lemma which states that given *two* close neurons whose weight vectors are correlated, we can merge them into a single neuron while incurring small square loss.

Lemma 3.13. Let $0 < \Delta \leq 1$. For $v_1, v_2, v \in \mathbb{R}^d$, suppose we have $v_1 = \gamma_1 v + v_1^\perp$ and $v_2 = \gamma_2 v + v_2^\perp$ for $1 \geq \gamma_1 \geq \gamma_2 \geq 0$ and v_1^\perp, v_2^\perp orthogonal to v . Suppose additionally that (v_1, b_1) and (v_2, b_2) are both (Δ, α) -close to (v, b) . For $s \in \{\pm 1\}$, we have that

$$\mathbb{E}_{x \sim \mathcal{N}(0, Id)} \left[(\sigma(\langle v_1, x \rangle - b_1) + s\sigma(\langle v_2, x \rangle - b_2) - (\gamma_1 + s\gamma_2)\sigma(\langle v, x \rangle - b))^2 \right] \leq O\left(\Delta^{2/5} + \alpha^2\right) \|v\|^2$$

Proof. For $i = 1, 2$, because $|\sin \angle(v_i, v)| \leq \Delta$, we find $\|v_i^\perp\| \leq \Delta \|v_i\| \leq O(\Delta \|v\|)$ for Δ sufficiently small. From Lemma 3.1 we have $\|\sigma(\langle v_i, \cdot \rangle - b_i) - \sigma(\langle \gamma_i v, \cdot \rangle - b_i)\| \leq O(\Delta^{1/5} \|v\|)$. Note that

$$(\gamma_i b - b_i) \|v\|^2 = b \langle v, v_i \rangle - b_i \|v\|^2 \leq \|v\| \|b v_i - b_i v\| \leq \alpha \|v\|^2 \|v_i\|,$$

i.e. $\gamma_i b - b_i \leq \alpha \|v_i\|$. So by Lemma 3.5, $\|\sigma(\langle \gamma_i v, \cdot \rangle - b_i) - \sigma(\langle \gamma_i v, \cdot \rangle - \gamma_i b)\| \leq \alpha \|v_i\|$. The lemma follows by triangle inequality and the fact that $\|v_i\| \leq \|v\| \sqrt{1 + \Delta^2} \leq 2\|v\|$. \square

Lemma 3.13 suggests the following binary operation.

Definition 3.14. Fix a vector $v^* \in \mathbb{R}^d$. Consider the set of all triples (s, v, b) for which $s \in \{\pm 1\}$, $b \in \mathbb{R}$, and v satisfies $0 \leq \langle v, v^* \rangle \leq \|v^*\|^2$. Define the binary operator \odot_{v^*} as follows. Suppose $v_1 = \gamma_1 v + v_1^\perp$ and $v_2 = \gamma_2 v + v_2^\perp$ as in Lemma 3.13, and define $\gamma = |s_1\gamma_1 + s_2\gamma_2|$. Then

$$(s_1, v_1, b_1) \odot_{v^*} (s_2, v_2, b_2) = (s_i, \gamma v, \gamma b) \text{ for } i = \arg \max_j \gamma_j$$

Note that s_i corresponds to the sign of $s_1\gamma_1 + s_2\gamma_2$, and $s_i\gamma = s_1\gamma_1 + s_2\gamma_2$.

In this notation we can restate Lemma 3.13 as follows:

Lemma 3.15. For v_1, v_2, b_1, b_2, v , satisfying the conditions of Lemma 3.13, if we define the tuple (s', v', b') by $(s', v', b') = (s_1, v_1, b_1) \odot_{v^*} (s_2, v_2, b_2)$ we have that

$$\mathbb{E}_{x \sim \mathcal{N}(0, Id)} \left[(s_1\sigma(\langle v_1, x \rangle - b_1) + s_2\sigma(\langle v_2, x \rangle - b_2) - s'\sigma(\langle v', x \rangle - b'))^2 \right] \leq O\left(\Delta^{2/5} + \alpha^2\right) \|v\|^2$$

It will be useful to record some basic properties of this binary operation:

Fact 3.16. \odot_{v^*} is associative and commutative. Moreover, if $(s_1, v_1, b_1) \odot_{v^*} \dots \odot_{v^*} (s_m, v_m, b_m) = (s, \gamma v, \gamma b)$ for s given by the sign of $\sum_i s_i \gamma_i$, where $v_i = \gamma_i v^* + v_i^\perp$ for v_i^\perp orthogonal to v^* , then s is the sign of $\sum s_i \gamma_i$, and $s\gamma = \sum s_i \gamma_i$.

Proof. That \odot_{v^*} is commutative is evident from the definition. For associativity, consider (s_1, v_1, b_1) , (s_2, v_2, b_2) , (s_3, v_3, b_3) . Recall that if $(s_1, v_1, b_1) \odot_{v^*} (s_2, v_2, b_2) = (s_i, \gamma_{12} v^*, \gamma_{12} b)$ for $\gamma_{12} = |s_1\gamma_1 + s_2\gamma_2|$, then s_i corresponds to the sign of $s_1\gamma_1 + s_2\gamma_2$, so $s_i\gamma = s_1\gamma_1 + s_2\gamma_2$. We conclude that

$$((s_1, v_1, b_1) \odot_{v^*} (s_2, v_2, b_2)) \odot_{v^*} (s_3, v_3, b_3) = (s_i, \gamma_{12} v^*, \gamma_{12} b) \odot_{v^*} (s_3, v_3, b_3) = (s_{i'}, \gamma_{123} v^*, \gamma_{123} b)$$

for $\gamma_{123} = |s_1\gamma_1 + s_2\gamma_2 + s_3\gamma_3|$ and $s_{i'}$ corresponding to the sign of $s_1\gamma_1 + s_2\gamma_2 + s_3\gamma_3$. It is therefore evident that \odot_{v^*} is associative. The last part of the claim follows by induction. \square

We show that merging many neurons which are all (Δ, α) close to some given neuron $\sigma(\langle v^*, \cdot \rangle - b^*)$ results in a neuron which is also close to $\sigma(\langle v^*, \cdot \rangle - b^*)$.

Lemma 3.17. *Let $m > 1$. Given $v_1, \dots, v_m, v^* \in \mathbb{R}^d$ and b_1, \dots, b_m, b^* for which every (v_i, b_i) is (Δ, α) -close to (v^*, b^*) and satisfies $\langle v_i, v^* \rangle \geq 0$, we have that for*

$$(s, v, b) \triangleq (s_1, v_1, b_1) \odot_{v^*} \dots \odot_{v^*} (s_m, v_m, b_m),$$

(v, b) is $(0, 0)$ -close to (v^*, b^*) and satisfies $\langle v, v^* \rangle \geq 0$. Furthermore, $\|v\| \leq \sum_i \|v_i\|$ and $|b| \leq \alpha \sum_i \|v_i\| + \sum_i |b_i|$.

Proof. Suppose first that $m = 2$. As usual, let $\Pi_{v^*} v_i = \gamma_i v^*$. Recall that $v = \gamma v^*$ and $b = \gamma b^*$ for $\gamma = |s_1 \gamma_1 + s_2 \gamma_2|$. As a result, we clearly have that $\langle v, v^* \rangle \geq 0$. Furthermore,

$$\|b v^* - b^* v\| = \|\gamma b^* v^* - \gamma b^* v^*\| = 0.$$

The first part of the claim then follows by induction. For the norm bound, note that $\|v\| = |\sum_i \gamma_i| \cdot \|v^*\| \leq \sum_i \|v_i\|$. For the bound on $|b|$, recall from Lemma 3.8 that for every i , $\|\gamma_i b^* - b_i\| \leq \alpha \|v_i\|$. So $|b| = |\sum_i \gamma_i b^*| \leq \sum_i (|b_i| + \alpha \|v_i\|)$ as claimed. \square

Putting everything from this subsection together, we are now ready to prove Lemma 3.12:

Proof of Lemma 3.12. Denote \odot_{v^*} by \odot . Let S^+ denote the set of $i \in [k]$ for which $\langle v^*, v_i \rangle \geq 0$, and let S^- denote the remaining indices $i \in [k]$. Define $F^+(x) \triangleq \sum_{i \in S^+} \sigma(\langle w_i, x \rangle - b_i)$ and $F^-(x) \triangleq \sum_{i \in S^-} \sigma(\langle w_i, x \rangle - b_i)$. By Lemma 3.13, Lemma 3.17, and triangle inequality, we have that for $(s^+, w^+, b^+) \triangleq \bigodot_{i \in S^+} (s_i, w_i, b_i)$ and $(s^-, w^-, b^-) \triangleq \bigodot_{i \in S^-} (s_i, w_i, b_i)$,

$$\|F^+ - s^+ \sigma(\langle w^+, \cdot \rangle - b^+)\|^2, \|F^- - s^- \sigma(\langle w^-, \cdot \rangle - b^-)\|^2 \leq O(k^2(\Delta^{2/5} + \alpha^2)) \|v^*\|^2.$$

Recalling that $(w^+, b^+) = (\gamma^+ v^*, \gamma^+ b^*)$ and $(w^-, b^-) = (\gamma^- v^*, \gamma^- b^*)$, we conclude the proof of (4) with one more application of triangle inequality. For the bounds in (5), we simply apply the last part of Lemma 3.17. \square

3.4 Constructing a Close Neuron

Note that Lemma 3.12 requires the existence of a neuron (v^*, b^*) which is close to all neurons $\{(v_i, b_i)\}$. In our algorithm, we will not have access to (v^*, b^*) but rather to some linear combination of the neurons $\{(v_i, b_i)\}$. We first show that provided this linear combination is not too small in norm, it will also be close to all the neurons $\{(v_i, b_i)\}$.

Lemma 3.18. *Suppose we have vectors $v_1, \dots, v_m, v^* \in \mathbb{R}^d$, biases $b_1, \dots, b_m, b^* \in \mathbb{R}$ for which every (v_i, b_i) is (Δ, α) -close to (v^*, b^*) . Then for any $s_1, \dots, s_m \in \{\pm 1\}$, if we define $v \triangleq \sum_{i=1}^m s_i v_i$ and $b \triangleq \sum_{i=1}^m s_i b_i$, then (v, b) is $(\Delta m, \alpha \sum_i \|v_i\| / \|v\|)$ -close to (v^*, b^*) .*

Proof. Note that $\angle(\sum_i s_i v_i, v) \leq \sum_i \angle(v_i, v)$. By Fact 2.1, we have that $\sin \angle(\sum_i s_i v_i, v) \leq \Delta m$.

The lemma then follows from noting that

$$\|b^* v - b v^*\| = \left\| \sum_i s_i (b_i v^* - b^* v_i) \right\| \leq \alpha \|v^*\| \cdot \sum_i \|v_i\| = \alpha \|v\| \|v^*\| \cdot \sum_i \|v_i\| / \|v\|.$$

\square

3.5 A Corner Case

This presents an issue: what if the linear combination of neurons that we get access to in our eventual algorithm has small norm, in which case Lemma 3.18 is not helpful? It turns out this linear combination takes a very specific form (see the vector in (6)), and we argue in this section that if it is indeed small, then the underlying network we are trying to approximate will be close to a linear function! The main result of this subsection is to show:

Lemma 3.19. *Suppose $(v_1, b_1), \dots, (v_m, b_m)$ are pairwise (Δ, α) -close, and let $[m] = S_1 \sqcup S_2$ denote the orientation induced by them (see Definition 3.10). If signs $s_1, \dots, s_m \in \{\pm 1\}$ satisfy*

$$\left\| \sum_{i \in S_1} s_i v_i - \sum_{i \in S_2} s_i v_i \right\| \leq (\Delta R)^{2/9}, \quad (6)$$

then for the network $F(x) \triangleq \sum_i s_i \sigma(\langle v_i, x \rangle - b_i)$, there exists an affine linear function $\ell(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ for which

$$\mathbb{E}_{x \sim \mathcal{N}(0, \text{Id})} \left[(F(x) - \ell(x))^2 \right] \leq \text{poly}(k, R, B) \cdot (\alpha^{1/2} + \Delta^{2/9}) \quad (7)$$

where $B \triangleq \max_i \|b_i\|$ and $R \triangleq \max_i \|v_i\|$, and $\ell \triangleq \langle w^*, \cdot \rangle - b^*$ satisfying

$$\|w^*\| \leq \sum_i \|v_i\| \quad \text{and} \quad |b^*| \leq \sum_i \|b_i\|. \quad (8)$$

Before proceeding to the proof, we will need the following stability result for affine linear threshold functions with possibly different thresholds.

Lemma 3.20. *Suppose (v, b) and (v', b') are (Δ, α) -close and $\|v\| \geq \|v'\|$. If $\langle v, v' \rangle \geq 0$ then*

$$\Pr[\langle v, x \rangle > b \wedge \langle v', x \rangle < b'] \leq O\left(\alpha + \sqrt{\Delta \|v\| / \|v'\|}\right). \quad (9)$$

Otherwise, if $\langle v, v' \rangle < 0$, then

$$\Pr[\langle v, x \rangle > b \wedge \langle v', x \rangle > b'] \leq O\left(\alpha + \sqrt{\Delta \|v\| / \|v'\|}\right).$$

Proof. Clearly it suffices to prove (9). Suppose $\|v'\| \leq \|v\|$ and write $v' = \gamma v + v^\perp$ for v^\perp orthogonal to v . Note that $\|v^\perp\| \leq \Delta \|v'\|$ and that $\|v^\perp\| \leq \gamma \|v\| \cdot \tan \angle(v, v') \leq O(\gamma \Delta \|v\|)$ for Δ sufficiently small.

Note that

$$\Pr[\text{sgn}(\langle v', x \rangle - \gamma b) \neq \text{sgn}(\langle v', x \rangle - b')] \leq \Pr_{g \sim \mathcal{N}(0, \|v'\|^2)}[g \in [\gamma b \wedge b', \gamma b \vee b']] \leq \frac{|b' - \gamma b|}{2\|v'\|}. \quad (10)$$

Because $\|bv' - b'v\| = \|(b\gamma - b')v + bv^\perp\| \leq \alpha \|v\| \|v'\|$, we have that $|b\gamma - b'| \leq \alpha \|v'\|$. We conclude that $\Pr[\text{sgn}(\langle v', x \rangle - \gamma b) \neq \text{sgn}(\langle v', x \rangle - b')] \leq \alpha/2$.

So by a union bound it suffices to bound $\Pr[\langle \gamma v, x \rangle > \gamma b \wedge \langle v', x \rangle < \gamma b]$. By Lemma 3.4, this is at most $\frac{\|v' - \gamma v\|}{\gamma b} = \frac{1}{\gamma b} \|v^\perp\| \leq O(\frac{\Delta \|v\|}{b})$.

We can also bound this in a different way. By a similar calculation to (10), we have $\Pr[\text{sgn}(\langle \gamma v, x \rangle - \gamma b) \neq \text{sgn}(\langle \gamma v, x \rangle)] \leq \frac{b}{2\|v\|}$ and $\Pr[\text{sgn}(\langle v', x \rangle - b) \neq \text{sgn}(\langle v', x \rangle)] \leq \frac{b}{2\|v'\|}$. And by Sheppard's formula, $\Pr[\text{sgn}(\langle v, x \rangle) \neq \text{sgn}(\langle v', x \rangle)] \leq \frac{\angle(v, v')}{\pi} \leq O(\Delta)$ for Δ sufficiently small.

We conclude that

$$\Pr[\langle \gamma v, x \rangle > \gamma b \wedge \langle v', x \rangle < \gamma b] \lesssim \frac{\Delta \|v\|}{b} \wedge \left(\frac{b}{\|v'\|} + \Delta \right) \lesssim \sqrt{\Delta \|v\| / \|v'\|},$$

from which the claim follows. \square

We can now prove Lemma 3.19.

Proof of Lemma 3.19. Define $\omega \triangleq \|\sum_{i \in S_1} s_i v_i - \sum_{i \in S_2} s_i v_i\|$. Let $S_0 \subseteq [m]$ denote the set of i for which $\|v_i\| \leq (\Delta R)^{1/9}$. For $i \in S_0$, note that by Lipschitz-ness of the ReLU function,

$$\|\sigma(\langle v_i, \cdot \rangle - b_i) - \sigma(-b_i)\|^2 \leq \|\langle v_i, \cdot \rangle\|^2 = \|v_i\|^2 \leq \Delta^{2/9} R^{2/9}.$$

So by triangle inequality it suffices to show that $\sum_{i \notin S_0} s_i \sigma(\langle v_i, x \rangle - b_i)$ is well-approximated by some affine linear function. We will thus assume without loss of generality that $S_0 = \emptyset$.

By Lemma 3.20 and a union bound over all pairs $i, j \in [m]$, we have that with probability at least $1 - O(m^2 \alpha + m^2 \Delta^{4/9} R^{4/9})$ over $x \sim \mathcal{N}(0, \text{Id})$, $\text{sgn}(\langle v_i, x \rangle - b_i) = \text{sgn}(\langle v_j, x \rangle - b_j)$ is the same for all $i, j \in S_1$ and for all $i, j \in S_2$, and $\text{sgn}(\langle v_i, x \rangle - b_i) \neq \text{sgn}(\langle v_j, x \rangle - b_j)$ for all $i \in S_1, j \in S_2$. Let $\mathbb{1}[x \in \mathcal{E}]$ denote the indicator for this event. In other words, with high probability all of the neurons in S_1 are activated and none in S_2 are, or vice versa; denote these two events by \mathcal{E}_1 and \mathcal{E}_2 respectively.

For $j = 1, 2$, note that when $x \in \mathcal{E}_j$, $F(x) = \left\langle \sum_{i \in S_j} s_i v_i, x \right\rangle - \sum_{i \in S_j} s_i b_i$. Define $\ell(x) = \left\langle \sum_{i \in S_1} s_i v_i, x \right\rangle - \sum_{i \in S_1} s_i b_i$. Obviously when $x \in S_1$, $F(x) = \ell(x)$. To handle $x \in S_2$, we need to bound $\delta \triangleq |\sum_{i \in S_1} s_i b_i - \sum_{i \in S_2} s_i b_i|$. Let $(v, b) = (v_1, b_1)$ and note that because (v_i, b_i) is (Δ, α) -close to (v, b) for all i ,

$$\alpha \|v\| \sum_i \|v_i\| \geq \left\| \left(\sum_{i \in S_1} s_i b_i - \sum_{i \in S_2} s_i b_i \right) v - b \left(\sum_{i \in S_1} s_i v_i - \sum_{i \in S_2} s_i v_i \right) \right\| \geq \delta \|v\| - |b| \omega.$$

In particular, $\delta \leq \alpha \sum_i \|v_i\| + |b| \omega / \|v\| \leq \alpha R + B \omega / (\Delta R)^{1/9}$.

We would like to apply Lemma 3.6 to $F(x) - \ell(x)$ (projected to the span of $\{v_i\}$). In that lemma, we can take $\epsilon(x) \leq |\langle \sum_{i \in S_1} s_i v_i - \sum_{i \in S_2} s_i v_i, x \rangle| + \delta$, for which we have $\mathbb{E}[\epsilon(x)^4]^{1/2} \leq O(\delta^2 + \omega^2)$. Additionally we can naively bound $F(0) - \ell(0) \leq 2 \sum_i |b_i|$ and therefore take M in that lemma to be $2 \sum_i |b_i| \leq 2mB$. In addition, we can take $\zeta = O(m^2 \alpha + m^2 \Delta^{4/9} R^{4/9})$, $L = 2mR$, and $d = .$

We conclude that

$$\begin{aligned} \mathbb{E}[(F(x) - \ell(x))^2] &= \mathbb{E}[(F(x) - \ell(x))^2 \mathbb{1}[x \in \mathcal{E}_2]] + \mathbb{E}[(F(x) - \ell(x))^2 \mathbb{1}[x \notin \mathcal{E}]] \\ &\lesssim (m \alpha^{1/2} + m \Delta^{2/9} R^{2/9}) \cdot (m^2 B^2 + m^4 R^2) + \alpha^2 R^2 + B^2 \omega^2 / (\Delta R)^{2/9} + \omega^2. \end{aligned}$$

Recalling that we paid an additional $m^2 (\Delta R)^{2/9}$ in square loss in reducing to the case where $S_0 = \emptyset$, we obtain the desired bound in (7). The bounds in (8) follow immediately from the definition of ℓ above. \square

3.6 Putting Everything Together

Putting Lemmas 3.12, 3.18, and 3.19 together, we conclude that networks whose hidden units are pairwise (Δ, α) -close can either be approximated by a particular size-two network, or by *some* affine linear function:

Lemma 3.21. *Suppose $(v_1, b_1), \dots, (v_k, b_k)$ are pairwise (Δ, α) -close, and let $[k] = S_1 \sqcup S_2$ denote the orientation induced by them (see Definition 3.10). Define $B \triangleq \max_i \|b_i\|$ and $R \triangleq \max_i \|v_i\|$. Let $s_1, \dots, s_m \in \{\pm 1\}$.*

Define $F(x) = \sum_i s_i \sigma(\langle v_i, x \rangle - b_i)$, $v^ = \sum_{i \in S_1} s_i v_i - \sum_{i \in S_2} s_i v_i$, and $b^* = \sum_{i \in S_1} s_i b_i - \sum_{i \in S_2} s_i b_i$. At least one of the following holds:*

1. There is an affine linear function $\ell : \mathbb{R}^d \rightarrow \mathbb{R}$ for which $\|F - \ell\|^2 \leq \text{poly}(k, R, B) \cdot (\alpha^{1/2} + \Delta^{2/9})$.
2. There exist coefficients $a^+, a^- \in \mathbb{R}$ for which $G(x) \triangleq a^+ \sigma(\langle v^*, x \rangle - b^*) - a^- \sigma(\langle -v^*, x \rangle + b^*)$ satisfies $\|F - G\|^2 \leq \text{poly}(k, R, B) \cdot (\Delta^{2/5} + \alpha^2 \Delta^{-4/9})$.

Proof. By assumption, every (v_i, b_i) is (Δ, α) -close to (v_1, b_1) . By Lemma 3.18 we get that for (v^*, b^*) defined in the lemma statement, (v_1, b_1) is $(\Delta k, \alpha m R / \|v^*\|)$ -close to (v^*, b^*) .

If $\|v^*\| \geq (\Delta R)^{2/9}$, then we conclude that (v_1, b_1) is $(\Delta k, \alpha m \Delta^{-2/9} R^{7/9})$ -close to (v^*, b^*) , and by Lemma 3.12 we find that there is a choice of a^+, a^- for which the function G defined in the lemma statement satisfies $\|F - G\|^2 \leq O(k^4 R^2 (\Delta^{2/5} k^{2/5} + \alpha^2 m^2 \Delta^{-4/9} R^{14/9}))$ (note that we used $\|v^*\| \leq \sum_i \|v_i\| \leq kR$).

If $\|v^*\| \leq (\Delta R)^{2/9}$, then by Lemma 3.19 we find that there is an affine linear ℓ for which $\|F - \ell\|^2 \leq \text{poly}(k, R, B) \cdot (\alpha^{1/2} + \Delta^{2/9})$. \square

4 Algorithm for Learning from Queries

In this section we give our algorithm for learning neural networks from queries. Throughout, we will suppose we have black-box query access to some unknown one-hidden layer neural network

$$F(x) \triangleq \sum_{i=1}^k s_i \sigma(\langle w_i, x \rangle - b_i), \quad (11)$$

where $s_i \in \{\pm 1\}$, $w_i \in \mathbb{R}^d$, $b_i \in \mathbb{R}$. Define the quantities $R \triangleq \max_i \|w_i\|$ and $B \triangleq \max_i |b_i|$; our bounds will be polynomial in these quantities, among others.

In Section 4.1, we give bounds on the separation among critical points of random restrictions of F . In Section 4.2 we prove our main *existence theorem* showing that by carefully searching along a random restriction of F , we are able to recover a collection of neurons that can be combined to approximate F . In Section 4.3 we show how to implement certain key steps in GETNEURONS involving querying the gradient and bias of F at certain points. Finally, in Section 4.4 we show to find an appropriate combination of these neurons.

4.1 Critical Points of One-Hidden Layer Networks

In this section, we compute the critical points of restrictions of F and argue that they are far apart along *random restrictions* unless if the corresponding neurons were close to begin with (in the sense of Definition 3.7).

First, we formalize the notion of a random restriction:

Definition 4.1. A Gaussian line L is a random line in \mathbb{R}^d formed as follows: sample $x_0 \sim \mathcal{N}(0, \text{Id})$ and Haar-random $v \in \mathbb{S}^{d-1}$ and form the line $L \triangleq \{x_0 + t \cdot v\}_{t \in \mathbb{R}}$.

Here we compute the critical points along a restriction of F .

Proposition 4.2. Given a line $L = \{x_0 + t \cdot v\}_{t \in \mathbb{R}}$, the restriction $F|_L(t) \triangleq F(x_0 + t \cdot v)$ is given by

$$F|_L(t) = \sum_{i=1}^k s_i \sigma(\langle w_i, x_0 \rangle - b_i + t \langle w_i, v \rangle).$$

This function has k critical points, namely $t = -\frac{\langle w_i, x_0 \rangle - b_i}{\langle w_i, v \rangle}$ for every $i \in [k]$.

Proof. The critical points of $F|_L$ are precisely the points t at which a neuron changes sign. So the critical point associated to the i -th neuron is the t for which $\langle w_i, x_0 \rangle - b_i + t \langle w_i, v \rangle = 0$, from which the claim follows. \square

We can show that these critical points are not too large, unless the norm of the corresponding weight vector is small. The reason for the latter caveat is that, e.g., if one took the one-dimensional neuron $\sigma(\epsilon z - b)$ for b fixed and $\epsilon \rightarrow 0$, the z at which it changes sign tends to ∞).

Lemma 4.3. *With probability at least $1 - \delta$ over the randomness of Gaussian line L , we have that $|t_i| \lesssim \frac{k(\sqrt{d} + \sqrt{\log(1/\delta)})}{\delta \|w_i\|} + k \left(\sqrt{d} + \sqrt{\log(1/\delta)} \right) \sqrt{\log(k/\delta)}$ for every critical point t_i of $F|_L$.*

Proof. By Lemma 2.9, with probability $1 - \delta$ we have that $|\langle w_i, v \rangle| \gtrsim \frac{\delta \|w_i\|}{k(\sqrt{d} + \sqrt{\log(1/\delta)})}$ for all $i \in [k]$.

Also note that $|\langle w_i, x_0 \rangle| \leq \|w_i\| \cdot \sqrt{\log(k/\delta)}$ for all $i \in [k]$ by Fact 2.6. By Proposition 4.2, the critical point corresponding to the i -th hidden unit satisfies

$$\begin{aligned} |t_i| &= \left| \frac{\langle w_i, x_0 \rangle - b_i}{\langle w_i, v \rangle} \right| \lesssim \frac{k(\sqrt{d} + \sqrt{\log(1/\delta)})}{\delta \|w_i\|} \left(B + \|w_i\| \sqrt{\log(k/\delta)} \right) \\ &\leq \frac{k(\sqrt{d} + \sqrt{\log(1/\delta)})}{\delta \|w_i\|} + k \left(\sqrt{d} + \sqrt{\log(1/\delta)} \right) \sqrt{\log(k/\delta)}. \end{aligned}$$

\square

Fix a separation parameter $\Delta > 0$ which we will tune in the sequel. We show that along Gaussian lines L , $F|_L$'s critical points are well-separated except for those corresponding to neurons which are (Δ, α) -close.

Lemma 4.4. *There is an absolute constant $c > 0$ for which the following holds. Given Gaussian line L , with probability at least $1 - \delta$ we have: for any pair of i, j for which (w_i, b_i) and (w_j, b_j) are not $(\Delta, c\Delta\sqrt{\log(k/\delta)})$ -close, the corresponding critical points are at least $\Omega\left(\frac{\Delta\delta^2}{k^4(\sqrt{d} + \sqrt{\log(k/\delta)})}\right)$ apart.*

Proof. For every $i \in [k]$, let $t_i \triangleq -\frac{\langle w_i, x_0 \rangle - b_i}{\langle w_i, v \rangle}$ denote the location of the critical point corresponding to neuron i . For any $i, j \in [k]$,

$$\begin{aligned} |t_j - t_i| &= \left| \frac{\langle w_j, v \rangle (\langle w_i, x_0 \rangle - b_i) - \langle w_i, v \rangle (\langle w_j, x_0 \rangle - b_j)}{\langle w_i, v \rangle \langle w_j, v \rangle} \right| \\ &\geq \frac{|\langle (\langle w_i, x_0 \rangle w_j - \langle w_j, x_0 \rangle w_i) - (b_i w_j - b_j w_i), v \rangle|}{\|w_i\| \|w_j\|} \triangleq |\langle z_{ij}, v \rangle|. \end{aligned}$$

Note that $(\langle w_i, x_0 \rangle w_j - \langle w_j, x_0 \rangle w_i) - (b_i w_j - b_j w_i)$ is distributed as $\mathcal{N}(\mu, \Sigma)$ for $\mu = -b_i w_j + b_j w_i$ and $\Sigma^{1/2} = w_j w_i^\top - w_i w_j^\top$. One can verify that

$$\|\Sigma\|_F^{1/2} = 2^{1/4} (\|w_i\|^2 \|w_j\|^2 - \langle w_i, w_j \rangle^2)^{1/2} = 2^{1/4} \|w_i\| \|w_j\| |\sin \angle(w_i, w_j)|$$

For the first part of the lemma, suppose $|\sin \angle(w_i, w_j)| \geq \Delta$ so that $\|\Sigma\|_F^{1/2} \geq \Omega(\Delta \|w_i\| \|w_j\|)$. Then by Lemma 2.8 we conclude that $\|z_{ij}\| \geq \Omega(\Delta \delta / k^2)$ with probability at least $1 - \delta/k^2$. Recall that v is a random unit vector drawn independently of x_0 , so the lemma follows by applying Lemma 2.9 and a union bound over all pairs i, j .

On the other hand, suppose $|\sin \angle(w_i, w_j)| \leq \Delta$ but $\|\mu\| \geq c\Delta\sqrt{\log(k/\delta)}\|w_i\|\|w_j\|$ for $c > 0$ sufficiently large. Note that Σ has rank 2, so by Fact 2.5, the norm of a sample from $\mathcal{N}(0, \Sigma)$ has norm at most $O(\|\Sigma^{1/2}\|_{\text{op}}(\sqrt{2} + \sqrt{\log(k/\delta)})) = O(\Delta\|w_i\|\|w_j\|\sqrt{\log(k/\delta)})$ with probability at least $1 - \delta/k^2$. So if we take c large enough that this is at least $\Omega\left(\frac{\Delta\delta^2}{k^4\sqrt{d}}\right)$ less than $c\Delta\|w_i\|\|w_j\|\sqrt{\log(k/\delta)}$, we conclude that $\|z_{ij}\| \geq \Omega(\Delta\delta/k^2)$ with probability at least $1 - \delta/k^2$. \square

4.2 Line Search and Existence Theorem

At a high level, our algorithm works by searching along $F|_L$, partitioning L into small intervals, and computing differences between the gradients/biases of F at the midpoints of these intervals. The primary structural result we must show is that there exists enough information in this set of differences to reconstruct F up to small error.

As we will be working with partitions of lines, it will be convenient to define the following notation:

Definition 4.5. *Given line $L \subset \mathbb{R}^d$ and finite interval $I \subseteq \mathbb{R}$ corresponding to a segment $\mathcal{I} \subset L$, let $\nabla_L(I)$ denote the gradient of F at the midpoint of \mathcal{I} . For $t_{\text{mid}} \in \mathbb{R}$ the midpoint of I , define $b_L(I) \triangleq F|_L(t_{\text{mid}}) - (F|_L)'(t_{\text{mid}}) \cdot t_{\text{mid}}$. Intuitively, this is the “y-intercept” of the linear piece of $F|_L$ that contains t_{mid} . When L is clear from context, we will drop subscripts and denote these objects by $\nabla(I)$ and $b(I)$.*

Definition 4.6. *Given line $L \subset \mathbb{R}^d$ and length $s > 0$, let $\{t_i\}$ denote the critical points of $F|_L$, and let $G_L^+(s) \subseteq [k]$ (resp. $G_L^-(s)$) denote the set of indices $a \leq i \leq b$ for which $t_{i+1} - t_i \geq s$ (resp. $t_i - t_{i-1} \geq s$). Let $G_L^*(r) \triangleq G_L^+(r) \cap G_L^-(r)$.*

The following observation motivates Definition 4.6:

Observation 4.7. *Given line $L \subset \mathbb{R}^d$, let $\{t_i\}$ denote the critical points of $F|_L$. Let I_1, \dots, I_m be a partition of some interval I into pieces of length r , and for $t_i \in I$ let $\ell(i)$ denote the index of the interval containing I .*

Then for any $i \in G_L^+(2r)$ (resp. $i \in G_L^-(2r)$), $I_{\ell(i)+1}$ is entirely contained within $[t_i, t_{i+1}]$ (resp. $I_{\ell(i)-1}$ is entirely contained within $[t_{i-1}, t_i]$). In particular, $I_{\ell(i)-1}$ and $I_{\ell(i)+1}$ are linear pieces of $F|_L$.

The following is the main result of this section. At a high level, it says if we partition a random line in \mathbb{R}^d into sufficiently small intervals and can compute the gradient of F at the midpoint of each interval, then we can produce a collection of neurons which can be used to approximate F .

Theorem 4.8. *For any $\epsilon, \delta > 0$, define*

$$r \triangleq O\left(\frac{\Delta\delta^2}{k^4(\sqrt{d} + \sqrt{\log(k/\delta)})}\right) \quad (12)$$

$$\tau \triangleq kr + \Theta\left(\frac{k(\sqrt{d} + \sqrt{\log(1/\delta)})}{\delta\|w_i\|} + k\left(\sqrt{d} + \sqrt{\log(1/\delta)}\right)\sqrt{\log(k/\delta)}\right). \quad (13)$$

Partition the interval $[-\tau, \tau]$ into intervals I_1, \dots, I_m of length r .

Let L be a Gaussian line, and let \mathcal{S} denote the set of all $m(m-1)$ pairs (w, b) obtained by taking distinct $i, j \in [k]$ and forming $(\nabla_L(I_i) - \nabla_L(I_j), b_L(I_i) - b_L(I_j))$. There exist $\{\pm 1\}$ -valued coefficients $\{a_{w,b}\}_{(w,b) \in \mathcal{S}}$, vector w^* , and $b^* \in \mathbb{R}$ for which

$$\left\| F - \sum_{(w,b) \in \mathcal{S}} a_{w,b} \cdot \sigma(\langle w, \cdot \rangle - b) - \langle w^*, \cdot \rangle - b^* \right\| \leq \epsilon + \mathfrak{P}_{k,R,B,\log(1/\delta)} \cdot \Delta^{2/9}.$$

for $\mathfrak{P}_{k,R,B,\log(1/\delta)}$ some absolute constant that is polynomially large in $k, R, B, \log(1/\delta)$. Furthermore, we have that

$$\|a_{w,b} \cdot w\| \leq kR \quad \text{and} \quad |a_{w,b} \cdot b| \leq c\Delta k^2 R \sqrt{\log(k/\delta)} + kB \quad (14)$$

$$\|w^*\| \leq kR \quad \text{and} \quad |b^*| \leq kB \quad (15)$$

Proof. Condition on the outcomes of Lemma 4.4 and Lemma 4.3 holding for L . Let t_1, \dots, t_k denote the critical points associated to neurons w_1, \dots, w_k , and for convenience we assume without loss of generality that $t_1 \leq \dots \leq t_k$. Let $a, b \in [k]$ denote the indices for which $|t_i| \leq \tau$ for $i \in [a, b]$. By Lemma 4.3 and the definition of τ , we have that for $i \notin [a, b]$, $\|w_i\| \leq \epsilon/k$.

By Lipschitzness of the ReLU function,

$$\begin{aligned} \left\| \sum_{i \notin [a,b]} s_i \sigma(\langle w_i, \cdot \rangle - b_i) - \sum_{i \notin [a,b]} s_i \sigma(-b_i) \right\| &\leq \sum_{i \notin [a,b]} \|\sigma(\langle w_i, \cdot \rangle - b_i) - \sigma(-b_i)\| \\ &\leq \sum_{i \notin [a,b]} \|w_i\| \leq (b - a + 1)\epsilon/k. \end{aligned} \quad (16)$$

Next, we handle the critical points $i \in [a, b]$. Given critical point t_i , let $\ell(i) \in [m]$ denote the index for which $t_i \in I_{\ell(i)}$. For convenience, denote $G_L^+(2r), G_L^-(2r), G_L^*(2r)$ by G^+, G^-, G^* . By Observation 4.7, we know that for $i \in G^*$, the linear piece of $F|_L$ immediately preceding critical point t_i contains $I_{\ell(i)-1}$, and the one immediately proceeding t_i contains $I_{\ell(i)+1}$. Therefore, $\nabla(I_{\ell(i)+1}) - \nabla(I_{\ell(i)-1})$ and $b(I_{\ell(i)-1}) - b(I_{\ell(i)+1})$ are equal to w_i and b_i up to a sign, so \mathcal{S} must contain the neurons (w_i, b_i) and $(-w_i, -b_i)$.

Now consider any neighboring $i_1 < i_2$ in $G^+ \Delta G^-$ for which $i_2 - i_1 > 1$; note that the latter condition implies that $i_1 \in G^- \setminus G^+$ and $i_2 \in G^+ \setminus G^-$, or else we would have a violation of the fact that i_1 and i_2 are neighboring. Furthermore, because i_1, i_2 are neighboring, for all $i_1 \leq i \leq i_2$ we have that $t_{i+1} - t_i \leq 2r$. By taking Δ in (the contrapositive of) Lemma 4.4 to be $\Delta \cdot k$, we conclude that for any $i_1 \leq i < j \leq i_2$, (w_i, b_i) and (w_j, b_j) are $(\Delta k, c\Delta k \sqrt{\log(k/\delta)})$ -close for all such i .

Let $\{i_1, \dots, i_2\} = S_1 \sqcup S_2$ denote the orientation induced by $(w_{i_1}, b_{i_1}), \dots, (w_{i_2}, b_{i_2})$. We would like to apply Lemma 3.21 to the subnetwork $\tilde{F}(x) \triangleq \sum_{j=i_1}^{i_2} s_j \sigma(\langle w_j, x \rangle - b_j)$. By another application of Observation 4.7, we know that $\nabla(I_{\ell(i_2)}) - \nabla(I_{\ell(i_1)})$ and $b(I_{\ell(i_1)}) - b(I_{\ell(i_2)})$ are, up to a common sign, precisely the vector v^* and bias b^* defined in Lemma 3.21, so we conclude that either there exists a network G consisting of neurons $\sigma(\langle v^*, x \rangle - b^*)$ and $\sigma(\langle -v^*, x \rangle + b^*)$ for which $\|\tilde{F} - G\|^2 \leq \text{poly}(k, R, B) \cdot (\Delta^{2/5} k^{2/5} + c^2 \Delta^{14/9} k^2 \log(k/\delta)) \leq \text{poly}(k, R, B) \Delta^{2/5} \log(1/\delta)$, or there is an affine linear function ℓ for which $\|\tilde{F} - \ell\|^2 \leq \text{poly}(k, R, B) \cdot (c^{1/2} \Delta^{1/2} \log(1/\delta)^{1/2} + \Delta^{2/9}) \leq \text{poly}(k, R, B) \cdot \Delta^{2/9} \log(1/\delta)^{1/2}$. Furthermore, the bounds in (14) and (15) follow from (5) in Lemma 3.12 (for $\alpha = c\Delta k \sqrt{\log(k/\delta)}$) and (8) in Lemma 3.19 respectively.

We have accounted for all critical points, except in the case where the smallest index a' in G^- is not a , or the largest index b' in G^+ is not b . In the former (resp. latter) case, note that

$t_a \leq \dots \leq t_{a'-1} \leq -\tau + kr$, (resp. $t_b \geq \dots \geq t_{b'+1} \geq \tau - kr$), so by Lemma 4.3, this implies that $\|w_{a'-1}\|, \dots, \|w_a\| \leq \epsilon/k$ (resp. $\|w_{b'+1}\|, \dots, \|w_b\| \leq \epsilon/k$). By Lipschitzness of the ReLU function, we can approximate these neurons by constants at a total cost of at most $(a' - a + b - b')\epsilon/k$ in L_2 using the same reasoning as (16). \square

4.3 Gradient and Bias Oracles

It remains to implement oracles to compute $b_L(I)$ and $\nabla_L(I)$ for prescribed line L and interval I . It is not clear how to do this for arbitrarily small intervals because for general networks there can be many arbitrarily close critical points, but we will only need to do so for certain “nice” I as suggested by Theorem 4.8.

To that end, first note that it is straightforward to form the quantities $b_L(I)$ for intervals I entirely contained within linear pieces of $F|_L$; we formalize this in Algorithm 1.

Algorithm 1: GETBIAS(L, I)

Input: Line $L \subset \mathbb{R}^d$, interval $I = [a, b] \subset \mathbb{R}$
Output: $b_L(I)$ if I is entirely contained within a linear piece of $F|_L$

- 1 $t_{\text{mid}} \leftarrow$ midpoint of I .
- 2 $y_0 \leftarrow F|_L(t_{\text{mid}})$.
- 3 $s \leftarrow \frac{F|_L(b) - F|_L(a)}{b - a}$.
- 4 **return** $y_0 - s \cdot t_{\text{mid}}$

It remains to demonstrate how to construct $\nabla_L(I)$. Intuitively one can accomplish this via “finite differencing,” i.e. the gradient of a piecewise linear function F at a point x can be computed from queries by computing $\frac{F(x+\delta) - F(x)}{\delta}$ several sufficiently small perturbations $\delta \in \mathbb{R}^d$ and solving the linear system.

With *a priori* precision estimates, we can similarly implement a gradient oracle, as formalized in Algorithm 2 and Lemma 4.9.

Algorithm 2: GETGRADIENT(x, α)

Input: $x \in \mathbb{R}^d$, $\alpha > 0$ for which (17) holds
Output: $\nabla F(x) \in \mathbb{R}^d$

- 1 **for** $j \in [d]$ **do**
- 2 Sample random unit vector $z_j \in \mathbb{S}^{d-1}$.
- 3 $v_j \leftarrow (F(x + \alpha z_j) - F(x)) / \alpha$.
- 4 Let w be the solution to the linear system $\{\langle w, z_j \rangle = v_j\}_{j \in [d]}$.
- 5 **return** w

Lemma 4.9. For any $\alpha > 0$ and any $x \in \mathbb{R}^d$ for which

$$|\langle w_i, x \rangle - b_i| \geq \alpha \|w_i\| \quad \forall i \in [k], \quad (17)$$

GETGRADIENT(x, α) makes d queries to F and outputs $\nabla F(x)$.

Proof. For any $z \in \mathbb{S}^{d-1}$, note that

$$\langle w_i, x + \alpha z \rangle - b_i = (\langle w_i, x \rangle - b_i) + \alpha \langle w_i, z \rangle,$$

and $\alpha|\langle w_i, z \rangle| \leq \alpha \cdot \|w_i\|$, so $\langle w_i, x + \alpha z \rangle - b_i$ and $\langle w_i, x \rangle + b_i$ have the same sign. As a result, if $S \subseteq [k]$ denotes the indices i for which $\langle w_i, x \rangle - b_i > 0$, then

$$\frac{F(x + \alpha z) - F(x)}{\alpha} = \left\langle \sum_{i \in S} s_i w_i, z \right\rangle = \langle \nabla F(x), z \rangle.$$

If $\{z_1, \dots, z_j\}$ are a collection of Haar-random unit vectors, they are linearly independent almost surely, in which case the linear system in Step 4 of GETGRADIENT has a unique solution, namely $\nabla F(x)$. \square

In order to use GETGRADIENT to construct the vectors $\nabla_L(I)$, we require estimates for α in Lemma 4.9. In the following lemma we show that with high probability over the randomness of L , if an interval I completely lies within a linear piece of $F|_L$, then we can bound how small we must take α to query the gradient of F at the midpoint of that interval.

Lemma 4.10. 4.10 *Let L be a Gaussian line. With probability at least $1 - \delta$ over the randomness of L , the following holds: in the partition $[-\tau, \tau] = I_1 \cup \dots \cup I_m$ in Theorem 4.8, for any I_ℓ which entirely lies within a linear piece of $F|_L$, GETGRADIENT(t_{mid}, α) correctly outputs $\nabla_L(I_\ell)$, where x_{mid} is the midpoint of the interval $\mathcal{I}_\ell \subset L$ that corresponds to interval $I_\ell \subset \mathbb{R}$ and $\alpha = \frac{\delta \cdot r}{4k\sqrt{d} + O(k\sqrt{\log(k/\delta)})}$ (where r is defined in (12)).*

Proof. Denote $L = \{x_0 + t \cdot v\}_{t \in \mathbb{R}}$. Let $t_{\text{mid}} \in \mathbb{R}$ denote the value corresponding to $x_{\text{mid}} \in \mathbb{R}^d$ on the line L . By Lemma 2.9 and a union bound over $[k]$, we have that

$$|\langle w_i, v \rangle| \geq \frac{\delta \|w_i\|}{2k\sqrt{d} + O(k\sqrt{\log(k/\delta)})} \quad \text{for all } i \in [k]$$

with probability at least $1 - \delta$ over the randomness of $v \in \mathbb{S}^{d-1}$. Now take any interval I_ℓ which entirely lies within a linear piece of $F|_L$. Because t_{mid} is the midpoint of I_ℓ , it is at least $r/2$ away from any critical point of $F|_L$. In particular, $|\langle w_i, x_{\text{mid}} \rangle - b| \geq (r/2) \cdot |\langle w_i, v \rangle| \geq (r/2) \cdot \frac{\delta \|w_i\|}{2k\sqrt{d} + O(k\sqrt{\log(k/\delta)})}$, so we can take $\alpha = \frac{\delta \cdot r}{4k\sqrt{d} + O(k\sqrt{\log(k/\delta)})}$ and invoke Lemma 4.9. \square

Putting these ingredients together, we obtain the following algorithm, GETNEURONS for producing a collection of neurons that can be used to approximate F .

We prove correctness of GETNEURONS in the following lemma:

Lemma 4.11. 4.11 *For any $\epsilon, \delta > 0$, GETNEURONS(ϵ, δ) makes $\text{poly}(k, d, R, B, 1/\epsilon, \log(1/\delta))$ queries and outputs a list \mathcal{S} of pairs (w, b) for which there exist $\{\pm 1\}$ -valued coefficients $\{a_{w,b}\}_{(w,b) \in \mathcal{S}}$ as well as a vector w^* and a scalar b^* such that*

$$\left\| F - \langle w^*, \cdot \rangle - b^* - \sum_{(w,b) \in \mathcal{S}} a_{w,b} \cdot \sigma(\langle w, \cdot \rangle - b) \right\| \leq \epsilon.$$

Proof. By Lemma 4.10, the choice of α in GETNEURONS is sufficiently small that for x_j the midpoint of any interval which is entirely contained within a linear piece of $F|_L$, GETGRADIENT(x_j, α) succeeds by Lemma 4.9. So the estimates $\hat{\nabla}$ and \hat{b} are exactly correct for all intervals that are entirely contained within a linear piece of $F|_L$. By the proof of Theorem 4.8, these are the only intervals for which we need $\nabla_L(I)$ and $b_L(I)$ in order for \mathcal{S} to contain enough neurons to approximate F by some linear combination to L_2 error ϵ . \square

Algorithm 3: GETNEURONS(ϵ, δ)

Input: Accuracy $\epsilon > 0$, confidence $\delta > 0$
Output: List \mathcal{S} of pairs (w, b) (see Theorem 4.8 for guarantee)

- 1 $\mathcal{S} \leftarrow \emptyset$.
- 2 Sample Gaussian line L .
- 3 $\Delta \leftarrow (\epsilon/\mathfrak{P}_{k,R,B,\log(1/\delta)})^{9/2}$. *// Theorem 4.8*
- 4 $\alpha \leftarrow \frac{\delta \cdot r}{4k\sqrt{d} + O(k\sqrt{\log(k/\delta)})}$. *// Lemma 4.10*
- 5 Define r, τ according to (12), (13).
- 6 Partition $[-\tau, \tau]$ into disjoint intervals I_1, \dots, I_m of length r .
- 7 **for** all $j \in [m]$ **do**
 - 8 $x_j \leftarrow$ midpoint of the interval $\mathcal{I}_j \subset L$ that corresponds to $I_j \subset \mathbb{R}$.
 - 9 $\hat{\nabla}_L(I_j) \leftarrow \text{GETGRADIENT}(x_j, \alpha)$.
 - 10 $\hat{b}_L(I_j) \leftarrow \text{GETBIAS}(L, I_j)$.
- 11 **for** all pairs of distinct $i, j \in [m]$ **do**
 - 12 $(v_j, b_j) \leftarrow (\hat{\nabla}_L(I_i) - \hat{\nabla}_L(I_j), \hat{b}_L(I_i) - \hat{b}_L(I_j))$.
 - 13 **if** (v_j, b_j) satisfies the bounds in (14) **then**
 - 14 Add (v_j, b_j) to \mathcal{S} .
- 15 **return** \mathcal{S} .

4.4 Linear Regression Over ReLU Features

It remains to show how to combine the neurons produced by GETNEURONS to obtain a good approximation to F . As Theorem 4.8 already ensures that some linear combination of them suffices, we can simply draw many samples $(x, F(x))$ for $x \sim \mathcal{N}(0, \text{Id})$, form the feature vectors computed by the neurons output by GETNEURONS, and run linear regression on these feature vectors.

Formally, let \mathcal{S} denote the set of pairs (w, b) guaranteed by Theorem 4.8. We will denote the w 's by $\{\hat{w}_j\}$ and the b 's by $\{\hat{b}_j\}$. Consider the following distribution over feature vectors computed by the neurons in \mathcal{S} :

Definition 4.12. Let \mathcal{D}' denote the distribution over $\mathbb{R}^{|\mathcal{S}|+d+1} \times \mathbb{R}$ of pairs (z, y) given by sampling $x \sim \mathcal{N}(0, \text{Id})$ and forming the vector z whose entries consist of all $\sigma(\langle \hat{w}_j, x \rangle - b_j)$ as well as the entries of x and the entry 1, and taking y to be $F(x)$ for the ground truth network F defined in (11).

We will also need to define a truncated version of \mathcal{D}' : let \mathcal{D} denote \mathcal{D}' conditioned on the norm of the $|\mathcal{S}|+1$ to $|\mathcal{S}|+d$ -th coordinates having norm at most $M \triangleq \sqrt{d} + O(\sqrt{\log(1/\delta)})$, which happens with probability at least $1 - \delta$ over \mathcal{D}' .

Our algorithm will be to sample sufficiently many pairs (z, y) from \mathcal{D}' (by querying F on random Gaussian inputs) and run ordinary least squares. This is outlined in LEARNFROMQUERIES below.

To show that regression-based algorithm successfully outputs a network that achieves low population loss with respect to F , we will use the following standard results on generalization.

Theorem 4.13. For \mathcal{D} a distribution over $\mathcal{X} \times \mathcal{Y}$ and $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ a loss function that is L -Lipschitz in its first argument and uniformly bounded above by c . Let \mathcal{F} be a class of functions $\mathcal{X} \rightarrow \mathcal{Y}$ such that for any $f \in \mathcal{F}$ and pairs $(x_1, y_1), \dots, (x_n, y_n)$ drawn independently from \mathcal{D} , with

Algorithm 4: LEARNFROMQUERIES(ϵ, δ)

Input: Accuracy $\epsilon > 0$, confidence $\delta > 0$
Output: One hidden-layer network $\tilde{F} : \mathbb{R}^d \rightarrow \mathbb{R}$ for which $\|F - \tilde{F}\| \leq O(\epsilon)$

- 1 $\mathcal{S} = \{(\hat{w}_j, \hat{b}_j)\} \leftarrow \text{GETNEURONS}(\epsilon, \delta)$.
- 2 Draw samples $(z_1, y_1), \dots, (z_n, y_n)$ from \mathcal{D} // Definition 4.12
- 3 Let \tilde{v} be the solution to the least-squares problem (19). Let \tilde{b} denote the last entry of \tilde{v} , and let \tilde{w} denote the vector given by the d entries of \tilde{v} prior to the last.
- 4 Form the network $\tilde{F}(x) \triangleq \sum_j \tilde{v}_j \sigma(\langle \hat{w}_j, x \rangle - \hat{b}_j) + \langle \tilde{w}, \cdot \rangle - \tilde{b}$.
- 5 **return** F .

probability at least $1 - \delta$,

$$\mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell(f(x), y)] \leq \frac{1}{n} \sum_i \ell(f(x_i), y_i) + 4L \cdot \mathcal{R}_n(\mathcal{F}) + 2c \cdot \sqrt{\frac{\log(1/\delta)}{2n}},$$

where $\mathcal{R}_n(\mathcal{F})$ denotes the Rademacher complexity of \mathcal{F} .

Theorem 4.14. If \mathcal{X} is the set of x satisfying $\|x\| \leq X$, and \mathcal{F} is the set of linear functions $\langle w, \cdot \rangle$ for $\|w\| \leq W$, then $\mathcal{R}_n(\mathcal{F}) \leq XW/\sqrt{n}$.

As these apply to bounded loss functions and covariates, we must first pass from \mathcal{D}' to \mathcal{D} and quantify the error in going from one to the other:

Lemma 4.15. For f satisfying $\mathbb{E}_{(z,y) \sim \mathcal{D}'} [(f(z) - y)^2] \leq \epsilon^2$, we have

$$\left| \mathbb{E}_{(z,y) \sim \mathcal{D}'} [(f(z) - y)^2] - \mathbb{E}_{(x,y) \sim \mathcal{D}} [(f(z) - y)^2] \right| \leq O(\epsilon^2). \quad (18)$$

Proof. Let Z denote the probability that a random draw from \mathcal{D}' lies in the support of \mathcal{D} so that $Z \geq 1 - \delta$; denote this event by \mathcal{E} . Then we can write $\mathbb{E}_{(z,y) \sim \mathcal{D}} [(f(z) - y)^2]$ as $\frac{1}{Z} \mathbb{E}_{(z,y) \sim \mathcal{D}'} [(f(z) - y)^2 \cdot \mathbf{1}[z \in \mathcal{E}]]$ and rewrite the left-hand side of (18) as

$$\left| \left(1 - \frac{1}{Z}\right) \cdot \mathbb{E}_{(z,y) \sim \mathcal{D}'} [(f(z) - y)^2 \cdot \mathbf{1}[z \in \mathcal{E}]] + \mathbb{E}_{(z,y) \sim \mathcal{D}'} [(f(z) - y)^2 \cdot \mathbf{1}[z \notin \mathcal{E}]] \right|.$$

Note that $|1 - 1/Z| \leq 2\delta \leq 1$ for δ sufficiently small, from which the claim follows. □

We are now ready to prove the main theorem of this section:

Theorem 4.16. 4.12 Let \mathcal{S} denote the list of pairs (\hat{w}_j, \hat{b}_j) output by GETNEURONS(ϵ, δ). Sample $(z_1, y_1), \dots, (z_n, y_n)$ from \mathcal{D} for $n = \text{poly}(k, R, B, 1/\epsilon, d, \log(1/\delta))$. With probability at least $1 - O(\delta)$ over the randomness of GETNEURONS and the samples, the following holds. Define

$$\tilde{v} \triangleq \arg \min_{\|v\| \leq W} \sum_{i=1}^n (\langle v, z_i \rangle - y_i)^2, \text{ for } W \triangleq \sqrt{\tau/r} + k(R + B), \quad (19)$$

let \tilde{b} denote the last entry of \tilde{v} , and let \tilde{w} denote the vector given by the d entries of \tilde{v} prior to the last. Then the one hidden-layer network $\tilde{F}(x) \triangleq \sum_j \tilde{v}_j \sigma(\langle \hat{w}_j, x \rangle - \hat{b}_j) + \langle \tilde{w}, \cdot \rangle - \tilde{b}$ satisfies $\|F - \tilde{F}\| \leq O(\epsilon)$.

Proof. Note that over the support of \mathcal{D} we have that the square loss $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ is uniformly bounded above by $(MkR + kB)^2$ and is $L = O(M \cdot k \cdot R + k \cdot B)$ -Lipschitz. Finally, note that for z in the support of \mathcal{D} ,

$$\begin{aligned} \|z\|^2 &= 1 + M^2 + 2M^2 \sum_j (\|\hat{w}_j\|^2 + \hat{b}_j^2) \\ &\lesssim (M^2\tau/r) \cdot (k^2R^2 + \Delta^2k^4R^2\log(k/\delta) + k^2B^2) \triangleq X^2. \end{aligned} \quad (20)$$

where τ, r are defined in Theorem 4.8 and we used (14) and Step 14 in GETNEURONS to bound $\|\hat{w}_j\|$ and $|\hat{b}_j|$.

By the guarantee on GETNEURONS given by Lemma 4.11, we know that there is a vector $v^* \in \{\pm 1\}^{|\mathcal{S}|} \times B^d(kR) \times [-kB, kB]$ which achieves ϵ^2 squared loss with respect to \mathcal{D}' . Note that

$$\|v^*\| \leq |\mathcal{S}|^{1/2} + k(R + B) = \sqrt{\tau/r} + k(R + B) \triangleq W. \quad (21)$$

By Lemma 4.15, v^* achieves $O(\epsilon^2)$ squared loss with respect to \mathcal{D} . As the random variable $(\langle v^*, z \rangle - y)^2$ for $(z, y) \sim \mathcal{D}$ is bounded above by

$$(\|v^*\| \|z\| + |y|)^2 \lesssim \text{poly}(k, R, B, 1/\epsilon, M),$$

for $n \geq \text{poly}(k, R, B, 1/\epsilon, M)$ we have that the empirical loss of v^* on $(z_1, y_1), \dots (z_n, y_n)$ is $O(\epsilon^2)$, and therefore that of the predictor \tilde{v} is $O(\epsilon^2)$.

By applying Theorem 4.14 with (20) and (21), we find that the Rademacher complexity $\mathcal{R}_n(\mathcal{F})$ of the family of linear predictors over $\|z\| \leq X$ and with norm bounded by W is C/\sqrt{n} for C which is polynomial in $k, R, B, 1/\epsilon, d, \log(1/\delta)$, from which the theorem follows by Theorem 4.13. \square

References

- [ATV21] Pranjal Awasthi, Alex Tang, and Aravindan Vijayaraghavan. Efficient algorithms for learning depth-2 neural networks with general relu activations. *arXiv preprint arXiv:2107.10209*, 2021.
- [BJW19] Ainesh Bakshi, Rajesh Jayaram, and David P Woodruff. Learning two layer rectified neural networks in polynomial time. In *Conference on Learning Theory*, pages 195–268. PMLR, 2019.
- [CJM20] Nicholas Carlini, Matthew Jagielski, and Ilya Mironov. Cryptanalytic extraction of neural network models. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 189–218. Springer, 2020.
- [CKM20] Sitan Chen, Adam Klivans, and Raghu Meka. Learning deep relu networks is fixed parameter tractable. *arXiv preprint arXiv:2009.13512*, 2020.
- [CM20] Sitan Chen and Raghu Meka. Learning polynomials of few relevant dimensions. *arXiv preprint arXiv:2004.13748*, 2020.
- [CS09] Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. In Yoshua Bengio, Dale Schuurmans, John D. Lafferty, Christopher K. I. Williams, and Aron Culotta, editors, *Advances in Neural Information Processing Systems 22: 23rd Annual*

Conference on Neural Information Processing Systems 2009. Proceedings of a meeting held 7-10 December 2009, Vancouver, British Columbia, Canada, pages 342–350. Curran Associates, Inc, 2009.

- [CW01] A. Carbery and J. Wright. Distributional and L^q norm inequalities for polynomials over convex bodies in R^n . *Mathematical Research Letters*, 8(3):233–248, 2001.
- [DG21] Amit Daniely and Elad Granot. An exact poly-time membership-queries algorithm for extraction a three-layer relu network. *arXiv preprint arXiv:2105.09673*, 2021.
- [DKKZ20] Ilias Diakonikolas, Daniel M Kane, Vasilis Kontonis, and Nikos Zarifis. Algorithms and sq lower bounds for pac learning one-hidden-layer relu networks. In *Conference on Learning Theory*, pages 1514–1539, 2020.
- [DV20] Amit Daniely and Gal Vardi. Hardness of learning neural networks with natural weights. *arXiv preprint arXiv:2006.03177*, 2020.
- [DV21] Amit Daniely and Gal Vardi. From local pseudorandom generators to hardness of learning. *CoRR*, abs/2101.08303, 2021.
- [GGJ⁺20] Surbhi Goel, Aravind Gollakota, Zhihan Jin, Sushrut Karmalkar, and Adam Klivans. Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. *arXiv preprint arXiv:2006.12011*, 2020.
- [GKLW18] Rong Ge, Rohith Kuditipudi, Zhize Li, and Xiang Wang. Learning two-layer neural networks with symmetric inputs. In *International Conference on Learning Representations*, 2018.
- [GKM18] Surbhi Goel, Adam R. Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. In Jennifer G. Dy and Andreas Krause 0001, editors, *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 1778–1786. PMLR, 2018.
- [GLM18] Rong Ge, Jason D Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.
- [Han09] Steve Hanneke. *Theoretical foundations of active learning*. Carnegie Mellon University, 2009.
- [JCB⁺20] Matthew Jagielski, Nicholas Carlini, David Berthelot, Alex Kurakin, and Nicolas Papernot. High accuracy and high fidelity extraction of neural networks. In Srdjan Capkun and Franziska Roesner, editors, *29th USENIX Security Symposium, USENIX Security 2020, August 12-14, 2020*, pages 1345–1362. USENIX Association, 2020.
- [JSA15] Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv*, pages arXiv–1506, 2015.
- [JWZ20] Rajesh Jayaram, David P. Woodruff, and Qiuyi Zhang. Span recovery for deep neural networks with applications to input obfuscation, 2020.
- [LMZ20] Yuanzhi Li, Tengyu Ma, and Hongyang R. Zhang. Learning over-parametrized two-layer relu neural networks beyond ntk. *CoRR*, abs/2007.04596, 2020.

[MSDH19] Smitha Milli, Ludwig Schmidt, Anca D. Dragan, and Moritz Hardt. Model reconstruction from model explanations. In *FAT*, pages 1–9. ACM, 2019.

[PMG⁺17] Nicolas Papernot, Patrick D. McDaniel, Ian J. Goodfellow, Somesh Jha, Z. Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In Ramesh Karri, Ozgur Sinanoglu, Ahmad-Reza Sadeghi, and Xun Yi, editors, *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security, AsiaCCS 2017, Abu Dhabi, United Arab Emirates, April 2-6, 2017*, pages 506–519. ACM, 2017.

[RK20] David Rolnick and Konrad P. Kording. Reverse-engineering deep relu networks. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 8178–8187. PMLR, 2020.

[TJ⁺16] Florian Tramèr, Fan Zhang 0022, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. *CoRR*, abs/1609.02943, 2016.

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

[ZSJ⁺17] Kai Zhong, Zhao Song, Prateek Jain, Peter L Bartlett, and Inderjit S Dhillon. Recovery guarantees for one-hidden-layer neural networks. In *Proceedings of the 34th International Conference on Machine Learning- Volume 70*, pages 4140–4149, 2017.

[ZYWG19] Xiao Zhang, Yaodong Yu, Lingxiao Wang, and Quanquan Gu. Learning one-hidden-layer relu networks via gradient descent. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1524–1534. PMLR, 2019.