

A Local Convergence Theory for Mildly Over-Parameterized Two-Layer Neural Network

Mo Zhou

Duke University

MOZHOU@CS.DUKE.EDU

Rong Ge

Duke University

RONGGE@CS.DUKE.EDU

Chi Jin

Princeton University

CHIJ@PRINCETON.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

While over-parameterization is widely believed to be crucial for the success of optimization for the neural networks, most existing theories on over-parameterization do not fully explain the reason—they either work in the Neural Tangent Kernel regime where neurons don't move much, or require an enormous number of neurons. In practice, when the data is generated using a teacher neural network, even mildly over-parameterized neural networks can achieve 0 loss and recover the directions of teacher neurons. In this paper we develop a local convergence theory for mildly over-parameterized two-layer neural net. We show that as long as the loss is already lower than a threshold (polynomial in relevant parameters), all student neurons in an over-parameterized two-layer neural network will converge to one of teacher neurons, and the loss will go to 0. Our result holds for any number of student neurons as long as it is at least as large as the number of teacher neurons, and our convergence rate is independent of the number of student neurons. A key component of our analysis is the new characterization of local optimization landscape—we show the gradient satisfies a special case of Lojasiewicz property which is different from local strong convexity or PL conditions used in previous work.

1. Introduction

Recent years, deep learning has achieved great empirical success in a wide range of applications including speech recognition, image detection, natural language processing, game playing, etc. In practice, simple optimization algorithms such as gradient descent (GD) and stochastic gradient descent (SGD) typically already achieve zero training loss. However, in theory, training deep neural networks remains a challenging problem, as it requires optimizing highly non-convex objective functions. Recent works suggest that over-parameterization is a key to the success of training for neural networks.

One line of work, known as the Neural Tangent Kernels (NTK) (Jacot et al., 2018; Chizat et al., 2019; Du et al., 2018; Allen-Zhu et al., 2018), shows that neural network training can get 0 training loss when the network is sufficiently over-parameterized. However, this theory also suggests that the neurons will not move very far from their initial positions, which is often not true for practical neural networks.

Another line of work uses a mean-field limit to analyze two-layer neural networks (Chizat and Bach, 2018; Mei et al., 2018). While this type of work would allow neurons to move far, the theoretical results often require the number of neurons to go to infinity, or be exponential in relevant

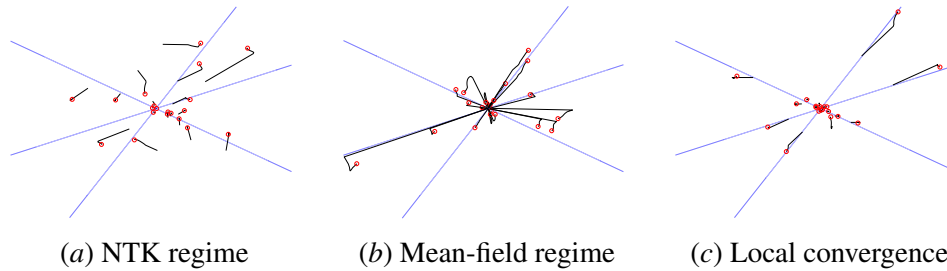


Figure 1: Training two-layer neural networks in 2 dimension with $m = 20$ student neurons and $r = 3$ teacher neurons. Blue lines represent the direction of teacher neurons, black curves represent the trajectories for each student neuron, and red points represent their end positions.

parameters. [Chizat et al. \(2019\)](#) unified the two lines of work by showing that NTK is equivalent to a lazy training regime (Figure 1(a)) where the initialization has a very large scale, while mean-field analysis can handle settings where the initialization is smaller.

In this paper we consider a simple teacher-student setting, where the training data (x, y) is generated by sampling x from a Gaussian and evaluating y using a ground truth two-layer teacher network. The goal is to train a student network that mimics the behavior of the teacher. Figure 1 illustrates the differences between two lines of work in the teacher student setting – in the NTK/lazy training regime (Figure 1(a)) student neurons do not move much, while in the mean field regime (Figure 1(b)) student neurons converges to one of the directions of teacher neurons¹. However, when the number of neurons is small, there are no analysis that shows why student neurons need to match the teacher neurons.

In fact, the phenomenon that student neurons converge locally and match teacher neurons is not even understood in a simpler over-parameterization setting, where initially there are already student neurons close to each teacher neuron (Figure 1(c)). [Safran et al. \(2020\)](#) observed that traditional techniques that rely on local strong convexity or PL conditions cannot be applied here. In this paper we focus on the following natural question:

When the initial loss is small, will student neurons always match teacher neurons for an over-parameterized two-layer student net?

We show that this is indeed true. In particular, we prove

Theorem 1 (Informal) *Given data generated by a two-layer teacher network² with r neurons that are Δ -separated (see Assumption 1). There exists a threshold $\tau = \text{poly}(\Delta/r)$ such that when loss is smaller than τ , gradient descent converges to a global optimum where all student neurons match (in direction) one of the teacher neurons.*

Note that the threshold τ is independent of the student network size – as long as the initial loss is low, even when the number of student neurons is equal or mildly larger than the number of

1. Similar empirical observations of student-teacher neuron matching were also known for deeper networks ([Tian, 2019](#)).

2. The specific architecture of the network is specified in Definition 2 in Section 2.

teacher neurons, gradient descent will still converge to the global optimal solution where all student neurons match teacher neurons. In low dimensions or for simple teacher neurons, we can also give initialization procedures that efficiently finds an initialization with loss smaller than τ .

1.1. Related Work

Neural Tangent Kernel (NTK) One line of the recent work connects the training of sufficiently over-parameterized neural networks with gradient descent to NTK (Jacot et al., 2018; Chizat et al., 2019; Du et al., 2018, 2019; Allen-Zhu et al., 2018; Cao and Gu, 2019; Zou et al., 2020; Li and Liang, 2018; Daniely et al., 2016; Arora et al., 2019a,b; Oymak and Soltanolkotabi, 2020; Ghorbani et al., 2021). In NTK regime, training neural networks with gradient descent is essentially solving kernel regression with NTK. A key technique in NTK analysis is to restrict the neurons to stay around initialization by choosing large enough width of the neural network. As pointed out in Chizat et al. (2019), neural networks essentially degenerate to linear function and makes the optimization become convex. Instead, our result allows neurons to move away from initialization and recover the ground truth neurons.

Mean-field analysis Another line of research uses mean-field approach to analyze the training of infinite-width neural networks (Chizat and Bach, 2018; Mei et al., 2018, 2019; Wei et al., 2019; Nguyen and Pham, 2020; Araújo et al., 2019; Nitanda and Suzuki, 2017; Sirignano and Spiliopoulos, 2020; Rotskoff and Vanden-Eijnden, 2018; Lu et al., 2020; Fang et al., 2020). In mean-field analysis, they focus on the dynamics of the distribution of neurons, and as the number of hidden neurons goes to infinite, gradient descent becomes Wasserstein gradient flow. Different from NTK regime, neurons can move away from initialization. However, these works often either require exponential (or infinite) number of neurons or only provide exponential convergence rate.

Local landscape analysis Several works have studied the local landscape property around the global minima in the teacher-student two-layer neural network setting. Zhong et al. (2017); Zhang et al. (2019) studied the problem in the exact-parameterization case and showed the Hessian around global minima is positive-definite. Chizat (2019) studied the over-parameterized neural network with a regularization term. They showed that when loss is small, it satisfies PL condition. Roughly speaking, their analysis relies on several kernels to be positive definite. However, they need to require some non-degenerate conditions that are difficult to verify in the case of neural networks. Further, these kernels become degenerate when the regularization term tends to zero. Safran et al. (2020) studied the over-parameterization case with orthogonal teacher neurons and showed that neither convexity nor PL condition could hold even in the local region of global minima. This indicates that the analysis discussed above cannot be applied in our over-parameterization setting. In contrast, we could show a slightly different version of PL condition holds when loss is small.

Student-teacher neuron matching and the lottery ticket hypothesis The lottery ticket hypothesis (Frankle and Carbin, 2018) showed that it is possible to prune a neural network such that even if training is done only on a small subset of randomly initialized neurons, the final network still achieves good accuracy. Our local convergence result gives a partial explanation of this phenomenon for two-layer teacher/student setting – as long as the initialization contains student neurons that are close to each teacher neuron, the training process can converge to a global optimal solution. See more discussions in Section E.4.

1.2. Outline

In Section 2 we formally define the neural network architecture that we work with. Then in Section 3 we summarize our main results, including the formal version of Theorem 1, potential initialization algorithms and generalizations in the setting of polynomial sample sizes. In Section 4, we illustrate the unique challenges in establishing local convergence results for overparameterized setting. In Section 5 we sketch the proof of a main lemma that lowerbounds the norm of the gradient, which is the main contribution of this paper. We also show how the main lemma can be used to prove Theorem 1. Finally we conclude in Section 6.

2. Preliminaries

Teacher/student setting for two-layer neural network We consider the standard teacher-student setting with Gaussian input $x \sim N(0, I_d)$. We parameterize the teacher/student networks according to the following definition:

Definition 2 (Teacher-student setup) *Teacher network is parameterized as $f^*(x) = \sum_{i=1}^r |w_i^{*\top} x|$, where $\{w_i^*\}_{i=1}^r$ ($w_i^* \in \mathbb{R}^d$) are the r teacher neurons. The student network is parameterized as $f(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x|$, where $\{w_i\}_{i=1}^m$ ($w_i \in \mathbb{R}^d$) are m student neurons ($m \geq r$). Denote $W = (w_1, \dots, w_m)$ as the weight matrix formed by student neurons. The loss function we optimize is the population square loss:*

$$\min_W L(W) = \mathbb{E}_{x \sim N(0, I)} \left[\frac{1}{2} \left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right]. \quad (1)$$

Choice of the neural network architecture Note that we use $\|w_i\|$ as the top layer weight in the student network. This parameterization of the student network ensures the smoothness of loss function, which helps our analysis as discussed in the later section. The same model was also used in Li et al. (2020). Since the two-layer neural network is 2-homogeneous, this restriction of the top layer weight is equivalent to requiring all top-layer weights to be nonnegative. Nonnegativity is important to our result as without this assumption, there might be two student neurons that completely cancel out each other and they may not converge to the direction of any teacher neuron.

We also remark that absolute value function is used as activation function in both teacher network and student network. We use absolute value function instead of ReLU (where $\text{ReLU}(x) = \max\{x, 0\}$) to ensure identifiability of our model. For ReLU activation, even at global minima, there may have student neuron that does not correspond to any teacher neuron. See the Claim below and more discussions in Section A.2.

Claim 1 *For problem (1) with ReLU activation, when loss is zero, there may exist a student neuron whose direction does not match direction of any teacher neuron.*

On the other hand, using absolute value function as activation we can achieve student-teacher matching at global minima. It is directly obtained by setting $\epsilon = 0$ in Lemma C.6.

Theorem 3 *For problem (1) with absolute value activation, when loss is zero, every student neuron's direction must match one of the teacher neuron's direction.*

We remark that we identify w and $-w$ as the same direction and measure the angle between directions up to a sign (i.e., $\angle(w, v) = \arccos[|w^\top v| / (\|w\| \|v\|)]$) when using the absolute value function as the activation. It is because absolute value function is an even function so there is no difference between $|w^\top x|$ and $|-w^\top x|$. Note that when ReLU is the activation, we cannot identify w and $-w$ as the same direction as $\text{ReLU}(w^\top x) \neq \text{ReLU}(-w^\top x)$.

Note that when the input data is Gaussian (or just symmetric), after first fitting the optimal linear function to the data, a teacher network with ReLU activation becomes a teacher network with absolute value activation. See more discussions in Section A.1.

Assumptions We make following two assumptions throughout this paper.

Assumption 1 (Δ -separation) *The teacher neurons $w_1^*, w_2^*, \dots, w_r^* \in \mathbb{R}^d$ are Δ -separated, i.e., $\angle(w_i^*, w_j^*) \geq \Delta$ for all $i \neq j \in [r]$, where $\angle(w_i^*, w_j^*) = \arccos[|w_i^{*\top} w_j^*| / (\|w_i^*\| \|w_j^*\|)]$.*

Intuitively, Δ indicates the level of difficulty to distinguish the weight vectors w_i^* separately. When Δ is large, all weight vectors w_i^* of the teacher network are well-separated. When Δ is small, there may exist two weight vectors that are close to each other. In this case, it may be hard to distinguish between them.

Assumption 2 (Norm Bounded) *The teacher neurons $w_1^*, w_2^*, \dots, w_r^* \in \mathbb{R}^d$ are norm bounded, i.e., $0 < w_{\min} \leq \|w_i^*\| \leq w_{\max}$ for all $i \in [r]$.*

Casual readers may assume that $\Delta, w_{\min}, w_{\max}, r$ are constants. In this case, teacher neurons are well-separated and their norms are approximately in the same order.

Notations We will use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For vector $x \in \mathbb{R}^d$, we use $\|x\| = (\sum_{i=1}^d x_i^2)^{1/2}$ to represent the 2-norm of x , and $\bar{x} = x / \|x\|$ as the unit vector in the direction of x . For matrix A , we use $\|A\|_F$ to denote the Frobenius norm of A . For vectors $w, v \in \mathbb{R}^d$, denote $\angle(w, v) = \arccos[|w^\top v| / (\|w\| \|v\|)] \in [0, \pi/2]$ as the angle between w and v (up to a sign). We will use \mathbb{I}_S as the indicator of the set S . Denote the inner product and norm between two function f and g on set S as $\langle f, g \rangle_S = \mathbb{E}_{x \sim N(0, I)}[f(x)g(x)\mathbb{I}_S]$, and $\|f - g\|_S^2 = \mathbb{E}_{x \sim N(0, I)}[(f(x) - g(x))^2 \mathbb{I}_S]$. When $S = \mathbb{R}^d$, we will omit S and denote it as $\|f - g\|$.

3. Main Results

In this section we give a formal version of Theorem 1 and talk about generalizations. We discuss proof ideas for the main results in the Section 5.

3.1. Gradient Lower Bound

We first present a result that characterizes the landscape of the loss function when the loss is small. We show that the gradient norm can be lower bounded by the value of the loss function, which is a special case of Łojasiewicz property (Łojasiewicz, 1963)³.

3. The general Łojasiewicz property is known as $\|\nabla f(x)\| \geq C(f - f^*)^\alpha$, where f^* is the optimal objective value. Theorem 4 corresponds to the case $\alpha = 1$. We remark that the well-known Polyak-Łojasiewicz (PL) property is also a special case of Łojasiewicz property with $\alpha = 1/2$, which has been discussed in many earlier works (see e.g. Karimi et al., 2016)

Theorem 4 (Gradient Lower Bound) For network and loss function $L(\cdot)$ defined in Definition 2, under Assumptions 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W such that loss $L(W) \leq \epsilon_0$, we have

$$\|\nabla_W L(W)\|_F \geq \kappa L(W),$$

where $\kappa = \Theta(r^{-1/2} w_{max}^{-1/2})$.

This result indicates that when the loss is smaller than certain threshold, gradient is zero only if the loss is zero. Hence there are no spurious stationary points.

3.2. Local Convergence

Now we are ready to state the formal version of Theorem 1, which gives local convergence for network defined in Definition 2.

Theorem 5 (Main Result) For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, suppose we run gradient descent on objective (1):

$$w_i^{(t+1)} = w_i^{(t)} - \eta \nabla_{w_i} L(W), \quad \text{for any } i \in [m]$$

there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ and $\eta_0 = O(r^{-1} w_{max}^{-1})$ such that for any $\epsilon > 0$, if initial loss $L(W^{(0)}) \leq \epsilon_0$ and step size $\eta \leq \eta_0$, then we have $L(W^{(T)}) \leq \epsilon$ in $T = O(r w_{max} / (\epsilon \eta))$ steps.

The proof relies on two geometric conditions about the landscape of loss function – the gradient lowerbound as in Theorem 4, and a “smoothness” type guarantee for the loss function (Lemma 13). See the proof in Section 5.5.

In Section 4 we show mild over-parameterization introduces significant difficulty for local convergence results, but the theorem shows that gradient descent can still achieve local convergence despite over-parameterization. The result only has mild requirement on the over-parameterization. As long as $m \geq r$ and the initial solution has low loss, our local convergence result holds. Moreover, neither the initial loss requirement or the convergence rate depends on the number of student neuron m . They only depend on the intrinsic quantities of the teacher network. This demonstrates that gradient descent could leverage the structure of a teacher network without explicitly knowing information like its number of neurons r or separation Δ .

3.3. Initialization

Next we talk about how one can find a good initialization to get into the local convergence regime. We first consider a simple random initialization (see Algorithm 1 in Section E), where the directions of neurons are chosen randomly, and the norm of neurons are fitted by least-squares. We also give a more complicated initialization algorithm (subspace initialization, see Algorithm 2 in Section E), which first estimates a subspace spanned by the teacher neurons and then randomly initialize student neurons in the subspace.

Theorem 6 For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, set $\epsilon_{init} = \epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$, to achieve a good initialization that has ϵ_{init} loss,

random initialization requires $O((rw_{max}/\sqrt{\epsilon_{init}})^d \cdot r \log(1/\delta))$ student neurons and subspace initialization requires $O((rw_{max}/\sqrt{\epsilon_{init}})^r \cdot r \log(1/\delta))$ student neurons with $\tilde{O}(dr^6 w_{max}^4 / \epsilon_{init}^2)$ samples. Suppose we run GD from either initialization, there exists a threshold $\eta_0 = O(r^{-1} w_{max}^{-1})$ such that for any $\epsilon > 0$, if step size $\eta \leq \eta_0$, with probability $1 - \delta$ we have $L(W^{(T)}) \leq \epsilon$ in $T = O(rw_{max}/(\epsilon\eta))$ steps.

When $\Delta = \Omega(1/r)$, $w_{min}, w_{max} = \Theta(1)$, random initialization gives global convergence in polynomial time for $d = O(1)$; while subspace initialization gives global convergence in polynomial time when $r \leq O(\log d / \log \log d)$.

3.4. Sample Complexity

Previous theorems require us to optimize the population least squares loss directly, which is impossible in practice with finitely many samples. Here we consider a setting where we have access to N data points $\{(x_k, y_k)\}_{k=1}^N$ where x_k are i.i.d. sampled from $N(0, I)$ and $y_k = \sum_{i=1}^r |w_i^{*\top} x_k|$. In this case, we can define an empirical loss:

$$\hat{L}(W) = \frac{1}{2N} \sum_{k=1}^N \left(\sum_{i=1}^m \|w_i\| |w_i^\top x_k| - \sum_{i=1}^r |w_i^{*\top} x_k| \right)^2.$$

We can show the gradient on the empirical loss is close to the gradient on the population loss when the number of data points N is large enough, which is formalized in Lemma F.1. Based on this, we can extend Theorem 5 to stochastic GD with mini-batch on training samples. See the proof and more discussions in Section F.

Theorem 7 *For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, suppose we run stochastic GD on W with N fresh samples within each mini-batch at every iteration, i.e., for any $i \in [m]$*

$$w_i^{(t+1)} = w_i^{(t)} - \eta \nabla_{w_i} \hat{L}_t(W),$$

where $\hat{L}_t(\cdot)$ is the empirical loss using the N samples at iteration t . Then, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ and $\eta_0 = O(r^{-1} w_{max}^{-1})$ such that for any $\epsilon > 0$, if initial loss $L(W^{(0)}) \leq \epsilon_0$, step size $\eta \leq \eta_0$ and batch size $N \geq O(r^5 w_{max}^5 d^2 \eta^{-1} \epsilon^{-3} \delta^{-1})$, then with probability $1 - \delta$ we have $L(W^{(T)}) \leq \epsilon$ in $T = O(rw_{max}/(\epsilon\eta))$ steps.

This shows that our local convergence results are robust even when we only use polynomially many samples.

4. Exact-Parameterization v.s. Over-Parameterization

Before talking about our proof ideas, we first highlight the key difference between the exact-parameterization setting and the over-parameterization setting, and the unique challenges in the latter setting. Consider a simple example where there is only one teacher neuron w^* . In the exact-parameterization setting, there is only one student neuron w . Suppose w is close to w^* within the distance δ in the sense $\angle(w, w^*) = \delta$ under the normalization $\|w\| = \|w^*\| = 1$. Then it is not

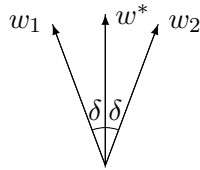


Figure 2: Warm-up Example: one teacher, two students

difficult to show the loss (1) is $\Theta(\delta^2)$. This illustrates that the Hessian around global minima is non-degenerate and positive definite (which is formally established in [Zhong et al. \(2017\)](#); [Zhang et al. \(2019\)](#)). Such properties immediately imply that the loss is locally strongly convex around the global minima, therefore gradient descent initialized in a small local region can converge to the global minima.

In sharp contrast, over-parameterization introduces a significantly more complicated geometry around the global minimum. Consider again the simple example in the over-parameterization setting where there are two student neurons and one teacher neuron (Figure 2). The claim below shows that when there are two student neurons that are δ -close to teacher neuron, the loss can be as small as $\Theta(\delta^3)$ instead.

Claim 2 *Suppose teacher neuron w^* and two student neurons w_1, w_2 belong to the same hyper-plane with position shown in Figure 2. If $\|w^*\| = 1$, $\angle(w_1, w^*) = \angle(w_2, w^*) = \delta$, $\|w_1\| = \|w_2\| = 1/\sqrt{2} \cos \delta$ (so that $\|w_1\| w_1 + \|w_2\| w_2 = w^*$) for a small enough constant δ , then we have loss $L(W) = \Theta(\delta^3)$.*

Claim 2 indicates that the over-parameterization setting is clearly different from the exact-parameterization setting. The scale $L(W) = \Theta(\delta^3)$ implies that the Hessian at the global minima must be degenerate, therefore the loss is not locally strongly convex. Intuitively, this is because that in the exact-parameterization case, we only have a set of isolated global minima. While in the over-parameterization setting, different global minima are connected and the loss could be small when the “average” student neuron ($\|w_1\| w_1 + \|w_2\| w_2$) matches the teacher neuron. The observation above showcases the unique challenge in the over-parameterization setting. A novel geometric characterization as well as corresponding analyses is necessary, which will be developed in Section 5.

5. Proof Overview: Local Geometry and Gradient Lower Bound

In this section, we provide proof sketch for Theorem 4. First, we explain how the gradient lower-bound can be reduced to constructing a descend direction (Section 5.1). Then in Section 5.2 we explain the intuitions for why the descent direction works. To prove our descent direction works, first we show that small loss implies that every teacher neuron has at least one nearby student neuron by constructing test functions (Section 5.3), then we exploit the notion of “average neuron” to decompose the residual into two terms (Section 5.4) and bound them separately. At the end (Section 5.5) we show how one can use Theorem 4 to prove the local convergence result in Theorem 5.

We also introduce following notations about the partition of student neurons that will be used in our analysis. For every teacher neuron w_i^* , denote T_i as the set of student neurons that are closer to w_i^* than other teacher neurons (break the tie arbitrarily), i.e., $T_i = \{j \in [m] | \angle(w_j, w_i^*) \leq \angle(w_j, w_k^*) \text{ for all } k \neq i\}$. It is easy to see that $\cup_{i \in [r]} T_i$ is the set of all student neurons, and $T_i \cap T_j =$

\emptyset for all $i \neq j \in [r]$. Also denote $T_i(\delta) = \{j \in T_i \mid \angle(w_j, w_i^*) \leq \delta\}$ as the set of student neurons w_j that are δ -close to teacher neuron w_i^* . Finally we use $\delta_j = \angle(w_i^*, w_j)$ to denote the angle between w_i^* and w_j for all $j \in T_i$.

5.1. Gradient Lowerbound (Theorem 4): Constructing Descent Direction

We discuss how to obtain the gradient lower bound (Theorem 4). We construct a ‘‘descent direction $g(W)$ ’’ that is correlated to the gradient and prove that $\langle \nabla L(W), g(W) \rangle \geq L(W) \geq 0$, this directly implies $\|\nabla L(W)\| \geq L(W) / \|g(W)\|$, which provides the result in form of Theorem 4. Intuitively, we group the neurons into $r + 1$ categories, where $T_i(\delta_{\max})$ represents student neurons that are within angle δ_{\max} with i -th teacher neuron (where δ_{\max} is a parameter chosen later), and the $r + 1$ -th group consists of neurons not close to any teacher neuron. The direction we construct will move student neurons closer to their corresponding teacher neurons, or 0 if it is in the last group. Formally, we prove the following lemma on descent direction.

Lemma 8 (Descent Direction) *For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{\max}^{-1}, w_{\min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have*

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \text{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq L(W),$$

where $\{q_{ij}\}_{i \in [r], j \in [m]}$ is any sequence that satisfies (1) $q_{ij} \geq 0$ for all (i, j) , (2) $q_{ij} = 0$ if $j \notin T_i(\delta_{\max})$; and (3) $\sum_{j \in T_i(\delta_{\max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$, where $\delta_{\max} = \Theta(r w_{\max} w_{\min}^{-5/3} \cdot \epsilon^{1/3})$.

In Lemma 8, the scalar q_{ij} describes how much fraction of teacher neuron w_i^* that student w_j should target to approximate. Recall that our activation is absolute value function, so there is a symmetry between w and $-w$ and we identify them as the same neuron. Therefore, $(w_j - q_{ij} \text{sgn}(w_j^\top w_i^*) w_i^*)$ can be understood as the difference between the current student neuron and its ‘‘optima’’. We further multiply it by matrix $(I + \bar{w}_j \bar{w}_j^\top)^{-1}$ due to technical reason raised by our parameterization of student neural network as $f(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x|$ instead of $f(x) = \sum_{i=1}^m |w_i^\top x|$ (see Section 2 for more details).

5.2. Descent Direction (Lemma 8): High-level Ideas

We briefly explain why the descent direction we constructed would reduce loss function in this section. First, we introduce a quantity which plays a key role in the remaining of Section 5:

$$\hat{w}_i \triangleq \sum_{j \in T_i} \|w_j\| w_j; \quad v_i \triangleq \hat{w}_i - w_i^* \quad \text{for } i \in [r] \quad (2)$$

Intuitively, \hat{w}_i represents the ‘‘average neuron’’ for student neurons close to teacher neuron w_i^* ; v_i represents the difference between a teacher neuron and its corresponding average student neuron. Note that in the above definition of ‘‘average neuron’’ \hat{w}_i , there is an ambiguity of the direction of the neurons w_j due to the symmetry of the absolute value activation. To break this ambiguity, we will assume the direction of w_j always has positive correlation with w_i^* , i.e., $w_j^\top w_i^* \geq 0$. This is without of generality because absolute value function is an even function so that intuitively we can

replace w_j with $-w_j$ when $w_j^\top w_i^* < 0$ without affecting anything else. In fact, we can show if there is a descent direction $(g(w_1), \dots, g(w_m))$ for (w_1, \dots, w_m) , then for arbitrary choice of $a_i \in \{\pm 1\}$ there is a descent direction $(a_1 g(w_1), \dots, a_m g(w_m))$ for $(a_1 w_1, \dots, a_m w_m)$. We defer the proof to the beginning of Section C.

In order to prove the Lemma 8, a main challenge is to properly exploit the precondition that loss $L(W)$ is small. Our strategy is to establish an intermediate result, and prove the lemma by the following three steps:

1. Show that when the loss $L(W) \leq \epsilon$, for each teacher neuron there is at least one student neuron that is $\delta_{\max} = \Omega(r w_{\max} w_{\min}^{-5/3} \cdot \epsilon^{1/3})$ close to the teacher neuron.
2. Show that loss $L(W)$ is small implies that the $\|v_i\|$ is small for all $i \in [r]$ (average neuron close to teacher neuron).
3. Show that 1 and 2 imply the conclusion of Lemma 8.

In fact, the third step directly follows from our choice of “descent direction”, and algebraic computation. We refer readers to Lemma C.2 in Appendix C.5 for details.

For the first step, we formalize it as the following lemma:

Lemma 9 *Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{\max}^{-1}, w_{\min})$ and an absolute constant C such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, if we choose $\delta_{\max} \geq C \cdot r w_{\max} w_{\min}^{-5/3} \cdot \epsilon^{1/3}$, then we have $|T_i(\delta_{\max})| \geq 1$ and $\sum_{j \in T_i(\delta_{\max})} \|w_j\|^2 \geq \frac{1}{2} \|w_i^*\|^2$ for all $i \in [r]$.*

The proof uses idea of test functions to lowerbound the loss, and will be discussed later in Section 5.3.

Now that each teacher neuron has at least one nearby student neuron, average neuron \hat{w}_i is never 0. Our second step shows that the average student neuron is always close to teacher neuron when the loss is small:

Lemma 10 (Average Student Is Close to Teacher) *Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{\max}^{-1}, w_{\min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\{v_i\}_{i \in [r]}$ (defined in (2)) satisfy $\|v_i\| \leq \text{poly}(r, w_{\max}, \Delta^{-1}) \epsilon^{3/8}$ for all $i \in [r]$.*

The proof of this relies on a characterization of the loss, which we describe in Section 5.4.

5.3. Local Property I (Lemma 9): Lowerbounding Loss using Test Functions

To prove Lemma 9, we show its contraposition – if there is a teacher neuron that does not have a nearby student neuron, then the loss must be large.

In order to lowerbound the loss, we use the idea of constructing a test function. The similar idea has also been used in maximum mean discrepancy (MMD) (Gretton et al., 2012), where they use test function to distinguish two distribution. Here, we try to explicitly construct a test function so that its correlation with the residual can be lower bounded. Formally, when there is a teacher neuron without any student neuron within angle δ , we construct a test function h such that $\langle R(x), h(x) \rangle / \|h(x)\| \geq \Omega(w_{\min}^{5/2} r^{-3/2} w_{\max}^{-3/2} \cdot \delta^{3/2})$, where $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^*{}^\top x|$ is the residual. This directly implies that $L^{1/2}(W) = \|R\| / \sqrt{2} \geq \Omega(w_{\min}^{5/2} r^{-3/2} w_{\max}^{-3/2} \cdot \delta^{3/2})$.

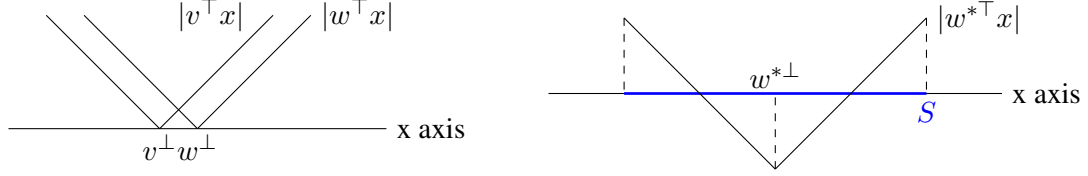


Figure 3: Left: Illustration for nonlinearity and nonlinear region. Here w^\perp and v^\perp represent the vector that is orthogonal to w and v . Right: Test function $h(x)$ in Lemma 9, where $S = \{x \mid |w^{*\top} x| \leq \tau\}$.

Our choice of test function crucially relies on the nonlinearity of absolute function. More specifically, absolute function $|w^\top x|$ is a linear function everywhere else except the hyperplane $w^\top x = 0$. We call the neighborhood of this hyperplane ($\{x \mid |w^\top x| \leq \tau\}$ for some small τ) the nonlinear region. The key observation here is that the loss in this nonlinear region would be small only if there is another neuron w' such that $w' \approx w$, and they cancel each other. See Figure 3 for an illustration. This gives the basic component for constructing test functions for Lemma 9.

Suppose w^* is a teacher neuron without any close-by student neurons. We focus on the nonlinear region of w^* , and construct test function $h(x) = (|w^{*\top} x| - \mathbb{E}_x[|w^{*\top} x| \mathbb{I}_S] / \mathbb{E}_x[\mathbb{I}_S]) \mathbb{I}_S$, where $S = \{x \mid |w^{*\top} x| \leq \tau\}$ with a small enough τ (see also Figure 3). Such a test function has almost zero correlation with any function that is linear in region S (i.e. terms contributed by student and teacher neurons far away from w^*). On the other hand, it has a large positive correlation with teacher neuron $|w^{*\top} x|$ term in the residual. See Section C.2 in appendix for detailed bounds.

5.4. Local Property II (Lemma 10): Residual Decomposition

We first introduce the notion of residual and its decomposition, which turns out to be the key in the proof of Lemma 10. For any fixed student neurons W , recall the residual function $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$. Intuitively, the residual is the difference between the outputs of the student network and the teacher network. Clearly our loss $L(W) = \mathbb{E}_x[R(x)^2]/2$. We also note the gradient of loss is closely related to the residual as $\nabla_{w_j} L(W) = \mathbb{E}_x[R(x) \|w_j\| (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x)]$. We decompose the residual as follows:

$$R_1(x) \triangleq \sum_{i=1}^r v_i^\top x \operatorname{sgn}(w_i^{*\top} x), \quad R_2(x) \triangleq \sum_{i=1}^r \sum_{j \in T_i} \|w_j\| w_j^\top x \left(\operatorname{sgn}(w_j^\top x) - \operatorname{sgn}(w_i^{*\top} x) \right). \quad (3)$$

where v_i is defined in (2). It is clear that $R(x) = R_1(x) + R_2(x)$. As discussed at the beginning of Section 5.2, we can fix the direction of student neurons to have positive correlation with their corresponding teacher neuron, so there is no ambiguity in the definition of ‘‘average neuron’’ \hat{w}_i and thus no ambiguity in the decomposition R_1 and R_2 defined above.

Intuitively, each neuron has a linear weight component and a nonlinear activation pattern. R_1 describes the difference in linear weights, while R_2 describes the difference in activation pattern. In some sense, if one focuses only on R_1 , then the model becomes *exactly parameterized* again because R_1 only depends on the average neurons \hat{w}_i and does not depend on individual neurons. Lemma 11 makes this observation more precise.

Using the decomposition and characterizations of R_1 and R_2 , we can now discuss the proof strategy for Lemma 10.

High-level proof strategy for Lemma 10 To prove Lemma 10, i.e., to show that $\|v_i\|$ is small for all $i \in [r]$, we do following two steps.

1. Show that $\|R_1\|^2$ is strongly convex in (v_1, \dots, v_r) . Thus to show $\|v_i\|$ is small for all i , it is sufficient to show that $\|R_1\|^2$ is small.
2. Show $\|R_2\|^2$ is small. Since $R_1 = R - R_2$, and we know $\|R\|^2 = L(W)$ is small, this implies $\|R_1\|^2$ is small as we need.

Step 1 let $v = (v_1, \dots, v_r)$, we observe that $\|R_1\|^2 = v^\top M v = \sum_{i,j} v_i^\top M_{ij} v_j$ where

$$M_{ij} = \mathbb{E}_{x \sim N(0, I)} \left[x x^\top \operatorname{sgn}(w_i^{*\top} x) \operatorname{sgn}(w_j^{*\top} x) \right]. \quad (4)$$

In fact, $M \in \mathbb{R}^{dr \times dr}$ also corresponds to the Hessian at global minima in the exact-parameterization case. We can show $\lambda_{\min}(M) = \Omega(\Delta^3/r^3)$, which ensures that $\|R_1\|^2$ is strongly convex in v . See the proof in Section C.4.1.

Lemma 11 Under Assumption 1, we have $\|R_1\|^2 \geq \Omega(\Delta^3/r^3) \|v\|^2$.

Step 2 we first observe following upper bound on R_2 .

$$\|R_2\|^2 = \mathbb{E}_x \left[\left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| w_j^\top x \left(\operatorname{sgn}(w_j^\top x) - \operatorname{sgn}(w_i^{*\top} x) \right) \right)^2 \right] \leq r \sum_{i=1}^r \left(\sum_{j \in T_i} \|w_j\|^2 \delta_j^{3/2} \right)^2$$

We obtain the upperbound for $\sum_{i=1}^m \|w_i\|^2$ and $\sum_{i=1}^m \|w_i\|^2 \delta_i^2$ in Lemma C.5 and Lemma C.6. The bound for the latter term uses similar ideas of constructing test functions as Lemma 9, but the test function is more complicated. Combining the two upperbounds we have the following result:

Lemma 12 Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \operatorname{poly}(\Delta, r^{-1}, w_{\max}^{-1}, w_{\min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\|R_2\|^2 = O(r^{5/2} w_{\max}^{1/2} \epsilon^{3/4})$.

5.5. Proof Sketch of Local Convergence Theorem (Theorem 5)

Theorem 4 proves that the gradient norm is lower bounded by the value of loss function when loss is small. In order to establish the convergence result, we need to additionally characterize the local “smoothness” of the loss function.

Lemma 13 (Smoothness) Under Assumption 2, if loss $L(W) = O(r^2 w_{\max}^2)$, then

$$\begin{aligned} L(W + U) \leq & L(W) + \langle \nabla_W L(W), U \rangle + O(r^{1/4} w_{\max}^{1/4}) L^{1/2}(W) \|U\|_F^{3/2} + O(r w_{\max}) \|U\|_F^2 \\ & + O(1) \|U\|_F^4. \end{aligned}$$

We note Lemma 13 is slightly different from the standard notion of smoothness in the optimization literature, but it turns out to be also sufficient to prove our result. See the proof in Section D.1. Furthermore, we can also prove that the function is Lipschitz when loss is small.

Lemma 14 (Lipschitz) *Under Assumption 2, if loss $L(W) = O(r^2 w_{max}^2)$, then $\|\nabla_W L(W)\|_F^2 = O(r^3 w_{max}^3)$.*

Finally, following the standard linear algebra calculations, Theorem 4, Lemma 13 and Lemma 14 jointly establish the convergence result (Theorem 5). We defer the detailed proof to Section D.3.

6. Conclusion

In this paper, we develop a local convergence theory for mildly over-parameterized two-layer neural networks. By characterizing the local landscape and showing gradient satisfies a special case of Łojasiewicz property, we prove that as long as initial loss is below a threshold that is polynomial in relevant parameters, gradient descent could converge to zero loss. Our result is different from NTK analysis and mean-field analysis, since student neurons converge to the ground-truth teacher neuron and we only have mild requirement on the over-parameterization. One immediate open question is when can gradient descent find such a good initialization. We hope our result could lead to stronger optimization results for mildly over-parameterized neural networks.

Acknowledgments

Rong Ge and Mo Zhou are supported in part by NSF-Simons Research Collaborations on the Mathematical and Scientific Foundations of Deep Learning (THEORINET), NSF Award CCF-1704656, CCF-1845171 (CAREER), CCF-1934964 (Tripods), a Sloan Research Fellowship, and a Google Faculty Research Award. Part of the work was done when Rong Ge and Chi Jin were visiting Instituted for Advanced Studies for “Special Year on Optimization, Statistics, and Theoretical Machine Learning” program.

References

- Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1948.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018.
- Dyego Araújo, Roberto I Oliveira, and Daniel Yukimura. A mean-field limit for certain deep neural networks. *arXiv preprint arXiv:1906.00193*, 2019.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019a.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019b.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *arXiv preprint arXiv:1905.13210*, 2019.
- Lenaic Chizat. Sparse optimization on measures with over-parameterized gradient descent. *arXiv preprint arXiv:1907.10300*, 2019.

- Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.
- Amit Daniely, Roy Frostig, and Yoram Singer. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. *arXiv preprint arXiv:1602.05897*, 2016.
- 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.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- Cong Fang, Jason D Lee, Pengkun Yang, and Tong Zhang. Modeling from features: a mean-field framework for over-parameterized deep neural networks. *arXiv preprint arXiv:2007.01452*, 2020.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Rong Ge, Rohith Kuditipudi, Zhize Li, and Xiang Wang. Learning two-layer neural networks with symmetric inputs. *arXiv preprint arXiv:1810.06793*, 2018.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021.
- Surbhi Goel, Adam Klivans, and Raghu Meka. Learning one convolutional layer with overlapping patches. In *International Conference on Machine Learning*, pages 1783–1791. PMLR, 2018.
- Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 795–811. Springer, 2016.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.

- Yuanzhi Li, Tengyu Ma, and Hongyang R Zhang. Learning over-parametrized two-layer relu neural networks beyond ntk. *arXiv preprint arXiv:2007.04596*, 2020.
- Stanislaw Lojasiewicz. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.
- Yiping Lu, Chao Ma, Yulong Lu, Jianfeng Lu, and Lexing Ying. A mean field analysis of deep resnet and beyond: Towards provably optimization via overparameterization from depth. In *International Conference on Machine Learning*, pages 6426–6436. PMLR, 2020.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR, 2019.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multi-layer neural networks. *arXiv preprint arXiv:2001.11443*, 2020.
- Atsushi Nitanda and Taiji Suzuki. Stochastic particle gradient descent for infinite ensembles. *arXiv preprint arXiv:1712.05438*, 2017.
- Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- Samet Oymak and Mahdi Soltanolkotabi. Towards moderate overparameterization: global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 2020.
- Grant M Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of neural networks: An interacting particle system approach. *arXiv preprint arXiv:1805.00915*, 2018.
- Itay Safran, Gilad Yehudai, and Ohad Shamir. The effects of mild over-parameterization on the optimization landscape of shallow relu neural networks. *arXiv preprint arXiv:2006.01005*, 2020.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.
- Yuangdong Tian. Student specialization in deep relu networks with finite width and input dimension. *arXiv*, page arXiv preprint arXiv:1909.13458, 2019.
- Colin Wei, Jason D Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. In *Advances in Neural Information Processing Systems*, pages 9712–9724, 2019.
- 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.

Kai Zhong, Zhao Song, Prateek Jain, Peter L Bartlett, and Inderjit S Dhillon. Recovery guarantees for one-hidden-layer neural networks. *arXiv preprint arXiv:1706.03175*, 2017.

Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep relu networks. *Machine Learning*, 109(3):467–492, 2020.

Appendix A. Discussions about ReLU Network and Absolute Network

A.1. Two-Layer ReLU network and Two-Layer Absolute network

We discuss two ways that reduce learning a two-layer teacher network with ReLU activation to a two-layer teacher network with absolute value activation with infinite number of data. Consider the following problem for learning two-layer ReLU net in the teacher-student setting with Gaussian input (or symmetric input⁴ as long as $\mathbb{E}_x[xx^\top]$ is full rank)

$$\min_{\{w_i\}_{i=1}^m} \mathbb{E}_x \left[\frac{1}{2} \left(\sum_{i=1}^m \|w_i\| \text{ReLU}(w_i^\top x) - \sum_{i=1}^r \text{ReLU}(w_i^{*\top} x) \right)^2 \right]. \quad (5)$$

The first way is to first fit the optimal linear function to the data, that is

$$\beta^* = \arg \min_{\beta} \mathbb{E}_x \left[\frac{1}{2} \left(\beta^\top x - \sum_{i=1}^r \text{ReLU}(w_i^{*\top} x) \right)^2 \right].$$

Since $\mathbb{E}_x[x \text{ReLU}(w^\top x)] = \frac{1}{2} \mathbb{E}_x[xx^\top] w$ (Goel et al., 2018; Ge et al., 2018), we know

$$\sum_{i=1}^r \text{ReLU}(w_i^{*\top} x) - \beta^{*\top} x = \frac{1}{2} \sum_{i=1}^r |w_i^{*\top} x|.$$

This implies once we first find the optimal linear function $\beta^{*\top} x$, we can then reduced the problem to learn a teacher network with absolute value activation.

The second way follows similar idea. We claim that the problem (5) is equivalent to problem (1) in the way described below.

$$\min_{\beta, \{w_i\}_{i=1}^m} \mathbb{E}_x \left[\frac{1}{2} \left(\beta^\top x + \sum_{i=1}^m \|w_i\| \text{ReLU}(w_i^\top x) - \sum_{i=1}^r \text{ReLU}(w_i^{*\top} x) \right)^2 \right], \quad (6)$$

where we have an additional linear term in the two-layer ReLU student net.

To see the equivalence, let us first optimize β which is essentially a least square problem, we will obtain $\beta = \frac{1}{2} \sum_{i=1}^r w_i^* - \frac{1}{2} \sum_{i=1}^m \|w_i\| w_i$. Then, plugging in the expression of β into (6), we have

$$\min_{\{w_i\}_{i=1}^m} \mathbb{E}_x \left[\frac{1}{2} \left(\frac{1}{2} \sum_{i=1}^m \|w_i\| |w_i^\top x| - \frac{1}{2} \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right],$$

which is exactly the same as the problem (1) where we use absolute value function as activation. Therefore, our results can also be extended to the ReLU activation case with some modifications as discussed in the above way.

4. We call x follows a symmetric distribution \mathcal{D} if for any x , the probability of observing $x \sim \mathcal{D}$ is the same as the probability of observing $-x \sim \mathcal{D}$.

A.2. Proof of Claim 1

Claim 1 For problem (1) with ReLU activation, when loss is zero, there may exist a student neuron whose direction does not match direction of any teacher neuron.

Proof We give one example to show Claim 1 is true. Consider a teacher network that is parameterized as $f^*(x) = \sum_{i=1}^r \text{ReLU}(w_i^{*\top} x)$ with $\sum_{i=1}^r w_i^* = 0$. For the student network $f(x) = \sum_{i=1}^m \|w_i\| \text{ReLU}(w_i^\top x)$, let $m = r$ and $\|w_i\| w_i = -w_i^*$ for $i \in [r]$. Using the fact that $\text{ReLU}(x) - \text{ReLU}(-x) = x$, we have

$$f(x) - f^*(x) = \sum_{i=1}^r \text{ReLU}(\|w_i\| w_i^\top x) - \sum_{i=1}^r \text{ReLU}(w_i^{*\top} x) = - \sum_{i=1}^r w_i^{*\top} x = 0.$$

This indicates that when loss is 0, it is possible that student neuron does not match the direction of any teacher neuron. \blacksquare

A.3. Proof of Claim 2

Claim 2 Suppose teacher neuron w^* and two student neurons w_1, w_2 belong to the same hyperplane with position shown in Figure 2. If $\|w^*\| = 1$, $\angle(w_1, w^*) = \angle(w_2, w^*) = \delta$, $\|w_1\| = \|w_2\| = 1/\sqrt{2} \cos \delta$ (so that $\|w_1\| w_1 + \|w_2\| w_2 = w^*$) for a small enough constant δ , then we have loss $L(W) = \Theta(\delta^3)$.

Proof Since $\|w_1\| w_1 + \|w_2\| w_2 = w^*$, we know $\|w_1\| |w_1^\top x| + \|w_2\| |w_2^\top x| = |w^{*\top} x|$ when $\text{sgn}(w_1^\top x) = \text{sgn}(w_2^\top x)$. Thus,

$$\begin{aligned} L(W) &= \mathbb{E}_x \left[\left(\|w_1\| |w_1^\top x| + \|w_2\| |w_2^\top x| - |w^{*\top} x| \right)^2 \mathbb{I}_{\text{sgn}(w_1^\top x) \neq \text{sgn}(w_2^\top x)} \right] \\ &= \mathbb{E}_{\tilde{x}} \left[\left(\|w_1\| |w_1^\top \tilde{x}| + \|w_2\| |w_2^\top \tilde{x}| - |w^{*\top} \tilde{x}| \right)^2 \mathbb{I}_{\text{sgn}(w_1^\top \tilde{x}) \neq \text{sgn}(w_2^\top \tilde{x})} \right], \end{aligned}$$

where \tilde{x} is a 3-dimensional Gaussian since the expectation only depends on three vectors w_1, w_2 and w^* . Note that when $\text{sgn}(w_1^\top \tilde{x}) \neq \text{sgn}(w_2^\top \tilde{x})$, $\|w_1\| |w_1^\top \tilde{x}| + \|w_2\| |w_2^\top \tilde{x}| - |w^{*\top} \tilde{x}| = O(\delta \|\tilde{x}\|)$. Therefore, we know $L(W) = O(\delta^3)$. \blacksquare

Appendix B. Proof of Gradient Lower Bound (Theorem 4)

Recall that in Lemma 8 we construct a descent direction and show that it has a positive correlation with the gradient. Then using Lemma 8 with $q_{ij} = \frac{\|w_{ij}\|}{\sum_{j \in T_i(\delta_{max})} \|w_{ij}\|^2}$ for $j \in T_i(\delta_{max})$ and $q_{ij} = 0$ otherwise, we can show a lower bound on gradient norm.

Theorem 4 (Gradient Lower Bound) For network and loss function $L(\cdot)$ defined in Definition 2, under Assumptions 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W such that loss $L(W) \leq \epsilon_0$, we have

$$\|\nabla_W L(W)\|_F \geq \kappa L(W),$$

where $\kappa = \Theta(r^{-1/2} w_{max}^{-1/2})$.

Proof By Lemma 8, we know

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq L(W),$$

where $\{q_{ij}\}_{i \in [r], j \in [m]}$ is any sequence that satisfies (1) $q_{ij} \geq 0$ for all (i, j) , (2) $q_{ij} = 0$ if $j \notin T_i(\delta_{max})$; and (3) $\sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$, where $\delta_{max} = \Theta(r w_{max} w_{min}^{-5/3} \cdot \epsilon^{1/3})$. Hence,

$$\begin{aligned} & \left(\sum_{i=1}^r \sum_{j \in T_i} \|\nabla_{w_j} L(W)\|^2 \right)^{1/2} \left(\sum_{i=1}^r \sum_{j \in T_i} \left\| (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \right\|^2 \right)^{1/2} \\ & \geq \sum_{i=1}^r \sum_{j \in T_i} \|\nabla_{w_j} L(W)\| \left\| (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \right\| \geq L(W). \end{aligned} \quad (7)$$

To lower bound the graident norm, we have to give a upper bound for the following term.

$$\begin{aligned} \sum_{i=1}^r \sum_{j \in T_i} \left\| (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \right\|^2 & \leq \sum_{i=1}^r \sum_{j \in T_i} \left\| w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^* \right\|^2 \\ & \leq 2 \sum_{i=1}^r \sum_{j \in T_i} \left(\|w_j\|^2 + \|q_{ij} w_i^*\|^2 \right) \\ & = 2 \sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 + 2 \sum_{i=1}^r \sum_{j \in T_i} q_{ij}^2 \|w_i^*\|^2. \end{aligned}$$

Let $q_{ij} = \frac{\|w_j\|}{\sum_{j \in T_i(\delta_{max})} \|w_j\|}$ for $j \in T_i(\delta_{max})$ and $q_{ij} = 0$ otherwise. We have

$$\sum_{i=1}^r \sum_{j \in T_i} q_{ij}^2 \|w_i^*\|^2 = \sum_{i=1}^r \sum_{j \in T_i(\delta_{max})} \frac{\|w_j\|^2 \|w_i^*\|^2}{\left(\sum_{j \in T_i(\delta_{max})} \|w_j\|^2 \right)^2} = \sum_{i=1}^r \frac{\|w_i^*\|^2}{\sum_{j \in T_i(\delta_{max})} \|w_j\|^2}.$$

Using Lemma 9 with $\delta_{max} = \Theta(r w_{max} w_{min}^{-5/3} \cdot \epsilon^{1/3})$, we have $\sum_{j \in T_i(\delta_{max})} \|w_j\|^2 \geq \frac{1}{2} \|w_i^*\|^2$.

Thus, with Lemma C.5, we have

$$\begin{aligned} \sum_{i=1}^r \sum_{j \in T_i} \left\| (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \right\|^2 & \leq 2 \sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 + 2 \sum_{i=1}^r \sum_{j \in T_i} q_{ij}^2 \|w_i^*\|^2 \\ & \leq O(r w_{max}) + 2 \sum_{i=1}^r 2 \|w_i^*\|^2 = O(r w_{max}). \end{aligned}$$

Together with (7), we have

$$\|\nabla_W L(W)\|_F = \left(\sum_{i=1}^r \sum_{j \in T_i} \|\nabla_{w_j} L(W)\|^2 \right)^{1/2} \geq \kappa L(W),$$

where $\kappa = \Theta\left(r^{-1/2} w_{max}^{-1/2}\right)$. ■

Appendix C. Proof of Descent Direction Lemma (Lemma 8)

We first give the proof of Lemma 8, then give the proof of each step (Lemma 9, Lemma 10, Lemma C.2) in the proof sketch accordingly in the following subsections.

Before presenting the proof of descent direction lemma, we first show that we can arbitrarily change the sign of student neurons and still have the direction of improvement. As discussed at the beginning of Section 5.2, this allows us to assume $w_j^\top w_i^* \geq 0$ for any $j \in T_i$. In the rest of this section, we assume the above always hold.

Lemma C.1 *Suppose there exists $\{g(w_j)\}_{j \in [m]}$ such that $W = (w_1, \dots, w_m)$ satisfies*

$$\sum_{j \in [m]} \langle \nabla_{w_j} L(W), g(w_j) \rangle \geq L(W).$$

Then, for any $a_i \in \{\pm 1\}$ if we replace W with $W' = (a_1 w_1, \dots, a_m w_m)$, we have

$$\sum_{j \in [m]} \langle \nabla_{w'_j} L(W'), a_j g(w_j) \rangle \geq L(W').$$

Proof To specify the dependency on W , let $R_W(x) \triangleq \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$. Denote $w'_i = a_i w_i$. Using $R_W = R_{W'}$, we have

$$\begin{aligned} \langle \nabla_{w_j} L(W), g(w_j) \rangle &= \mathbb{E}_x \left[R_W(x) \|w_j\| g(w_j)^\top (I + \bar{w}_j \bar{w}_j) x \operatorname{sgn}(w_j^\top x) \right] \\ &= \mathbb{E}_x \left[R_{W'}(x) \|w'_j\| g(w_j)^\top (I + \bar{w}'_j \bar{w}'_j) x \operatorname{sgn}(w_j^\top x a_j) \right] \\ &= \mathbb{E}_x \left[R_{W'}(x) \|w'_j\| g(w_j)^\top (I + \bar{w}'_j \bar{w}'_j) a_j x \operatorname{sgn}(w_j^\top x) \right] \\ &= \langle \nabla_{w'_j} L(W'), a_j g(w_j) \rangle \end{aligned}$$

Since $L(W) = L(W')$, we know the result holds. ■

C.1. Proof of Descent Direction Lemma (Lemma 8)

We need the follow lemma to prove Lemma 8. In fact, this lemma corresponds to the third step of proof sketch as described in Section 5.2. It shows that as long as every teacher neuron has close-by student neuron and the difference between average neuron \hat{w}_i and w_i^* is small, the direction we construct is indeed a descent direction. The proof is provided in Section C.5.

Lemma C.2 *Under Assumption 2, if $\|v_i\| \leq \alpha$ and $|T_i(\delta_{max})| \geq 1$ for all $i \in [r]$, then*

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq \|R\|^2 - O(r^2 w_{max} \alpha \delta_{max}^2),$$

where $\{q_{ij}\}_{i \in [r], j \in [m]}$ is any sequence that satisfies (1) $q_{ij} \geq 0$ for all (i, j) , (2) $q_{ij} = 0$ if $j \notin T_i(\delta_{max})$; and (3) $\sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$.

Now we are ready to prove Lemma 8 by following the three-step proof sketch in Section 5.2.

Lemma 8 (Descent Direction) For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \text{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq L(W),$$

where $\{q_{ij}\}_{i \in [r], j \in [m]}$ is any sequence that satisfies (1) $q_{ij} \geq 0$ for all (i, j) , (2) $q_{ij} = 0$ if $j \notin T_i(\delta_{max})$; and (3) $\sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$, where $\delta_{max} = \Theta(r w_{max} w_{min}^{-5/3} \cdot \epsilon^{1/3})$.

Proof Using Lemma 9, we know there exists such $\delta_{max} = \Theta(r w_{max} w_{min}^{-5/3} \cdot \epsilon^{1/3})$ satisfies $|T_i(\delta_{max})| \geq$

1. From Lemma 10, we know $\|v_i\| \leq \alpha = O(r^{11/4} w_{max}^{1/4} \Delta^{-3/2} \cdot \epsilon^{3/8})$ for all $i \in [r]$.

Then, by Lemma C.2 we have

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \text{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq \|R\|^2 - O(r^2 w_{max} \alpha \delta_{max}^2).$$

By the choice of δ_{max} , we have $O(r^2 w_{max} \alpha \delta_{max}^2) = O(r^{27/4} w_{max}^{13/4} \Delta^{-3/2} w_{min}^{-10/3} \cdot \epsilon^{25/24})$. Then with the fact that $\|R\|^2 = 2L(W) = 2\epsilon$, when $\epsilon \leq \epsilon_0 = O(\Delta^{36} w_{min}^{80} r^{-162} w_{max}^{-78})$ we have

$$\begin{aligned} \sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \text{sgn}(w_j^\top w_i^*) w_i^*) \rangle &\geq 2\epsilon - O\left(\frac{r^{27/4} w_{max}^{13/4}}{\Delta^{3/2} w_{min}^{10/3}} \epsilon^{25/24}\right) \\ &\geq \epsilon. \end{aligned}$$

■

C.2. Proof of Lemma 9

We give the proof of Lemma 9, which is the first step of the proving Lemma 8 as described in Section 5.2. It shows every teacher neuron has at least one close-by student neuron when loss is small. The proof follows the idea of test function as described in Section 5.3.

Lemma 9 Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ and an absolute constant C such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, if we choose $\delta_{max} \geq C \cdot r w_{max} w_{min}^{-5/3} \cdot \epsilon^{1/3}$, then we have $|T_i(\delta_{max})| \geq 1$ and $\sum_{j \in T_i(\delta_{max})} \|w_j\|^2 \geq \frac{1}{2} \|w_i^*\|^2$ for all $i \in [r]$.

Proof The following proof has two parts: In Part 1, we aim to show $|T_i(\delta_{max})| \geq 1$; In Part 2, we aim to show $\sum_{j \in T_i(\delta_{max})} \|w_j\|^2 \geq \frac{1}{2} \|w_i^*\|^2$. Denote $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ as the error function. We will use δ instead of δ_{max} to simplify the notation in the proof. WLOG, assume $\delta \leq \Delta$ as we can choose small enough $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$.

Part 1: We aim to show every teacher neuron has at least one close student neuron. Assume toward contradiction that there is a teacher neuron w_i^* such that for all student neurons w_j , we have $\angle(w_i^*, w_j) \geq \delta$. Denote set $S = \{x \in \mathbb{R}^d \mid |w_i^{*\top} x| \leq \tau\}$. To simplify notation, we will use w^* to represent w_i^* . Let w be a normalized vector satisfying $\phi \triangleq \angle(w^*, w) \geq \delta$. In following part (i)-(iii), WLOG, assume $\|w\| = \|w^*\| = 1$, $w^* = (1, 0, \dots, 0)^\top$ and $w = (\cos \phi, \sin \phi, 0, \dots, 0)^\top$. We are going to focus on $|w^\top x|$, so we can further assume $\phi \in [\delta, \frac{\pi}{2}]$. Also, we will assume $\tau, \frac{\tau}{\phi} \leq c$ for a sufficiently small constant c .

(i) Estimate $\langle |w^{*\top} x|, |w^\top x| \rangle_S$. We have

$$\langle |w^{*\top} x|, |w^\top x| \rangle_S = \frac{1}{2\pi} \int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2} \int_{-\infty}^{\infty} |x_1 \cos \phi + x_2 \sin \phi| e^{-x_2^2/2} dx_2 dx_1.$$

Note that since $\phi \in (0, \frac{\pi}{2}]$, we know $x_1 \cos \phi + x_2 \sin \phi \geq 0$ is equivalent to $x_2 \geq -x_1 \cot \phi$. Thus,

$$\begin{aligned} & \int_{-\infty}^{\infty} |x_1 \cos \phi + x_2 \sin \phi| e^{-x_2^2/2} dx_2 \\ &= \int_{-x_1 \cot \phi}^{\infty} (x_1 \cos \phi + x_2 \sin \phi) e^{-x_2^2/2} dx_2 - \int_{-\infty}^{-x_1 \cot \phi} (x_1 \cos \phi + x_2 \sin \phi) e^{-x_2^2/2} dx_2 \\ &= x_1 \cos \phi \int_{-x_1 \cot \phi}^{x_1 \cot \phi} e^{-x_2^2/2} dx_2 + 2 \sin \phi \int_{|x_1 \cot \phi|}^{\infty} x_2 e^{-x_2^2/2} dx_2 \\ &= \sqrt{2\pi} x_1 \cos \phi \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) + 2 \sin \phi e^{-x_1^2 \cot^2 \phi / 2}. \end{aligned} \tag{8}$$

This leads to

$$\begin{aligned} \langle |w^{*\top} x|, |w^\top x| \rangle_S &= \frac{1}{2\pi} \int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2} \left(\sqrt{2\pi} x_1 \cos \phi \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) + 2 \sin \phi e^{-x_1^2 \cot^2 \phi / 2} \right) dx_1 \\ &= \frac{\cos \phi}{\sqrt{2\pi}} \int_{-\tau}^{\tau} |x_1| x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 + \frac{\sin \phi}{\pi} \int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2 \sin^2 \phi} dx_1, \end{aligned} \tag{9}$$

For the first term in (9), we have

$$\int_{-\tau}^{\tau} |x_1| x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 = 2 \int_0^{\tau} x_1^2 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1.$$

To get the upper bound, using the fact that $x_1 \cot \phi \geq 0$ and $\operatorname{erf}(x) \leq \frac{2x}{\sqrt{\pi}}$ for $x \geq 0$, we have

$$\begin{aligned} \int_{-\tau}^{\tau} |x_1| x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 &\leq 2 \sqrt{\frac{2}{\pi}} \cot \phi \int_0^{\tau} x_1^3 e^{-x_1^2/2} dx_1 \\ &= 2 \sqrt{\frac{2}{\pi}} \cot \phi (2 - (2 + \tau^2) e^{-\tau^2/2}). \end{aligned}$$

To get the lower bound, using the fact that $x_1 \cot \phi \leq \frac{\tau}{\phi} \leq c$ and $\operatorname{erf}(x) \geq \frac{x}{\sqrt{\pi}}$ for $x \in [0, 1]$, we have

$$\begin{aligned} \int_{-\tau}^{\tau} |x_1| x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 &\geq \sqrt{\frac{2}{\pi}} \cot \phi \int_0^{\tau} x_1^3 e^{-x_1^2/2} dx_1 \\ &= \sqrt{\frac{2}{\pi}} \cot \phi (2 - (2 + \tau^2) e^{-\tau^2/2}). \end{aligned}$$

For the second term in (9), we have

$$\int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2 \sin^2 \phi} dx_1 = 2 \int_0^{\tau} x_1 e^{-x_1^2/2 \sin^2 \phi} dx_1 = 2 \sin^2 \phi (1 - e^{-\tau^2/2 \sin^2 \phi}).$$

Therefore, combining the above two term, we have the upper bound

$$\begin{aligned} \langle |w^{*\top} x|, |w^\top x| \rangle_S &\leq \frac{\cos \phi}{\sqrt{2\pi}} 2 \sqrt{\frac{2}{\pi}} \cot \phi (2 - (2 + \tau^2) e^{-\tau^2/2}) + \frac{2}{\pi} \sin^3 \phi (1 - e^{-\tau^2/2 \sin^2 \phi}) \\ &\leq \frac{2 \cos^2 \phi}{\pi \sin \phi} \left(2 - (2 + \tau^2) \left(1 - \frac{\tau^2}{2}\right) \right) + \frac{2}{\pi} \sin^3 \phi \left(\frac{\tau^2}{2 \sin^2 \phi} - \frac{\tau^4}{16 \sin^4 \phi} \right) \\ &= \frac{\cos^2 \phi}{\pi \sin \phi} \tau^4 + \frac{\sin \phi}{\pi} \tau^2 - \frac{\tau^4}{8\pi \sin \phi}, \end{aligned}$$

where the second inequality is because $\tau \leq c$, $\frac{\tau}{\sin \phi} \leq \frac{2\tau}{\phi} \leq c$ and $1 - e^{-x^2} \leq x^2 - \frac{x^4}{4} \leq x^2$ for $x \in [0, 1]$.

Also, we have the lower bound

$$\begin{aligned} \langle |w^{*\top} x|, |w^\top x| \rangle_S &\geq \frac{\cos \phi}{\sqrt{2\pi}} \sqrt{\frac{2}{\pi}} \cot \phi (2 - (2 + \tau^2) e^{-\tau^2/2}) + \frac{2}{\pi} \sin^3 \phi (1 - e^{-\tau^2/2 \sin^2 \phi}) \\ &\geq \frac{\cos^2 \phi}{\pi \sin \phi} \left(2 - (2 + \tau^2) \left(1 - \frac{\tau^2}{2} + \frac{\tau^4}{8}\right) \right) + \frac{2}{\pi} \sin^3 \phi \left(\frac{\tau^2}{2 \sin^2 \phi} - \frac{\tau^4}{8 \sin^4 \phi} \right) \\ &= \frac{\cos^2 \phi}{\pi \sin \phi} \left(\frac{\tau^4}{4} - \frac{\tau^6}{8} \right) + \frac{\sin \phi}{\pi} \tau^2 - \frac{\tau^4}{4\pi \sin \phi}, \end{aligned}$$

where the second inequality is because $\tau, \frac{\tau}{\sin \phi} \geq 0$ and $1 - e^{-x^2} \geq x^2 - \frac{x^4}{2}$ for $x \geq 0$.

Thus, we have the estimation

$$\langle |w^{*\top} x|, |w^\top x| \rangle_S = \frac{\sin \phi}{\pi} \tau^2 \pm \Theta \left(\frac{\tau^4}{\phi} \right).$$

(ii) Estimate $\langle 1, |w^\top x| \rangle_S$. We have

$$\begin{aligned} \langle 1, |w^\top x| \rangle_S &= \frac{1}{2\pi} \int_{-\tau}^{\tau} e^{-x_1^2/2} \int_{-\infty}^{\infty} |x_1 \cos \phi + x_2 \sin \phi| e^{-x_2^2/2} dx_2 dx_1 \\ &= \frac{1}{2\pi} \int_{-\tau}^{\tau} e^{-x_1^2/2} \left(\sqrt{2\pi} x_1 \cos \phi \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) + 2 \sin \phi e^{-x_1^2 \cot^2 \phi / 2} \right) dx_1 \quad (10) \\ &= \frac{\cos \phi}{\sqrt{2\pi}} \int_{-\tau}^{\tau} x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 + \frac{\sin \phi}{\pi} \int_{-\tau}^{\tau} e^{-x_1^2/2 \sin^2 \phi} dx_1, \end{aligned}$$

where we use (8) in the second line.

For the first term in (10), we have

$$\int_{-\tau}^{\tau} x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1 = 2 \int_0^{\tau} x_1 e^{-x_1^2/2} \operatorname{erf} \left(\frac{x_1 \cot \phi}{\sqrt{2}} \right) dx_1.$$

To get the upper bound, using the fact that $x_1 \cot \phi \geq 0$ and $\operatorname{erf}(x) \leq \frac{2x}{\sqrt{\pi}}$ for $x \geq 0$, we have

$$\begin{aligned} \int_{-\tau}^{\tau} x_1 e^{-x_1^2/2} \operatorname{erf}\left(\frac{x_1 \cot \phi}{\sqrt{2}}\right) dx_1 &\leq 2\sqrt{\frac{2}{\pi}} \cot \phi \int_0^{\tau} x_1^2 e^{-x_1^2/2} dx_1 \\ &= 2\sqrt{\frac{2}{\pi}} \cot \phi \left(-\tau e^{-\tau^2/2} + \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)\right). \end{aligned}$$

To get the lower bound, using the fact that $x_1 \cot \phi \leq \frac{\tau}{\phi} \leq c$ and $\operatorname{erf}(x) \geq \frac{x}{\sqrt{\pi}}$ for $x \in [0, 1]$, we have

$$\begin{aligned} \int_{-\tau}^{\tau} x_1 e^{-x_1^2/2} \operatorname{erf}\left(\frac{x_1 \cot \phi}{\sqrt{2}}\right) dx_1 &\geq \sqrt{\frac{2}{\pi}} \cot \phi \int_0^{\tau} x_1^2 e^{-x_1^2/2} dx_1 \\ &= \sqrt{\frac{2}{\pi}} \cot \phi \left(-\tau e^{-\tau^2/2} + \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)\right). \end{aligned}$$

For the second term in (10), we have

$$\int_{-\tau}^{\tau} e^{-x_1^2/2 \sin^2 \phi} dx_1 = 2 \int_0^{\tau} e^{-x_1^2/2 \sin^2 \phi} dx_1 = \sqrt{2\pi} \sin \phi \operatorname{erf}\left(\frac{\tau}{\sqrt{2} \sin \phi}\right).$$

Therefore, combining the two terms above, we have the upper bound

$$\begin{aligned} \langle 1, |w^\top x| \rangle_S &\leq \frac{\cos \phi}{\sqrt{2\pi}} 2\sqrt{\frac{2}{\pi}} \cot \phi \left(-\tau e^{-\tau^2/2} + \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)\right) + \frac{\sin \phi}{\pi} \sqrt{2\pi} \sin \phi \operatorname{erf}\left(\frac{\tau}{\sqrt{2} \sin \phi}\right) \\ &\leq \frac{2 \cos^2 \phi}{\pi \sin \phi} \left(-\tau \left(1 - \frac{\tau^2}{2}\right) + \tau\right) + \sqrt{\frac{2}{\pi}} \sin^2 \phi \sqrt{\frac{2}{\pi}} \frac{\tau}{\sin \phi} \\ &= \frac{\cos^2 \phi}{\pi \sin \phi} \tau^3 + \frac{2 \sin \phi}{\pi} \tau, \end{aligned}$$

where the second inequality is because $\tau \leq c$, $\frac{\tau}{\sin \phi} \leq \frac{2\tau}{\phi} \leq c$, $1 - e^{-x^2} \leq x^2 - \frac{x^4}{4} \leq x^2$ for $x \in [0, 1]$ and $\operatorname{erf}(x) \leq \frac{2x}{\sqrt{\pi}}$ for $x \geq 0$.

Also, we have the lower bound,

$$\begin{aligned} \langle 1, |w^\top x| \rangle_S &\geq \frac{\cos \phi}{\sqrt{2\pi}} \sqrt{\frac{2}{\pi}} \cot \phi \left(-\tau e^{-\tau^2/2} + \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)\right) + \frac{\sin \phi}{\pi} \sqrt{2\pi} \sin \phi \operatorname{erf}\left(\frac{\tau}{\sqrt{2} \sin \phi}\right) \\ &\geq \frac{\cos^2 \phi}{\pi \sin \phi} \left(-\tau \left(1 - \frac{\tau^2}{2} + \frac{\tau^4}{8}\right) + \tau - \frac{\tau^3}{6}\right) + \sqrt{\frac{2}{\pi}} \sin^2 \phi \left(\sqrt{\frac{2}{\pi}} \frac{\tau}{\sin \phi} - \frac{2}{3\sqrt{\pi}} \frac{\tau^3}{2\sqrt{2} \sin^3 \phi}\right) \\ &= \frac{\cos^2 \phi}{\pi \sin \phi} \left(\frac{\tau^3}{3} - \frac{\tau^5}{8}\right) + \frac{2 \sin \phi}{\pi} \tau - \frac{\tau^3}{3\pi \sin \phi}, \end{aligned}$$

where the second inequality is because $1 - e^{-x^2} \geq x^2 - \frac{x^4}{2}$ for $x \geq 0$ and $\operatorname{erf}(x) \geq \frac{2x}{\sqrt{\pi}} - \frac{2}{3\sqrt{\pi}} x^3$ for $x \geq 0$.

Thus, we have the estimation

$$\langle 1, |w^\top x| \rangle_S = \frac{2 \sin \phi}{\pi} \tau \pm \Theta\left(\frac{\tau^3}{\phi}\right).$$

(iii) **Determine g .** Let $g = \frac{\langle |w^{*\top} x|, 1 \rangle_S}{\langle 1, 1 \rangle_S}$, which implies that

$$\langle |w^{*\top} x| - g, g \rangle_S = 0.$$

Note that

$$\begin{aligned} \langle |w^{*\top} x|, 1 \rangle_S &= \frac{1}{\sqrt{2\pi}} \int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2} dx_1 = \sqrt{\frac{2}{\pi}} (1 - e^{-\tau^2/2}), \\ \langle 1, 1 \rangle_S &= \frac{1}{\sqrt{2\pi}} \int_{-\tau}^{\tau} e^{-x_1^2/2} dx_1 = \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right). \end{aligned}$$

Thus,

$$g = \frac{\langle |w^{*\top} x|, 1 \rangle_S}{\langle 1, 1 \rangle_S} = \sqrt{\frac{2}{\pi}} \frac{1 - e^{-\tau^2/2}}{\operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)} = \frac{\tau}{2} + O(\tau^3).$$

(iv) **Complete the proof.** Note that

$$\begin{aligned} \langle |w^{*\top} x|, |w^{*\top} x| \rangle_S &= \frac{1}{\sqrt{2\pi}} \int_{-\tau}^{\tau} x_1^2 e^{-x_1^2/2} dx_1 = \sqrt{\frac{2}{\pi}} \left(-\tau e^{-\tau^2/2} + \sqrt{\frac{\pi}{2}} \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) \right), \\ \langle 1, |w^{*\top} x| \rangle_S &= \frac{1}{\sqrt{2\pi}} \int_{-\tau}^{\tau} |x_1| e^{-x_1^2/2} dx_1 = \sqrt{\frac{2}{\pi}} (1 - e^{-\tau^2/2}). \end{aligned}$$

Now, combining with the results in (i)(ii)(iii), we have

$$\begin{aligned} \langle |w^{*\top} x| - g, g \rangle_S &= 0, \\ \langle |w^{*\top} x| - g, |w^{*\top} x| \rangle_S &= \langle |w^{*\top} x|, |w^{*\top} x| \rangle_S - g \langle 1, |w^{*\top} x| \rangle_S = \Theta(\tau^3), \\ \left| \langle |w^{*\top} x| - g, |w^{*\top} x| \rangle_S \right| &= \left| \frac{\sin \phi}{\pi} \tau^2 \pm \Theta\left(\frac{\tau^4}{\phi}\right) - \left(\frac{\tau}{2} + O(\tau^3)\right) \left(\frac{2 \sin \phi}{\pi} \tau \pm \Theta\left(\frac{\tau^3}{\phi}\right)\right) \right| \leq \Theta\left(\frac{\tau^4}{\phi}\right). \end{aligned}$$

In the following, we will no longer assume $\|w^*\| = 1$, and will directly use $\|w^*\|$. Recall that $\angle(w_i, w^*) \geq \delta$ and $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$. We have

$$\begin{aligned} \langle |w^{*\top} x| - g, R(x) \rangle_S &\geq \Theta(\tau^3) \|w^*\|^2 - \sum_{i=1}^m \|w_i\|^2 \|w^*\| \Theta\left(\frac{\tau^4}{\delta}\right) - \sum_{j \neq i} \|w_j^*\| \|w^*\| \Theta\left(\frac{\tau^4}{\Delta}\right) \\ &= \left(\Theta(\tau^3) \|w^*\| - \Theta\left(\frac{r w_{\max} \tau^4}{\delta}\right) \right) \|w^*\|, \end{aligned}$$

where we use $\delta \leq \phi$, $\delta \leq \Delta$ and Lemma C.5.

Hence, when $\tau \leq \frac{c_1 w_{\min}}{r w_{\max}} \delta$ with a sufficiently small constant c_1 , we have $\langle |w^{*\top} x| - g, R(x) \rangle_S \geq \Theta(\tau^3) \|w^*\|^2$. Thus,

$$\|R\|_S^2 \geq \frac{\langle |w^{*\top} x| - g, R(x) \rangle_S^2}{\langle |w^{*\top} x| - g, |w^{*\top} x| - g \rangle_S} = \Theta(\tau^3) \|w^*\|^2 = \Theta\left(\frac{w_{\min}^3}{r^3 w_{\max}^3} \delta^3\right) \|w^*\|^2,$$

where we choose $\tau = \frac{c_1 w_{\min}}{r w_{\max}} \delta$. Since $\delta^3 \geq \frac{C_1 r^3 w_{\max}^3}{w_{\min}^5} \epsilon$ with a sufficiently large constant C_1 , we have

$$\epsilon > \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] = \|R\|^2 \geq \|R\|_S^2 \geq \epsilon,$$

which is a contradiction.

Part 2: We aim to show that the total mass of student neuron is large. Assume toward contradiction that there exists a teacher neuron w_i^* such that $\sum_{j \in \mathcal{I}_i(\delta)} \|w_j\| < \frac{1}{2} \|w_i^*\|$. We follow the same notations and assumptions as mentioned at the beginning of Part 1. We also assume $\phi \leq c$ for a sufficient small constant c .

We are going to show that when $\phi \leq \delta$, $h(\phi) \triangleq \langle |\bar{w}^{*\top} x| - g, |w^\top x| \rangle_S$ is non-increasing, i.e., $h'(\phi) \leq 0$. With (9)(10), we have

$$\begin{aligned} h'(\phi) &= \left(\left(2e^{-\tau^2/2}(e^{-\tau^2/2} - 1) + \sqrt{2\pi}e^{-\tau^2/2}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) \right) \frac{\operatorname{erf}\left(\frac{\tau \cot \phi}{\sqrt{2}}\right)}{\pi \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)} \right. \\ &\quad \left. + 4 \operatorname{OwenT}(\tau, \cot \phi) - 1 + \frac{2\phi}{\pi} \right) \sin \phi \\ &= \underbrace{\left(2e^{-\tau^2/2}(e^{-\tau^2/2} - 1) + \sqrt{2\pi}e^{-\tau^2/2}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) - \sqrt{\frac{\pi}{2}}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) \right)}_{h_1(\tau)} \frac{\operatorname{erf}\left(\frac{\tau \cot \phi}{\sqrt{2}}\right)}{\pi \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right)} \sin \phi \\ &\quad + \underbrace{\left(4 \operatorname{OwenT}(\tau, \cot \phi) - 1 + \frac{2\phi}{\pi} + \frac{\tau}{\sqrt{2\pi}} \operatorname{erf}\left(\frac{\tau \cot \phi}{\sqrt{2}}\right) \right)}_{h_2(\tau, \phi)} \sin \phi, \end{aligned}$$

where $\operatorname{OwenT}(x, a) = \frac{1}{2\pi} \int_0^a e^{-x^2(1+t^2)/2} \frac{1}{1+t^2} dt$.

For $h_1(\tau)$, we have

$$\begin{aligned} h_1(\tau) &= 2e^{-\tau^2/2}(e^{-\tau^2/2} - 1) + \sqrt{2\pi}e^{-\tau^2/2}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) - \sqrt{\frac{\pi}{2}}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) \\ &= 2e^{-\tau^2/2}(e^{-\tau^2/2} - 1) + \sqrt{\frac{\pi}{2}}e^{-\tau^2/2}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) + \sqrt{\frac{\pi}{2}}\tau \operatorname{erf}\left(\frac{\tau}{\sqrt{2}}\right) (e^{-\tau^2/2} - 1) \\ &\leq 2 \left(-\frac{\tau^2}{2} + \frac{\tau^4}{8} \right) + \sqrt{\frac{\pi}{2}}\tau \frac{2}{\sqrt{\pi}} \frac{\tau}{\sqrt{2}} + \sqrt{\frac{\pi}{2}}\tau \frac{2}{\sqrt{\pi}} \frac{\tau}{\sqrt{2}} \left(-\frac{\tau^2}{2} + \frac{\tau^4}{8} \right) \\ &= -\frac{\tau^4}{4} + \frac{\tau^6}{8} \leq 0, \end{aligned}$$

where in the first inequality we use $\operatorname{erf}(x) \leq \frac{2x}{\sqrt{\pi}}$, $e^{-x} - 1 \leq -x + \frac{x^2}{2}$ for $x \geq 0$ and $\tau \leq c$ in the last line.

For $h_2(\tau, \phi)$, by Lemma C.8 we have $h_2(\tau, \phi) \leq 0$ when $\tau, \phi \leq c$.

Therefore, we have $h'(\phi) \leq 0$ when $\tau, \phi \leq c$. This implies that when $\tau, \phi \leq c$,

$$h(\phi) = \langle |\bar{w}^{*\top} x| - g, |w^\top x| \rangle_S \leq \langle |\bar{w}^{*\top} x| - g, |\bar{w}^{*\top} x| \rangle_S = h(0).$$

In Part 1, we know that when $\tau \leq c$, $\frac{\tau}{\phi} \leq c$, $|\langle |\bar{w}^{*\top} x| - g, |w^\top x| \rangle_S| \leq \Theta\left(\frac{\tau^4}{\phi}\right)$. Thus, if $\tau \leq \min\{c, c^2\}$ and $\phi \geq c$, we have

$$\left| \langle |\bar{w}^{*\top} x| - g, |w^\top x| \rangle_S \right| \leq \Theta\left(\frac{\tau^4}{\phi}\right) \leq \langle |\bar{w}^{*\top} x| - g, |\bar{w}^{*\top} x| \rangle_S = \Theta(\tau^3).$$

In summary, we have when $\tau \leq \min\{c, c^2\}$,

$$\langle |\bar{w}^{*\top} x| - g, |w^\top x| \rangle_S \leq \langle |\bar{w}^{*\top} x| - g, |\bar{w}^{*\top} x| \rangle_S.$$

Recall that $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$, we have

$$\begin{aligned} \langle |w^{*\top} x| - g, R(x) \rangle_S &\geq \|w^*\| \left(\|w^*\| - \sum_{j \in T_i(\delta)} \|w_j\|^2 \right) \langle |\bar{w}^{*\top} x| - g, |\bar{w}^{*\top} x| \rangle_S \\ &\quad - \sum_{j \in [m] \setminus T_i(\delta)} \|w_j\|^2 \|w^*\| \Theta \left(\frac{\tau^4}{\delta} \right) - \sum_{j \neq i} \|w_j^*\| \|w^*\| \Theta \left(\frac{\tau^4}{\Delta} \right) \\ &= \left(\Theta(\tau^3) \|w^*\| - \Theta \left(\frac{r w_{\max} \tau^4}{\delta} \right) \right) \|w^*\|, \end{aligned}$$

where we use $\delta \leq \Delta$ and Lemma C.5.

Then, following the same argument at the end of Part 1, we get the contradiction

$$\epsilon > \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] = \|R\|^2 \geq \|R\|_S^2 \geq \epsilon,$$

which finishes the proof. ■

C.3. Lemma C.3: Property of The Residual

Before we give the proof of second step (Lemma 10) and third step (Lemma C.2) of the proof sketch, we first present the following result that characterize the property of the residual. This lemma will be useful in the later analysis. Recall that we have the residual decomposition that $R(x) = R_1(x) + R_2(x)$. The two terms have different properties – one can further show that R_1 is “flat” (whose value is uniformly bounded for all x) while R_2 is “spiky” (whose value is large only in a local region), see also Figure 4. More precisely, we have

Lemma C.3 (Property of The Residual) *Under Assumption 2, if $\|v_i\| \leq \alpha$ for all $i \in [r]$, then R_1, R_2 as defined in (3) satisfy $|R_1(x)| \leq \|x\| r\alpha$ and $R_2(x) \geq 0$, for all x .*

Proof It is easy to verify $R(x) = R_1(x) + R_2(x)$ for all x .

For $R_1(x)$, we have

$$|R_1(x)| = \left| \sum_{i=1}^r v_i^\top x \operatorname{sgn}(w_i^{*\top} x) \right| \leq \sum_{i=1}^r \left| v_i^\top x \operatorname{sgn}(w_i^{*\top} x) \right| \leq \sum_{i=1}^r \|v_i\| \|x\| \leq r\alpha \|x\|.$$

For $R_2(x)$, note that for any $x, i \in [r]$ and $j \in T_i$, we have

$$w_j^\top x (\operatorname{sgn}(w_j^\top x) - \operatorname{sgn}(w_i^{*\top} x)) \geq 0.$$

Therefore,

$$R_2(x) = \sum_{i=1}^r \sum_{j \in T_i} \|w_j\| w_j^\top x \left(\operatorname{sgn}(w_j^\top x) - \operatorname{sgn}(w_i^{*\top} x) \right) \geq 0. \quad \blacksquare$$

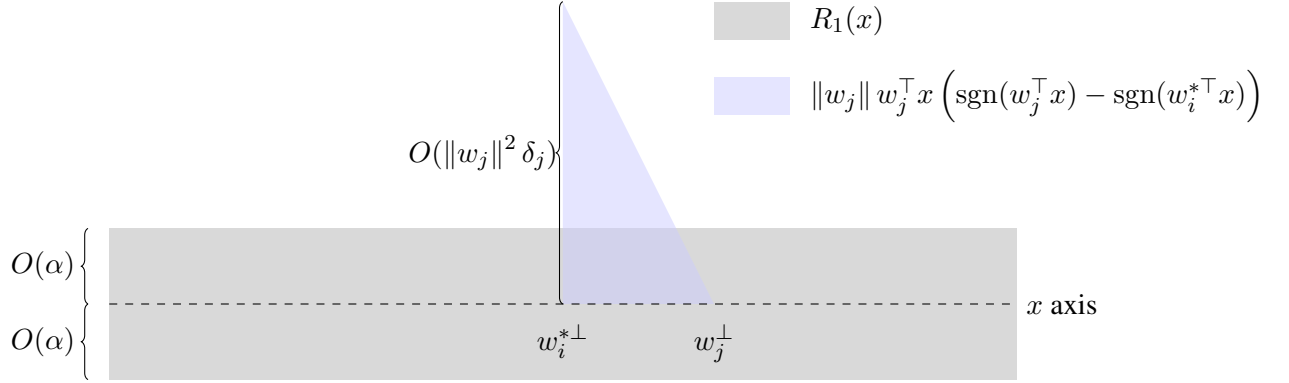


Figure 4: Illustration for the decomposition of residual $R(x)$ in Lemma C.3 (assume $\|x\| = 1$). Here w^\perp represents a vector that is orthogonal to w , and $\left\| \sum_{j \in T_i} \|w_j\| w_j - w_i^* \right\| \leq \alpha$ for all $i \in [r]$.

C.4. Proof of Lemma 10

In this subsection, we first following the proof sketch in Section 5.4 to give a proof of Lemma 10 (the second step of proving Lemma 8). Then, we give the proof of Lemma 11 and Lemma 12 respectively. In this subsection, we focus on the residual which is invariant with the sign change of student neuron. Recall the discussion at the beginning of Section 5.2, we can assume $w_j^\top w_i^* \geq 0$ for all $j \in T_i$ due to the symmetry of the absolute value function.

Lemma 10 (Average Student Is Close to Teacher) *Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{\max}^{-1}, w_{\min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\{v_i\}_{i \in [r]}$ (defined in (2)) satisfy $\|v_i\| \leq \text{poly}(r, w_{\max}, \Delta^{-1}) \epsilon^{3/8}$ for all $i \in [r]$.*

Proof

For R_1 , from Lemma 11 we have $\|R_1\|^2 = \Omega\left(\frac{\Delta^3}{r^3}\right) \|v\|^2$.

For R_2 , from Lemma 12 we have $\|R_2\|^2 = O(r^{5/2} w_{\max}^{1/2} \epsilon^{3/4})$.

With $\|R\|^2 = 2L(W) \leq 2\epsilon$, we have

$$\Omega\left(\frac{\Delta^{3/2}}{r^{3/2}}\right) \|v\| \leq \|R_1\| \leq \|R\| + \|R_2\| = O(\epsilon^{1/2} + r^{5/4} w_{\max}^{1/4} \epsilon^{3/8}),$$

which leads to $\left\| \sum_{j \in T_i} \|w_j\| w_j - w_i^* \right\|^2 = \|v_i\|^2 \leq \|v\|^2 = O\left(r^{11/2} w_{\max}^{1/2} \Delta^{-3} \cdot \epsilon^{3/4}\right)$ with our choice of ϵ . \blacksquare

C.4.1. PROOF OF LEMMA 11

Before proving Lemma 11 (the first step of proving Lemma 10), we need the following lemma that shows M is positive definite. The proof relies on the fact that teacher neurons are Δ -separated. Intuitively, this matrix M corresponds to the Hessian matrix at global minima in the exact-parameterization case. When teacher neurons are well-separated, one can imagine this matrix is positive definite.

Lemma C.4 Under Assumption 1, for matrix $M \in \mathbb{R}^{dr \times dr}$ defined in (4), we have $\lambda_{\min}(M) = \Omega(\Delta^3/r^3)$.

Now given the relation between $\|R_1\|$ and matrix M , we are ready to prove Lemma 11.

Lemma 11 Under Assumption 1, we have $\|R_1\|^2 \geq \Omega(\Delta^3/r^3) \|v\|^2$.

Proof Recall that $R_1(x) = \sum_{i=1}^r v_i^\top x \operatorname{sgn}(w_i^{*\top} x)$. We have

$$\|R_1\|^2 = \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^r v_i^\top x \operatorname{sgn}(w_i^{*\top} x) \right)^2 \right] = v^\top M v,$$

where $M \in \mathbb{R}^{dr \times dr}$ is a matrix with block entry $M_{ij} = \mathbb{E}_x [x x^\top \operatorname{sgn}(w_i^{*\top} x) \operatorname{sgn}(w_j^{*\top} x)]$ as defined in (4) and $v = (v_1, \dots, v_r) \in \mathbb{R}^{dr}$. Therefore, we have $\|R_1\|^2 \geq \Omega(\Delta^3/r^3) \|v\|^2$. \blacksquare

We now give the proof of Lemma C.4. The high-level idea of this proof is that when teacher neurons are well-separated, it is hard to cancel the nonlinearity induced by each teacher neuron.

Proof [Proof of Lemma C.4]

WLOG, assume $\|w_i^*\| = 1$. Denote vector $v \in \mathbb{R}^{dr}$ with $\|v\| = 1$ as $v = (v_1, \dots, v_r)^\top$, where $v_i \in \mathbb{R}^d$. It suffices to give lower bound on $v^\top M v$.

$$v^\top M v = \mathbb{E}_x \left[\left(\sum_{j=1}^r v_j^\top x \operatorname{sgn}(w_j^{*\top} x) \right)^2 \right] \geq \sum_{i=1}^r \mathbb{E}_x \left[\left(\sum_{j=1}^r v_j^\top x \operatorname{sgn}(w_j^{*\top} x) \right)^2 \mathbb{I}_{S_i} \right], \quad (11)$$

where set S_i will be determined later and satisfies $S_i \cap S_j = \emptyset$ for $i \neq j$.

Denote $S(w, \delta) = \{x \in \mathbb{R}^d \mid |w^\top x| \leq \delta\}$ as the nonlinear region of neuron w . We will use $S_i^*(\delta)$ to represent $S(w_i^*, \delta)$. Also denote

$$\begin{aligned} & A(\alpha, \beta, \delta_1, \delta_2) \\ &= \{x \mid |\alpha^\top x| \leq \delta_1, \text{ there exists } y \text{ such that } |\alpha^\top y| \leq \delta_1, |\beta^\top y| \leq \delta_2, (I - \alpha \alpha^\top)x = (I - \alpha \alpha^\top)y\} \end{aligned}$$

as the projection of $S(\alpha, \delta_1) \cap S(\beta, \delta_2)$ onto $S(\alpha, \delta_1)$ in the direction of α . We will use $A_i^*(\beta, \delta_1, \delta_2)$ to represent $A(w_i^*, \beta, \delta_1, \delta_2)$.

Let set

$$S_i = S_i^*(\delta_1) \setminus \left(\cup_{j \in [r], j \neq i} A_i^*(w_j^*, \delta_1, \delta_1) \right)$$

with $\delta_1 = O(\frac{\Delta}{r})$ and $\delta_1 \leq 1$. It is easy to see that $S_i \cap S_j = \emptyset$ for $i \neq j$, since $S_i^*(\delta_1) \cap S_j^*(\delta_1) \subseteq S_i^*(\delta_1) \cap A_i^*(w_j^*, \delta_1, \delta_1)$. By Lemma C.7, we know

$$\mathbb{E}_x [\mathbb{I}_{S_i^*(\delta_1)}] \geq \sqrt{\frac{2}{\pi}} e^{-\frac{1}{2}} \delta_1, \quad \mathbb{E}_x [\mathbb{I}_{A_i^*(w_j^*, \delta_1, \delta_1)}] \leq \frac{\delta_1^2 (1 + \cos \Delta)}{\pi \sin \Delta}.$$

Together with $\sin x \geq \frac{2x}{\pi}$ for $0 \leq x \leq \frac{\pi}{2}$, we have

$$\mathbb{E}_x [\mathbb{I}_{S_i}] \geq \Omega \left(\delta_1 - \frac{r \delta_1^2}{\Delta} \right) = c_1 \delta_1,$$

where c_1 is a constant.

In the following, we focus on one term in (11), and show it can be lower bounded. That is,

$$\mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{j=1}^r v_j^\top x \operatorname{sgn}(w_j^{*\top} x) \right)^2 \mathbb{I}_{S_i} \right].$$

For any $x^+ \in S_i^+ \triangleq S_i \cap S_i^*(\frac{\delta_1}{2}) \cap \{x | w_i^{*\top} x \geq 0\}$, let

$$x^- = x^+ - \frac{\delta_1}{2} w_i^*, \quad x^{++} = x^+ + \frac{\delta_1}{2} w_i^*, \quad x^{--} = x^+ - \delta_1 w_i^*.$$

We know that when $x^+ \in S_i$, $x^-, x^{++}, x^{--} \in S_i$. Further, we have

$$\begin{aligned} x^- &\in S_i^- \triangleq S_i \cap S_i^*(\frac{\delta_1}{2}) \cap \{x | w_i^{*\top} x \leq 0\}, \\ x^{++} &\in S_i^{++} \triangleq S_i \cap \left(S_i^*(\delta_1) \setminus S_i^*(\frac{\delta_1}{2}) \right) \cap \{x | w_i^{*\top} x \geq 0\}, \\ x^{--} &\in S_i^{--} \triangleq S_i \cap \left(S_i^*(\delta_1) \setminus S_i^*(\frac{\delta_1}{2}) \right) \cap \{x | w_i^{*\top} x \leq 0\}, \end{aligned}$$

and $S_i^+ \cup S_i^{++} \cup S_i^- \cup S_i^{--} = S_i$.

Consider $g(x) = \sum_{j=1}^r v_j^\top x \operatorname{sgn}(w_j^{*\top} x)$ on these four data points x^+, x^-, x^{++}, x^{--} . Note that when data has the form $x = x^+ + \gamma w_i^*$ for any fixed $x^+ \in S_i$, we have $g(x) = \alpha + \beta^\top x + v_i^\top x \operatorname{sgn}(w_i^{*\top} x)$, by the choice of S_i . Here α, β only depend on the choice of x^+ . Thus,

$$\begin{aligned} &0 \cdot g(x^{++}) - g(x^{--}) - g(x^+) + 2g(x^-) \\ &= 0 \cdot \left(\alpha + \beta^\top \left(x^+ + \frac{\delta_1}{2} w_i^* \right) + v_i^\top \left(x^+ + \frac{\delta_1}{2} w_i^* \right) \right) - \left(\alpha + \beta^\top \left(x^+ - \delta_1 w_i^* \right) - v_i^\top \left(x^+ - \delta_1 w_i^* \right) \right) \\ &\quad - \left(\alpha + \beta^\top x^+ + v_i^\top x^+ \right) + 2 \left(\alpha + \beta^\top \left(x^+ - \frac{\delta_1}{2} w_i^* \right) - v_i^\top \left(x^+ - \frac{\delta_1}{2} w_i^* \right) \right) \\ &= -2v_i^\top x^+, \end{aligned}$$

which indicates that $\max\{|g(x^{++})|, |g(x^{--})|, |g(x^+)|, |g(x^-)|\} \geq \frac{1}{2} |v_i^\top x^+|$.

Note that

$$\begin{aligned} \mathbb{E}_x [g^2(x) \mathbb{I}_{S_i^-}] &= \frac{1}{(2\pi)^{d/2}} \int_{S_i^-} g^2(x) e^{-\|x\|^2/2} dx \\ &= \frac{1}{(2\pi)^{d/2}} \int_{S_i^+} g^2 \left(x^+ - \frac{\delta_1}{2} w_i^* \right) e^{-\left\| x^+ - \frac{\delta_1}{2} w_i^* \right\|^2/2} dx^+ \\ &\geq \frac{1}{(2\pi)^{d/2}} \int_{S_i^+} g^2 \left(x^+ - \frac{\delta_1}{2} w_i^* \right) e^{-\left(\|x^+\|^2 + \frac{\delta_1^2}{4} \right)/2} dx^+ \\ &\geq e^{-1/8} \mathbb{E}_{x^+} [g^2(x^-) \mathbb{I}_{S_i^+}], \end{aligned}$$

where the second line is because we change x to $x^+ - \frac{\delta_1}{2}w_i^*$, the third line is because $w_i^{*\top}x^+ \geq 0$, and the last line is due to $\delta_1 \leq 1$ and $x^- = x^+ - \frac{\delta_1}{2}w_i^*$. Similarly, we have $\mathbb{E}_x[g^2(x)\mathbb{I}_{S_i^{--}}] \geq e^{-1/2}\mathbb{E}_{x^+}[g^2(x^-)\mathbb{I}_{S_i^+}]$ and $\mathbb{E}_x[g^2(x)\mathbb{I}_{S_i^{++}}] \geq e^{-3/8}\mathbb{E}_{x^+}[g^2(x^{++})\mathbb{I}_{S_i^+}]$.

Therefore,

$$\begin{aligned}\mathbb{E}_x[g^2(x)\mathbb{I}_{S_i}] &= \mathbb{E}_x[g^2(x)(\mathbb{I}_{S_i^+} + \mathbb{I}_{S_i^-} + \mathbb{I}_{S_i^{++}} + \mathbb{I}_{S_i^{--}})] \\ &\geq e^{-1/2}\mathbb{E}_{x^+}[(g^2(x^+) + g^2(x^-) + g^2(x^{++})g^2(x^{--}))\mathbb{I}_{S_i^+}] \\ &\geq \frac{1}{4\sqrt{e}}\mathbb{E}_{x^+}[|v_i^\top x^+|^2\mathbb{I}_{S_i^+}],\end{aligned}$$

where we use $\max\{|g(x^{++})|, |g(x^{--})|, |g(x^+)|, |g(x^-)|\} \geq \frac{1}{2}|v_i^\top x^+|$.

Recall that

$$S_i^+ = S_i \cap S_i^*(\frac{\delta_1}{2}) \cap \{x|w_i^{*\top}x \geq 0\} = \left(S_i^*(\frac{\delta_1}{2}) \cap \{x|w_i^{*\top}x \geq 0\}\right) \setminus (\cup_{j \in [r], j \neq i} A_i^*(w_j^*, \delta_1, \delta_1)).$$

By Lemma C.7, we know

$$\mathbb{E}_x[\mathbb{I}_{S_i^+}] \geq \frac{1}{2} \left(\mathbb{E}_x[\mathbb{I}_{S_i(\frac{\delta_1}{2})}] - \sum_{j \in [r], j \neq i} \mathbb{E}_x[\mathbb{I}_{A_i^*(w_j^*, \delta_1, \delta_1)}] \right) \geq \Omega\left(\delta_1 - \frac{r\delta_1^2}{\Delta}\right) = c_2\delta_1.$$

Let $S_i^{+'} = S(\bar{v}_i, \sqrt{\frac{\pi}{8}}c_2\delta_1)$. By Lemma C.7, we know $\mathbb{E}_x[\mathbb{I}_{S_i^{+'}}] \leq \frac{c_2}{2}\delta_1$. Hence, $\mathbb{E}[\mathbb{I}_{S_i^+ \setminus S_i^{+'}}] \geq \frac{c_2}{2}\delta_1$ and $|v_i^\top x^+| \geq \sqrt{\frac{\pi}{8}}c_2\delta_1 \|v_i\|$ for $x^+ \in S_i^+ \setminus S_i^{+'}$. Therefore,

$$\begin{aligned}\mathbb{E}_x[g^2(x)\mathbb{I}_{S_i}] &\geq \frac{1}{4\sqrt{e}}\mathbb{E}_{x^+}[|v_i^\top x^+|^2\mathbb{I}_{S_i^+ \setminus S_i^{+'}}] \geq \frac{1}{4\sqrt{e}}\frac{\pi}{8}c_2^2\delta_1^2 \|v_i\|^2 \mathbb{E}_{x^+}[\mathbb{I}_{S^+ \setminus S^{+'}}] \\ &\geq \frac{\pi}{64\sqrt{e}}c_2^3\delta_1^3 \|v_i\|^2 = \frac{c_3\Delta^3}{r^3} \|v_i\|^2.\end{aligned}$$

Therefore, by (11), we have

$$v^\top Mv \geq \sum_{i=1}^r \mathbb{E}_x[g^2(x)\mathbb{I}_{S_i}] \geq \sum_{i=1}^r \frac{c_3\Delta^3}{r^3} \|v_i\|^2 = \Omega\left(\frac{\Delta^3}{r^3}\right) \|v\|^2,$$

which implies that $\lambda_{\min}(M) = \Omega\left(\frac{\Delta^3}{r^3}\right)$. ■

C.4.2. PROOF OF LEMMA 12

We need the following two lemmas to prove Lemma 12 (the second step of proving Lemma 10). Lemma C.5 claims that when loss is small, the sum of norm of the student neurons would be bounded by a fixed quantity. Lemma C.6 claims student neurons should not be far away from teacher neuron when loss is small. See more discussions and corresponding proofs in the next two subsections.

Lemma C.5 Under Assumption 2, if loss $L(W) = O(r^2 w_{max}^2)$, then $\sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 = O(r w_{max})$.

Lemma C.6 Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 \delta_j^2 = O(\epsilon^{1/2})$.

Now we are ready to prove Lemma 12 by combining above two lemmas and Holder's inequality.

Lemma 12 Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\|R_2\|^2 = O(r^{5/2} w_{max}^{1/2} \epsilon^{3/4})$.

Proof Recall that $R_2(x) = \sum_{i=1}^r \sum_{j \in T_i} \|w_j\| w_j^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w_i^{*\top} x) \right)$. We have

$$\begin{aligned} \|R_2\|^2 &= \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| w_j^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w_i^{*\top} x) \right) \right)^2 \right] \\ &\leq r \sum_{i=1}^r \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{j \in T_i} \|w_j\| w_j^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w_i^{*\top} x) \right) \right)^2 \right]. \end{aligned}$$

In the following, we focus on one term in the above sum. Since we use absolute value function as activation, WLOG assume $w_j^\top w_i^* \geq 0$ for all $j \in T_i$ as discussed at the beginning of Section 5.2. We have

$$\begin{aligned} &\mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{j \in T_i} \|w_j\| w_j^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w_i^{*\top} x) \right) \right)^2 \right] \\ &= \sum_{j, k \in T_i} \mathbb{E}_{x \sim N(0, I)} \left[\|w_j\| w_j^\top x \cdot \|w_k\| w_k^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w_i^{*\top} x) \right) \left(\text{sgn}(w_k^\top x) - \text{sgn}(w_i^{*\top} x) \right) \right] \\ &\leq \sum_{j, k \in T_i} c_1 \|w_j\|^2 \|w_k\|^2 \delta_j \delta_k \min\{\delta_j, \delta_k\} \\ &\leq c_1 \left(\sum_{j \in T_i} \|w_j\|^2 \delta_j^{3/2} \right)^2 \\ &\leq c_1 \left(\left(\sum_{j \in T_i} \|w_j\|^2 \right)^{1/4} \left(\sum_{j \in T_i} \|w_j\|^2 \delta_j^2 \right)^{3/4} \right)^2 \\ &= O(r^{1/2} w_{max}^{1/2} \epsilon^{3/4}), \end{aligned}$$

where we use Lemma C.10 in the first inequality, Holder's inequality in the last inequality, and Lemma C.5 and Lemma C.6 in the last line.

Thus, we have $\|R_2\|^2 = O(r^{5/2} w_{max}^{1/2} \epsilon^{3/4})$. ■

C.4.3. PROOF OF LEMMA C.5

Intuitively, when loss is small, the norm of student neurons are bounded. The proof follows by simple calculations.

Lemma C.5 *Under Assumption 2, if loss $L(W) = O(r^2 w_{max}^2)$, then $\sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 = O(r w_{max})$.*

Proof We have

$$\begin{aligned} \mathbb{E}_x \left[\left(\sum_{i=1}^r \sum_{j \in T_i} |w_j^\top x| \right)^2 \right] &\leq 2 \mathbb{E}_x \left[\left(\sum_{i=1}^r \sum_{j \in T_i} |w_j^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] + 2 \mathbb{E}_x \left[\left(\sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] \\ &= O(r^2 w_{max}^2). \end{aligned}$$

Also, we have

$$\mathbb{E}_x \left[\left(\sum_{i=1}^r \sum_{j \in T_i} |w_j^\top x| \right)^2 \right] \geq c \left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| \right)^2,$$

where c is a constant. Thus, $\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| = O(r w_{max})$. ■

C.4.4. PROOF OF LEMMA C.6

Recall that $\delta_j = \angle(w_i^*, w_j)$ as the angle between w_i^* and w_j for all $j \in T_i$. Lemma C.6 upper-bounds the weighted norm of student neurons by the value of the loss, where the weights depend on the angle between student and teacher neuron. Since more weights are assigned to the student neurons that are far away from any teacher neuron, this lemma implies that if there are many far-away neurons with large weights, then the loss is also large. The proof idea is to use test function as described in Section 5.3. We construct a test function such that it has large correlation with far-away student neurons and almost zero correlation with teacher neurons and near-by student neurons.

Lemma C.6 *Under Assumption 1, 2, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ such that for any W satisfying loss $\epsilon \triangleq L(W) \leq \epsilon_0$, we have $\sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 \delta_j^2 = O(\epsilon^{1/2})$.*

Proof Since we use absolute value function as activation, WLOG assume $w_j^\top w_i^* \geq 0$ for all $j \in T_i$ as discussed at the beginning of Section 5.2. We will first choose our test function in (i) and then complete the proof in (ii).

(i) Choose test function. Let $h(x) = \sqrt{\frac{\pi}{2}} - \sum_{k=1}^r \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_k^\top x)$ be the test function, where $h_l(x)$ is l -th Hermite polynomial, $\hat{\sigma}_l$ is l -th Hermite coefficient of absolute value function and $l = 2 \max\{\lceil \log_{\frac{1}{\cos(\Delta/2)}} \frac{1}{\epsilon} \rceil, 1\}$. See Section C.7 for the definition of Hermite polynomial and some properties. By Lemma C.11, we know $\hat{\sigma}_l^2 = \Theta(l^{-3/2})$. We are going to estimate $\langle h(x), R(x) \rangle$, where $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$ is the residual.

By Lemma C.11 and Claim C.1, for any w with $\|w\| = 1$, we have

$$\begin{aligned} \langle h_l(\bar{w}_i^{*\top} x), |w^\top x| \rangle &= \sum_{k=0}^{\infty} \hat{\sigma}_k \mathbb{E}[h_l(\bar{w}_i^{*\top} x) h_k(w^\top x)] = \hat{\sigma}_l \langle \bar{w}_i^*, w \rangle^l, \\ \langle 1, |w^\top x| \rangle &= \sqrt{\frac{2}{\pi}}. \end{aligned}$$

Denote $\phi = \angle(w_i^*, w)$. When $\sin^2 \phi \leq 2/l$, we have

$$\left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_i^{*\top} x), |w^\top x| \right\rangle = \cos^l \phi = (1 - \sin^2 \phi)^{l/2} \leq 1 - \frac{l}{4} \sin^2 \phi \leq 1 - \frac{\sin^2 \phi}{4},$$

where we use $(1-x)^n \leq 1 - \frac{nx}{2}$ for $x \in [0, 1/n]$. When $\sin^2 \phi > 2/l$, we have

$$\left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_i^{*\top} x), |w^\top x| \right\rangle = \cos^l \phi = (1 - \sin^2 \phi)^{l/2} \leq \left(1 - \frac{2}{l}\right)^{l/2} \leq \frac{1}{e} \leq 1 - \frac{\sin^2 \phi}{4}.$$

This implies that $\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_i^{*\top} x), |w^\top x| \rangle \leq 1 - \frac{\sin^2 \phi}{4}$ holds for any ϕ .

Further, when $\angle(w_i^*, w) \geq \frac{\Delta}{2}$, by the choice of l , we have

$$0 \leq \left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_i^{*\top} x), |w^\top x| \right\rangle \leq \cos^l \frac{\Delta}{2} \leq \epsilon.$$

Note that for $j \in T_i$ and $k \neq i$, we have $\Delta \leq \angle(w_k^*, w_i^*) \leq \angle(w_k^*, w_j) + \angle(w_i^*, w_j) \leq 2\angle(w_k^*, w_j)$. This indicates that $\angle(w_k^*, w_j) \geq \Delta/2$ for all $k \neq i$. Thus, for $j \in T_i$, we have

$$\begin{aligned} \langle h(x), |w_j^\top x| \rangle &= \sqrt{\frac{\pi}{2}} \langle 1, |w_j^\top x| \rangle - \sum_{k=1}^r \left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_k^{*\top} x), |w_j^\top x| \right\rangle \\ &\geq \left(1 - \left(1 - \frac{\sin^2 \delta_j}{4}\right) - (r-1)\epsilon\right) \|w_j\| \\ &= \frac{1}{4} \|w_j\| \sin^2 \delta_j - (r-1) \|w_j\| \epsilon. \end{aligned}$$

We also have for every teacher neuron w_j^* ,

$$\begin{aligned} \langle h(x), |w_j^{*\top} x| \rangle &= \sqrt{\frac{\pi}{2}} \langle 1, |w_j^{*\top} x| \rangle - \sum_{k=1}^r \left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_k^{*\top} x), |w_j^{*\top} x| \right\rangle \\ &= \left(1 - 1 - \sum_{k \neq j} \left\langle \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_k^{*\top} x), |w_j^{*\top} x| \right\rangle\right) \|w_j^*\| \leq 0. \end{aligned}$$

Therefore, for residual $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$, we have

$$\langle h(x), R(x) \rangle \geq \frac{1}{4} \sum_{i=1}^m \|w_i\| \sin^2 \delta_i - (r-1) \sum_{i=1}^m \|w_i\| \epsilon \geq \frac{1}{4} \sum_{i=1}^m \|w_i\| \sin^2 \delta_i - O(r^2 w_{\max} \epsilon),$$

where we use Lemma C.5 in the last inequality.

(ii) Complete Proof. Recall that $\|R\|^2 = 2L(W) \leq 2\epsilon$. From (i) we could have

$$\|h\| \sqrt{2\epsilon} \geq \|h\| \|R\| \geq \langle h(x), R(x) \rangle \geq \frac{1}{4} \sum_{i=1}^m \|w_i\| \sin^2 \delta_i - O(r^2 w_{max} \epsilon),$$

We now estimate the norm of h . Recall that $\hat{\sigma}_l^2 = \Theta(l^{-3/2})$, we have

$$\begin{aligned} \|h\|^2 &= \mathbb{E} \left[\left(\sqrt{\frac{\pi}{2}} - \sum_{k=1}^r \frac{1}{\hat{\sigma}_l} h_l(\bar{w}_k^* \top x) \right)^2 \right] = \frac{\pi}{2} + \frac{1}{\hat{\sigma}_l^2} \sum_{k,j=1}^r \langle \bar{w}_k^*, \bar{w}_j^* \rangle^l \\ &= \frac{\pi}{2} + O(l^{3/2} r^2 \epsilon) = O \left(1 + r^2 \epsilon \log^{\frac{3/2}{\cos(\Delta/2)}} \frac{1}{\epsilon} \right). \end{aligned}$$

Therefore, we have

$$\begin{aligned} \frac{1}{\pi^2} \sum_{i=1}^m \|w_i\| \delta_i^2 &\leq \frac{1}{4} \sum_{i=1}^m \|w_i\| \sin^2 \delta_i \leq \|h\| \sqrt{2\epsilon} + O(r^2 w_{max} \epsilon) \\ &= O \left(\epsilon^{1/2} + r^2 w_{max} \epsilon + r^2 \epsilon^{3/2} \log^{\frac{3/2}{\cos(\Delta/2)}} \frac{1}{\epsilon} \right), \end{aligned}$$

which finishes the proof. \blacksquare

C.5. Proof of Lemma C.2

Finally, we give the proof of Lemma C.2, which is the third step of proving Lemma 8. It shows that when the first two steps (Lemma 9 and Lemma 10) hold, the inner product between the gradient and descent direction can be lower bounded. The proof is involved and relies on algebraic computations.

Lemma C.2 *Under Assumption 2, if $\|v_i\| \leq \alpha$ and $|T_i(\delta_{max})| \geq 1$ for all $i \in [r]$, then*

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} \operatorname{sgn}(w_j^\top w_i^*) w_i^*) \rangle \geq \|R\|^2 - O(r^2 w_{max} \alpha \delta_{max}^2),$$

where $\{q_{ij}\}_{i \in [r], j \in [m]}$ is any sequence that satisfies (1) $q_{ij} \geq 0$ for all (i, j) , (2) $q_{ij} = 0$ if $j \notin T_i(\delta_{max})$; and (3) $\sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$.

Proof We focus on the case where $w_j^\top w_i^* \geq 0$ for all $j \in T_i$. The general case directly follows from Lemma C.1. Recall that residual $R(x) = \sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$. For any student neuron w_j where $j \in T_i$, we have

$$\begin{aligned} &\langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j^\top)^{-1} (w_j - q_{ij} w_i^*) \rangle \\ &= \mathbb{E}_x \left[R(x) \|w_j\| (w_j - q_{ij} w_i^*)^\top x \operatorname{sgn}(w_j^\top x) \right] \\ &= \mathbb{E}_x \left[R(x) \left(\|w_j\| |w_j^\top x| - q_{ij} \|w_j\| |w_i^{*\top} x| \right) \right] \\ &\quad + \mathbb{E}_x \left[R(x) q_{ij} \|w_j\| w_i^{*\top} x \left(\operatorname{sgn}(w_i^{*\top} x) - \operatorname{sgn}(w_j^\top x) \right) \right]. \end{aligned}$$

Sum over all student neurons w_j , we have

$$\begin{aligned}
 & \sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} w_i^*) \rangle \\
 &= \mathbb{E}_x \left[R^2(x) + R(x) \sum_{i=1}^r \sum_{j \in T_i} q_{ij} \|w_j\| w_i^{*\top} x \left(\text{sgn}(w_i^{*\top} x) - \text{sgn}(w_j^\top x) \right) \right] \\
 &= \|R\|^2 + \mathbb{E}_x \left[\underbrace{R(x) \sum_{i=1}^r \sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| w_i^{*\top} x \left(\text{sgn}(w_i^{*\top} x) - \text{sgn}(w_j^\top x) \right)}_{I_1(x)} \right], \tag{12}
 \end{aligned}$$

where the last line is because we set $q_{ij} = 0$ if $j \notin T_i(\delta_{max})$.

We are going to lower bound the second term $\mathbb{E}_x[R(x)I_1(x)]$ in the last line. Note that when $\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)$, we have $w_i^{*\top} x \left(\text{sgn}(w_i^{*\top} x) - \text{sgn}(w_j^\top x) \right) = 2|w_i^{*\top} x|$. Hence,

$$\begin{aligned}
 & \mathbb{E}_x \left[R(x) q_{ij} \|w_j\| w_i^{*\top} x \left(\text{sgn}(w_i^{*\top} x) - \text{sgn}(w_j^\top x) \right) \right] \\
 &= 2q_{ij} \|w_j\| \mathbb{E}_x \left[R(x) |w_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right] \\
 &= 2q_{ij} \|w_j\| \underbrace{\mathbb{E}_x \left[R_1(x) \cdot |w_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right]}_{I_2} \\
 & \quad + 2q_{ij} \|w_j\| \underbrace{\mathbb{E}_x \left[R_2(x) \cdot |w_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right]}_{I_3},
 \end{aligned}$$

where $R_1(x)$ and $R_2(x)$ are the residual decomposition as defined in (3).

According to Lemma C.3, we have the residual $R(x) = R_1(x) + R_2(x)$ satisfies $|R_1(x)| \leq \|x\| r\alpha$, $R_2(x) \geq 0$ for all x .

For the first term I_2 , recall that $R_1(x) = \sum_{i=1}^r v_i^\top x \text{sgn}(w_i^{*\top} x)$ and $\|v_i\| \leq \alpha$. We have

$$\begin{aligned}
 \mathbb{E}_x \left[R_1(x) |w_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right] &= \sum_{i=1}^r \mathbb{E}_x \left[v_i^\top x \text{sgn}(w_i^{*\top} x) |w_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right] \\
 &\geq - \sum_{i=1}^r \|v_i\| \|w_i^*\| \mathbb{E}_x \left[|\bar{v}_i^\top x| |\bar{w}_i^{*\top} x| \mathbb{I}_{\text{sgn}(w_i^{*\top} x) \neq \text{sgn}(w_j^\top x)} \right] \\
 &\geq - \sum_{i=1}^r \alpha \|w_i^*\| \delta_j \mathbb{E}_{\tilde{x}} \left[\|\tilde{x}\|^2 \mathbb{I}_{\text{sgn}(w_i^{*\top} \tilde{x}) \neq \text{sgn}(w_j^\top \tilde{x})} \right] \\
 &= -c_1 r w_{max} \alpha \delta_j^2,
 \end{aligned}$$

where in the second to last line \tilde{x} represents a 3-dimensional Gaussian since the expression only depends on three vectors v_i, w_j, w_i^* , and we use Lemma C.9 in the last line and c_1 is a constant.

For the second term I_3 , recall that $R_2(x) \geq 0$. Hence,

$$\mathbb{E}_x \left[R_2(x) |w_i^*{}^\top x| \mathbb{I}_{\text{sgn}(w_i^*{}^\top x) \neq \text{sgn}(w_j^\top x)} \right] \geq 0.$$

Therefore, we have

$$\begin{aligned} \mathbb{E}_x [R(x)I_1(x)] &\geq - \sum_{i=1}^r \sum_{j \in T_i(\delta_{max})} 2c_1 q_{ij} \|w_j\| r w_{max} \alpha \delta_j^2 \\ &\geq -2c_1 r w_{max} \alpha \delta_{max}^2 \sum_{i=1}^r \sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = -2c_1 r^2 w_{max} \alpha \delta_{max}^2, \end{aligned}$$

where in the last line is because $\sum_{j \in T_i(\delta_{max})} q_{ij} \|w_j\| = 1$ for all $i \in [r]$.

Thus, with (12) we have

$$\sum_{i=1}^r \sum_{j \in T_i} \langle \nabla_{w_j} L(W), (I + \bar{w}_j \bar{w}_j)^{-1} (w_j - q_{ij} w_i^*) \rangle \geq \|R\|^2 - O(r^2 w_{max} \alpha \delta_{max}^2).$$

■

C.6. Technical Lemmas

Lemma C.7 Consider $\alpha, \beta \in \mathbb{R}^d$ with $\|\alpha\| = \|\beta\| = 1$ and $x \sim N(0, I)$. Denote $\angle(\alpha, \beta) = \phi$, then if $\sin \phi > 0$, we have

$$\begin{aligned} \sqrt{\frac{2}{\pi}} \delta_1 e^{-\frac{\delta_1^2}{2}} &\leq \mathbb{E}_x [\mathbb{I}_{|\alpha^\top x| \leq \delta_1}] \leq \sqrt{\frac{2}{\pi}} \delta_1, \\ \frac{\delta_1(\delta_2 + \delta_1 \cos \phi)}{\pi \sin \phi} e^{-\frac{\delta_1^2 + \delta_2^2 + 2\delta_1 \delta_2 \cos \phi}{2 \sin^2 \phi}} &\leq \mathbb{E}_x [\mathbb{I}_A] \leq \frac{\delta_1(\delta_2 + \delta_1 \cos \phi)}{\pi \sin \phi}, \end{aligned}$$

where $A = \{x | |\alpha^\top x| \leq \delta_1, \text{ there exists } y \text{ such that } |\alpha^\top y| \leq \delta_1, |\beta^\top y| \leq \delta_2, (I - \alpha \alpha^\top)x = (I - \alpha \alpha^\top)y\}$

Proof WLOG, assume $\alpha = (1, 0, \dots, 0)^\top$ and $\beta = (\cos \phi, \sin \phi, 0, \dots, 0)^\top$ with $\phi \in (0, \frac{\pi}{2}]$. We first prove the first inequality. We have

$$\mathbb{E}_x [\mathbb{I}_{|\alpha^\top x| \leq \delta_1}] = \mathbb{E}_x [\mathbb{I}_{|x_1| \leq \delta_1}] = \frac{1}{\sqrt{2\pi}} \int_{|x_1| \leq \delta_1} e^{-\frac{x_1^2}{2}} dx_1 = \sqrt{\frac{2}{\pi}} \int_{0 \leq x_1 \leq \delta_1} e^{-\frac{x_1^2}{2}} dx_1.$$

Therefore, we have the upper bound

$$\mathbb{E}_x [\mathbb{I}_{|\alpha^\top x| \leq \delta_1}] \leq \sqrt{\frac{2}{\pi}} \delta_1,$$

and the lower bound

$$\mathbb{E}_x [\mathbb{I}_{|\alpha^\top x| \leq \delta_1}] \geq \sqrt{\frac{2}{\pi}} \delta_1 e^{-\frac{\delta_1^2}{2}}.$$

Now, we prove the second inequality. We have the set A is

$$\begin{aligned} A &= \{x \mid |x_1| \leq \delta_1, \text{ there exists } y \text{ such that } |y_1| \leq \delta_1, |y_1 \cos \phi + y_2 \sin \phi| \leq \delta_2, x_2 = y_2, \dots, x_d = y_d\} \\ &= \{x \mid |x_1| \leq \delta_1, \text{ there exists } y_1 \text{ such that } |y_1| \leq \delta_1, |y_1 \cos \phi + x_2 \sin \phi| \leq \delta_2\}. \end{aligned}$$

When $\cos \phi = 0$, we have $A = \{x \mid |x_1| \leq \delta_1, |x_2 \sin \phi| \leq \delta_2\}$. Thus,

$$\mathbb{E}_x[\mathbb{I}_A] = \frac{1}{4\pi} \int_{|x_1| \leq \delta_1} e^{-\frac{x_1^2}{2}} dx_1 \int_{|x_2| \leq \delta_2} e^{-\frac{x_2^2}{2}} dx_2.$$

We could have the upper bound

$$\mathbb{E}_x[\mathbb{I}_A] \leq \frac{1}{4\pi} \cdot 2\delta_1 \cdot 2\delta_2 = \frac{\delta_1 \delta_2}{\pi},$$

and lower bound

$$\mathbb{E}_x[\mathbb{I}_A] \geq \frac{1}{4\pi} \cdot 2\delta_1 e^{-\frac{\delta_1^2}{2}} \cdot 2\delta_2 e^{-\frac{\delta_2^2}{2}} = \frac{\delta_1 \delta_2}{\pi} e^{-\frac{\delta_1^2 + \delta_2^2}{2}}.$$

When $\cos \phi \neq 0$, we have

$$\begin{aligned} A &= \{x \mid |x_1| \leq \delta_1, \text{ there exists } y_1 \text{ such that } |y_1| \leq \delta_1, |y_1 \cos \phi + x_2 \sin \phi| \leq \delta_2\} \\ &= \{x \mid |x_1| \leq \delta_1, \frac{\delta_2 - x_2 \sin \phi}{\cos \phi} \geq -\delta_1, \frac{-\delta_2 - x_2 \sin \phi}{\cos \phi} \leq \delta_1\} \\ &= \{x \mid |x_1| \leq \delta_1, |x_2| \leq \frac{\delta_2 + \delta_1 \cos \phi}{\sin \phi}\}. \end{aligned}$$

Thus,

$$\mathbb{E}_x[\mathbb{I}_A] = \frac{1}{4\pi} \int_{|x_1| \leq \delta_1} e^{-\frac{x_1^2}{2}} dx_1 \int_{|x_2| \leq \frac{\delta_2 + \delta_1 \cos \phi}{\sin \phi}} e^{-\frac{x_2^2}{2}} dx_2.$$

We could have the upper bound

$$\mathbb{E}_x[\mathbb{I}_A] \leq \frac{1}{4\pi} \cdot 2\delta_1 \cdot \frac{2(\delta_2 + \delta_1 \cos \phi)}{\sin \phi} = \frac{\delta_1(\delta_2 + \delta_1 \cos \phi)}{\pi \sin \phi},$$

and lower bound

$$\mathbb{E}_x[\mathbb{I}_A] \geq \frac{1}{4\pi} \cdot 2\delta_1 e^{-\frac{\delta_1^2}{2}} \cdot \frac{2(\delta_2 + \delta_1 \cos \phi)}{\sin \phi} e^{-\frac{(\delta_2 + \delta_1 \cos \phi)^2}{2 \sin^2 \phi}} = \frac{\delta_1(\delta_2 + \delta_1 \cos \phi)}{\pi \sin \phi} e^{-\frac{\delta_1^2 + \delta_2^2 + 2\delta_1 \delta_2 \cos \phi}{2 \sin^2 \phi}}.$$

■

Lemma C.8 Suppose $0 \leq \tau, \phi \leq c$ for a sufficiently small constant. Then, we have

$$h(\tau, \phi) = 4 \text{OwenT}(\tau, \cot \phi) - 1 + \frac{2\phi}{\pi} + \frac{\tau}{\sqrt{2\pi}} \operatorname{erf}\left(\frac{\tau \cot \phi}{\sqrt{2}}\right) \leq 0,$$

where $\text{OwenT}(x, a) = \frac{1}{2\pi} \int_0^a e^{-x^2(1+t^2)/2} \frac{1}{1+t^2} dt$ and $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$.

Proof We have

$$\frac{dh(\tau, \phi)}{d\phi} = \frac{1}{\pi} \left(2 - 2e^{-\frac{\tau^2}{2\sin^2\phi}} - \frac{\tau^2}{\sin^2\phi} e^{-\frac{\tau^2 \cot^2\phi}{2}} \right) \triangleq \frac{1}{\pi} h_1(\tau, \phi).$$

Further,

$$\begin{aligned} \frac{dh_1(\tau, \phi)}{d\tau} &= \frac{\tau}{\sin^2\phi} \left(2e^{-\frac{\tau^2}{2\sin^2\phi}} - 2e^{-\frac{\tau^2 \cot^2\phi}{2}} + e^{-\frac{\tau^2 \cot^2\phi}{2}} \tau^2 \cot^2\phi \right) \\ &= \frac{\tau}{\sin^2\phi} e^{-\frac{\tau^2 \cot^2\phi}{2}} \left(2e^{-\frac{\tau^2}{2}} - 2 + \tau^2 \cot^2\phi \right) \\ &\geq \frac{\tau}{\sin^2\phi} e^{-\frac{\tau^2 \cot^2\phi}{2}} \left(2\left(1 - \frac{\tau^2}{2}\right) - 2 + \tau^2 \cot^2\phi \right) \\ &= \frac{\tau}{\sin^2\phi} e^{-\frac{\tau^2 \cot^2\phi}{2}} \tau^2 (\cot^2\phi - 1) \geq 0, \end{aligned}$$

where we use $e^{-x} \geq 1 - x$ when $x \geq 0$ in the first inequality and $\phi \leq \frac{\pi}{4}$ in the last line.

Hence, when $\phi \leq \frac{\pi}{4}$, $h_1(\tau, \phi) \geq h_1(0, \phi) = 0$. This implies that if $h(\tau, \phi_0) \leq 0$ for some $\phi_0 \leq \frac{\pi}{4}$, then we have $h(\tau, \phi) \leq h(\tau, \phi_0) \leq 0$.

We are going to show that when $\phi_0 = \text{arccot } 3 \leq \frac{\pi}{4}$ and $\tau \leq c$, we have $h(\tau, \phi_0) \leq 0$.

$$\begin{aligned} \frac{dh(\tau, \phi_0)}{d\tau} &= e^{-\frac{\tau^2 \cot^2\phi_0}{2}} \frac{\tau \cot\phi_0}{\pi} + \frac{\text{erf}\left(\frac{\tau \cot\phi_0}{\sqrt{2}}\right)}{\sqrt{2\pi}} - \sqrt{\frac{2}{\pi}} e^{-\frac{\tau^2}{2}} \text{erf}\left(\frac{\tau \cot\phi_0}{\sqrt{2}}\right) \\ &\leq \left(1 - \frac{\tau^2 \cot^2\phi_0}{2} + \frac{\tau^4 \cot^4\phi_0}{8} \right) \frac{\tau \cot\phi_0}{\pi} + \frac{1}{\sqrt{2\pi}} \frac{2}{\sqrt{\pi}} \frac{\tau \cot\phi_0}{\sqrt{2}} \\ &\quad - \sqrt{\frac{2}{\pi}} \left(1 - \frac{\tau^2}{2} \right) \left(\frac{2}{\sqrt{\pi}} \frac{\tau \cot\phi_0}{\sqrt{2}} - \frac{2}{3\sqrt{\pi}} \frac{\tau^3 \cot^3\phi_0}{2\sqrt{2}} \right) \\ &= \left(\frac{\cot\phi_0}{\pi} - \frac{\cot^3\phi_0}{6\pi} \right) \tau^3 + \left(\frac{\cot^5\phi_0}{8\pi} - \frac{\cot^3\phi_0}{6\pi} \right) \tau^5 = -\frac{3}{2\pi} \tau^3 + \frac{207}{8\pi} \tau^5 \leq 0, \end{aligned}$$

where we use $e^{-x} \leq 1 - x + \frac{x^2}{2}$, $e^{-x} \geq 1 - x$ and $\text{erf}(x) \geq \frac{2x}{\sqrt{\pi}} - \frac{2x^3}{3\sqrt{\pi}}$ for $x \geq 0$ in the first inequality, and $\tau \leq c$, $\phi_0 = \text{arccot } 3$ in the last line.

Therefore, we have $h(\tau, \phi_0) \leq h(0, \phi_0) = 0$, when $\tau \leq c$. Together with $h_1(\tau, \phi) \geq 0$ when $\phi \leq \frac{\pi}{4}$, we know $h(\tau, \phi) \leq h(\tau, \phi_0) \leq 0$, when $\tau \leq c$ and $\phi \leq \frac{\pi}{4}$. \blacksquare

Lemma C.9 Consider $\alpha, \beta \in \mathbb{R}^3$ with $\angle(\alpha, \beta) = \phi$ and $\alpha^\top \beta \geq 0$. We have

$$\mathbb{E}_x[\|x\|^2 \mathbb{I}_{\text{sgn}(\alpha^\top x) \neq \text{sgn}(\beta^\top x)}] = O(\phi).$$

Proof WLOG, assume $\alpha = (1, 0, 0)^\top$ and $\beta = (\cos \phi, \sin \phi, 0)^\top$. We have

$$\begin{aligned}
 & \mathbb{E}_x [\|x\|^2 \mathbb{I}_{\text{sgn}(\alpha^\top x) \neq \text{sgn}(\beta^\top x)}] \\
 &= \mathbb{E}_{x_1, x_2, x_3} [(x_1^2 + x_2^2 + x_3^2) \mathbb{I}_{\text{sgn}(x_1) \neq \text{sgn}(x_1 \cos \phi + x_2 \sin \phi)}] \\
 &= \mathbb{E}_{x_1, x_2} [(x_1^2 + x_2^2) \mathbb{I}_{\text{sgn}(x_1) \neq \text{sgn}(x_1 \cos \phi + x_2 \sin \phi)}] + \mathbb{E}_{x_1, x_2} [\mathbb{I}_{\text{sgn}(x_1) \neq \text{sgn}(x_1 \cos \phi + x_2 \sin \phi)}] \\
 &= \frac{1}{2\pi} \int_0^\infty r^3 e^{-r^2/2} dr \int_0^{2\pi} \mathbb{I}_{\text{sgn}(\cos \theta) \neq \text{sgn}(\cos \phi \cos \theta + \sin \phi \sin \theta)} d\theta \\
 &\quad + \frac{1}{2\pi} \int_0^\infty r e^{-r^2/2} dr \int_0^{2\pi} \mathbb{I}_{\text{sgn}(\cos \theta) \neq \text{sgn}(\cos \phi \cos \theta + \sin \phi \sin \theta)} d\theta \\
 &= O(\phi).
 \end{aligned}$$

■

Lemma C.10 Consider $w_j, w_k, w^* \in \mathbb{R}^d$ with $\|w_j\| = \|w_k\| = \|w^*\| = 1$ and $w_j^\top w^*, w_k^\top w^* \geq 0$. Denote $\phi_j = \angle(w_j, w^*)$ and $\phi_k = \angle(w_k, w^*)$. We have

$$0 \leq \mathbb{E}_x \left[w_j^\top x \cdot w_k^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w^{\top} x) \right) \left(\text{sgn}(w_k^\top x) - \text{sgn}(w^{\top} x) \right) \right] = O(\phi_j \phi_k \min\{\phi_j, \phi_k\}).$$

Proof Given that there are only three vectors w_j, w_k, w^* and $x \sim N(0, I)$, it is equivalent to consider a three dimensional Gaussian $\tilde{x} \in \mathbb{R}^3$ that lies in the span of w_j, w_k, w^* . In the following, we slightly abuse the notation and use x to denote this three dimensional Gaussian.

For any x , it is easy to see that

$$w_j^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w^{\top} x) \right) \geq 0, \quad w_k^\top x \left(\text{sgn}(w_k^\top x) - \text{sgn}(w^{\top} x) \right) \geq 0.$$

Therefore, we have the lower bound

$$\mathbb{E}_{x \sim N(0, I)} \left[w_j^\top x \cdot w_k^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w^{\top} x) \right) \left(\text{sgn}(w_k^\top x) - \text{sgn}(w^{\top} x) \right) \right] \geq 0.$$

For upper bound, note that when $\text{sgn}(w_j^\top x) \neq \text{sgn}(w^{\top} x)$, we have $w_j^\top x = O(\phi_j \|x\|)$. Similarly, we have $w_k^\top x = O(\phi_k \|x\|)$ when $\text{sgn}(w_k^\top x) \neq \text{sgn}(w^{\top} x)$. Thus, we have

$$\begin{aligned}
 & \mathbb{E}_{x \sim N(0, I)} \left[w_j^\top x \cdot w_k^\top x \left(\text{sgn}(w_j^\top x) - \text{sgn}(w^{\top} x) \right) \left(\text{sgn}(w_k^\top x) - \text{sgn}(w^{\top} x) \right) \right] \\
 & \leq O(\phi_j \phi_k) \mathbb{E}_x [\|x\|^2 \mathbb{I}_{\text{sgn}(w_j^\top x) \neq \text{sgn}(w^{\top} x), \text{sgn}(w_k^\top x) \neq \text{sgn}(w^{\top} x)}] \\
 & \leq O(\phi_j \phi_k) \min\{\mathbb{E}_x [\|x\|^2 \mathbb{I}_{\text{sgn}(w_j^\top x) \neq \text{sgn}(w^{\top} x)}], \mathbb{E}_x [\|x\|^2 \mathbb{I}_{\text{sgn}(w_k^\top x) \neq \text{sgn}(w^{\top} x)}]\} \\
 & = O(\phi_j \phi_k \min\{\phi_j, \phi_k\}),
 \end{aligned}$$

where we use Lemma C.9 in the last line. ■

C.7. Some Properties of Hermite Polynomials

In this section, we give several properties of Hermite Polynomials that are useful in our analysis. Let H_k be the probabilists' Hermite polynomial where

$$H_k(x) = (-1)^k e^{x^2/2} \frac{d^k}{dx^k} (e^{-x^2/2})$$

and $h_k = \frac{1}{\sqrt{k!}} H_k$ be the normalized Hermite polynomials.

The following lemma gives the Hermite coefficients for absolute value function.

Lemma C.11 *Let $\sigma(x) = |x|$. Then, $\sigma(x) = \sum_{k=0}^{\infty} \hat{\sigma}_k h_k(x)$, where $\{h_k\}_{k=0}^{\infty}$ are Hermite polynomials and*

$$\hat{\sigma}_k = \begin{cases} 0 & , k \text{ is odd} \\ \sqrt{\frac{2}{\pi}} & , k = 0 \\ \sqrt{\frac{1}{\pi}} & , k = 2 \\ (-1)^{\frac{k}{2}-1} \sqrt{\frac{2}{\pi \cdot k!}} (k-3)!! & , k \text{ is even and } n \geq 4 \end{cases}$$

are Hermite coefficients of $|x|$. Here, $\hat{\sigma}_k^2 = \Theta(k^{-3/2})$.

Proof We calculate $\hat{\sigma}_k$ for the following cases. **Case 1:** k is odd.

Since $h_k(x)$ is odd when k is odd, we have

$$\hat{\sigma}_k = \langle \sigma(x), h_k(x) \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |x| h_k(x) e^{-x^2/2} dx = 0.$$

Case 2: $k = 0$.

$$\hat{\sigma}_0 = \langle \sigma(x), h_0(x) \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |x| e^{-x^2/2} dx = \sqrt{\frac{2}{\pi}}.$$

Case 3: $k = 2$.

$$\hat{\sigma}_2 = \langle \sigma(x), h_2(x) \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |x| \cdot \frac{x^2 - 1}{\sqrt{2}} \cdot e^{-x^2/2} dx = \frac{1}{\sqrt{\pi}}$$

Case 4: k is even and $k \geq 4$.

When k is even, we know h_k is even. Thus,

$$\begin{aligned}
 \hat{\sigma}_k &= \langle \sigma(x), h_k(x) \rangle = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |x| \cdot \frac{H_k}{\sqrt{k!}} \cdot e^{-x^2/2} dx \\
 &= \frac{1}{\sqrt{2\pi} \cdot k!} \int_{-\infty}^{\infty} |x| \cdot (-1)^k e^{x^2/2} \frac{d^k}{dx^k} (e^{-x^2/2}) \cdot e^{-x^2/2} dx \\
 &= \frac{(-1)^k}{\sqrt{2\pi} \cdot k!} \cdot 2 \int_0^{\infty} x \cdot \frac{d^k}{dx^k} (e^{-x^2/2}) dx \\
 &= \frac{2(-1)^k}{\sqrt{2\pi} \cdot k!} \cdot \left(x \cdot \frac{d^{k-1}}{dx^{k-1}} (e^{-x^2/2}) \Big|_0^{\infty} - \int_0^{\infty} \frac{d^{k-1}}{dx^{k-1}} (e^{-x^2/2}) dx \right) \\
 &= \frac{2(-1)^k}{\sqrt{2\pi} \cdot k!} \cdot \left(- \frac{d^{k-2}}{dx^{k-2}} (e^{-x^2/2}) \Big|_0^{\infty} \right) \\
 &= \sqrt{\frac{2}{\pi \cdot k!}} \cdot H_{k-2}(0) \\
 &= \sqrt{\frac{2}{\pi \cdot k!}} (-1)^{\frac{k}{2}-1} (k-3)!!,
 \end{aligned}$$

where we use $H_k(x) = (-1)^k e^{x^2/2} \frac{d^k}{dx^k} (e^{-x^2/2})$ and $H_k(0) = (-1)^{k/2} (k-1)!!$ when k is even (Abramowitz and Stegun, 1948). ■

The following is another helpful property of Hermite polynomial.

Claim C.1 ((O'Donnell, 2014), Section 11.2) *Let (x, y) be ρ -correlated standard normal variables (that is, both x, y have marginal distribution $N(0, 1)$ and $\mathbb{E}[xy] = \rho$). Then, $\mathbb{E}[h_m(x)h_n(y)] = \rho^n \delta_{mn}$.*

Appendix D. Proof of Local Convergence (Theorem 5)

In this section, we first show that the loss function is ‘‘smooth’’ (Lemma 13) and Lipschitz (Lemma 14). At the end, we give the proof of local convergence theorem (Theorem 5).

D.1. Proof of Theorem 13

As the first step to prove local convergence theorem, we show that the loss function satisfies ‘‘smoothness’’-type conditions as below. The proof is involved and relies on algebraic computations to carefully bound each term.

Lemma 13 (Smoothness) *Under Assumption 2, if loss $L(W) = O(r^2 w_{max}^2)$, then*

$$\begin{aligned}
 L(W + U) &\leq L(W) + \langle \nabla_W L(W), U \rangle + O(r^{1/4} w_{max}^{1/4}) L^{1/2}(W) \|U\|_F^{3/2} + O(r w_{max}) \|U\|_F^2 \\
 &\quad + O(1) \|U\|_F^4.
 \end{aligned}$$

Proof Denote the residual as

$$R_W(x) = \sum_{j=1}^m \|w_j\| |w_j^\top x| - \sum_{i=1}^r |w_i^{*\top} x|.$$

Then, we have

$$\begin{aligned}
 & L(W + U) - L(W) - \langle \nabla_W L(W), U \rangle \\
 &= \frac{1}{2} \mathbb{E}_x [R_{W+U}^2(x)] - \frac{1}{2} \mathbb{E}_x [R_W^2(x)] - \mathbb{E}_x \left[R_W(x) \sum_{j=1}^m \|w_j\| u_j^\top (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right] \\
 &= \frac{1}{2} \underbrace{\mathbb{E}_x [(R_{W+U}(x) - R_W(x))^2]}_{I_1} \\
 &\quad + \underbrace{\mathbb{E}_x \left[R_W(x) \left(R_{W+U}(x) - R_W(x) - \sum_{j=1}^m \|w_j\| u_j^\top (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right) \right]}_{I_2}.
 \end{aligned}$$

For term I_1 , let $\gamma_j = \|w_j + u_j\| (w_j + u_j) - \|w_j\| w_j$. Note that

$$\begin{aligned}
 \|\gamma_j\|^2 &\leq 2 \left(\|\|w_j + u_j\| w_j - \|w_j\| w_j\|^2 + \|\|w_j + u_j\| u_j\|^2 \right) \\
 &\leq 2 \left(\|w_j\|^2 \|u_j\|^2 + \|w_j + u_j\|^2 \|u_j\|^2 \right) \\
 &\leq 2 \left(3 \|w_j\|^2 + 2 \|u_j\|^2 \right) \|u_j\|^2.
 \end{aligned}$$

Hence, we have

$$\begin{aligned}
 I_1 &= \mathbb{E}_x [(R_{W+U}(x) - R_W(x))^2] = \mathbb{E}_x \left[\left(\sum_{j=1}^m \left(\|\|w_j + u_j\| (w_j + u_j)^\top x\| - \|\|w_j\| w_j^\top x\| \right) \right)^2 \right] \\
 &\leq \mathbb{E}_x \left[\left(\sum_{j=1}^m |\gamma_j^\top x| \right)^2 \right] \\
 &\leq c_0 \left(\sum_{j=1}^m \|\gamma_j\| \right)^2 \\
 &\leq c_0 \left(\sum_{j=1}^m \left(6 \|w_j\|^2 + 4 \|u_j\|^2 \right)^{1/2} \|u_j\| \right)^2 \\
 &\leq c_0 \sum_{j=1}^m \left(6 \|w_j\|^2 + 4 \|u_j\|^2 \right) \cdot \sum_{j=1}^m \|u_j\|^2 \\
 &= O(rw_{max}) \sum_{j=1}^m \|u_j\|^2 + O(1) \left(\sum_{j=1}^m \|u_j\|^2 \right)^2,
 \end{aligned}$$

where we use Lemma D.4 in the third line, and the last line is because of Lemma C.5.

For term I_2 , we have

$$\begin{aligned}
 I_2 &= \mathbb{E} \left[R_W(x) \left(R_{W+U}(x) - R_W(x) - \sum_{j=1}^m \|w_j\| u_j^\top (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right) \right] \\
 &= \mathbb{E} \left[R_W(x) \sum_{j=1}^m \underbrace{\left(\|w_j + u_j\| (w_j + u_j)^\top x - \|w_j\| w_j^\top x - \|w_j\| u_j^\top (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right)}_{I_j(x)} \right] \\
 &\leq \sum_{j=1}^m \|R_W\| \|I_j\|.
 \end{aligned}$$

Denote $\alpha = \|w_j + u_j\| (w_j + u_j)$ and $\beta = \|w_j\| (w_j + (I + \bar{w}_j \bar{w}_j^\top) u_j)$. In the following, we drop the subscript j for simplicity. We have

$$\begin{aligned}
 \|I_j\|^2 &= \mathbb{E}_x \left[\left(|\alpha^\top x| - \beta^\top x \operatorname{sgn}(w^\top x) \right)^2 \right] \\
 &\leq 2\mathbb{E}_x \left[\left(|\alpha^\top x| - |\beta^\top x| \right)^2 + \left(\beta^\top x (\operatorname{sgn}(\beta^\top x) - \operatorname{sgn}(w^\top x)) \right)^2 \right] \\
 &\leq 2\mathbb{E}_x \left[\left((\alpha - \beta)^\top x \right)^2 \right] + 2\mathbb{E}_x \left[\left(\beta^\top x \right)^2 \mathbb{I}_{\operatorname{sgn}(\beta^\top x) \neq \operatorname{sgn}(w^\top x)} \right] \\
 &= 2\|\alpha - \beta\|^2 + 2 \underbrace{\mathbb{E}_x \left[\left(\beta^\top x \right)^2 \mathbb{I}_{\operatorname{sgn}(\beta^\top x) \neq \operatorname{sgn}(w^\top x)} \right]}_{I_3}.
 \end{aligned}$$

For first term $\|\alpha - \beta\|$, let $g(w) = \|w\| w$. By Lemma D.1, we know

$$g(w + u) = g(w) + \langle \nabla_w g(w), u \rangle + \Delta,$$

where $\|\Delta\| \leq c_1 \|u\|^2$ for some constant c_1 . Note that $g(w + u) = \alpha$ and $g(w) + \langle \nabla_w g(w), u \rangle = \beta$. Hence, we have $\|\alpha - \beta\| \leq c_1 \|u\|^2$.

For second term I_3 , by Lemma D.2, we have

$$I_3 = \mathbb{E}_x \left[\left(\beta^\top x \right)^2 \mathbb{I}_{\operatorname{sgn}(\beta^\top x) \neq \operatorname{sgn}(w^\top x)} \right] = \frac{\|\beta\|^2}{\pi} (\phi - \sin \phi \cos \phi) \leq \frac{\|\beta\|^2}{\pi} \phi^3,$$

where $\phi = \arccos(\bar{\beta}^\top \bar{w})$ and we use $\phi - \sin \phi \cos \phi \leq \phi^3$ for $\phi \geq 0$.

Denote $\theta = \arccos(\bar{w}^\top \bar{u})$. We have

$$\begin{aligned}
 \cos \phi &= \frac{\beta^\top w}{\|\beta\| \|w\|} = \frac{\|w\|^3 + 2\|w\| w^\top u}{\|\beta\| \|w\|} = \frac{\|w\| + 2\|u\| \cos \theta}{\sqrt{\|w\|^2 + \|u\|^2 + 3\|u\|^2 \cos^2 \theta + 4\|w\| \|u\| \cos \theta}} \\
 &= \frac{\|w\| + 2\|u\| \cos \theta}{\sqrt{(\|w\| + 2\|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta}}.
 \end{aligned}$$

Note that $\cos \phi \leq 1 - \frac{\phi^2}{5}$ for $0 \leq \phi \leq \pi$. We have

$$I_3 \leq \frac{\|\beta\|^2}{\pi} 5^{3/2} \left(1 - \frac{\|w\| + 2\|u\| \cos \theta}{\sqrt{(\|w\| + 2\|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta}} \right)^{3/2}. \quad (13)$$

Note that when $\|u\| = 0$, it is easy to verify that $I_3 = 0$. In the following, we assume $\|u\| > 0$.

Case 1: $\|w\| \geq 4\|u\|$.

In this case, we have $\|w\| + 2\|u\| \cos \theta \geq \frac{\|w\|}{2} \geq 2\|u\| > 0$. Using $1 - \frac{1}{\sqrt{1+x^2}} \leq \frac{x^2}{2}$ with (13), we have

$$\begin{aligned} I_3 &\leq \frac{5^{3/2} \|\beta\|^2}{\pi} \left(\frac{\|u\|^2 \sin^2 \theta}{2(\|w\| + 2\|u\| \cos \theta)^2} \right)^{3/2} \\ &= \frac{5^{3/2} \|w\|^2 \left((\|w\| + 2\|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta \right)}{2^{3/2} \pi \|w\| + 2\|u\| \cos \theta^3} \|u\|^3 |\sin \theta|^3 \\ &= \frac{5^{3/2} \|w\|}{2^{3/2} \pi \|w\| + 2\|u\| \cos \theta} \frac{(\|w\| + 2\|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta}{\|w\| + 2\|u\| \cos \theta^2} \|w\| \|u\|^3 |\sin \theta|^3 \\ &\leq \frac{5^{3/2}}{2^{3/2} \pi} 2 \left(1 + \frac{\sin^2 \theta}{4} \right) \|w\| \|u\|^3 |\sin \theta|^3 \\ &\leq 4 \|w\| \|u\|^3 \end{aligned}$$

Case 2: $\|w\| < 4\|u\|$.

In this case, we have $\|\beta\|^2 = \|w\|^2 \left((\|w\| + 2\|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta \right) \leq 592 \|u\|^4$. With (13), we have $I_3 \leq \frac{592 \|u\|^4}{\pi} 5^{3/2} 2^{3/2}$.

Combine above two cases, we have $I_3 \leq c_2 \max\{\|w\| \|u\|^3, \|u\|^4\}$, where c_2 is a constant. Therefore, we have $\|I_j\| \leq c_3 \max\{\|w_j\|^{1/2} \|u_j\|^{3/2}, \|u_j\|^2\}$, where c_3 is a constant. This leads to

$$\begin{aligned} I_2 &\leq \sum_{j=1}^m \|R_W\| \|I_{ij}\| \leq \|R_W\| \sum_{j=1}^m c_3 \max\{\|w_j\|^{1/2} \|u_j\|^{3/2}, \|u_j\|^2\} \\ &\leq \|R_W\| \sum_{j=1}^m c_3 \|w_j\|^{1/2} \|u_j\|^{3/2} + \|R_W\| \sum_{j=1}^m c_3 \|u_j\|^2 \\ &\leq c_3 \|R_W\| \left(\sum_{j=1}^r \|w_j\|^2 \right)^{1/4} \left(\sum_{j=1}^m \|u_{ij}\|^2 \right)^{3/4} + c_3 \|R_W\| \sum_{j=1}^m \|u_{ij}\|^2, \end{aligned}$$

where we use Holder's inequality in the last line.

Denote $\|U\|_F^2 = \sum_{j=1}^m \|u_j\|^2$. With Lemma C.5, we have

$$\begin{aligned} &L(W + U) - L(W) - \langle \nabla_W L(W), U \rangle \\ &\leq I_1 + I_2 \\ &\leq O(rw_{max}) \|U\|_F^2 + O(1) \|U\|_F^4 + O(r^{1/4} w_{max}^{1/4}) L^{1/2}(W) \|U\|_F^{3/2} + O(1) L^{1/2}(W) \|U\|_F^2 \\ &= O(r^{1/4} w_{max}^{1/4}) L_2^{1/2}(W) \|U\|_F^{3/2} + O(rw_{max}) \|U\|_F^2 + O(1) \|U\|_F^4. \end{aligned}$$

■

D.2. Proof of Lemma 14

The second step is to show the loss function is Lipschitz, i.e., the gradient norm is upperbound. The proof follows from simple computations.

Lemma 14 (Lipschitz) *Under Assumption 2, if loss $L(W) = O(r^2 w_{max}^2)$, then $\|\nabla_W L(W)\|_F^2 = O(r^3 w_{max}^3)$.*

Proof Denote $R(x) = \sum_{i=1}^m \sum_{j \in T_i} \|w_j\| |w_j^\top x| - \sum_{i=1}^r |w_i^{*\top} x|$ as the residual. For any $k \in [m]$, we have

$$\begin{aligned}
 \|\nabla_{w_k} L(W)\| &= \left\| \mathbb{E}_x \left[R(x) \|w_k\| (I + \bar{w}_k \bar{w}_k^\top) x \operatorname{sgn}(w_k^\top x) \right] \right\| \\
 &= \|w_k\| \left\| \mathbb{E}_x \left[\left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| |w_j^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right) (I + \bar{w}_k \bar{w}_k^\top) x \operatorname{sgn}(w_k^\top x) \right] \right\| \\
 &\leq \|w_k\| \left(\sum_{i=1}^r \sum_{j \in T_i} \left\| \mathbb{E}_x \left[\|w_j\| |w_j^\top x| (I + \bar{w}_k \bar{w}_k^\top) x \operatorname{sgn}(w_k^\top x) \right] \right\| \right. \\
 &\quad \left. + \sum_{i=1}^r \left\| \mathbb{E}_x \left[|w_i^{*\top} x| (I + \bar{w}_k \bar{w}_k^\top) x \operatorname{sgn}(w_k^\top x) \right] \right\| \right) \\
 &\leq \|w_k\| \left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\|^2 \cdot O(1) + \sum_{i=1}^r \|w_i^*\| \cdot O(1) \right) \\
 &\leq O(r w_{max}) \|w_k\|,
 \end{aligned}$$

where we use Lemma D.3 in the second to last line, and Lemma C.5 in the last line. This leads to

$$\|\nabla_W L(W)\|_F^2 = \sum_{j=1}^m \|\nabla_{w_j} L_2(W)\|^2 \leq O(r^2 w_{max}^2) \sum_{j=1}^m \|w_j\|^2 = O(r^3 w_{max}^3),$$

where we again use Lemma C.5. ■

D.3. Proof of Theorem 5

Now we are ready to prove the local convergence result given the smoothness and Lipschitz of loss function.

Theorem 5 (Main Result) *For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, suppose we run gradient descent on objective (1):*

$$w_i^{(t+1)} = w_i^{(t)} - \eta \nabla_{w_i} L(W), \quad \text{for any } i \in [m]$$

there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ and $\eta_0 = O(r^{-1} w_{max}^{-1})$ such that for any $\epsilon > 0$, if initial loss $L(W^{(0)}) \leq \epsilon_0$ and step size $\eta \leq \eta_0$, then we have $L(W^{(T)}) \leq \epsilon$ in $T = O(r w_{max} / (\epsilon \eta))$ steps.

Proof Using Theorem 13 with $u_j = -\eta \nabla_{w_j} L(W)$, we have

$$\begin{aligned} L(W + U) &\leq L(W) - \eta \|\nabla_W L(W)\|_F^2 + O(\eta^{3/2} r^{1/4} w_{max}^{1/4}) L^{1/2}(W) \|\nabla_W L(W)\|_F^{3/2} \\ &\quad + O(\eta^2 r w_{max}) \|\nabla_W L(W)\|_F^2 + O(\eta^4) \|\nabla_W L(W)\|_F^4 \\ &\leq L(W) - \left(\eta - O\left(\eta^{3/2} r^{1/4} w_{max}^{1/4} \kappa^{-1/2} + \eta^2 r w_{max} + \eta^4 r^3 w_{max}^3 \right) \right) \|\nabla_W L(W)\|_F^2 \\ &\leq L(W) - \frac{\eta \kappa^2}{4} L^2(W), \end{aligned}$$

where in the second line we use gradient norm upper bound (Lemma 14) and gradient norm lower bound (Theorem 4), in the last line we use Theorem 4 again and $\eta \leq \eta_0 = O(r^{-1} w_{max}^{-1})$. This implies that

$$L(W^{(t+1)}) \leq L(W^{(t)}) - \frac{\eta \kappa^2}{4} L^2(W^{(t)}).$$

Then, note that $0 < L(W^{(t+1)}) \leq L(W^{(t)})$, we could have

$$\frac{1}{L(W^{(t)})} \leq \frac{1}{L(W^{(t+1)})} - \frac{\eta \kappa^2}{4} \frac{L(W^{(t)})}{L(W^{(t+1)})} \leq \frac{1}{L(W^{(t+1)})} - \frac{\eta \kappa^2}{4}.$$

This implies that

$$0 < \frac{1}{L(W^{(0)})} \leq \frac{1}{L(W^{(t)})} - \frac{t\eta \kappa}{4},$$

which leads to

$$L(W^{(t)}) \leq \min \left\{ \frac{4}{t\eta \kappa^2}, L(W^{(0)}) \right\} = \min \left\{ O\left(\frac{r w_{max}}{t\eta} \right), \epsilon_0 \right\}.$$

Therefore, we have $L(W^{(T)}) \leq \epsilon$ for some $T = O\left(\frac{r w_{max}}{\epsilon \eta} \right)$. ■

D.4. Technical Lemmas

Lemma D.1 *Let $g(w) = \|w\| w$. Then, g is $(1 + \sqrt{3})$ -smooth.*

Proof The gradient is

$$\nabla_w g(w) = \|w\| (I + \bar{w} \bar{w}^\top).$$

Note that when $w = 0$, $\nabla_w g(w) = 0$.

For any u , we have

$$\begin{aligned} &\|\nabla_w g(w + u) - \nabla_w g(w)\|_F \\ &= \left\| \|w + u\| (I + \overline{w + u} \overline{w + u}^\top) - \|w\| (I + \bar{w} \bar{w}^\top) \right\|_F \\ &\leq \|(\|w + u\| - \|w\|)I\|_F + \left\| \|w + u\| \overline{w + u} \overline{w + u}^\top - \|w\| \bar{w} \bar{w}^\top \right\|_F \\ &\leq \|u\| + \underbrace{\left\| \|w + u\| \overline{w + u} \overline{w + u}^\top - \|w\| \bar{w} \bar{w}^\top \right\|_F}_{I_1}. \end{aligned}$$

Since $\nabla_w g(0) = 0$, when $w = 0$ or $w + u = 0$, it is clear that $\|\nabla_w g(w + u) - \nabla_w g(w)\|_F \leq 2 \|u\|$.

For I_1 , we have

$$\begin{aligned}
 I_1^2 &= \left\| \|w + u\| \overline{w + u} \overline{w + u}^\top - \|w\| \overline{w} \overline{w}^\top \right\|_F^2 \\
 &= \text{tr} \left(\left(\|w + u\| \overline{w + u} \overline{w + u}^\top - \|w\| \overline{w} \overline{w}^\top \right)^2 \right) \\
 &= \|w + u\|^2 + \|w\|^2 - 2 \|w + u\| \|w\| \left(\overline{w + u}^\top \overline{w} \right)^2 \\
 &= (\|w + u\| - \|w\|)^2 + 2 \|w + u\| \|w\| \left(1 - \left(\overline{w + u}^\top \overline{w} \right)^2 \right) \\
 &\leq \|u\|^2 + 2 \|w + u\| \|w\| \left(1 - \left(\overline{w + u}^\top \overline{w} \right)^2 \right).
 \end{aligned}$$

Denote $\theta = \arccos(\overline{w}^\top \overline{u})$. Then,

$$\begin{aligned}
 \overline{w + u}^\top \overline{w} &= \frac{\|w\|^2 + \|w\| \|u\| \cos \theta}{\|w\| \|w + u\|} = \frac{\|w\| + \|u\| \cos \theta}{\sqrt{\|w\|^2 + \|u\|^2 + 2 \|w\| \|u\| \cos \theta}} \\
 &= \frac{\|w\| + \|u\| \cos \theta}{\sqrt{(\|w\| + \|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta}}.
 \end{aligned}$$

Hence,

$$\|w + u\| \|w\| \left(1 - \left(\overline{w + u}^\top \overline{w} \right)^2 \right) = \frac{\|w + u\| \|w\| \|u\|^2 \sin^2 \theta}{(\|w\| + \|u\| \cos \theta)^2 + \|u\|^2 \sin^2 \theta} = \frac{\|w\| \sin^2 \theta}{\|w + u\|} \|u\|^2.$$

It is easy to check $\frac{\|w\| \sin \theta}{\|w + u\|} \leq 1$. Thus, we have $I_1^2 \leq (1 + 2 \sin \theta) \|u\|^2 \leq 3 \|u\|^2$. Therefore, for any u

$$\|\nabla_w g(w + u) - \nabla_w g(w)\|_F \leq (1 + \sqrt{3}) \|u\|.$$

■

Lemma D.2 Consider $\alpha, \beta \in \mathbb{R}^d$ with $\|\alpha\| = \|\beta\| = 1$. Denote $\phi = \arccos(\alpha^\top \beta)$, then

$$\mathbb{E}_{x \sim N(0, I)} [(\beta^\top x)^2 \mathbb{I}_{\text{sgn}(\alpha^\top x) \neq \text{sgn}(\beta^\top x)}] = \frac{1}{\pi} (\phi - \sin \phi \cos \phi) = \Theta(\phi^3).$$

Proof WLOG, assume $\beta = (1, 0, \dots, 0)^\top$ and $\alpha = (\cos \phi, \sin \phi, 0, \dots, 0)^\top$. We have

$$\begin{aligned}
 &\mathbb{E}_{x \sim N(0, I)} [(\beta^\top x)^2 \mathbb{I}_{\text{sgn}(\alpha^\top x) \neq \text{sgn}(\beta^\top x)}] \\
 &= \mathbb{E}_{x \sim N(0, I)} [x_1^2 \mathbb{I}_{\text{sgn}(x_1 \cos \phi + x_2 \sin \phi) \neq \text{sgn}(x_1)}] \\
 &= \frac{1}{2\pi} \int_0^\infty r^3 e^{-r^2/2} dr \int_0^{2\pi} \cos^2 \theta \mathbb{I}_{\text{sgn}(\cos \theta \cos \phi + \sin \theta \sin \phi) \neq \text{sgn}(\cos \theta)} d\theta \\
 &= \frac{1}{2\pi} \cdot 2 \cdot 2 \int_{\pi/2}^{\pi/2 + \phi} \cos^2 \theta d\theta \\
 &= \frac{2}{\pi} \cdot \frac{1}{2} (\phi - \sin \phi \cos \phi) = \frac{1}{\pi} (\phi - \sin \phi \cos \phi) = \Theta(\phi^3).
 \end{aligned}$$

■

Lemma D.3 Consider $\alpha, \beta \in \mathbb{R}^d$ with $\|\alpha\| = \|\beta\| = 1$. We have

$$\left\| \mathbb{E}_{x \sim N(0, I)} \left[|\alpha^\top x| (I + \beta\beta^\top) x \operatorname{sgn}(\beta^\top x) \right] \right\|^2 = O(1).$$

Proof WLOG, assume $\beta = (1, 0, \dots, 0)^\top$ and $\alpha = (\alpha_1, \alpha_2, 0, \dots, 0)^\top$.

We have

$$\begin{aligned} & \left\| \mathbb{E}_{x \sim N(0, I)} \left[|\alpha^\top x| (I + \beta\beta^\top) x \operatorname{sgn}(\beta^\top x) \right] \right\|^2 \\ &= \left\| \mathbb{E}_{x \sim N(0, I)} \left[|\alpha_1 x_1 + \alpha_2 x_2| (I + \beta\beta^\top) x \operatorname{sgn}(x_1) \right] \right\|^2 \\ &= \left(\mathbb{E}_{x \sim N(0, I)} [|\alpha_1 x_1 + \alpha_2 x_2| \cdot 2x_1 \cdot \operatorname{sgn}(x_1)] \right)^2 + \left(\mathbb{E}_{x \sim N(0, I)} [|\alpha_1 x_1 + \alpha_2 x_2| \cdot x_2 \cdot \operatorname{sgn}(x_1)] \right)^2 \\ &\leq \mathbb{E}_{x \sim N(0, I)} \left[(|\alpha_1 x_1 + \alpha_2 x_2| \cdot 2x_1 \cdot \operatorname{sgn}(x_1))^2 \right] + \mathbb{E}_{x \sim N(0, I)} \left[(|\alpha_1 x_1 + \alpha_2 x_2| \cdot x_2 \cdot \operatorname{sgn}(x_1))^2 \right] \\ &\leq 4 \mathbb{E}_{x \sim N(0, I)} [|\alpha_1 x_1 + \alpha_2 x_2|^2 \cdot (x_1^2 + x_2^2)] \\ &\leq 4 \mathbb{E}_{x \sim N(0, I)} [(x_1^2 + x_2^2)^2] \\ &= O(1), \end{aligned}$$

where we use Jensen's inequality in the first inequality. ■

Lemma D.4 Consider $\alpha_i \in \mathbb{R}^d$ for $i \in [n]$. We have

$$\mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^n |\alpha_i^\top x| \right)^2 \right] \leq c_0 \left(\sum_{i=1}^n \|\alpha_i\| \right)^2,$$

where c_0 is a constant.

Proof We have

$$\mathbb{E}_x \left[\left(\sum_{i=1}^n |\alpha_i^\top x| \right)^2 \right] = \mathbb{E}_x \left[\sum_{i,j=1}^n |\alpha_i^\top x| |\alpha_j^\top x| \right] \leq \sum_{i,j=1}^n c_0 \|\alpha_i\| \|\alpha_j\| = c_0 \left(\sum_{i=1}^n \|\alpha_i\| \right)^2.$$

To prove the first inequality above, it suffices to prove the following

$$\mathbb{E}_{x \sim N(0, I)} \left[|\alpha_i^\top x| |\alpha_j^\top x| \right] \leq c_0 \|\alpha_i\| \|\alpha_j\|.$$

Note that the LHS above only depends on two vectors, so it is equivalent to consider the expectation over a 2-dimensional Gaussian \tilde{x} . We have

$$\mathbb{E}_{x \sim N(0, I)} \left[|\alpha_i^\top x| |\alpha_j^\top x| \right] = \mathbb{E}_{\tilde{x} \sim N(0, I)} \left[|\alpha_i^\top \tilde{x}| |\alpha_j^\top \tilde{x}| \right] \leq \|\alpha_i\| \|\alpha_j\| \mathbb{E}_{\tilde{x}} \left[\|\tilde{x}\|^2 \right] = c_0 \|\alpha_i\| \|\alpha_j\|,$$

where c_0 is a constant. ■

Appendix E. Initialization

We present the details of two initialization algorithms: (1) random initialization (Algorithm 1) and (2) subspace initialization (Algorithm 2) and prove their correctness respectively in the following two subsections. At the end of this section, we give the proof of Theorem 6.

E.1. Random Initialization (Algorithm 1)

Random Initialization (Algorithm 1) initializes the direction of neurons randomly and adjust the norm of neurons by least-squares. In the proof we show that as long as every teacher neuron has at least one close student neuron in direction, the solution returned by least-squares will be small.

Algorithm 1 Random Initialization

Input: number of student neurons m .

Initialize student neurons as $w_1, w_2, \dots, w_m \sim N(0, I_d)$

Set new student neurons $W' = (w'_1, \dots, w'_m)$ as $w'_i = \sqrt{z_i^*/\|w_i\|} \cdot w_i$ where $\{z_i^*\}_{i=1}^m$ is the solution of the following least-squares problem

$$\min_{\{z_i\}_{i=1}^m: \forall i \in [m], z_i \geq 0} \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right]$$

Output: W'

Lemma E.1 (Random Initialization) *Under Assumption 2, for any $\epsilon_{init} > 0$, if we use Algorithm 1 with $m \geq m_0 = O((r w_{max}/\sqrt{\epsilon_{init}})^d \cdot r \log(1/\delta))$, then with probability $1 - \delta$, we have $L(W') \leq \epsilon_{init}$.*

Proof For simplicity, we use ϵ instead of ϵ_{init} in the proof. Recall that after random initialization, we adjust the norm of student neurons by setting new student neurons $W' = (w'_1, \dots, w'_m)$ as $w'_i = \sqrt{z_i^*/\|w_i\|} \cdot w_i$ where $\{z_i^*\}_{i=1}^m$ is the solution of the following least-squares problem

$$\min_{\{z_i\}_{i=1}^m: \forall i \in [m], z_i \geq 0} \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right].$$

Therefore, it is equivalent to consider $w_1, w_2, \dots, w_m \sim \text{Uniform}(S_{d-1})$. First, we are going to show that for every teacher neuron w_i^* , with probability $1 - \delta$ and with $O(\gamma^{-d} \log \frac{1}{\delta})$ neurons, at least one of these neurons satisfy $\angle(w, w_i^*) \leq \gamma$.

Note that for a single neuron w ,

$$\mathbb{P}(\angle(w, w_i^*) \leq \gamma) = \frac{2 \int_0^\gamma \sin^{d-2} \theta_1 d\theta_1}{\int_0^\pi \sin^{d-2} \theta_1 d\theta_1} = \Omega(\gamma^d).$$

Hence, with $O(\gamma^{-d} \log \frac{1}{\delta})$ neurons, the probability of none of these neurons satisfy $\angle(w, w_i^*) \leq \gamma$ is at most $(1 - \Omega(\gamma^d))^{O(\gamma^{-d} \log \frac{1}{\delta})} \leq \delta$. Therefore, with $m \geq m_0 = O(r \gamma^{-d} \log \frac{1}{\delta})$ neurons, we

guarantee that with probability $1 - \delta$, every teacher neuron has at least one student neuron satisfy $\angle(w_i, w_i^*) \leq \gamma$.

Then, we show that if every teacher neuron has at least one student neuron satisfy $\angle(w_i, w_i^*) \leq \gamma$, then after doing least square to fit the norms $\|w_i\|$, the loss $L(W) \leq O(r^2 w_{max}^2 \gamma^2)$. We prove this by constructing a feasible solution that satisfy the constraints.

Denote student neuron w_i as the one that satisfies $\angle(w_i, w_i^*) \leq \gamma$. Let $z_i = \|w_i^*\| / \|w_i\|$ for $i \in [r]$ and $z_i = 0$ otherwise. Consider the least-squares objective under this case, we have

$$\begin{aligned} L(W) &= \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] = \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^r \|w_i^*\| (|\bar{w}_i^\top x| - |\bar{w}_i^{*\top} x|) \right)^2 \right] \\ &\leq \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^r \|w_i^*\| |(\bar{w}_i - \bar{w}_i^*)^\top x| \right)^2 \right] \\ &\leq \frac{c_0}{2} \left(\sum_{i=1}^r \|w_i^*\| \|\bar{w}_i - \bar{w}_i^*\| \right)^2 = O(r^2 w_{max}^2 \gamma^2), \end{aligned}$$

where we use Lemma D.4 and $\angle(w_i, w_i^*) \leq \gamma$ in the last line.

Therefore, let $\gamma = O(r^{-1} w_{max}^{-1} \epsilon^{1/2})$, we know after least square $L(W) \leq \epsilon$. In this case, we need $m \geq m_0 = O\left(r \left(\frac{r w_{max}}{\epsilon^{1/2}}\right)^d \log \frac{1}{\delta}\right)$ neurons at initialization. \blacksquare

E.2. Subspace Initialization (Algorithm 2)

Subspace Initialization (Algorithm 2) first uses samples to estimate the subspace space spanned by the teacher neuron. Then following the same argument in Random Initialization (Algorithm 1), if we random initializes the direction of neurons in this subspace and adjust the norms by least-squares, we could have a small initial loss.

Algorithm 2 Subspace Initialization

Input: number of student neurons m , number of teacher neuron r , N samples (x_i, y_i) where $x_i \sim$

$$N(0, I) \text{ and } y_i = \sum_{i=1}^r |w_i^{*\top} x_i|.$$

$$\text{Let } \widehat{M} = \frac{1}{N} \sum_{i=1}^N y_i (x_i x_i^\top - I).$$

Let S be the subspace spanned by the top- r eigenvectors of \widehat{M} and $Q \in \mathbb{R}^{d \times r}$ be a matrix formed by an orthonormal basis of S .

Initialize student neurons as $w_1, w_2, \dots, w_m \sim N(0, QQ^\top)$.

Set new student neurons $W' = (w'_1, \dots, w'_m)$ as $w'_i = \sqrt{z_i^* / \|w_i\|} \cdot w_i$ where $\{z_i^*\}_{i=1}^m$ is the solution of the following least-squares problem

$$\min_{\{z_i\}_{i=1}^m: \forall i \in [m], z_i \geq 0} \frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right]$$

Output: W'

We need the following results for our second initialization that show the span of top eigenvectors of M is approximately the span of teacher neurons. Recall that $f^*(x) = \sum_{i=1}^r |w_i^{*\top} x|$.

Lemma E.2 (Claim 5.2, (Zhong et al., 2017), with our notation)

$$M \triangleq \mathbb{E}_{x \sim N(0, I)} \left[f^*(x) \left(x x^\top - I \right) \right] = c \sum_{i=1}^r \|w_i^*\| \bar{w}_i^* \bar{w}_i^{*\top} \quad (14)$$

Lemma E.3 (Lemma E.2, (Zhong et al., 2017), with our notation) For M defined in (14), denote $\widehat{M} = \frac{1}{N} \sum_{i=1}^N f^*(x_i) (x_i x_i^\top - I)$. If there are $N = \tilde{O}(d/\epsilon^2)$ samples (x_i, y_i) , where $y_i = f^*(x_i)$, then with probability $1 - \delta$ we have

$$\left\| \widehat{M} - M \right\|_2 = O(r w_{\max} \epsilon).$$

Lemma E.4 (Lemma 3.6, (Diakonikolas et al., 2020), with our notation) For M defined in (14), let $\widehat{M} \in \mathbb{R}^{d \times d}$ be a matrix such that $\left\| M - \widehat{M} \right\|_2 \leq \epsilon$ and let \mathcal{V} be the subspace of \mathbb{R}^d that is spanned by the top- r eigenvectors of \widehat{M} . There exist r vectors $v^{(i)} \in \mathcal{V}$ such that

$$\mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^r |v^{(i)\top} x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] = O(r^2 w_{\max} \epsilon).$$

Now we are ready to prove the lemma for Subspace Initialization.

Lemma E.5 (Subspace Initialization)

Under Assumption 2, for any $\epsilon_{\text{init}} > 0$, if we use Algorithm 2 with $\epsilon_{\text{init}} \leq O(r^2 w_{\max}^2)$, $N \geq N_0 = \tilde{O}(dr^6 w_{\max}^4 / \epsilon_{\text{init}}^2)$ and $m \geq m_0 = O((r w_{\max} / \sqrt{\epsilon_{\text{init}}})^r \cdot r \log(1/\delta))$, then with probability $1 - \delta$, we have $L(W') \leq \epsilon_{\text{init}}$.

Proof For simplicity, we use ϵ instead of ϵ_{init} in the proof. We use \tilde{d} to represent an upperbound of r . From Lemma E.3 and Lemma E.4 and $N \geq N_0 = \tilde{O}(dr^6 w_{\max}^4 / \epsilon^2)$, we know with probability $1 - \delta$ there exists r vectors $v^{(i)} \in S$ such that

$$\mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^r |v^{(i)\top} x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] \leq \epsilon/2.$$

Since $w_1, w_2, \dots, w_m \sim N(0, QQ^\top)$ and $Q \in \mathbb{R}^{d \times \tilde{d}}$ is a matrix formed by an orthonormal basis of S , we are effectively sampling in S . That is, $w_i = Qu_i$ with $u_i \sim N(0, I_{\tilde{d}})$. Then, following the similar arguments in proof of Lemma E.1, we know using $m \geq m_0 = O(r \gamma^{-d} \log \frac{1}{\delta})$ neurons, with probability $1 - \delta$ every $v^{(i)}$ has at least one close enough neuron w_i in the sense of $\angle(v^{(i)}, w_i) \leq \gamma$ and when $z_i = \|v^{(i)}\| / \|w_i\|$ for $i \in [r]$ and $z_i = 0$ otherwise, we have

$$\frac{1}{2} \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |v^{(i)\top} x| \right)^2 \right] = \left(\sum_{i=1}^r \|v^{(i)}\| \right)^2 O(\gamma^2).$$

By Lemma C.5, we know $\sum_{i=1}^r \|v^{(i)}\| = O(rw_{max})$. Therefore, let $\gamma = O(r^{-1}w_{max}^{-1}\epsilon^{1/2})$, we have

$$\mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |v^{(i)\top} x| \right)^2 \right] = O(r^2 w_{max}^2 \gamma^2) \leq \epsilon/2.$$

Hence, the loss after least square is bounded by

$$L(W) \leq \mathbb{E}_{x \sim N(0, I)} \left[\left(\sum_{i=1}^r |v^{(i)\top} x| - \sum_{i=1}^r |w_i^{*\top} x| \right)^2 \right] + \mathbb{E}_x \left[\left(\sum_{i=1}^m z_i |w_i^\top x| - \sum_{i=1}^r |v^{(i)\top} x| \right)^2 \right] \leq \epsilon.$$

This implies when $m \geq m_0 = O\left(r \left(\frac{rw_{max}}{\epsilon^{1/2}}\right)^{\tilde{d}} \log \frac{1}{\delta}\right)$, the loss after least square is at most ϵ . ■

E.3. Proof of Theorem 6

Now we are ready to prove Theorem 6 given the above two initialization procedures.

Theorem 6 *For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, set $\epsilon_{init} = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$, to achieve a good initialization random initialization requires $O\left((rw_{max}/\sqrt{\epsilon_{init}})^d \cdot r \log(1/\delta)\right)$ student neurons and subspace initialization requires $O\left((rw_{max}/\sqrt{\epsilon_{init}})^r \cdot r \log(1/\delta)\right)$ student neurons with $\tilde{O}(dr^6 w_{max}^4 / \epsilon_{init}^2)$ samples. Suppose we run GD from either initialization, there exists a threshold $\eta_0 = O(r^{-1}w_{max}^{-1})$ such that for any $\epsilon > 0$, if step size $\eta \leq \eta_0$, with probability $1 - \delta$ we have $L(W^{(T)}) \leq \epsilon$ in $T = O(rw_{max}/(\epsilon\eta))$ steps.*

Proof Set $\epsilon_{init} = \epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$. By Lemma E.1 and Lemma E.5, we know the initial loss returned by Algorithm 1 and Algorithm 2 satisfies $L(W^l) \leq \epsilon_0$. Then, combining Theorem 5 we finishes the proof. ■

E.4. Student-Teacher Neuron Matching and the Lottery Ticket Hypothesis

In this subsection, we give a more detailed discussion about the connection between student-teacher neuron matching as indicated in our local convergence result and the lottery ticket hypothesis (Frankle and Carbin, 2018). The lottery ticket hypothesis suggests that one can prune a neural network such that even if we only train the neural network with a small subset of randomly initialized neurons (the pruned network), we can still have a model with good performance. Our local convergence result shows if one can maintain at least one student neuron close to every teacher neuron at every step, gradient descent will eventually converge to the global optimal solution and recover the teacher neurons. This gives a partial explanation for lottery ticket hypothesis in the two-layer teacher/student setting – as long as the initialization contains student neurons that are close to each teacher neuron and one can prune away neurons that were not close to teacher neurons at initialization, then the training process will converge to the global optima. Intuitively, the two initialization algorithms in Section E use least-squares to prune away the useless neurons given enough number of student neurons at initialization.

Appendix F. Sample Complexity

Recall we have the empirical loss

$$\widehat{L}(W) = \frac{1}{2N} \sum_{k=1}^N \left(\sum_{i=1}^r \sum_{j \in T_i} \|w_j\| |w_j^\top x_k| - \sum_{i=1}^r |w_i^{*\top} x_k| \right)^2,$$

where N sample $\{(x_k, y_k)\}$ are i.i.d. sampled from $x_k \sim N(0, I)$ and $y_k = f^*(x_k) = \sum_{i=1}^r |w_i^{*\top} x_k|$.

We first give the following concentration result that shows when the number of sample is large enough, the gradient on empirical loss is close to the gradient on population loss.

Lemma F.1 *Under Assumption 2, for any fixed W , if loss $L(W) = O(r^2 w_{max}^2)$ and the number of data $N \geq O(r^3 w_{max}^3 d^2 \epsilon^{-2} \delta^{-1})$, with probability $1 - \delta$ we have*

$$\left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F \leq \epsilon.$$

Proof We have

$$\begin{aligned} & \mathbb{P} \left(\left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F \geq \epsilon \right) \\ & \leq \epsilon^{-2} \mathbb{E}_{x_1, x_2, \dots, x_N} \left[\left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F^2 \right] = \epsilon^{-2} \sum_{j=1}^m \mathbb{E}_{x_1, x_2, \dots, x_N} \left[\left\| \nabla_{w_j} \widehat{L}(W) - \nabla_{w_j} L(W) \right\|^2 \right] \\ & = \epsilon^{-2} N^{-1} \sum_{j=1}^m \mathbb{E}_x \left[\left\| \left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right) \|w_j\| (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) - \nabla_{w_j} L(W) \right\|^2 \right] \\ & = 2\epsilon^{-2} N^{-1} \sum_{j=1}^m \mathbb{E}_x \left[\left\| \left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right) \|w_j\| (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right\|^2 \right] \\ & \quad + 2\epsilon^{-2} N^{-1} \sum_{j=1}^m \left\| \nabla_{w_j} L(W) \right\|^2. \end{aligned}$$

We now bound the first term above. We have

$$\begin{aligned} & \mathbb{E}_x \left[\left\| \left(\sum_{i=1}^m \|w_i\| |w_i^\top x| - \sum_{i=1}^r |w_i^{*\top} x| \right) \|w_j\| (I + \bar{w}_j \bar{w}_j^\top) x \operatorname{sgn}(w_j^\top x) \right\|^2 \right] \\ & = 4 \|w_j\|^2 \mathbb{E}_x \left[\left\| x \right\|^4 \left(\sum_{i=1}^m \|w_i\| |w_i^\top \bar{x}| - \sum_{i=1}^r |w_i^{*\top} \bar{x}| \right)^2 \right] \\ & \leq 4 \|w_j\|^2 \mathbb{E}_x \left[\left\| x \right\|^4 \left(\sum_{i=1}^m \|w_i\|^2 + \sum_{i=1}^r \|w_i^*\|^2 \right)^2 \right] \\ & = O(r^2 w_{max}^2 d^2) \|w_j\|^2, \end{aligned}$$

where in the last line we use Lemma C.5, which gives $\sum_{i=1}^m \|w_i\|^2 = O(rw_{max})$. Hence, use Lemma C.5 again, we upper bound the first term by $O(\epsilon^{-2}N^{-1}r^3w_{max}^3d^2)$.

For the second term, from Lemma 14, we know

$$\sum_{j=1}^m \|\nabla_{w_j} L(W)\|^2 = O(r^3w_{max}^3).$$

Therefore, we have

$$\mathbb{P}\left(\left\|\nabla_W \widehat{L}_2(W) - \nabla_W L_2(W)\right\| \geq \epsilon\right) = O\left(\frac{r^3w_{max}^3d^2}{\epsilon^2N}\right).$$

By our choice of N , this finishes the proof. \blacksquare

Given the above concentration result, we are ready to prove Theorem 7. The proof is very similar with the proof of Theorem 5.

Theorem 7 *For network and loss function $L(\cdot)$ defined in Definition 2, under Assumption 1, 2, suppose we run stochastic GD on W with N fresh samples within each mini-batch at every iteration, i.e., for any $i \in [m]$*

$$w_i^{(t+1)} = w_i^{(t)} - \eta \nabla_{w_i} \widehat{L}_t(W),$$

where $\widehat{L}_t(\cdot)$ is the empirical loss using the N samples at iteration t . Then, there exists a threshold $\epsilon_0 = \text{poly}(\Delta, r^{-1}, w_{max}^{-1}, w_{min})$ and $\eta_0 = O(r^{-1}w_{max}^{-1})$ such that for any $\epsilon > 0$, if initial loss $L(W^{(0)}) \leq \epsilon_0$, step size $\eta \leq \eta_0$ and batch size $N \geq O(r^5w_{max}^5d^2\eta^{-1}\epsilon^{-3}\delta^{-1})$, then with probability $1 - \delta$ we have $L(W^{(T)}) \leq \epsilon$ in $T = O(rw_{max}/(\epsilon\eta))$ steps.

Proof By Lemma 14, we know

$$\|\nabla_W L(W)\|_F^2 = O(r^3w_{max}^3).$$

By Theorem 4, we know

$$\|\nabla_W L(W)\|_F \geq \kappa L(W),$$

where $\kappa = \Theta(r^{-1/2}w_{max}^{-1/2})$.

Now, using Theorem 13 with $u_j = -\eta \nabla_{w_j} \widehat{L}(W)$, we have

$$\begin{aligned} L(W + U) &\leq L(W) - \eta \langle \nabla_W L(W), \nabla_W \widehat{L}(W) \rangle + O(\eta^{3/2}r^{1/4}w_{max}^{1/4})L^{1/2}(W) \left\| \nabla_W \widehat{L}(W) \right\|_F^{3/2} \\ &\quad + O(\eta^2rw_{max}) \left\| \nabla_W \widehat{L}(W) \right\|_F^2 + O(\eta^4) \left\| \nabla_W \widehat{L}(W) \right\|_F^4. \end{aligned}$$

Using Lemma F.1 with $N \geq O(r^3w_{max}^3d^2\epsilon_g^{-2}\delta_g^{-1})$ samples where $\epsilon_g^2 \leq \frac{1}{4} \|\nabla_W L(W)\|^2$, then with probability $1 - \delta_g$ we have $\left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F \leq \epsilon_g$. Hence,

$$\begin{aligned} \langle \nabla_W L(W), \nabla_W \widehat{L}(W) \rangle &= \|\nabla_W L(W)\|_F^2 + \langle \nabla_W L(W), \nabla_W \widehat{L}(W) - \nabla_W L(W) \rangle \\ &\geq \|\nabla_W L(W)\|_F^2 - \|\nabla_W L(W)\|_F \left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F \\ &\geq \frac{1}{2} \|\nabla_W L(W)\|_F^2, \end{aligned}$$

$$\begin{aligned} \left\| \nabla_W \widehat{L}(W) \right\|_F^{3/2} &\leq \left(\left\| \nabla_W L(W) \right\|_F + \left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F \right)^{3/2} \leq O(1) \left\| \nabla_W L(W) \right\|_F^{3/2}, \\ \left\| \nabla_W \widehat{L}(W) \right\|_F^2 &\leq 2 \left\| \nabla_W L(W) \right\|_F^2 + 2 \left\| \nabla_W \widehat{L}(W) - \nabla_W L(W) \right\|_F^2 \leq O(1) \left\| \nabla_W L(W) \right\|_F^2. \end{aligned}$$

With the above bounds, we have

$$\begin{aligned} L(W + U) &\leq L(W) - \left(\eta - O \left(\eta^{3/2} r^{1/4} w_{max}^{1/4} \kappa^{-1/2} + \eta^2 r w_{max} + \eta^4 r^3 w_{max}^3 \right) \right) \left\| \nabla_W L(W) \right\|_F^2 \\ &\leq L(W) - \frac{\eta \kappa^2}{4} L^2(W), \end{aligned}$$

where in the last line we use $\eta \leq \eta_0 = O(\min\{r^{-1/2} w_{max}^{-1/2} \kappa, r^{-1} w_{max}^{-1}\})$. This implies that with probability $1 - \delta_g$

$$L(W^{(t+1)}) \leq L(W^{(t)}) - \frac{\eta \kappa^2}{4} L^2(W^{(t)}).$$

Then, note that $0 < L(W^{(t+1)}) \leq L(W^{(t)})$, we could have

$$\frac{1}{L(W^{(t)})} \leq \frac{1}{L(W^{(t+1)})} - \frac{\eta \kappa^2}{4} \frac{L(W^{(t)})}{L(W^{(t+1)})} \leq \frac{1}{L(W^{(t+1)})} - \frac{\eta \kappa^2}{4}.$$

This implies that

$$0 < \frac{1}{L(W^{(0)})} \leq \frac{1}{L(W^{(t)})} - \frac{t\eta \kappa^2}{4},$$

which leads to

$$L(W^{(t)}) \leq \min \left\{ \frac{4}{t\eta \kappa^2}, L(W^0) \right\} = \min \left\{ O \left(\frac{r w_{max}}{t\eta} \right), \epsilon_0 \right\}.$$

Therefore, with probability $1 - T\delta_g$, we have $L(W^{(T)}) \leq \epsilon$ in $T = O \left(\frac{r w_{max}}{\epsilon \eta} \right)$. Setting $\delta_g = \delta/T$ and $\epsilon_g = \frac{1}{2} \kappa \epsilon \leq \frac{1}{2} \left\| \nabla_W L(W) \right\|_F$ finishes the proof. \blacksquare