

ROBUSTNESS GUARANTEES FOR ADVERSARIAL TRAINING ON NON-SEPARABLE DATA

Anonymous authors

Paper under double-blind review

ABSTRACT

Adversarial training has emerged as a popular approach for training models that are robust to inference time attacks. However, our theoretical understanding of why and when it works remains limited. Prior work has offered convergence analysis of adversarial training, but they are either restricted to the Neural Tangent Kernel (NTK) regime or make restrictive assumptions about data such as linearly realizability. In this work, we provide convergence and generalization guarantees for adversarial training of two-layer networks of any width on non-separable data. Our analysis goes beyond the NTK regime and holds for both smooth and non-smooth activation functions. We support our theoretical findings with an empirical study on synthetic and real-world data.

1 INTRODUCTION

Machine learning models are ubiquitous in real-world applications, achieving state-of-the-art performance on various tasks such as image classification and speech recognition. However, several recent studies have shown that these models, especially those based on deep neural networks, are highly vulnerable to small, nearly imperceptible, albeit strategic, perturbation of data. These perturbations, called adversarial examples, are abundant and easy to find computationally (Bubeck et al., 2021; Wang et al., 2022). The potential of such adversarial attacks to substantially degrade the performance of an otherwise well-performing model has been a source of significant concern regarding deploying machine learning models in real-world systems. It is no surprise, then, that developing algorithms that can provably defend against such attacks and are guaranteed to improve the robustness of machine learning has gained tremendous traction in recent years.

One of the most prominent empirical defense algorithms against inference-time attacks is the adversarial training method of Madry et al. (2018). Adversarial training proceeds by simulating attacks as part of training – generating adversarial examples from (clean) training examples and using them to train a neural network. We can view adversarial training as a two-player game, wherein the learner seeks to minimize their error on the training set while an adversary strives to maximize the error by crafting small strategic corruptions of the input training examples. Several empirical studies show that by using adversarial training, the learner returns a model that is more resilient to perturbations in the input space (Madry et al., 2018; Shafahi et al., 2019b; Dong et al., 2020; Pang et al., 2021).

Despite the empirical success of adversarial training, our understanding of its theoretical underpinnings is far from complete. Several prior works study statistical and computational aspects of adversarial training but in somewhat restrictive settings; e.g., assuming linear separability of data (Mianjy and Arora, 2022), or essentially assuming away nonconvexity of neural networks by considering an overly parametrized regime wherein the trajectory dynamics are in the lazy regime (aka, the neural tangent kernel or the NTK setting) (Gao et al., 2019; Zhang et al., 2020; Li and Telgarsky, 2023). In this paper, we forego these simplifying assumptions and present theoretical convergence and generalization guarantees for adversarial training on two-layer neural networks, of any width, on non-separable data. Our key contributions are as follows.

1. We establish convergence guarantees for adversarial training of two-layer neural networks. We allow the network to be of *arbitrary width* thereby extending our results to networks beyond the NTK regime. Furthermore, we do not make any assumptions about the separability or robust realizability of data.

2. We provide generalization guarantees on both the clean test error and the robust test error. For a moderately large network, we show that for norm-bounded additive adversarial attacks, if the perturbation budget is not too large, the robust test error approximates the label noise rate. For adversarial attacks with a large perturbation budget, we show that the robust test error is bounded from below by a constant.
3. We validate our theoretical results with experiments on both synthetic and real-world datasets.

1.1 RELATED WORK

Convergence Analysis of Standard Training. Several recent works study the convergence of (stochastic) gradient descent for training neural networks (Arora et al., 2019; Allen-Zhu et al., 2019; Cao and Gu, 2019). Most of these works focus on a lazy training regime wherein the network weights remain close to initialization through the run of the algorithm (owing to an extreme over-parametrization); this is also referred to as the neural tangent kernel (NTK) setting. While interesting from a theoretical perspective (we essentially end up with a convex learning problem), this assumption is typically violated in practice. Analyzing SGD beyond the NTK setting is much more challenging owing to the non-convexity of learning problems associated with training neural networks of arbitrary widths. There has been some progress toward addressing this challenge – Frei et al. (2022) provide a first guarantee for finite-width neural networks trained on logistic loss for data drawn from a Gaussian mixture model. Concurrently, Cao et al. (2022) characterize the generalization guarantees of two-layer convolutional neural networks, assuming that the input data is a sum of a label-dependent signal patch and a label-independent noise patch. While both of the works above consider a smooth activation function, follow-up works by Kou et al. (2023); Xu and Gu (2023) extend the result to SGD for training neural networks with non-smooth activation functions (e.g., ReLU networks).

Convergence Analysis of Adversarial Training. Adversarial training, introduced by Madry et al. (2018), is one of the most popular algorithms for training models that are robust to adversarial attacks. Subsequent works have explored variants, including the TRADES (Zhang et al., 2019) and MART (Wang et al., 2020) algorithms. Despite their success, a theoretical understanding of why and when adversarial training succeeds remains elusive. Much of the recent work (Charles et al., 2019; Li et al., 2020; Zou et al., 2021; Chen et al., 2021) has focused on studying adversarial training of linear models wherein the adversarial examples are given in a simple closed-form expression – this simplifies the problem greatly reducing it to standard training. Adversarial training of neural networks was analyzed by Gao et al. (2019) and further improved by Zhang et al. (2020); however, both of these works focus on ensuring convergence of the training procedure and do not provide generalization guarantees on robust loss. This gap has been addressed in very recent work by Li and Telgarsky (2023). However, the work of Li and Telgarsky (2023), and the prior work all focus on the lazy training regime, which, unfortunately, has been proven to be at odds with robustness Wang et al. (2022). Finally, Mianjy and Arora (2022) provide an end-to-end analysis of adversarial training beyond the NTK setting with a variant of adversarial training that involves using a slightly different (reflected) loss for the inner loop maximization problem (for finding an attack vector as part of adversarial training). The results of Mianjy and Arora (2022) are limited to distributions that are robustly realizable.

Our work builds on that of Frei et al. (2022) and considers a high-dimensional setting for a class-conditional model; the data model, as well as various other data assumptions we need, were first introduced and studied in Chatterji and Long (2021). While our proof techniques are inspired by Frei et al. (2022), we differ in many respects. To the best of our knowledge, ours is the first work that provides the convergence and generalization guarantees for *adversarial training* for a *non-separable data distribution*. We consider neural networks with both smooth and non-smooth activation functions, e.g., ReLU networks; the analysis of Frei et al. (2022) is limited to smooth activation functions. Additionally, unlike prior works (Gao et al., 2019; Zhang et al., 2020; Li and Telgarsky, 2023) that are limited to the NTK setting, our guarantees hold for neural networks of *arbitrary width* and analyze GD-based adversarial training in the rich regime (i.e., *beyond the lazy regime*).

2 PRELIMINARIES

Notation Throughout the paper, we denote scalars, vectors, and matrices with lowercase italics, lowercase bold, and uppercase bold Roman letters, respectively; e.g., u , \mathbf{u} , and \mathbf{U} . We use $[m]$

to denote the set $\{1, 2, \dots, m\}$ and use both $\|\cdot\|$ and $\|\cdot\|_2$ for ℓ_2 -norm. Given a matrix $U = [u_1, \dots, u_m] \in \mathbb{R}^{d \times m}$, we use $\|U\|_F$ and $\|U\|_2$ to represent the Frobenius norm and spectral norm, respectively. We use $\mathcal{B}_2(u, \alpha)$ to denote the ℓ_2 ball centered at $u \in \mathbb{R}^d$ of radius α . We use the standard O-notation (\mathcal{O} , Θ and Ω).

2.1 PROBLEM SETUP

We focus on binary classification and denote the input space and label space as $\mathcal{X} = \mathbb{R}^d$, $\mathcal{Y} = \{\pm 1\}$, respectively. We assume that the data are drawn from a noisy mixture data distribution \mathcal{D} on $\mathcal{X} \times \mathcal{Y}$ that, along with its variants, has been studied in several recent works (Chatterji and Long, 2021; Cao et al., 2021; Frei et al., 2022). Formally, we consider the following data distribution.

Definition (Data Distribution). Let $\mathcal{D}_{\text{clust}}$ be a λ -strongly log-concave distribution over \mathbb{R}^d for some $\lambda > 0$. We assume that $\mathcal{D}_{\text{clust}} = \mathcal{D}_{\text{clust}}^{(1)} \times \dots \times \mathcal{D}_{\text{clust}}^{(d)}$ is a product distribution whose marginals are all mean-zero with the sub-Gaussian norm at most one. We further assume that $\mathbb{E}_{\xi \sim \mathcal{D}_{\text{clust}}} [\|\xi\|^2] \geq \kappa d$ holds for some $0 < \kappa < 1$. Let \mathcal{D}_c be a distribution over $\mathcal{X} \times \mathcal{Y}$. We first draw a sample $(x_c, y_c) \sim \mathcal{D}_c$ by sampling $y_c \in \{\pm 1\}$ uniformly at random, sampling $\xi \sim \mathcal{D}_{\text{clust}}$, and setting $x_c = y_c \mu + \xi$. Given a noise rate $\beta > 0$, we define our true data distribution \mathcal{D} to be any distribution over $\mathcal{X} \times \mathcal{Y}$ such that the marginal distribution of \mathcal{D} and \mathcal{D}_c on \mathcal{X} are the same, and the total variation distance between the two distributions is bounded by β , i.e., $d_{\text{TV}}(\mathcal{D}_c, \mathcal{D}) \leq \beta$.

The standard coupling lemma states that given two distributions \mathcal{D} and \mathcal{D}_c over the same domain $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, there exists a joint distribution over $\mathcal{Z} \times \mathcal{Z}$ such that the marginals along the projections $(z, z') \mapsto z$ and $(z, z') \mapsto z'$ are \mathcal{D} and \mathcal{D}_c , respectively. Given that the marginal on \mathcal{X} for \mathcal{D} and \mathcal{D}_c are the same (see the definition above), this implies that for $(x, y) \sim \mathcal{D}$, $(x_c, y_c) \sim \mathcal{D}_c$, $\mathbb{P}(x = x_c) = 1$ and $\mathbb{P}(y \neq y_c) \leq \beta$. The definition above includes two settings: 1) Independent label flip, where for each sample, label y is obtained by flipping y_c with probability at most β , independent of how other labels are generated; 2) Non-independent label flip, where there exists potential correlations between labels y . A yet another special instance that has been studied extensively in the adversarial learning literature is that of Gaussian distribution (Javanmard et al., 2020; Dobriban et al., 2020; Dan et al., 2020) which is a special case of the data generative model above for $\beta = 0$.

Hypothesis Class. We focus on learning two-layer neural networks defined as: $f(x; \mathbf{W}) := \frac{1}{\sqrt{m}} \sum_{s=1}^m a_s \phi(w_s, x)$ where m is an even integer representing the number of hidden nodes and $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an activation function. The weight matrix at the bottom layer is denoted as $\mathbf{W} = [w_1, \dots, w_m] \in \mathbb{R}^{d \times m}$ and the weight vector at the top layer by $a = [a_1, \dots, a_m] = [1, \dots, 1, -1, \dots, -1] \in \mathbb{R}^m$. The top layer weight vector a is kept fixed throughout the training process. The weight vectors at the bottom layer are initialized randomly as $w_s^0 \sim \mathcal{N}(0, \omega_{\text{init}}^2 \mathbf{I})$, for $s \in \{1, \dots, \frac{m}{2}\}$, and setting $w_s^0 = w_{s-\frac{m}{2}}^0$ for $s \in \{\frac{m}{2} + 1, \dots, m\}$. This ensures symmetry at initialization and yields $f(x; \mathbf{W}^0) = 0$ for all x . This symmetric initialization technique is commonly used in related work (Langer, 2021; Bartlett et al., 2021; Montanari and Zhong, 2022) and we employ here for analytical purposes.

Training Data. We are given a training data of size n sampled i.i.d. from the noisy data distribution, $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}$. Let \mathcal{C} denote the set of indices of training data corresponding to the clean labels; i.e., for $i \in \mathcal{C}$, we have that $(x_i, y_i) \sim \mathcal{D}_c$; similarly, let \mathcal{N} to denote the set of indices corresponding to noisy labels; i.e., $(x_i, -y_i) \sim \mathcal{D}_c$ for all $i \in \mathcal{N}$.

Loss Function. The 0-1 loss of a predictor $f(\cdot, \mathbf{W})$ on a data point (x, y) is defined as $\ell^{0/1}((x, y); \mathbf{W}) = \mathbb{1}(yf(x; \mathbf{W}) \leq 0)$, where $\mathbb{1}(\cdot)$ is the indicator function. For computational reasons, as is typical, we use the logistic loss, denoted $\ell(z) = \log(1 + \exp(-z))$, to train the two-layer neural networks. The population and the empirical loss w.r.t. $\ell(\cdot)$ are denoted as:

$$L(\mathbf{W}) := \mathbb{E}_{(x,y) \sim \mathcal{D}} \ell(yf(x; \mathbf{W})), \text{ and } \widehat{L}(\mathbf{W}) := \frac{1}{n} \sum_{i=1}^n \ell(y_i f(x_i; \mathbf{W})).$$

Robust Loss. We consider ℓ_2 norm-bounded adversarial attacks with a perturbation budget of size $\alpha > 0$. The set of all such perturbations for an input example $x \in \mathcal{X}$ is represented by $\mathcal{B}_2(x, \alpha)$. This threat model motivates minimizing the robust 0-1 loss defined as $\ell_{\text{rob}}^{0/1}((x, y); \mathbf{W}) =$

$\max_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} \mathbb{1}(yf(\tilde{x}; \mathbf{W}) \leq 0)$. The population and empirical risk w.r.t. the 0-1 loss and the robust 0-1 loss, respectively, are denoted as $L^{0/1}$, $\widehat{L}^{0/1}$, $L_{\text{rob}}^{0/1}$, and $\widehat{L}_{\text{rob}}^{0/1}$. Analogously, the population and empirical risk w.r.t. the (surrogate) logistic loss $\ell(\cdot)$ are defined as:

$$L_{\text{rob}}(\mathbf{W}) := \mathbb{E}_{(x, y) \sim \mathcal{D}} \max_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} \ell(yf(\tilde{x}; \mathbf{W})), \text{ and } \widehat{L}_{\text{rob}}(\mathbf{W}) := \frac{1}{n} \sum_{i=1}^n \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W})).$$

Note that we are ultimately interested in bounding the 0-1 loss and its robust variant.

Adversarial Training. The gradient descent-based adversarial training algorithm is presented in Algorithm 1. We denote the adversarial training example for some input x_i given model parameter \mathbf{W}^t , at round t as $\tilde{x}_i^t = \arg \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W}^t)) = \arg \min_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} y_i f(\tilde{x}_i; \mathbf{W}^t)$.

3 MAIN RESULT

3.1 SMOOTH ACTIVATION FUNCTION

In this section, we consider a strictly increasing, 1-Lipschitz, H -smooth activation function that is approximately homogeneous with $\phi(0) = 0$. Formally, there exists $\gamma, H > 0, 0 \leq \zeta < 1, c_1 \geq 0, c_2 \geq 0$ such that

$$0 < \gamma \leq \phi'(z) \leq 1, \phi'(z) \text{ is } H\text{-Lipschitz}, \text{ and } |\phi'(z) \cdot z - \phi(z)| \leq c_1 + c_2 |z|^\zeta, \forall z \in \mathbb{R}.$$

Smooth activation functions have been extensively studied both theoretically and empirically (Liu and Di, 2021; Biswas et al., 2022). One example of such an activation function that satisfies our condition is the smoothed Leaky ReLU activation (Frei et al., 2022) defined as follows:

$$\phi_{\text{SLReLU}}(z) = \begin{cases} z - \frac{1-\gamma}{4H}, & z \geq \frac{1}{H} \\ \frac{1-\gamma}{4} H z^2 + \frac{1+\gamma}{2} z, & |z| \leq \frac{1}{H} \\ \gamma z - \frac{1-\gamma}{4H}, & z \leq -\frac{1}{H} \end{cases}. \quad (1)$$

However, we do need an additional assumption on top of what Frei et al. (2022) require. In particular, we assume that $\phi'(z)z$ and $\phi(z)$ are close to each other. We argue that this is a mild assumption, and holds trivially for standard ReLU and Leaky ReLU, with $c_1 = c_2 = 0$. For $\phi_{\text{SLReLU}}(z)$, of Frei et al. (2022), the assumption holds with $\zeta = 0$ with $c_1 = \frac{1-\gamma}{4H}$, and $c_2 = 0$. The reason we need this additional assumption is because the neural networks with $\phi_{\text{SLReLU}}(z)$ activation function are no longer homogeneous. Consequently, without the assumption we end up with terms in the upper bound on the empirical robust risk that depends on the Frobenius norm of the weight matrix (see Section 4.2 for more details).

We make the following set of assumptions about our problem setup. Specifically, we consider a high dimensional setting where the dimension d is much larger than the number of training samples n , as stated below in Assumption (A2). Such a regime is popular in biomedical settings where the data comes from limited patient information such as MRI or DNA sequence. Assumption (A6) requires a small initialization to ensure that the first step of adversarial training dominates the behavior of the neural network, pushing it beyond the lazy training regime. Such initialization technique has also been introduced in previous work (Ba et al., 2019; Xing et al., 2021). Given that the objective of adversarial training is to achieve a classifier that is robust against small input perturbations imperceptible to human eyes, Assumption (A7) is reasonable as it imposes a mild constraint on the attack strength. Finally, we note that when $\alpha = 0$, these assumptions are essentially the same as in Frei et al. (2022).

Assumption 1. Let $\delta \in (0, 1/2)$. We assume that there exists a positive constant C such that the following holds: (A1) The number of samples satisfies $n > C \log(1/\delta)$. (A2) The dimension satisfies $C \max\{\|\mu\|^2 n, n^2 (\log(n/\delta) + \alpha^2)\} \leq d \leq \|\mu\|^4 / C$. (A3) The signal size satisfies $\|\mu\|^2 \geq C \log(n/\delta)$. (A4) noise rate $\beta \in [0, 1/C]$. (A5) Step size $\eta \leq (Cd^2(1 + \frac{H}{\sqrt{m}})^2)^{-1}$. (A6) Initialization variance satisfies $\omega_{\text{init}} \sqrt{md} \leq \eta$. (A7) Adversarial perturbation $\alpha \leq 0.99 \|\mu\|$.

Next we present our main result of this section that describes the effects of adversarial training on a neural network with smooth activation functions trained on samples from the noisy distribution \mathcal{D} (see Section 2.1). Our findings suggest that, as we run adversarial training for more epochs, the robust training loss goes to zero. Furthermore, the clean test error and the robust test error is approximately equal to the noise rate, provided that the attack strength, α , is small.

Theorem 3.1. Let $0 < \varepsilon \leq \frac{1}{2^n}$, $\delta \in (0, 1/2)$. Let ϕ be a γ -leaky H -smooth activation with $0 \leq \zeta < 1$. Let $\kappa \in (0, 1)$, $\lambda > 0$. Then, given that Assumption 1 holds with some constant $C > 0$, there exists a constant $c > 0$ such that after running Algorithm 1 for $T \geq \Omega\left(\left(\frac{1 + \sqrt{m/d^3}}{(199\|\mu\| - 200\alpha)\gamma\eta\varepsilon}\right)^{\frac{2}{1-\zeta}}\right)$ iterations, we have that with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample of size n , the following holds:

1. The robust training loss satisfies $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon$.
2. The clean test error satisfies $L^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(0.99 - \frac{\alpha}{\|\mu\|}\right)^2\right)$.
3. For $\frac{\alpha}{\|\mu\|} \leq \frac{0.99\sqrt{n}\|\mu\|}{\sqrt{n}\|\mu\| + C\sqrt{d}}$, the robust test error satisfies

$$L_{\text{rob}}^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-c\lambda \|\mu\|^2 \left(\frac{\sqrt{n}\|\mu\|}{C\sqrt{d}} \left(0.99 - \frac{\alpha}{\|\mu\|}\right) - \frac{\alpha}{\|\mu\|}\right)^2\right).$$

For the smooth Leaky ReLU activation function of Frei et al. (2022), we have the following result.

Corollary 3.2. For any γ -leaky H -smooth ReLU activation ϕ_{SLReLU} defined in Equation (1), and for all $\kappa \in (0, 1)$, $\lambda > 0$, given Assumption 1 holds, we have that with probability at least $1 - 2\delta$ over the random initialization and the draws of the samples, the robust training loss satisfies

$$\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \mathcal{O}\left(\frac{1 + \sqrt{(1-\gamma)/H}m^{1/4}}{(199\|\mu\| - 200\alpha)\gamma\sqrt{\eta}\sqrt{T}}\right).$$

3.2 NON-SMOOTH ACTIVATION FUNCTION

Here, we consider a more practical setting where the activation function is no longer smooth. We consider a homogeneous non-smooth activation function that satisfies the following properties.

$$\phi(0) = 0, \phi'(z)z = \phi(z), z \in \mathbb{R}; \quad 0 \leq \phi'(z) \leq 1, z \in \mathbb{R}; \quad \phi'(z) \geq \gamma, z \geq 0, \gamma \in (0, 1].$$

This includes ReLU and Leaky ReLU activation functions. Additionally, we assume the following.

Assumption 2. Let $\delta \in (0, 1/2)$. We assume that there exists a positive constant C such that the following holds: (B1) The network width satisfies $m \geq C \log(n/\delta)$. (B2) The signal size satisfies $\|\mu\| \geq C \max\left\{\left(\frac{d}{n} \log(md/n\delta)\right)^{1/4}, \sqrt{\log(n/\delta)}\right\}$. (B3) The dimension satisfies $d \geq C \max\{\|\mu\|^2 n, n^2 (\log(n/\delta) + \alpha^2)\}$. (B4) noise rate $\beta \in [0, 1/C]$. (B5) Initialization variance satisfies $\omega_{\text{init}}\sqrt{md} \leq \eta$. (B6) Step size $\eta \leq (Cd^2)^{-1}$. (B7) The number of samples satisfies $n \geq C \log(m/\delta)$. (B8) Adversarial perturbation $\alpha \leq \sqrt{n/d}\|\mu\|$.

Assumption (B1) is a relatively mild constraint on the network width. Assumption (B2) is slightly more stringent compared to Assumption (A3). However, it is worth noting that in the clean setting, the minimax generalization error is at least $\mathcal{O}\left(\exp\left(-\min\left(\|\mu\|^2, n\|\mu\|^4/d\right)\right)\right)$ (Giraud and Verzelen, 2019), implying that Assumption(B2) is unavoidable up to logarithmic factors if we desire a classifier with good generalization. Assumptions (B7) and (B8) are also more restrictive compared to Assumptions (A1) and (A7), respectively. These assumptions ensure the presence of sufficient number of neurons to have positive activation at the initial stage of adversarial training, which is a crucial aspect of our analysis in terms of relaxing the requirement of a smooth activation function, as opposed to Section 3.1. The analogous result to Theorem 3.1 is presented below.

Theorem 3.3. Let $0 < \varepsilon \leq \frac{1}{2^n}$, $\delta \in (0, 1/2)$. Let ϕ be a non-smooth activation with $\gamma \in (0, 1]$. Let $\kappa \in (0, 1)$, $\lambda > 0$. Then, given that Assumption 2 holds with some constant $C > 0$, there exists a constant $c > 0$ such that after running Algorithm 1 for $T \geq \Omega\left(\left((199\|\mu\| - 200\alpha)\gamma\sqrt{\eta}\varepsilon\right)^{-2}\right)$ iterations, we have that with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample of size n , the following holds:

1. The robust training loss satisfies $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon$.
2. The clean test error satisfies $L^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(1 - \frac{\alpha}{\|\mu\|}\right)^2\right)$.
3. For $\frac{\alpha}{\|\mu\|} \leq \frac{\sqrt{n}\|\mu\|}{\sqrt{n}\|\mu\| + C\sqrt{d}}$, the robust test error satisfies

$$L_{\text{rob}}^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-c\lambda \|\mu\|^2 \left(\frac{\sqrt{n}\|\mu\|}{C\sqrt{d}} \left(1 - \frac{\alpha}{\|\mu\|}\right) - \frac{\alpha}{\|\mu\|}\right)^2\right).$$

3.3 DISCUSSION

Theorems 3.1 and 3.3 suggest an interesting interplay between the parameters d , n , and $\|\mu\|$ as described in Assumptions 1 and 2. Importantly, when $n \geq \tilde{\Omega}\left(\frac{d \max(1, \alpha^2)}{\|\mu\|^2 (\|\mu\| - \alpha)^2}\right)$, it ensures a small robust test error. Furthermore, when $n \geq \tilde{\Omega}\left(\frac{d}{\|\mu\|^2 (\|\mu\| - \alpha)^2}\right)$, the clean test error is also guaranteed to be small. In cases where $\alpha = 0$, Theorem 3.1 and 3.3 recover the results in the standard setting (Frei et al., 2022; Xu and Gu, 2023). Compared to the standard setting, we pay an additional price proportional to $\frac{\max(1, \alpha^2)}{(1 - \alpha/\|\mu\|)^2}$ in terms of the sample size. It is worth noting that both the clean test error and the robust test error decrease as n/d increases or the attack strength $\frac{\alpha}{\|\mu\|}$ decreases, which is consistent with the findings in previous literature (Schmidt et al., 2018; Shafahi et al., 2019a).

Next, we provide a lower bound on the robust test error that is independent of the algorithm as well as the hypothesis class.

Theorem 3.4. We consider independent label flip with probability β . Let $p(x)$ be the density function of $\mathcal{D}_{\text{clust}}$. For any given classifier $f(\cdot; \mathbf{W})$, when $\alpha < \|\mu\|$, we have $L_{\text{rob}}^{0/1}(\mathbf{W}) \geq \beta + \frac{1-2\beta}{4} \int_{\mathbb{R}^d} \min\{p(\xi), p(\xi + \nu)\} d\xi$, where $\nu = 2(1 - \alpha/\|\mu\|)\mu$. When $\alpha \geq \|\mu\|$, the robust test error satisfies $L_{\text{rob}}^{0/1}(\mathbf{W}) \geq 0.5$.

Consider the special instance of when $\mathcal{D}_{\text{clust}}$ is a standard Gaussian distribution. Theorem 3.4 recovers the optimal risk in Dobriban et al. (2020) up to a scaling factor when $\beta = 0$. Moreover, the upper bound on the robust test error (denoted as UBD) that we provide in Theorems 3.1 and 3.3 and the lower bound (denoted as LBD) in 3.4 satisfy the following relationship: $(\text{UBD} - \beta) = (\text{LBD} - \beta)^{\mathcal{O}(n\|\mu\|^2/d)}$. When $\frac{n\|\mu\|^2}{d} = \Omega(1)$, $\alpha \leq \mathcal{O}(\|\mu\|)$, our upper bound roughly matches the lower bound.

Overfitting with Adversarial Training. Recent empirical studies have observed overfitting with adversarial training, wherein the robust training loss continues to decrease with the number of epochs, whereas the robust test error first decreases and then starts increasing (Rice et al., 2020). While our result may, at first, seem in conflict with this empirical observation, we note that there is actually no contradiction since we consider a specific data-generative model and a bound on the size of the adversarial perturbation during adversarial training. Indeed, recent empirical studies by Dong et al. (2021) and Yu et al. (2022) confirm that small α prevents adversarial training from overfitting. Furthermore, Xing et al. (2022) explored the phase transition between standard training and adversarial training and showed that the optimization trajectories in the two settings are close to each other when α is small. One interesting future direction is to justify the generalization guarantee for moderately large attack strength $\frac{\alpha}{\|\mu\|}$.

Comparison with Theoretical Works Several recent works focus on giving convergence and generalization guarantees for adversarial training (Gao et al., 2019; Zhang et al., 2020; Mianjy and Arora, 2022; Li and Telgarsky, 2023); here we compare and contrast our work with each of these.

The work of Gao et al. (2019) prove convergence for a modified algorithm for adversarial training wherein the iterates are projected onto a norm ball to ensure that the network weights stay close to initialization. However, they further need to assume that a robust network exists in the vicinity of the initialization. Such an assumption has been shown to be invalid in a recent work (Wang et al., 2022). In a related work, Zhang et al. (2020) provide a fine-grained convergence analysis for datasets that are well-separated. More recently, Li and Telgarsky (2023) give convergence and generalization guarantees for adversarial training of shallow networks with early stopping. Unfortunately, all of the aforementioned works are limited to the lazy regime (aka, the NTK setting) which has been shown to be at odds with adversarial robustness (Wang et al., 2022). Mianjy and Arora (2022) were the first to

provide both convergence and generalization guarantees beyond the NTK regime, yet their analysis was restricted to robust realizable data distributions.

Our work stands out from prior work in several ways. First, we study the standard adversarial training algorithm commonly used in practice. Second, we do not make restrictive assumptions regarding data separability; our generative model allows for the data to be non-separable. Finally, our results hold for neural networks of arbitrary width and can be trained for arbitrary many iterations allowing $\|\mathbf{W}^t\|$ to go to infinity, i.e., beyond the NTK regime. The following result shows that for certain step sizes and initialization, the neural network weights move far from the initialization after the first step of adversarial training based on gradient descent.

Proposition 3.5. Consider the same setting as in Theorem 3.1. Then, for some absolute constant $C > 1$, with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample, we have that $\frac{\|\mathbf{W}^1 - \mathbf{W}^0\|_F}{\|\mathbf{W}^0\|_F} \geq \frac{\gamma(199\|\mu\| - 200\alpha)}{1000}$.

Finally, we note that Dan et al. (2020) establish a minimax-type lower bound for the classification excess risk in the conditional Gaussian model, with a bound of $\Omega_P \left(\exp\left(-\left(\frac{1}{8} + o(1)\right)r^2\right)\frac{d}{n}\right)^1$, where r is the adversarial signal-to-noise ratio; this bound is shown to be achieved by a plug-in linear estimator. While useful, their result does not elucidate why adversarial training helps train robust networks. It also remains to be seen if adversarial training can achieve a matching upper bound.

4 PROOF SKETCH

We begin by providing some intuition for our proof. We show that when the perturbation size is relatively small, the trajectory of the adversarial training remains close to that of the standard training. Furthermore, given a good initialization of the neural network the dynamics of the training algorithm can be shown to be nearly linear. We also leverage a result from high dimensional probability, that the training data we draw is (nearly) separable even though the underlying data distribution is non-separable. We show that both of these events happen with high probability and establish what we refer to as a “good” run of the algorithm and are central to our proof.

Next, we formalize this intuition and provide a brief proof sketch of our main result. We focus primarily on neural networks with smooth activation function (i.e., Theorem 3.1) and note the differences in the analysis when extending the result to the non-smooth activation functions. In our analysis, we borrow many ideas from Frei et al. (2022) and Xu and Gu (2023). However, the extension is not straightforward and our focus in this section is on highlighting the technical challenges we overcome and the key insights we utilized in our analysis. For detailed proofs, we refer the reader to the Appendix.

4.1 GENERALIZATION GUARANTEE

As a proof strategy we seek to get an upper bound on the robust test error in terms of a lower bound on the normalized expected conditional margin. This follows using a concentration argument given that $\mathcal{D}_{\text{clust}}$ is λ -strongly log-concave.

Lemma 4.1. Suppose that $\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = \bar{y}] - \|\mathbf{W}\|_2 \alpha \geq 0$ holds for both $\bar{y} = 1$ and $\bar{y} = -1$. Then, there exists a universal constant $c > 0$ such that

$$L_{\text{rob}}^{0/1}(\mathbf{W}) \leq \beta + \sum_{\bar{y} \in \{-1, +1\}} \exp\left(-c\lambda \left(\frac{\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = \bar{y}] - \alpha}{\|\mathbf{W}\|_2}\right)^2\right)$$

Next, we need to show that the assumption in Lemma 4.1 does indeed hold for our setting. Here, we leverage the smoothness property of the activation function to derive a lower bound on the increment in the un-normalized margin for an independent test example (x, y) .

Lemma 4.2 (Informal). For some constant C_2 , with high probability, we have for any $t \geq 0$ and $(x, y) \in \mathbb{R}^d \times \{\pm 1\}$, there exist $\tilde{\rho}_i^t = \rho(\mathbf{W}^t, \tilde{\mathbf{x}}_i^t, x) \in [\gamma^2, 1]$ such that

$$y [f(x; \mathbf{W}^{t+1}) - f(x; \mathbf{W}^t)] \geq \frac{\eta}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \left(\tilde{\xi}_i^t \langle y_i \tilde{\mathbf{x}}_i^t, y \mathbf{x} \rangle - \frac{H \|\mathbf{x}\|^2 C_2^2 d \eta}{2\sqrt{mn}} \right).$$

¹For a sequence of random variables, X_n , and corresponding constants c_n , $X_n = \Omega_P(c_n)$ denotes that c_n/X_n converges to zero in probability as $n \rightarrow \infty$.

where $\tilde{g}_i(\mathbf{W}^t) = -\ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) = 1/(1 + \exp(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)))$.

For the non-smooth activation function, we get a similar result which we defer to the Appendix due to space constraints. Finally, we seek a positive lower bound on un-normalized expected conditional margin for model \mathbf{W}^t by expressing it in terms of the cumulative increments of margin; i.e., showing $\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c | y_c=1} [y_c f(x; \mathbf{W}^t)] = \sum_{t=1}^T \mathbb{E}_{(x, y_c) \sim \mathcal{D}_c | y_c=1} [y_c f(x; \mathbf{W}^t) - y_c f(x; \mathbf{W}^{t-1})] + \mathbb{E}_{(x, y_c) \sim \mathcal{D}_c | y_c=1} [y_c f(x; \mathbf{W}^0)]$. A positive lower bound holds trivially positive if $\langle y_i \tilde{\mathbf{x}}_i^t, y_c \mathbf{x} \rangle$ is always bounded below by some positive constant. However, due to the presence of noisy labels y_i and adversarial examples $\tilde{\mathbf{x}}_i$, $\langle y_i \tilde{\mathbf{x}}_i^t, y_c \mathbf{x} \rangle$ may be negative. Note, though, that the term $\langle y_i \tilde{\mathbf{x}}_i^t, y_c \mathbf{x} \rangle$ scales with $\tilde{g}_i(\mathbf{W}^t)$. If we can show that $\tilde{g}_i(\mathbf{W}^t)$ is of the same order across all training examples, and assume a small perturbation budget and that only a small fraction of labels are noisy, then we can mitigate the effect of the negative terms. The key lemma providing such a result by bounding the loss ratio is as follows.

Lemma 4.3 (Informal). Given Assumption 1, there is an absolute constant $C_r > 0$ such that with high probability, we have for all $t \geq 0$, $\max_{i, j \in [n]} \frac{\tilde{g}_i(\mathbf{W}^t)}{\tilde{g}_j(\mathbf{W}^t)} \leq C_r$.

To see why the above holds, note that for any given $i, j \in [n]$, we have that $\frac{\tilde{g}_i(\mathbf{W}^t)}{\tilde{g}_j(\mathbf{W}^t)} \leq \max \left\{ 2, \frac{2 \exp(-y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \right\}$, where $\tilde{\mathbf{x}}_i^t = \arg \min_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(x_i; \alpha)} y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)$. For successive iterates we get that $\frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^{t+1}; \mathbf{W}^{t+1}))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^{t+1}; \mathbf{W}^{t+1}))} \leq \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \cdot \frac{\exp(y_i f(\tilde{\mathbf{x}}_i^{t+1}; \mathbf{W}^t) - y_i f(\tilde{\mathbf{x}}_i^{t+1}; \mathbf{W}^{t+1}))}{\exp(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t) - y_j f(\tilde{\mathbf{x}}_j^{t+1}; \mathbf{W}^{t+1}))}$. Finally, we use induction to complete the proof.

For smooth activation function, the proof of Lemmas 4.2 and 4.3, follows by controlling the term $y [f(x; \mathbf{W}^{t+1}) - f(x; \mathbf{W}^t)]$ via Taylor approximation. For non-smooth activation functions, we need to ensure that there exist enough neurons have positive activations at initialization.

Remark 4.4. We can modify Assumption 2 by allowing the network initialization to depend on the training data: $a_s \mathbf{w}_s^0 = \frac{\hat{\mu}}{\|\hat{\mu}\|} \omega_{\text{init}} \sqrt{d}$ where $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n y_i x_i$. Then, Assumption (B8) can be relaxed to allow $\alpha \leq \mathcal{O}(\|\mu\|)$. Under Assumption 2 with the above modifications, Lemma 4.5 is still applicable and therefore Theorem 3.3 continues to hold.

Lemma 4.5 (Informal). Given Assumption 2, with high probability, for all $s \in [m]$, we have $|\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| = \Theta(n)$; for all $i \in [n]$, we have $|\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| = \Theta(m)$.

We further show that the number of positive neurons remains large throughout the training process.

4.2 CONVERGENCE GUARANTEE

In order to control the robust training loss, a naive approach would be to decouple the increment of the robust training loss, from iterate t to $t + 1$, into two terms as follows:

$$\begin{aligned} \widehat{L}_{\text{rob}}(\mathbf{W}^{t+1}) - \widehat{L}_{\text{rob}}(\mathbf{W}^t) &= \frac{1}{n} \sum_{i=1}^n [(\ell(y_i f(\tilde{\mathbf{x}}_i^{t+1}; \mathbf{W}^{t+1})) - \ell(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^{t+1}))) + (\ell(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^{t+1})) - \ell(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)))] . \end{aligned}$$

The second term can be controlled by the smoothness property of the loss function. The first term, unfortunately, is upper bounded by $\|\mathbf{W}^{t+1}\| \|\tilde{\mathbf{x}}_i^{t+1} - \tilde{\mathbf{x}}_i^t\|$, and the robust training loss hence inevitably depends on the norm of iterates $\|\mathbf{W}^{t+1}\|$ if no additional assumptions are made. This poses a problem if we do not constrain the model weights within a bounded domain, as $\|\mathbf{W}^t\|$ may tend to infinity as the number of epochs increases. To mitigate this issue, we instead control the robust training loss via the norm of the iterates. Specifically, we first show that $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \frac{2}{T} \sum_{t=0}^{T-1} G_{\text{rob}}(\mathbf{W}^t)$ where $G_{\text{rob}}(\mathbf{W}) := \frac{1}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(x_i, \alpha)} -\ell'(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}))$; this holds due to a property of the loss $\ell(\cdot)$ (see the Appendix for more details). We then bound $G_{\text{rob}}(\mathbf{W}^t)$ by a constant scaling of $\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle$, where $\mathbf{V} \in \mathbb{R}^{m \times d}$ is a matrix with row $\mathbf{v}_s = a_s \mu / \|\mu\|$. We achieve this result using Lemma 4.3 and the fact that only a small fraction labels are noisy. Given $\sum_{t=0}^{T-1} \langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle = \langle \mathbf{W}^T, \mathbf{V} \rangle - \langle \mathbf{W}^0, \mathbf{V} \rangle \leq \|\mathbf{W}^T\|_F + \|\mathbf{W}^0\|$, the only thing we need to

prove is that the growth rate of $\|W_T\|$ is smaller than $\mathcal{O}(T)$. This property holds for both smooth activation functions that satisfy our construction and non-smooth activation functions such as ReLU and Leaky ReLU.

5 EXPERIMENTS

In this section, we present a simple empirical study on a synthetic dataset to support our theoretical results. We follow the generative model in Section 2 to synthesize a dataset with independent label flips when generating y from y_c . We set $\mu = \|\mu\|_2 [1, 0, 0, \dots, 0]^\top$, $\beta = 0.1$, and generate $n = 100$ training samples and 2K test samples with the noise vector sampled from the standard multivariate Gaussian distribution, $\xi \sim \mathcal{N}(0, I)$. We train a two-layer ReLU network with width 1K. We use the default initialization in PyTorch and train the network applying full-batch gradient-descent based adversarial training using logistic loss for 1K iterations. We use PGD attack to generate adversarial examples with attack strength $\alpha / \|\mu\|$ and attack stepsize $\alpha/5 \|\mu\|$ for 20 iterations. The outer minimization is trained using an initial learning rate of 0.1 with decay by 10 after training for every 500 iterations. We note that adversarial training achieves 100% robust training accuracy. We estimate the robust test accuracy using the same PGD attack. We consider settings with varying dimension d and attack strength $\frac{\alpha}{\|\mu\|}$.




Figure 1: Clean test accuracy (left) / robust test accuracy (right) as a function of signal size $\|\mu\|$ and dimension d , for a fixed perturbation ratio $\alpha / \|\mu\| = 0.1$.

Figure 2: Robust test accuracy as a function of d and $\frac{\alpha}{\|\mu\|}$ for a fixed $\|\mu\| = 5$.

For our first experiment, we fix the perturbation ratio $\frac{\alpha}{\|\mu\|} = 0.1$, and vary the value of the signal strength $\|\mu\|$ from 1 to 10 and the dimension d from 1K to 18K. We show the results in Figure 1 as a heat map of clean accuracy and robust accuracy averaged over ten independent random runs. We observe a phase transition for both clean accuracy and robust accuracy at the value of dimension d around $\mathcal{O}(\|\mu\|^4)$ for clean accuracy and $\mathcal{O}(\|\mu\|^2)$ for robust accuracy. This is consistent with the main theorems (see discussion in Section 3.3).

For our next experiment, we fix the signal size $\|\mu\| = 5.0$, vary dimension d from 500 to 6K and perturbation ratio $\frac{\alpha}{\|\mu\|}$ from 0.05 to 0.45. Figure 2 plots the robust accuracy as a heat map averaged over ten independent runs. Our findings indicate that, increasing the dimension leads to a smaller perturbation ratio required to achieve the same level of robust test accuracy.

We observe the same trends on the MNIST dataset even though the data generative assumptions are no longer valid. We defer a detailed discussion of experiments on MNIST to the Appendix.

6 CONCLUSION

We presented the convergence and generalization guarantees for adversarial training of two-layer neural networks of arbitrary width under a non-separable data distribution. Our work suggests several promising future directions. Our results assume a generative model with a structured log-concave data distribution. It is natural to explore whether our findings can be extended to more general data distributions. Another interesting direction is to investigate whether our results generalize to the setting where the data dimension and the number of training samples have the same scale. Finally, we note that our main result only partially characterizes the phase transition from small to large test errors for small and large attack strengths, respectively. An important next step is to provide generalization guarantees for attacks of moderate strength and to explore the relationship between the perturbation size, signal size, dimension, and the number of training samples.

REFERENCES

- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR, 2019.
- Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu. A convergence analysis of gradient descent for deep linear neural networks. In *International Conference on Learning Representations*, 2019.
- Jimmy Ba, Murat Erdogdu, Taiji Suzuki, Denny Wu, and Tianzong Zhang. Generalization of two-layer neural networks: An asymptotic viewpoint. In *International conference on learning representations*, 2019.
- Peter L Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- Koushik Biswas, Sandeep Kumar, Shilpak Banerjee, and Ashish Kumar Pandey. Smooth Maximum Unit: Smooth Activation Function for Deep Networks using Smoothing Maximum Technique. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 794–803, 2022.
- Sébastien Bubeck, Yeshwanth Cherapanamjeri, Gauthier Gidel, and Remi Tachet des Combes. A single gradient step finds adversarial examples on random two-layers neural networks. *Advances in Neural Information Processing Systems*, 34:10081–10091, 2021.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. *Advances in neural information processing systems*, 32, 2019.
- Yuan Cao, Quanquan Gu, and Mikhail Belkin. Risk bounds for over-parameterized maximum margin classification on sub-gaussian mixtures. *Advances in Neural Information Processing Systems*, 34: 8407–8418, 2021.
- Yuan Cao, Zixiang Chen, Mikhail Belkin, and Quanquan Gu. Benign Overfitting in Two-layer Convolutional Neural Networks. *arXiv preprint arXiv:2202.06526*, 2022.
- Zachary Charles, Shashank Rajput, Stephen Wright, and Dimitris Papailiopoulos. Convergence and margin of adversarial training on separable data. *arXiv preprint arXiv:1905.09209*, 2019.
- Niladri S Chatterji and Philip M Long. Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *The Journal of Machine Learning Research*, 22(1):5721–5750, 2021.
- Jinghui Chen, Yuan Cao, and Quanquan Gu. Benign Overfitting in Adversarially Robust Linear Classification. *arXiv preprint arXiv:2112.15250*, 2021.
- Chen Dan, Yuting Wei, and Pradeep Ravikumar. Sharp statistical guarantees for adversarially robust gaussian classification. In *International Conference on Machine Learning*, pages 2345–2355. PMLR, 2020.
- Edgar Dobriban, Hamed Hassani, David Hong, and Alexander Robey. Provable tradeoffs in adversarially robust classification. *arXiv preprint arXiv:2006.05161*, 2020.
- Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 321–331, 2020.
- Yinpeng Dong, Ke Xu, Xiao Yang, Tianyu Pang, Zhijie Deng, Hang Su, and Jun Zhu. Exploring memorization in adversarial training. In *International Conference on Learning Representations*, 2021.
- Spencer Frei, Niladri S Chatterji, and Peter Bartlett. Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. In *Conference on Learning Theory*, pages 2668–2703. PMLR, 2022.

- Ruiqi Gao, Tianle Cai, Haochuan Li, Cho-Jui Hsieh, Liwei Wang, and Jason D Lee. Convergence of adversarial training in overparametrized neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Christophe Giraud and Nicolas Verzelen. Partial recovery bounds for clustering with the relaxed k -means. *Mathematical Statistics and Learning*, 1(3):317–374, 2019.
- Adel Javanmard, Mahdi Soltanolkotabi, and Hamed Hassani. Precise tradeoffs in adversarial training for linear regression. In *Conference on Learning Theory*, pages 2034–2078. PMLR, 2020.
- Yiwen Kou, Zixiang Chen, Yuanzhou Chen, and Quanquan Gu. Benign Overfitting for Two-layer ReLU Networks. *arXiv preprint arXiv:2303.04145*, 2023.
- Sophie Langer. Analysis of the rate of convergence of fully connected deep neural network regression estimates with smooth activation function. *Journal of Multivariate Analysis*, 182:104695, 2021.
- Justin D Li and Matus Telgarsky. On achieving optimal adversarial test error. In *International Conference on Learning Representations*, 2023.
- Yan Li, Ethan Fang, Huan Xu, and Tuo Zhao. Implicit bias of gradient descent based adversarial training on separable data. In *International Conference on Learning Representations*, 2020.
- Xinyu Liu and Xiaoguang Di. TanhExp: A Smooth Activation Function with High Convergence Speed for Lightweight Neural Networks. *IET Computer Vision*, 15(2):136–150, 2021.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.
- Poorya Mianjy and Raman Arora. Robustness guarantees for adversarially trained neural networks. 2022.
- Andrea Montanari and Yiqiao Zhong. The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *The Annals of Statistics*, 50(5):2816–2847, 2022.
- Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *International Conference on Learning Representations*, 2021.
- Leslie Rice, Eric Wong, and Zico Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*, pages 8093–8104. PMLR, 2020.
- Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. *Advances in neural information processing systems*, 31, 2018.
- Ali Shafahi, W Ronny Huang, Christoph Studer, Soheil Feizi, and Tom Goldstein. Are adversarial examples inevitable? In *International Conference on Learning Representations*, 2019a.
- Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *Advances in Neural Information Processing Systems*, 32, 2019b.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2020.
- Yunjuan Wang, Enayat Ullah, Poorya Mianjy, and Raman Arora. Adversarial robustness is at odds with lazy training. *Advances in Neural Information Processing Systems*, 35:6505–6516, 2022.
- Yue Xing, Qifan Song, and Guang Cheng. On the generalization properties of adversarial training. In *International Conference on Artificial Intelligence and Statistics*, pages 505–513. PMLR, 2021.

- Yue Xing, Qifan Song, and Guang Cheng. Phase transition from clean training to adversarial training. In *Advances in Neural Information Processing Systems*, 2022.
- Xingyu Xu and Yuantao Gu. Benign overfitting of non-smooth neural networks beyond lazy training. *International Conference on Artificial Intelligence and Statistics*, pages 11094–11117, 2023.
- Chaojian Yu, Bo Han, Li Shen, Jun Yu, Chen Gong, Mingming Gong, and Tongliang Liu. Understanding robust overfitting of adversarial training and beyond. In *International Conference on Machine Learning*, pages 25595–25610. PMLR, 2022.
- Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019.
- Yi Zhang, Orestis Plevrakis, Simon S Du, Xingguo Li, Zhao Song, and Sanjeev Arora. Over-parameterized adversarial training: An analysis overcoming the curse of dimensionality. *Advances in Neural Information Processing Systems*, 33:679–688, 2020.
- Difan Zou, Spencer Frei, and Quanquan Gu. Provable robustness of adversarial training for learning halfspaces with noise. In *International Conference on Machine Learning*, pages 13002–13011. PMLR, 2021.

Supplementary Material

A ADDITIONAL EXPERIMENTS

In order to see if our results extend beyond the generative data model we consider in this paper, we run the same set of experiments as above on the MNIST dataset. MNIST is a dataset of 28×28 greyscale handwritten digits. We extract examples corresponding to images of the digits ‘0’ and ‘1’, resulting in 12,665 training examples and 2,115 test examples. We view the input image as a vector input of size d and normalize the data to ensure that the ℓ_2 norm of each input vector is equal to $\|\mu\|$. A random subset of size $n = 100$ is used for training. We do not introduce any label noise; i.e., $\beta = 0$. We train a two-layer ReLU network with width 1K using the same training procedure as for the experiments on the synthetic data.




Figure 3: Clean test accuracy (left) / robust test accuracy (right) on MNIST dataset as a function of signal size $\|\mu\|$ and dimension d , for a fixed perturbation ratio $\alpha / \|\mu\| = 0.3$

The perturbation ratio is set to $\frac{\alpha}{\|\mu\|} = 0.3$, the signal size $\|\mu\|$ is varied from 0.1 to 2.0. We downsample the images by different factors to simulate data with dimension d ranging between 25 and 784. We plot the heat map for both the clean accuracy and the robust accuracy averaged over ten independent random runs in Figure 3. We observe a phase transition in both subplots at the value of dimension d around $\mathcal{O}(\|\mu\|^4)$ for clean accuracy and $\mathcal{O}(\|\mu\|^2)$ for robust accuracy. This confirms that even when the data distribution deviates from a Gaussian mixture model, the result in Theorem 3.1 is still indicative of an interesting relationship between d , $\|\mu\|$, $\frac{\alpha}{\|\mu\|}$ and n .

As for the synthetic data, for a second set of experiments, we fix the signal size to $\|\mu\| = 5.0$ and vary the dimension d from 25 to 784 and perturbation ratios $\frac{\alpha}{\|\mu\|}$ from 0.05 to 0.45. The resulting heat map of robust accuracy averaged over five independent runs is presented in Figure 4. We see a similar trend as in Figure 2.




Figure 4: Robust test accuracy (right) as a function of dimension d and perturbation ratio $\frac{\alpha}{\|\mu\|}$, for a fixed signal size $\|\mu\| = 5.0$ on MNIST

B MISSING PROOFS

We start by introducing some important notations that will be used throughout our proof. We find the negative derivative of the logistic loss to be useful in our discussion; we denote it as $g(z) := -\ell'(z) = 1/(1 + \exp(z))$. Note that $g(\cdot)$ is non-negative and decreasing and can serve as a surrogate for the 0-1 loss. More importantly, we can check that finding adversarial examples that maximize $\ell(\cdot)$ is equivalent to maximizing $g(\cdot)$, i.e., $\operatorname{argmax}_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W})) = \operatorname{argmax}_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} g(y_i f(\tilde{x}_i; \mathbf{W}))$. For simplicity, we denote $\tilde{\ell}_i(\mathbf{W})$ and $\tilde{g}_i(\mathbf{W})$ to represent

$\max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W}))$ and $\max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} g(y_i f(\tilde{x}_i; \mathbf{W}))$, respectively. The empirical risk and the robust empirical risk w.r.t. the surrogate loss $g(\cdot)$ are denoted as

$$\widehat{G}(\mathbf{W}) := \frac{1}{n} \sum_{i=1}^n g(y_i f(x_i; \mathbf{W})), \quad \widehat{G}_{\text{rob}}(\mathbf{W}) := \frac{1}{n} \sum_{i=1}^n \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} g(y_i f(\tilde{x}_i; \mathbf{W})).$$

B.1 MISSING PROOFS IN SECTION 3.1

Improvements over Frei et al. (2022). We have identified two non-rigorous arguments in the proof of Frei et al. (2022) and addressed them in our analysis. The first issue arises in the Lemma 4.1 of Frei et al. (2022), where the concentration inequality for the Lipschitz function class (Equation (2) in Frei et al. (2022)) is applied. However, the expectation should be taken with respect to x instead of (x, y) . To resolve this, we introduce Lemma B.1, conditioning on the label y_c , and apply the concentration argument twice. The second issue is found in the proof of Lemma 4.11, Equation 24 in Frei et al. (2022), specifically in the calculation of $\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c} [\xi_i \langle y_i x_i, y_c x \rangle]$. In their analysis, the expectation is taken only over $\langle y_i x_i, y_c x \rangle$, but it should also consider the dependence of ξ_i on x . In our analysis, presented in Lemma B.10, we provide a careful treatment of this expression, incorporating the additional assumption that $d \leq \frac{\|\mu\|^4}{C}$.

Theorem 3.1. Let $0 < \varepsilon \leq \frac{1}{2n}$, $\delta \in (0, 1/2)$. Let ϕ be a γ -leaky H -smooth activation with $0 \leq \zeta < 1$. Let $\kappa \in (0, 1)$, $\lambda > 0$. Then, given that Assumption 1 holds with some constant $C > 0$, there exists a constant $c > 0$ such that after running Algorithm 1 for $T \geq \Omega\left(\left(\frac{1 + \sqrt{m/d^3}}{(199\|\mu\| - 200\alpha)\gamma\eta\varepsilon}\right)^{\frac{2}{1-\zeta}}\right)$ iterations, we have that with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample of size n , the following holds:

1. The robust training loss satisfies $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon$.
2. The clean test error satisfies $L^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(0.99 - \frac{\alpha}{\|\mu\|}\right)^2\right)$.
3. For $\frac{\alpha}{\|\mu\|} \leq \frac{0.99\sqrt{n}\|\mu\|}{\sqrt{n}\|\mu\| + C\sqrt{d}}$, the robust test error satisfies

$$L_{\text{rob}}^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-c\lambda \|\mu\|^2 \left(\frac{\sqrt{n}\|\mu\|}{C\sqrt{d}} \left(0.99 - \frac{\alpha}{\|\mu\|}\right) - \frac{\alpha}{\|\mu\|}\right)^2\right).$$

Proof of Theorem 3.1. By Lemma B.2 and Lemma B.3, a good run occurs with probability at least $1 - 2\delta$. The robust training loss bound is proved in Lemma B.11. For the generalization guarantee, we apply Lemma B.1 with Lemma B.10, which give us with probability at least $1 - 2\delta$,

$$\begin{aligned} L_{\text{rob}}^{0/1}(\mathbf{W}^T) &= \mathbf{P}_{(x, y) \sim \mathcal{D}}[\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y \neq \text{sign}(f(\tilde{x}; \mathbf{W}^T))] \\ &\leq \beta + 2 \exp\left(-c\lambda \left(\frac{\gamma^2 \sqrt{n}}{8C_2 \sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right) - \alpha\right)^2\right) \\ &\leq \beta + 2 \exp\left(-c\lambda \left(\frac{\sqrt{n}}{C\sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right) - \alpha\right)^2\right), \quad (\text{Choose } C \geq \frac{8C_2}{\gamma^2}) \end{aligned}$$

where the last line holds for $\frac{\alpha}{\|\mu\|} \leq \frac{0.99\|\mu\|\sqrt{n}}{\sqrt{n}\|\mu\| + C\sqrt{d}}$, so that $\frac{\sqrt{n}}{C\sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right) - \alpha \geq 0$.

Applying Lemma B.1 gives us

$$\begin{aligned} L^{0/1}(\mathbf{W}^T) &\leq \beta + \exp\left(-c\lambda \left(\frac{\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c}[y_c f(x; \mathbf{W}^T) | y_c = 1]}{\|\mathbf{W}^T\|_2}\right)^2\right) \\ &\quad + \exp\left(-c\lambda \left(\frac{\mathbb{E}_{(x, y_c) \sim \mathcal{D}_c}[y_c f(x; \mathbf{W}^T) | y_c = -1]}{\|\mathbf{W}^T\|_2}\right)^2\right) \\ &\leq \beta + 2 \exp\left(-c\lambda \left(\frac{\gamma^2 \sqrt{n}}{8C_2 \sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right)\right)^2\right) \quad (\text{Lemma B.10}) \end{aligned}$$

$$\begin{aligned}
&\leq \beta + 2 \exp \left(-c\lambda \left(\frac{\sqrt{n}}{C\sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha \right) \right)^2 \right) \quad (\text{Choose } C \geq \frac{8C_2}{\gamma^2}) \\
&= \beta + 2 \exp \left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(0.99 - \frac{\alpha}{\|\mu\|} \right)^2 \right).
\end{aligned}$$

□

The proof of Theorem 3.1 builds upon a sequence of Lemmas, which we show below. Lemma B.1 bound the robust test error by the normalized expected conditional margin via a concentration argument.

Lemma B.1. Suppose that $\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = \bar{y}] - \|\mathbf{W}\|_2 \alpha \geq 0$ holds for both $\bar{y} = 1$ and $\bar{y} = -1$. Then, there exists a universal constant $c > 0$ such that

$$\begin{aligned}
L_{\text{rob}}^{0/1}(\mathbf{W}) &\leq \beta + \sum_{\bar{y} \in \{-1, +1\}} \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = \bar{y}]}{\|\mathbf{W}\|_2} - \alpha \right)^2 \right), \\
L^{0/1}(\mathbf{W}) &\leq \beta + \sum_{\bar{y} \in \{-1, +1\}} \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = \bar{y}]}{\|\mathbf{W}\|_2} \right)^2 \right).
\end{aligned}$$

Proof of Lemma B.1. We have

$$\begin{aligned}
&\mathbf{P}_{(x,y) \sim \mathcal{D}} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y \neq \text{sign}(f(\tilde{x}; \mathbf{W}))] \\
&= \mathbf{P}_{(x,y) \sim \mathcal{D}} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y f(\tilde{x}; \mathbf{W}) \leq 0] \\
&\leq \beta + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y_c f(\tilde{x}; \mathbf{W}) \leq 0] \\
&= \beta + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right].
\end{aligned}$$

For any $\tilde{x} \in \mathcal{B}_2(x, \alpha)$, we have

$$\begin{aligned}
|y_c f(x; \mathbf{W}) - y_c f(\tilde{x}; \mathbf{W})| &= \frac{1}{\sqrt{m}} \left| \sum_{s=1}^m a_s [\phi(\langle \mathbf{w}_s, x \rangle) - \phi(\langle \mathbf{w}_s, \tilde{x} \rangle)] \right| \\
&\leq \frac{1}{\sqrt{m}} \sum_{s=1}^m |a_s| |\langle \mathbf{w}_s, x - \tilde{x} \rangle| \quad (\phi \text{ is 1-Lipschitz}) \\
&\leq \frac{1}{\sqrt{m}} \sqrt{\sum_{s=1}^m a_s^2} \sqrt{\sum_{s=1}^m \langle \mathbf{w}_s, x - \tilde{x} \rangle^2} \quad (\text{Cauchy-Schwartz}) \\
&= \|\mathbf{W}(x - \tilde{x})\| \\
&\leq \|\mathbf{W}\|_2 \alpha. \quad (\text{By the definition of the spectral norm})
\end{aligned}$$

Since $\mathcal{D}_{\text{clust}}$ is λ -strongly log concave, and $y_c f(x; \mathbf{W})$ is $\|\mathbf{W}\|_2$ -Lipschitz, there is an absolute constant $\bar{c} > 0$ such that for any $q \geq 1$, $\|y_c f(x; \mathbf{W}) - \mathbb{E}[y_c f(x; \mathbf{W})]\|_{L^q} \leq \bar{c} \|\mathbf{W}\|_2 \sqrt{q/\lambda}$. Therefore, there is an absolute constant $c > 0$ such that for any $t \geq 0$, for fixed $y_c = 1$ (same for $y_c = -1$), we have

$$\mathbf{P}(y_c f(x; \mathbf{W}) - \mathbb{E}[y_c f(x; \mathbf{W})] \leq -t) \leq \exp \left(-c\lambda \left(\frac{t}{\|\mathbf{W}\|_2} \right)^2 \right). \quad (2)$$

where the expectation is w.r.t. x . Choose $t = \mathbb{E}[y_c f(x; \mathbf{W})] - \|\mathbf{W}\|_2 \alpha \geq 0$, we have

$$\begin{aligned}
&\mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \mid y_c = 1 \right) \\
&= \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(y_c f(x; \mathbf{W}) - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) | y_c = 1] \leq y_c f(x; \mathbf{W}) \right)
\end{aligned}$$

$$\begin{aligned}
& - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) \middle| y_c = 1 \right] - \min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \middle| y_c = 1 \Big) \\
& \leq \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(y_c f(x; \mathbf{W}) - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) \middle| y_c = 1 \right] \leq \|\mathbf{W}\|_2 \alpha \right. \\
& \quad \left. - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) \middle| y_c = 1 \right] \middle| y_c = 1 \right) \\
& \leq \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) \middle| y_c = 1 \right] - \|\mathbf{W}\|_2 \alpha}{\|\mathbf{W}\|_2} \right)^2 \right).
\end{aligned}$$

Consider both $y_c = 1$ and $y_c = -1$ gives us

$$\begin{aligned}
& \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right) \\
& = \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \middle| y_c = -1 \right) \cdot \mathbf{P}(y_c = 1) \\
& \quad + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \middle| y_c = -1 \right) \cdot \mathbf{P}(y_c = -1) \\
& \leq \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = 1] - \|\mathbf{W}\|_2 \alpha}{\|\mathbf{W}\|_2} \right)^2 \right) \\
& \quad + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = -1] - \|\mathbf{W}\|_2 \alpha}{\|\mathbf{W}\|_2} \right)^2 \right).
\end{aligned}$$

Similarly,

$$\begin{aligned}
& L^{0/1}(\mathbf{W}) \\
& = \mathbf{P}_{(x,y) \sim \mathcal{D}} [y f(x; \mathbf{W}) \leq 0] \\
& \leq \beta + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \leq 0] \\
& \leq \beta + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \leq 0 \middle| y_c = 1] + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \leq 0 \middle| y_c = -1] \\
& \leq \beta + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = 1] \leq -\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = 1] \middle| y_c = 1 \right] \\
& \quad + \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[y_c f(x; \mathbf{W}) - \mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = -1] \leq -\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = -1] \middle| y_c = -1 \right] \\
& \leq \beta + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = 1]}{\|\mathbf{W}\|_2} \right)^2 \right) + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}) \middle| y_c = -1]}{\|\mathbf{W}\|_2} \right)^2 \right).
\end{aligned}$$

□

Now we only need to derive a lower bound on the normalized expected conditional margin. Below is a series of structural results that leads us to our destination. Lemma B.2 and B.3 are the properties of initialized network weights as well as the generated data.

Lemma B.2 (Lemma 4.2 in Frei et al. (2022)). There is a universal constant $C_0 > 1$ such that with probability at least $1 - \delta/2$ over the random initialization,

$$\frac{1}{2} \omega_{\text{init}}^2 m d \leq \|\mathbf{W}^0\|_F^2 \leq \frac{3}{2} \omega_{\text{init}}^2 m d, \quad \|\mathbf{W}^0\|_2 \leq C_0 \omega_{\text{init}} (\sqrt{m} + \sqrt{d})$$

Proof of lemma B.2. For the Forbenius norm statement, note that $\sum_{s=1}^{\frac{m}{2}} \|\mathbf{w}_s^0\|_2^2$ is a ω_{init}^2 -multiple of a chi-squared random variable with $\frac{md}{2}$ degrees of freedom. By concentration of the χ^2 distribution, for any $t \in (0, 1]$,

$$\mathbb{P} \left(\left| \frac{1}{\frac{md}{2}\omega_{\text{init}}^2} \sum_{s=1}^{\frac{m}{2}} \|\mathbf{w}_s^0\|_2^2 - 1 \right| \geq t \right) \leq 2 \exp \left(-\frac{mdt^2}{8} \right).$$

In particular, if we choose $t = 1/2$, and notice that $\mathbf{w}_s = \mathbf{w}_{s+\frac{m}{2}}$,

$$\|\mathbf{W}^0\|_F^2 = 2 \sum_{s=1}^{\frac{m}{2}} \|\mathbf{w}_s^0\|_2^2 \in \left[\frac{1}{2}\omega_{\text{init}}^2 md, \frac{3}{2}\omega_{\text{init}}^2 md \right]$$

holds with probability at least

$$1 - 2 \exp(-md/64) \geq 1 - \delta/4. \quad (d \geq C \log(1/\delta) \text{ from Assumption (A2)})$$

For the spectral norm, since the entries of $[\mathbf{w}_1, \dots, \mathbf{w}_{\frac{m}{2}}]/\omega_{\text{init}}$ are i.i.d. standard normal variables, by Theorem 4.4.5 in Vershynin (2018), there exists a universal constant $c > 0$ such that for any $u \geq 0$, with probability at least $1 - 2 \exp(-u^2)$, we have

$$\|[\mathbf{w}_1, \dots, \mathbf{w}_{\frac{m}{2}}]\|_2 \leq c\omega_{\text{init}}(\sqrt{m/2} + \sqrt{d} + u).$$

In particular, taking $u = \sqrt{\log(8/\delta)}$, we have with probability at least $1 - \delta/4$, $\|[\mathbf{w}_1, \dots, \mathbf{w}_{\frac{m}{2}}]\|_2 \leq c\omega_{\text{init}}(\sqrt{m/2} + \sqrt{d} + \sqrt{\log(8/\delta)})$. Since $\|\mathbf{W}^0\|_2 = \sqrt{2} \|[\mathbf{w}_1, \dots, \mathbf{w}_{\frac{m}{2}}]\|_2$ holds by symmetric initialization, and $\sqrt{d} \geq \sqrt{C \log(1/\delta)} \geq \sqrt{\log(8/\delta)}$ by Assumption (A2) and C sufficiently large, we are done with the spectral norm. \square

Lemma B.3. Let $(\mathbf{x}_i, y_i) \sim \mathcal{D}, \forall i \in [n]$, where $\mathbf{x}_i = y_i^c \mu + \xi_i$, $\mathbb{P}(y_i^c \neq y_i) \leq \beta$. For all $0 < \kappa < 1$, there exists $C_1 = \frac{10}{\kappa} > 1$ such that for large enough C , with probability at least $1 - \delta$ over \mathcal{D}^n , the following hold

$$(C1) \quad \forall i \in [n], \frac{\kappa d}{2} \leq \|\xi_i\|^2 \leq (3 + \frac{\kappa}{2})d, \quad d/C_1 \leq \|\mathbf{x}_i\|^2 \leq C_1 d; \quad \forall \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}'_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha), \\ \left(\sqrt{d/C_1} - \alpha \right)^2 \leq \|\tilde{\mathbf{x}}_i\|^2 \leq (\sqrt{C_1 d} + \alpha)^2, \quad \langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}'_i \rangle \geq \left(\sqrt{d/C_1} - \alpha \right)^2.$$

$$(C2) \quad \forall i \neq j \in [n], |\langle \xi_i, \xi_j \rangle| \leq C_1 \left(\sqrt{d \log(n/\delta)} \right), \quad |\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \leq C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right).$$

$$\forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha), \tilde{\mathbf{x}}_j \in \mathcal{B}_2(\mathbf{x}_j, \alpha), |\langle \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j \rangle| \leq C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha\sqrt{C_1 d} + \alpha^2,$$

$$|\langle \mathbf{x}_i, \tilde{\mathbf{x}}_j \rangle| \leq C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + \alpha\sqrt{C_1 d}.$$

$$(C3) \quad \forall z_1, z_2, \dots, z_n \in \mathbb{R}, \left\| \sum_{i=1}^n z_i \xi_i \right\|^2 \leq 4d \sum_{i=1}^n z_i^2.$$

$$(C4) \quad \forall i \in \mathcal{C}, \left| \langle \mu, y_i \mathbf{x}_i \rangle - \|\mu\|^2 \right| \leq \|\mu\|^2 / 200, \quad \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha), \frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \leq \langle \mu, y_i \tilde{\mathbf{x}}_i \rangle \leq \\ \frac{201}{200} \|\mu\|^2 + \|\mu\| \alpha.$$

$$(C5) \quad \forall i \in \mathcal{N}, \left| \langle \mu, y_i \mathbf{x}_i \rangle + \|\mu\|^2 \right| \leq \|\mu\|^2 / 200, \quad \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha), -\frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \leq \langle \mu, y_i \tilde{\mathbf{x}}_i \rangle \leq \\ -\frac{199}{200} \|\mu\|^2 + \|\mu\| \alpha.$$

$$(C6) \quad \text{The number of noisy samples satisfies } |\mathcal{N}|/n \leq \beta + \sqrt{\frac{2}{C}}.$$

Proof of Lemma B.3. The proof is a simple extension of Lemma 13 in Chatterji and Long (2021). For statement (C1), $\frac{\kappa d}{2} \leq \|\xi_i\|^2 \leq (3 + \frac{\kappa}{2})d$ and $d/C_1 \leq \|\mathbf{x}_i\|^2 \leq C_1 d$ follows directly from the

proof of Lemma 19 in Chatterji and Long (2021). Since

$$\begin{aligned} \|x_i - \tilde{x}_i\| &\leq \alpha \leq 0.5 \|\mu\| && \text{(Assumption (A7))} \\ &\leq \frac{1}{2} \sqrt{\frac{d}{C_1}} && (d \geq C \|\mu\|^2 n \text{ from Assumption (A2) with sufficiently large } C) \\ &\leq \frac{1}{2} \|x_i\|, \end{aligned}$$

$(\sqrt{d/C_1} - \alpha)^2 \leq \|\tilde{x}_i\|^2 \leq (\sqrt{C_1 d} + \alpha)^2$ holds. Because $\|x_i - \tilde{x}'_i\| \leq \alpha \leq \frac{1}{2} \|x_i\|$ also holds, through some simple calculation, we have $\langle \tilde{x}_i, \tilde{x}'_i \rangle \geq (\|x_i\| - \alpha)^2 \geq (\sqrt{d/C_1} - \alpha)^2$.

For statement (C2), $|\langle \xi_i, \xi_j \rangle| \leq C_1 (\sqrt{d \log(n/\delta)})$ and $|\langle x_i, x_j \rangle| \leq C_1 (\|\mu\|^2 + \sqrt{d \log(n/\delta)})$ follows directly from the proof of Lemma 20 in Chatterji and Long (2021).

$$\begin{aligned} |\langle \tilde{x}_i, \tilde{x}_j \rangle| &= |\langle x_i, x_j \rangle + \langle \tilde{x}_i - x_i, x_j \rangle + \langle x_i, \tilde{x}_j - x_j \rangle + \langle \tilde{x}_i - x_i, \tilde{x}_j - x_j \rangle| \\ &\leq |\langle x_i, x_j \rangle| + \alpha \|x_j\| + \alpha \|x_i\| + \|\tilde{x}_i - x_i\| \cdot \|\tilde{x}_j - x_j\| \\ &\leq C_1 (\|\mu\|^2 + \sqrt{d \log(n/\delta)}) + 2\alpha \sqrt{C_1 d} + \alpha^2. \end{aligned}$$

$$\begin{aligned} |\langle x_i, \tilde{x}_j \rangle| &= |\langle x_i, x_j \rangle + \langle x_i, \tilde{x}_j - x_j \rangle| \\ &\leq |\langle x_i, x_j \rangle| + \alpha \|x_i\| \\ &\leq C_1 (\|\mu\|^2 + \sqrt{d \log(n/\delta)}) + \alpha \sqrt{C_1 d}. \end{aligned}$$

The statement (C3) holds since

$$\begin{aligned} \left\| \sum_{i=1}^n z_i \xi_i \right\|^2 &= \sum_{i=1}^n z_i^2 \|\xi_i\|^2 + 2 \sum_{i<j} z_i z_j \langle \xi_i, \xi_j \rangle \\ &\leq \sum_{i=1}^n z_i^2 (3 + \frac{\kappa}{2}) d + 2C_1 \sum_{i<j} \frac{z_i^2 + z_j^2}{2} \sqrt{d \log(n/\delta)} && \text{((C1), (C2))} \\ &\leq \sum_{i=1}^n z_i^2 (3 + \frac{\kappa}{2}) d + 2C_1 \sum_{i<j} \frac{z_i^2 + z_j^2}{2} \frac{d}{\sqrt{C}n} && \text{(Assumption (A2))} \\ &\leq (3 + \frac{\kappa}{2} + \frac{C_1}{\sqrt{C}}) d \sum_{i=1}^n z_i^2 \leq 4d \sum_{i=1}^n z_i^2. && \text{(C sufficiently large)} \end{aligned}$$

Statement (C4) and statement (C5) follow similarly from the proof of Lemma 21 and 22 in Chatterji and Long (2021) and combining the fact that

$$|\langle \mu, y_i x_i \rangle - \langle \mu, y_i \tilde{x}_i \rangle| \leq \|\mu\| \cdot \|x_i - \tilde{x}_i\| \leq \|\mu\| \alpha.$$

The last statement follows from Hoeffding's inequality:

$$\begin{aligned} &\mathbb{P}(|\mathcal{N}|/n - \beta > \sqrt{\frac{2}{C}}) \\ &\leq e^{-2n(\sqrt{\frac{2}{C}})^2} \\ &\leq e^{-2C \log(1/\delta)(\frac{2}{C})} && \text{(Assumption (A1))} \\ &= \delta^4 \\ &\leq \delta/6. && (\delta < 0.5) \end{aligned}$$

□

Definition. If the events in Lemma B.2 and Lemma B.3 occur, let us say that we have a good run.

Lemma B.2 and Lemma B.3 show that a good run occurs with probability at least $1 - 2\delta$. In the following we assume a good run occurs. Lemma B.4 leverages the smoothness property of activation function and derive the result via Taylor approximation. Lemma B.5 characterizes the relationship between $\widehat{L}_{\text{rob}}(\cdot)$ and $\widehat{G}_{\text{rob}}(\cdot)$. Lemma B.6 further derives the bounds on the gradient norm given adversarial training example, as well as the pairwise correlations of the gradients given different adversarial training examples. These are standard results that have been derived by Frei et al. (2022), and we simply extend them for adversarial training scenario.

Lemma B.4 (Lemma 4.5 in Frei et al. (2022)). For an H -smooth activation ϕ and any $\mathbf{W}, \mathbf{V} \in \mathbb{R}^{m \times d}$, and $\mathbf{x} \in \mathbb{R}^d$, we have

$$|f(\mathbf{x}; \mathbf{W}) - f(\mathbf{x}; \mathbf{V}) - \langle \nabla f(\mathbf{x}; \mathbf{V}), \mathbf{W} - \mathbf{V} \rangle| \leq \frac{H \|\mathbf{x}\|^2}{2\sqrt{m}} \|\mathbf{W} - \mathbf{V}\|_2^2.$$

Lemma B.5. Let $C_1 > 1$ be the constant from Lemma B.3. For an H -smooth activation ϕ and any $\mathbf{W}, \mathbf{V} \in \mathbb{R}^{m \times d}$, on a good run it holds that

$$\frac{1}{\sqrt{C_1 d} + \alpha} \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}) \right\|_F \leq \widehat{G}_{\text{rob}}(\mathbf{W}) \leq \widehat{L}_{\text{rob}}(\mathbf{W}) \wedge 1.$$

Proof of Lemma B.5. Since ϕ is 1-Lipschitz, we have $\forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$,

$$\|\nabla f(\tilde{\mathbf{x}}_i; \mathbf{W})\|_F^2 = \frac{1}{m} \sum_{s=1}^m \|a_s \phi'(\langle \mathbf{w}_s, \tilde{\mathbf{x}}_i \rangle) \tilde{\mathbf{x}}_i\|^2 \leq \left(\sqrt{C_1 d} + \alpha \right)^2. \quad (3)$$

For $\forall i \in [n]$, choose $\tilde{\mathbf{x}}_i = \operatorname{argmax}_{\tilde{\mathbf{x}} \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}; \mathbf{W}))$ so that $\tilde{\ell}_i(\mathbf{W}) = \ell(y_i f(\tilde{\mathbf{x}}; \mathbf{W}))$, $\tilde{g}_i(\mathbf{W}) = g(y_i f(\tilde{\mathbf{x}}; \mathbf{W}))$, we have

$$\begin{aligned} \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}) \right\|_F &= \left\| \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}) y_i \nabla f(\tilde{\mathbf{x}}_i; \mathbf{W}) \right\|_F \\ &\leq \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}) \|\nabla f(\tilde{\mathbf{x}}_i; \mathbf{W})\|_F && \text{(Jensen's inequality)} \\ &\leq \frac{\sqrt{C_1 d} + \alpha}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}) && \text{(Equation (3))} \\ &\leq \frac{\sqrt{C_1 d} + \alpha}{n} \sum_{i=1}^n \min(\tilde{\ell}_i(\mathbf{W}), 1) && \text{(By the definition of } \tilde{g}_i(\mathbf{W}) \text{ and } \tilde{\ell}_i(\mathbf{W})) \\ &\leq \left(\sqrt{C_1 d} + \alpha \right) \left(\widehat{L}_{\text{rob}}(\mathbf{W}) \wedge 1 \right). && \text{(Jensen's inequality)} \end{aligned}$$

□

Lemma B.6. Let $C_1 > 1$ be the constant from Lemma B.3. For a γ -leaky, H -smooth activation ϕ , on a good run, for any $i, j \in [n]$, $i \neq j$, $\forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$, $\forall \tilde{\mathbf{x}}'_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$, $\forall \tilde{\mathbf{x}}_j \in \mathcal{B}_2(\mathbf{x}_j, \alpha)$, we have

$$\begin{aligned} |\langle \nabla f(\tilde{\mathbf{x}}_i, \mathbf{W}), \nabla f(\tilde{\mathbf{x}}_j, \mathbf{W}) \rangle| &\leq C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + \alpha \sqrt{C_1 d}, \\ |\langle \nabla f(\tilde{\mathbf{x}}_i, \mathbf{W}), \nabla f(\tilde{\mathbf{x}}'_i, \mathbf{W}) \rangle| &\leq C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2. \end{aligned}$$

Moreover, for any $i \in [n]$ and any $\mathbf{W} \in \mathbb{R}^{m \times d}$, we have

$$\begin{aligned} \left(\sqrt{d/C_1} - \alpha \right)^2 \gamma^2 &\leq \|\nabla f(\tilde{\mathbf{x}}_i; \mathbf{W})\|_F^2 \leq \left(\sqrt{C_1 d} + \alpha \right)^2, \\ \left(\sqrt{d/C_1} - \alpha \right)^2 \gamma^2 &\leq |\langle \nabla f(\tilde{\mathbf{x}}_i, \mathbf{W}), \nabla f(\tilde{\mathbf{x}}'_i, \mathbf{W}) \rangle| \leq \left(\sqrt{C_1 d} + \alpha \right)^2. \end{aligned}$$

Proof of Lemma B.6. The proof is similar as Lemma 4.7 in Frei et al. (2022).

$$\langle \nabla f(\mathbf{x}, \mathbf{W}), \nabla f(\mathbf{y}, \mathbf{W}) \rangle = \frac{1}{m} \langle \mathbf{x}, \mathbf{y} \rangle \sum_{s=1}^m \phi'(\langle \mathbf{w}_s, \mathbf{x} \rangle) \phi'(\langle \mathbf{w}_s, \mathbf{y} \rangle).$$

Therefore,

$$|\langle \nabla f(\mathbf{x}, \mathbf{W}), \nabla f(\mathbf{y}, \mathbf{W}) \rangle| = \frac{1}{m} |\langle \mathbf{x}, \mathbf{y} \rangle| \sum_{s=1}^m \phi'(\langle \mathbf{w}_s, \mathbf{x} \rangle) \phi'(\langle \mathbf{w}_s, \mathbf{y} \rangle) \in [\gamma^2 |\langle \mathbf{x}, \mathbf{y} \rangle|, |\langle \mathbf{x}, \mathbf{y} \rangle|].$$

Thus, the first two inequalities follow from Lemma B.3 (C2). The last two inequalities follow from Lemma B.3 (C1). \square

Lemma B.7 plays a crucial role in our analysis. It demonstrates that the margin increases with each epoch of adversarial training, given any adversarial examples. More importantly, it proves the loss g is at the same scale across all adversarial training examples.

Lemma B.7. For a γ -leaky, H -smooth activation ϕ , there is a constant $C_r = \frac{64C_1(\sqrt{C_1+0.5}\sqrt{\frac{1}{C_1}})^2}{\gamma^2}$ such that on a good run, provided $C > 1$ is sufficiently large, we have for all $t \geq 0$,

$$\begin{aligned} y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^{t+1}) &\geq y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^t) \geq 0, \forall \tilde{\mathbf{x}}_k \in \mathcal{B}_2(\mathbf{x}_k, \alpha), \forall k \in [n], \\ \max_{i,j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} &\leq \frac{16(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 (\sqrt{d/C_1} - \alpha)^2} \leq C_r, \end{aligned}$$

where $\tilde{\mathbf{x}}_i^t = \operatorname{argmax}_{\tilde{\mathbf{x}} \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}; \mathbf{W}^t))$, $\tilde{\mathbf{x}}_j^t = \operatorname{argmax}_{\tilde{\mathbf{x}} \in \mathcal{B}_2(\mathbf{x}_j, \alpha)} \ell(y_j f(\tilde{\mathbf{x}}; \mathbf{W}^t))$.

Proof of Lemma B.7. By Fact A.2 in Frei et al. (2022), we have

$$\frac{g(x)}{g(y)} \leq \max \left(2, 2 \frac{\exp(-x)}{\exp(-y)} \right)$$

holds for any $x, y \in \mathbb{R}$, so

$$\max_{i,j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq \max \left(2, 2 \cdot \max_{i,j \in [n]} \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \right).$$

In the remainder of the proof we will show that the ratio of the exponential losses is bounded. We will prove it by induction. Since a good run occurs, all the events in Lemma B.2 and Lemma B.3 occurs. In particular, we have $\|\mathbf{W}^0\|_2 \leq C_0 \omega_{\text{init}}(\sqrt{m} + \sqrt{d})$ and $\|\tilde{\mathbf{x}}_i^0\| \leq \sqrt{C_1 d} + \alpha$. Note that at initialization, we have $|f(\tilde{\mathbf{x}}_i^0; \mathbf{W}^0)| = 0$. For any $\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$, $\tilde{\mathbf{x}}_j \in \mathcal{B}_2(\mathbf{x}_j, \alpha)$, consider $t = 0$, we have

$$\max_{i,j \in [n]} \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^0; \mathbf{W}^0))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^0; \mathbf{W}^0))} = 1 \leq \frac{8(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 (\sqrt{d/C_1} - \alpha)^2}.$$

Assume the result holds at time t and consider the case $t + 1$. For simplicity we only consider the exponential ratio for the first sample and the second sample, and denote $A_t := \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))}$. Then $A_t \leq \frac{8(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 (\sqrt{d/C_1} - \alpha)^2}$. Fix $k \in [n]$, consider $\forall \tilde{\mathbf{x}}_k \in \mathcal{B}(\mathbf{x}_k, \alpha)$, de-

fine $\tilde{\rho}_i^t = \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^{(t)}, \tilde{\mathbf{x}}_i^t \rangle) \phi'(\langle \mathbf{w}_s^{(t)}, \tilde{\mathbf{x}}_i^t \rangle) \in [\gamma^2, 1]$. We first need to show that $y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^{t+1}) \geq y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^t)$.

$$y_k [f(\tilde{\mathbf{x}}_k; \mathbf{W}^{t+1}) - f(\tilde{\mathbf{x}}_k; \mathbf{W}^t)]$$

$$\geq y_k [\langle \nabla f(\tilde{\mathbf{x}}_k; \mathbf{W}^t), \mathbf{W}^{t+1} - \mathbf{W}^t \rangle] - \frac{H \|\tilde{\mathbf{x}}_k\|^2}{2\sqrt{m}} \|\mathbf{W}^{t+1} - \mathbf{W}^t\|_2^2$$

(for $y \in \{\pm 1\}$, apply Lemma B.4)

$$\begin{aligned}
&= y_k \eta \left[\left\langle \nabla f(\tilde{\mathbf{x}}_k; \mathbf{W}^t), \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) y_i \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \right\rangle \right] - \frac{H \|\tilde{\mathbf{x}}_k\|^2 \eta^2}{2\sqrt{m}} \left\| \nabla \hat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_2^2 \\
&\geq \eta \left[\frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \langle y_k \nabla f(\tilde{\mathbf{x}}_k; \mathbf{W}^t), y_i \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \rangle \right] - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta^2}{2\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\hspace{15em} \text{(Lemma B.5)} \\
&= \frac{\eta}{n} \left(\tilde{g}_k(\mathbf{W}^t) \tilde{\rho}_k^t \langle \tilde{\mathbf{x}}_k^t, \tilde{\mathbf{x}}_k \rangle + \sum_{i \neq k} \tilde{g}_i(\mathbf{W}^t) \tilde{\rho}_i^t \langle y_i \tilde{\mathbf{x}}_i^t, y_k \tilde{\mathbf{x}}_k \rangle \right) - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta^2}{2\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\geq \frac{\eta}{n} \left[\tilde{g}_k(\mathbf{W}^t) \left(\gamma^2 (\sqrt{d/C_1} - \alpha)^2 - \frac{\max_j \tilde{g}_j(\mathbf{W}^t)}{\tilde{g}_k(\mathbf{W}^t)} \sum_{i \neq k} \left(C_1 (\|\mu\|^2 + \sqrt{d \log(n/\delta)}) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right) \right] \\
&\hspace{15em} \text{(Lemma B.3)} \\
&\quad - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta^2}{2\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\geq \frac{\eta}{n} \left[\tilde{g}_k(\mathbf{W}^t) \left(\gamma^2 (\sqrt{d/C_1} - \alpha)^2 - C_r n \left(C_1 (\|\mu\|^2 + \sqrt{d \log(n/\delta)}) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right) \right] \\
&\hspace{15em} \text{(By induction, } \frac{\max_j \tilde{g}_j(\mathbf{W}^t)}{\tilde{g}_k(\mathbf{W}^t)} \leq C_r \text{)} \\
&\quad - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta^2}{2\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\geq \frac{\eta \gamma^2 (\sqrt{d/C_1} - \alpha)^2}{2n} \tilde{g}_k(\mathbf{W}^t) - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta^2}{2\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\hspace{15em} \text{(this line holds with large enough } C \text{ via Assumption (A2))} \\
&\geq \eta \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(\frac{\gamma^2 (\sqrt{d/C_1} - \alpha)^2}{2nC_r} - \frac{H (\sqrt{C_1 d} + \alpha)^4 \eta}{2\sqrt{m}} \right) \\
&\hspace{15em} \text{(By induction, } \tilde{g}_k(\mathbf{W}^t) \geq \frac{1}{C_r} \hat{G}_{\text{rob}}(\mathbf{W}^t) \text{)} \\
&\geq \eta \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(\frac{\gamma^2 \left(\frac{1}{2} \sqrt{d/C_1} \right)^2}{2nC_r} - \frac{H (2\sqrt{C_1 d})^4 \eta}{2\sqrt{m}} \right) \text{ (Assumption (A7) and Assumption (A2))} \\
&\geq \eta \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(\frac{\gamma^2 \left(\frac{1}{2} \sqrt{d/C_1} \right)^2}{2nC_r} - \frac{(2\sqrt{C_1 d})^4}{4Cd^2} \right) \hspace{10em} \text{(Assumption (A5))} \\
&\geq 0, \hspace{15em} \text{(4)}
\end{aligned}$$

where the last line holds from Assumption (A2) with sufficiently large C .

Now we are back to prove the upper bound of the exponential ratio A_t . We have

$$\begin{aligned}
A_{t+1} &= \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\
&= \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \cdot \frac{\exp(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t) - y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t) - y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\
&\leq A_t \cdot \frac{\exp(y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^t) - y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t) - y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^{t+1}))} \\
&= A_t \cdot \frac{\exp(y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^t) - y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^t - \frac{\eta}{n} \sum_{i=1}^n \nabla \ell(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)))}{\exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t) - y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t - \frac{\eta}{n} \sum_{i=1}^n \nabla \ell(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)))}
\end{aligned}$$

$$\begin{aligned}
&\leq A_t \cdot \frac{\exp\left(-\frac{\eta}{n} \sum_{i=1}^n y_1 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)}{\exp\left(-\frac{\eta}{n} \sum_{i=1}^n y_2 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)} \quad (\text{Lemma B.4}) \\
&\quad \cdot \exp\left(\frac{2H(\sqrt{C_1 d} + \alpha)^2 \eta^2}{\sqrt{m}} \left\| \frac{1}{n} \sum_{i=1}^n \nabla \ell(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \right\|^2\right) \\
&= A_t \cdot \exp\left(-\frac{\eta}{n} g(y_1 f(\tilde{x}_1^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_1^t; \mathbf{W}^t) \rangle\right) \\
&\quad + \frac{\eta}{n} g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_2^t; \mathbf{W}^t) \rangle \\
&\quad \cdot \frac{\exp\left(-\frac{\eta}{n} \sum_{i \neq 1} y_1 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)}{\exp\left(-\frac{\eta}{n} \sum_{i \neq 2} y_2 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)} \\
&\quad \cdot \exp\left(\frac{2H(\sqrt{C_1 d} + \alpha)^2 \eta^2}{\sqrt{m}} \left\| \frac{1}{n} \sum_{i=1}^n \nabla \ell(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \right\|^2\right),
\end{aligned}$$

where the first inequality holds since $\exp(y_1 f(\tilde{x}_1^{t+1}; \mathbf{W}^t)) \geq \exp(y_1 f(\tilde{x}_1^t; \mathbf{W}^t))$, $\exp(y_2 f(\tilde{x}_2^{t+1}; \mathbf{W}^{t+1})) \leq \exp(y_2 f(\tilde{x}_2^t; \mathbf{W}^{t+1}))$ by the definition of $\tilde{x}_1^t, \tilde{x}_2^{t+1}$.

We next bound each of the above term separately. For the first term, we have

$$\begin{aligned}
&\exp\left(-\frac{\eta}{n} g(y_1 f(\tilde{x}_1^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_1^t; \mathbf{W}^t) \rangle + \frac{\eta}{n} g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_2^t; \mathbf{W}^t) \rangle\right) \\
&= \exp\left(-\frac{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t)) \eta}{n} \left(\frac{g(y_1 f(\tilde{x}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t))} \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_1^t; \mathbf{W}^t) \rangle - \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_2^t; \mathbf{W}^t) \rangle\right)\right) \\
&\leq \exp\left(-\frac{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t)) \gamma^2 \eta}{n} \left(\frac{g(y_1 f(\tilde{x}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t))} (\sqrt{d/C_1} - \alpha)^2 - \frac{(\sqrt{C_1 d} + \alpha)^2}{\gamma^2}\right)\right) \quad (\text{Lemma B.6}) \\
&= \exp\left(-\frac{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t)) \eta \gamma^2 (\sqrt{d/C_1} - \alpha)^2}{n} \left(\frac{g(y_1 f(\tilde{x}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{x}_2^t; \mathbf{W}^t))} - \frac{(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 (\sqrt{d/C_1} - \alpha)^2}\right)\right).
\end{aligned}$$

For the second term, we have

$$\begin{aligned}
&\frac{\exp\left(-\frac{\eta}{n} \sum_{i \neq 1} y_1 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)}{\exp\left(-\frac{\eta}{n} \sum_{i \neq 2} y_2 y_i g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle\right)} \\
&\leq \exp\left(\frac{\eta}{n} \sum_{i \neq 1} g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) |\langle \nabla f(\tilde{x}_1^{t+1}; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle| \right. \\
&\quad \left. + \frac{\eta}{n} \sum_{i \neq 2} g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) |\langle \nabla f(\tilde{x}_2^t; \mathbf{W}^t), \nabla f(\tilde{x}_i^t; \mathbf{W}^t) \rangle| \right) \\
&\leq \exp\left(\frac{2\eta}{n} \sum_{i=1}^n g(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \left(C_1 \|\mu\|^2 + C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2\right)\right). \quad (\text{Lemma B.6})
\end{aligned}$$

For the third term, we have

$$\begin{aligned}
&\exp\left(\frac{2H(\sqrt{C_1 d} + \alpha)^2 \eta^2}{\sqrt{m}} \left\| \frac{1}{n} \sum_{i=1}^n \nabla \ell(y_i f(\tilde{x}_i^t; \mathbf{W}^t)) \right\|^2\right) \\
&\leq \exp\left(\frac{2H(\sqrt{C_1 d} + \alpha)^4 \eta^2}{\sqrt{m}} \cdot \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{x}_i^t; \mathbf{W}^t))\right) \quad (\text{Lemma B.5})
\end{aligned}$$

$$\leq \exp\left(\frac{\eta}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))\right). \quad (\text{Large enough } C \text{ for assumption (A5)})$$

Combining the above results gives us that

$$\begin{aligned} A_{t+1} &\leq A_t \exp\left(-\frac{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))\eta\gamma^2\left(\sqrt{\frac{d}{C_1}} - \alpha\right)^2}{n} \left(\frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} - \frac{(\sqrt{C_1 d} + \alpha)^2}{\gamma^2\left(\sqrt{\frac{d}{C_1}} - \alpha\right)^2}\right)\right) \\ &\quad \cdot \exp\left(\frac{2\eta}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2\right)\right). \end{aligned}$$

Now consider the following two cases. If $\frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \leq \frac{2(\sqrt{C_1 d} + \alpha)^2}{\gamma^2(\sqrt{d/C_1} - \alpha)^2}$, then we have

$$\begin{aligned} A_{t+1} &\leq A_t \cdot \exp\left(\frac{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))\eta(\sqrt{C_1 d} + \alpha)^2}{n}\right) \quad (g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t)) \geq 0) \\ &\quad \cdot \exp\left(\frac{2\eta}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2\right)\right) \\ &\leq A_t \cdot \exp\left(\frac{\eta(\sqrt{C_1 d} + \alpha)^2}{n}\right) \exp\left(2\eta\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2\right)\right) \\ &\quad (0 \leq g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) \leq 1) \\ &\leq \frac{2g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \exp\left(\frac{\eta(\sqrt{C_1 d} + \alpha)^2}{n}\right) \exp\left(2\eta\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2\right)\right) \\ &\quad (\frac{1}{2} \exp(-z) \leq g(z) \leq \exp(-z), \forall z \geq 0; \text{Equation (4)}) \\ &= \frac{4(\sqrt{C_1 d} + \alpha)^2}{\gamma^2\left(\sqrt{d/C_1} - \alpha\right)^2} \exp\left(2\eta\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2 + \frac{(\sqrt{C_1 d} + \alpha)^2}{2n}\right)\right) \\ &\leq \frac{8(\sqrt{C_1 d} + \alpha)^2}{\gamma^2\left(\sqrt{d/C_1} - \alpha\right)^2}, \end{aligned}$$

where the last line holds by assumption (A5) with sufficiently large C so that the following holds

$$\begin{aligned} &2\eta\left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha\sqrt{C_1 d} + \alpha^2 + \frac{(\sqrt{C_1 d} + \alpha)^2}{2n}\right) \\ &= \frac{\eta(\sqrt{C_1 d} + \alpha)^2}{n} + 2\eta\left(C_1 \|\mu\|^2 + \alpha^2\right) + 4\eta\left(C_1 \sqrt{d \log(n/\delta)} + \alpha\sqrt{C_1 d}\right) \\ &\leq \frac{(2\sqrt{C_1 d})^2}{Cd^2 n} + \frac{2(2C_1 \|\mu\|^2)}{Cd^2} + \frac{4(C_1 \sqrt{\frac{d^2}{Cn^2}} + C_1 d)}{Cd^2} \\ &\quad (\text{Assumption (A2), (A5), (A7) with sufficiently large } C) \\ &\leq \frac{1}{8}. \end{aligned}$$

Otherwise, $\frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} > \frac{2(\sqrt{C_1 d} + \alpha)^2}{\gamma^2(\sqrt{d/C_1} - \alpha)^2}$, then we have

$$A_{t+1} \leq A_t \cdot \exp\left(-\frac{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))\eta\gamma^2\left(\sqrt{\frac{d}{C_1}} - \alpha\right)^2}{n} \left(\frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} - \frac{(\sqrt{C_1 d} + \alpha)^2}{\gamma^2\left(\sqrt{\frac{d}{C_1}} - \alpha\right)^2}\right)\right)$$

$$\begin{aligned}
& \cdot \exp \left(\frac{2\eta g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))}{n} \sum_{i=1}^n \frac{g(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2; \mathbf{W}^t))} \left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right) \\
& \leq A_t \cdot \exp \left(- \frac{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t)) \eta \gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2}{n} \left(\frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} - \frac{(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \right) \right) \\
& \quad \cdot \exp \left(2\eta g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t)) \left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \max \left\{ 2, \frac{16(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \right\} \right) \\
& \leq A_t \cdot \exp \left(- g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t)) \eta \left(\frac{(\sqrt{C_1 d} + \alpha)^2}{n} \right. \right. \\
& \quad \left. \left. - \frac{32(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right) \right) \\
& \quad \quad \quad \left(\text{Assumption that } \frac{g(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{g(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} > \frac{2(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \right) \\
& \leq A_t \leq \frac{8(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2},
\end{aligned}$$

where the last line holds from Assumption (A7), Assumption (A2) with C being sufficiently large that

$$\begin{aligned}
& \frac{32(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \left(C_1 \|\mu\|^2 + 2C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \\
& \leq \frac{32(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\frac{1}{2} \sqrt{d/C_1} \right)^2} \left(C_1 \frac{d}{Cn} + 4C_1 \sqrt{d} \sqrt{\log(n/\delta)} + \alpha^2 + \|\mu\|^2 \right) \\
& \leq \frac{32(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\frac{1}{2} \sqrt{d/C_1} \right)^2} \left(C_1 \frac{d}{Cn} + 4C_1 \sqrt{\frac{d^2}{Cn^2} + \frac{d}{Cn}} \right) \\
& \leq \frac{(\sqrt{C_1 d} + \alpha)^2}{n}.
\end{aligned}$$

We have shown through induction that

$$\max_{i,j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq \frac{16(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2}.$$

By Assumption (A7) and Assumption (A2), we know $\alpha \leq \|\mu\| \leq \frac{1}{2} \sqrt{d/C_1}$. Therefore,

$$\max_{i,j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq \frac{16(\sqrt{C_1 d} + \alpha)^2}{\gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2} \leq \frac{64C_1 \left(\sqrt{C_1} + 0.5 \sqrt{\frac{1}{C_1}} \right)^2}{\gamma^2} = C_r.$$

□

With Lemma B.7, we are able to give a tighter bound on the norm of \mathbf{W}^t .

Lemma B.8. There is absolute constant $C_2 > 1$ such that for $C > 1$ sufficiently large, on a good run we have that

$$\|\mathbf{W}^t\|_F \leq \|\mathbf{W}^0\|_F + C_2 \eta \sqrt{\frac{d}{n}} \sum_{s=0}^{t-1} \hat{G}_{\text{rob}}(\mathbf{W}^s).$$

Lemma B.8. By triangle inequality we have that

$$\begin{aligned} \|\mathbf{W}^t\|_F &= \left\| \mathbf{W}^{t-1} - \eta \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^{t-1}) \right\|_F \\ &\leq \|\mathbf{W}^{t-1}\|_F + \left\| \eta \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^{t-1}) \right\|_F \\ &\leq \|\mathbf{W}^0\|_F + \eta \sum_{s=0}^{t-1} \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^s) \right\|_F. \end{aligned} \quad (\text{Telescope})$$

Consider $\tilde{\mathbf{x}}_i^s = \operatorname{argmax}_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^s))$. Then we have the following

$$\begin{aligned} &\left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^s) \right\|_F^2 \\ &= \frac{1}{n^2} \left\| \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^s) y_i \nabla f(\tilde{\mathbf{x}}_i^s; \mathbf{W}^s) \right\|_F^2 \\ &= \frac{1}{n^2} \left[\sum_{i=1}^n (\tilde{g}_i(\mathbf{W}^s))^2 \|\nabla f(\tilde{\mathbf{x}}_i^s; \mathbf{W}^s)\|_F^2 + \sum_{i \neq j} y_i y_j \tilde{g}_i(\mathbf{W}^s) \tilde{g}_j(\mathbf{W}^s) \langle \nabla f(\tilde{\mathbf{x}}_i^s; \mathbf{W}^s), \nabla f(\tilde{\mathbf{x}}_j^s; \mathbf{W}^s) \rangle \right] \\ &\leq \frac{1}{n^2} \left[\sum_{i=1}^n (\tilde{g}_i(\mathbf{W}^s))^2 \left(\sqrt{C_1 d} + \alpha \right)^2 \right. \\ &\quad \left. + \sum_{i \neq j} \tilde{g}_i(\mathbf{W}^s) \tilde{g}_j(\mathbf{W}^s) \left(C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right] \\ &\leq \frac{1}{n^2} \max_{k \in [n]} \tilde{g}_k(\mathbf{W}^s) \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^s) \left(2C_1 d + 2\alpha^2 + n \left(C_1 \|\mu\|^2 + C_1 \sqrt{d \log(n/\delta)} + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \right) \\ &\leq \frac{5C_1 d}{n} \max_{k \in [n]} \tilde{g}_k(\mathbf{W}^s) \widehat{G}_{\text{rob}}(\mathbf{W}^s), \end{aligned} \quad (\text{Lemma B.6})$$

where the last line follows Assumption (A2) and (A7).

Applying Lemma B.7 gives us

$$\max_{k \in [n]} \tilde{g}_k(\mathbf{W}^s) \leq \frac{C_r}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^s) = C_r \widehat{G}_{\text{rob}}(\mathbf{W}^s).$$

Define $C_2 := \sqrt{5C_1 C_r}$, then we have

$$\left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^s) \right\|_F \leq \sqrt{\frac{5C_1 C_r d}{n}} \widehat{G}_{\text{rob}}(\mathbf{W}^s) = C_2 \sqrt{\frac{d}{n}} \widehat{G}_{\text{rob}}(\mathbf{W}^s). \quad (5)$$

As a result, we have

$$\|\mathbf{W}^t\|_F \leq \|\mathbf{W}^0\|_F + C_2 \eta \sqrt{\frac{d}{n}} \sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s). \quad (6)$$

□

Recall that our goal is to give a lower bound on the normalized expected conditional margin. We start by giving a lower bound in terms of the cumulative increments of margin given any independent test example (x, y) , shown in Lemma B.9.

Lemma B.9. Let $C_2 > 1$ be the constant from Lemma B.8. For a γ -leaky, H -smooth activation ϕ , on a good run, we have for any $t \geq 0$ and $(x, y) \in \mathbb{R}^d \times \{\pm 1\}$, there exist $\tilde{\rho}_i^t = \rho(\mathbf{W}^t, \tilde{\mathbf{x}}_i^t, x) \in [\gamma^2, 1]$ such that

$$y [f(x; \mathbf{W}^{t+1}) - f(x; \mathbf{W}^t)] \geq \frac{\eta}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \left(\tilde{\rho}_i^t \langle y_i \tilde{\mathbf{x}}_i^t, yx \rangle - \frac{H \|\mathbf{x}\|^2 C_2^2 d \eta}{2\sqrt{mn}} \right)$$

Proof of Lemma B.9. Note that since a good run occurs, Lemma B.4 implies

$$|f(\mathbf{x}; \mathbf{W}^{t+1}) - f(\mathbf{x}; \mathbf{W}^t) - \langle \nabla f(\mathbf{x}; \mathbf{W}^t), \mathbf{W}^{t+1} - \mathbf{W}^t \rangle| \leq \frac{H \|\mathbf{x}\|^2}{2\sqrt{m}} \|\mathbf{W}^{t+1} - \mathbf{W}^t\|_2^2 \quad (7)$$

Therefore, we have

$$\begin{aligned} & y [f(\mathbf{x}; \mathbf{W}^{t+1}) - f(\mathbf{x}; \mathbf{W}^t)] \\ & \geq y [\langle \nabla f(\mathbf{x}; \mathbf{W}^t), \mathbf{W}^{t+1} - \mathbf{W}^t \rangle] - \frac{H \|\mathbf{x}\|^2}{2\sqrt{m}} \|\mathbf{W}^{t+1} - \mathbf{W}^t\|_2^2 \quad (\text{for } y \in \{\pm 1\}, \text{ apply (7)}) \\ & = y \eta \left[\left\langle \nabla f(\mathbf{x}; \mathbf{W}^t), \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) y_i \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \right\rangle \right] - \frac{H \|\mathbf{x}\|^2 \eta^2}{2\sqrt{m}} \|\nabla \hat{L}_{\text{rob}}(\mathbf{W}^t)\|_2^2 \\ & \geq \eta \left[\frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \langle y \nabla f(\mathbf{x}; \mathbf{W}^t), y_i \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \rangle \right] - \frac{H \|\mathbf{x}\|^2 C_2^2 d \eta^2}{2\sqrt{mn}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \\ & \quad (\text{Equation (5), } \|\nabla \hat{L}_{\text{rob}}(\mathbf{W}^t)\|_2 \leq \|\nabla \hat{L}_{\text{rob}}(\mathbf{W}^t)\|_F, \hat{G}_{\text{rob}}(\mathbf{W}^t) \leq 1) \\ & = \frac{\eta}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \left(\tilde{\rho}_i^t \langle y_i \tilde{\mathbf{x}}_i^t, y \mathbf{x} \rangle - \frac{H \|\mathbf{x}\|^2 C_2^2 d \eta}{2\sqrt{mn}} \right) \quad (\hat{G}_{\text{rob}}(\mathbf{W}^t) = \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t)) \end{aligned}$$

where the last equality follows by defining

$$\tilde{\rho}_i^t = \rho(\mathbf{W}^t, \tilde{\mathbf{x}}_i^t, \mathbf{x}) = \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^{(t)}, \mathbf{x} \rangle) \phi'(\langle \mathbf{w}_s^{(t)}, \tilde{\mathbf{x}}_i^t \rangle) \in [\gamma^2, 1].$$

□

Leveraging Lemma B.7 and B.9, we now formally derive a lower bound on the normalized expected conditional margin.

Lemma B.10. For a γ -leaky H -smooth activation ϕ , and for all $C > 1$ sufficiently large, on a good run, for any $t \geq 1$, we have

$$\begin{aligned} \frac{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} [y f(\mathbf{x}; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} & \geq \frac{\gamma^2 \sqrt{n}}{8C_2 \sqrt{d}} (0.99 \|\mu\|^2 - \|\mu\| \alpha); \\ \frac{\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} [y f(\mathbf{x}; \mathbf{W}^t) | y = -1]}{\|\mathbf{W}^t\|_2} & \geq \frac{\gamma^2 \sqrt{n}}{8C_2 \sqrt{d}} (0.99 \|\mu\|^2 - \|\mu\| \alpha). \end{aligned}$$

where C_2 is the constant from Lemma B.8.

Lemma B.10. From Lemma B.8, we have

$$\|\mathbf{W}^t\|_F \leq \|\mathbf{W}^0\|_F + C_2 \eta \sqrt{\frac{d}{n}} \sum_{s=0}^{t-1} \hat{G}_{\text{rob}}(\mathbf{W}^s).$$

Recall the following definition

$$\tilde{\rho}_i^t = \rho(\mathbf{W}^t, \tilde{\mathbf{x}}_i^t, \mathbf{x}) = \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^{(t)}, \mathbf{x} \rangle) \phi'(\langle \mathbf{w}_s^{(t)}, \tilde{\mathbf{x}}_i^t \rangle) \in [\gamma^2, 1].$$

By Lemma B.3, we have

$$\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} [\tilde{\rho}_i^t \langle y_i \tilde{\mathbf{x}}_i^t, \mu \rangle | y = 1] \geq \begin{cases} \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right), & i \in \mathcal{C} \\ -\frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha, & i \in \mathcal{N} \end{cases}$$

If $i \in \mathcal{C}$:

$$\begin{aligned}
& \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} \rangle \middle| y = 1 \right] \\
&= \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle + \tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, \mu \rangle \middle| y = 1 \right] \\
&\geq \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \geq 0) \middle| y = 1 \right] \\
&\quad + \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle < 0) \middle| y = 1 \right] + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&\geq \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\gamma^2 \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \geq 0) \middle| y = 1 \right] \\
&\quad + \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle < 0) \middle| y = 1 \right] + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&= -\frac{1-\gamma^2}{2} \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[|\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle| \right] + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&\geq -\frac{1-\gamma^2}{2} c_3 \|\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle\|_{\psi_2} + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \quad (c_3 \text{ is an absolute constant}) \\
&\geq -\frac{1-\gamma^2}{2} c_4 \|y_i \tilde{x}_i^t\|_2 + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \quad (c_4 \text{ is an absolute constant}) \\
&\geq -\frac{1-\gamma^2}{2} c_4 (\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right).
\end{aligned}$$

If $i \in \mathcal{N}$:

$$\begin{aligned}
& \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} \rangle \middle| y = 1 \right] \\
&= \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle + \tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, \mu \rangle \middle| y = 1 \right] \\
&\geq \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \geq 0) \middle| y = 1 \right] \\
&\quad + \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle < 0) \middle| y = 1 \right] - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \\
&\geq \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\gamma^2 \langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \geq 0) \middle| y = 1 \right] \\
&\quad + \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle \cdot \mathbf{1}(\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle < 0) \middle| y = 1 \right] - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \\
&= -\frac{1-\gamma^2}{2} \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[|\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle| \middle| y = 1 \right] - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \\
&\geq -\frac{1-\gamma^2}{2} c_3 \|\langle y_i \tilde{x}_i^t, y\mathbf{x} - \mu \rangle\|_{\psi_2} - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \quad (c_3 \text{ is an absolute constant}) \\
&\geq -\frac{1-\gamma^2}{2} c_4 \|y_i \tilde{x}_i^t\|_2 - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \quad (c_4 \text{ is an absolute constant}) \\
&\geq -\frac{1-\gamma^2}{2} c_4 (\sqrt{C_1 d} + \alpha) - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha.
\end{aligned}$$

$$\begin{aligned}
& \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[yf(\mathbf{x}; \mathbf{W}^{s+1}) - yf(\mathbf{x}; \mathbf{W}^s) \middle| y = 1 \right] \\
&\geq \frac{\eta}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^s) \mathbb{E}_{(x,y) \sim \mathcal{D}_c} \left[\tilde{\rho}_i^t \langle y_i \tilde{x}_i^t, y\mathbf{x} \rangle - \frac{H \|\mathbf{x}\|^2 C_2^2 d \eta}{2\sqrt{mn}} \middle| y = 1 \right] \quad (\text{Lemma B.9}) \\
&\geq \eta \left(\frac{1}{n} \sum_{i \in \mathcal{C}} \tilde{g}_i(\mathbf{W}^s) \left(-\frac{1-\gamma^2}{2} c_4 (\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \right) \right)
\end{aligned}$$

$$\begin{aligned}
& + \frac{1}{n} \sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^s) \left(-\frac{1-\gamma^2}{2} c_4(\sqrt{C_1 d} + \alpha) - \frac{201}{200} \|\mu\|^2 - \|\mu\| \alpha \right) - \frac{H c_5 d C_2^2 d \eta}{2\sqrt{mn}} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \\
& \hspace{15em} (c_5 \text{ is an absolute constant}) \\
& = \eta \left(-\frac{1-\gamma^2}{2} c_4(\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \right) \\
& \quad \cdot \left(\left(1 - \frac{H c_5 C_2^2 d^2 \eta}{(-1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^s) \right. \\
& \quad \left. - \left(1 + \frac{(1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{201}{100} \|\mu\|^2 + 2\alpha \|\mu\|}{(-1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \right) \frac{1}{n} \sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^s) \right) \\
& \geq \eta \left(-\frac{1-\gamma^2}{2} c_4(\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \right) \\
& \quad \cdot \left(\left(1 - \frac{H c_5 C_2^2 d^2 \eta}{(-1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^s) \right. \\
& \quad \left. - (\beta + \sqrt{\frac{2}{C}}) C_r \left(1 + \frac{(1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{201}{100} \|\mu\|^2 + 2\alpha \|\mu\|}{(-1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^s) \right) \\
& \geq \frac{\eta \widehat{G}_{\text{rob}}(\mathbf{W}^s)}{4} \left(-\frac{1-\gamma^2}{2} c_4(\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \right) \\
& \geq \frac{\eta \widehat{G}_{\text{rob}}(\mathbf{W}^s) \gamma^2}{4} \left(0.99 \|\mu\|^2 - \alpha \|\mu\| \right),
\end{aligned}$$

where the second last inequality follows from $\eta \leq \frac{\sqrt{m}}{C d^2 H^2} \leq \frac{0.005 \gamma^4 \|\mu\|^2 \sqrt{mn}}{H c_5 C_2^2 d^2}$, $\beta \leq 1/C$, $d \leq \frac{\|\mu\|^4}{C} \leq \frac{10^{-4} \gamma^4 \|\mu\|^4}{4c_4^2 C_1}$, $\alpha \leq \|\mu\| \leq \sqrt{C_1 d}$ and C sufficiently large so that

$$\begin{aligned}
& \frac{H c_5 C_2^2 d^2 \eta}{\left(-1-\gamma^2 \right) c_4 \left(\sqrt{C_1 d} + \alpha \right) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \sqrt{mn} \\
& \leq \frac{0.005 \gamma^4 \|\mu\|^2}{-1-\gamma^2 \left) c_4 \left(2\sqrt{C_1 d} \right) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 1.98 \gamma^2 \|\mu\|^2} \\
& \leq \frac{0.005 \gamma^4 \|\mu\|^2}{-0.01 \gamma^2 (1-\gamma^2) \|\mu\|^2 + \frac{199}{100} \gamma^2 \|\mu\|^2 - 1.98 \gamma^2 \|\mu\|^2} = 0.5,
\end{aligned}$$

and

$$\begin{aligned}
& (\beta + \sqrt{\frac{2}{C}}) C_r \left(1 + \frac{(1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{201}{100} \|\mu\|^2 + 2\alpha \|\mu\|}{(-1-\gamma^2)c_4(\sqrt{C_1 d} + \alpha) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 2\alpha \gamma^2 \|\mu\|} \right) \\
& \leq \left(\frac{1}{C} + \sqrt{\frac{2}{C}} \right) C_r \left(1 + \frac{(1-\gamma^2)c_4(2\sqrt{C_1 d}) + \frac{201}{100} \|\mu\|^2 + 2\|\mu\|^2}{(-1-\gamma^2)c_4(2\sqrt{C_1 d}) + \frac{199}{100} \gamma^2 \|\mu\|^2 - 1.98 \gamma^2 \|\mu\|^2} \right) \\
& \leq \left(\frac{1}{C} + \sqrt{\frac{2}{C}} \right) C_r \left(1 + \frac{(1-\gamma^2)0.01 \gamma^2 \|\mu\|^2 + \frac{201}{100} \|\mu\|^2 + 2\|\mu\|^2}{(-1-\gamma^2)0.01 \gamma^2 \|\mu\|^2 + \frac{199}{100} \gamma^2 \|\mu\|^2 - 1.98 \gamma^2 \|\mu\|^2} \right) \\
& = \left(\frac{1}{C} + \sqrt{\frac{2}{C}} \right) C_r \left(1 + \frac{(1-\gamma^2)0.01 \gamma^2 + \frac{201}{100} + 2}{0.01 \gamma^4} \right) \\
& \leq 0.25.
\end{aligned}$$

The third last inequality follows from Lemma B.7 that

$$\begin{aligned}
\sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^s) &\leq |\mathcal{N}| \cdot \max_i \tilde{g}_i(\mathbf{W}^s) \\
&\leq \frac{|\mathcal{N}|}{n} \sum_{k=1}^n \max_i \tilde{g}_i(\mathbf{W}^s) \\
&\leq C_r \cdot |\mathcal{N}| \cdot \widehat{G}_{\text{rob}}(\mathbf{W}^s) \\
&\leq C_r (\beta + \sqrt{\frac{2}{C}}) n \widehat{G}_{\text{rob}}(\mathbf{W}^s). \tag{8}
\end{aligned}$$

The last inequality follows from $d \leq \frac{\|\mu\|^4}{C} \leq \frac{10^{-4} \gamma^4 \|\mu\|^4}{4c_4^2 C_1}$, and $\alpha \leq \|\mu\| \leq \sqrt{C_1 d}$ so that

$$\begin{aligned}
&-\frac{1-\gamma^2}{2} c_4 (\sqrt{C_1 d} + \alpha) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&\geq -\frac{1-\gamma^2}{2} c_4 (2\sqrt{C_1 d}) + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&\geq -\frac{1-\gamma^2}{200} \gamma^2 \|\mu\|^2 + \gamma^2 \left(\frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha \right) \\
&\geq 0.99 \gamma^2 \|\mu\|^2 - \gamma^2 \alpha \|\mu\|.
\end{aligned}$$

Applying the above result gives us the following

$$\begin{aligned}
&\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} \\
&= \frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^0) | y = 1] + \sum_{s=0}^{t-1} \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^{s+1}) - yf(x; \mathbf{W}^s) | y = 1]}{\|\mathbf{W}^t\|_2} \\
&\geq \frac{\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \eta \gamma^2}{4 \|\mathbf{W}^t\|_F} (0.99 \|\mu\|^2 - \|\mu\| \alpha). \quad (f(x; \mathbf{W}^0) = 0 \text{ via symmetric initialization})
\end{aligned}$$

Note that we have

$$\widehat{G}_{\text{rob}}(\mathbf{W}^0) = \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^0) = -\frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\bar{x}_i^0; \mathbf{W}^0)) = \frac{1}{2}. \tag{9}$$

Along with Lemma B.2, Assumption (A6) and Assumption (A2) gives us

$$\|\mathbf{W}^0\|_F \leq 2\omega_{\text{init}} \sqrt{md} \leq 2\eta \leq \eta \sqrt{d/n} \widehat{G}_{\text{rob}}(\mathbf{W}^0). \tag{10}$$

Then if $\|\mathbf{W}^t\|_F \leq 2 \|\mathbf{W}^0\|_F$, we have

$$\begin{aligned}
\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} &\geq \frac{\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \eta \gamma^2}{8 \|\mathbf{W}^0\|_F} (0.99 \|\mu\|^2 - \|\mu\| \alpha) \\
&\geq \frac{\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \eta \gamma^2}{8 \eta \sqrt{d/n} \widehat{G}_{\text{rob}}(\mathbf{W}^0)} (0.99 \|\mu\|^2 - \|\mu\| \alpha) \quad (\text{Equation (10)}) \\
&\geq \frac{\sqrt{n} \gamma^2}{8 \sqrt{d}} (0.99 \|\mu\|^2 - \|\mu\| \alpha). \quad (\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \geq \widehat{G}_{\text{rob}}(\mathbf{W}^0))
\end{aligned}$$

If $\|\mathbf{W}^t\|_F > 2 \|\mathbf{W}^0\|_F$, by Lemma B.8, we have

$$2 \|\mathbf{W}^0\|_F \leq \|\mathbf{W}^t\|_F \leq \|\mathbf{W}^0\|_F + C_2 \eta \sqrt{d/n} \sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s).$$

Thus we have

$$\begin{aligned} \frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} &\geq \frac{\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \eta \gamma^2}{8C_2 \eta \sqrt{d/n} \sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s)} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right) \\ &\geq \frac{\sqrt{n} \gamma^2}{8C_2 \sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right). \end{aligned}$$

Similarly, we can get

$$\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = -1]}{\|\mathbf{W}^t\|_2} \geq \frac{\gamma^2 \sqrt{n}}{8C_2 \sqrt{d}} \left(0.99 \|\mu\|^2 - \|\mu\| \alpha\right).$$

□

We finally provide the convergence guarantees of robust training loss in Lemma B.11.

Lemma B.11. For a γ -leaky, H -smooth activation ϕ , provided $C > 1$ is sufficiently large, then on a good run we have that,

$$\left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F \geq \frac{\gamma (199 \|\mu\| - 200\alpha)}{400} \widehat{G}_{\text{rob}}(\mathbf{W}^t)$$

Moreover, the robust training loss satisfies

$$\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \frac{3500 + 800\sqrt{m/d^3}}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}}$$

Proof of Lemma B.11. Consider \tilde{x}_i^t as the adversarial examples given model parameter \mathbf{W}^t ; i.e., $\tilde{x}_i^t = \operatorname{argmax}_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W}^t))$. We first need to show a lower bound for $\left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F = \sup_{\mathbf{U}: \|\mathbf{U}\|_F=1} \left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{U} \right\rangle$, and it suffices to construct a matrix \mathbf{V} with Frobenius norm at most one such that $\left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle$ is bounded from below by a positive constant. To this end, choose $\mathbf{V} \in \mathbb{R}^{m \times d}$ be the matrix with rows $\mathbf{v}_s = \frac{a_s \mu}{\|\mu\| \sqrt{m}}, \forall s \in [m]$. Then $\|\mathbf{V}\|_F = 1$ since $a_s = \pm 1$, and we have for any $\mathbf{W} \in \mathbb{R}^{m \times d}$,

$$\left\langle \nabla f(x_i; \mathbf{W}), \mathbf{V} \right\rangle = \frac{1}{\sqrt{m}} \sum_{s=1}^m a_s \phi'(\langle \mathbf{w}_s, \mathbf{x}_i \rangle) \langle \mathbf{v}_s, \mathbf{x}_i \rangle = \left\langle \frac{\mu}{\|\mu\|}, \mathbf{x}_i \right\rangle \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s, \mathbf{x}_i \rangle) \quad (11)$$

By Lemma B.3, we have

$$\begin{cases} y_i \langle \mu, \mathbf{x}_i \rangle \geq \frac{199}{200} \|\mu\|^2, & i \in \mathcal{C} \\ |\langle \mu, \tilde{\mathbf{x}}_i \rangle| \leq \frac{201}{200} \|\mu\|^2, & i \in \mathcal{N} \end{cases}, \begin{cases} y_i \langle \mu, \tilde{\mathbf{x}}_i \rangle \geq \frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha, & i \in \mathcal{C} \\ |\langle \mu, \tilde{\mathbf{x}}_i \rangle| \leq \frac{201}{200} \|\mu\|^2 + \|\mu\| \alpha, & i \in \mathcal{N} \end{cases}$$

And $\forall z, \phi'(z) \geq \gamma > 0$, equation (11) implies that we have the following lower bound for any $\mathbf{W} \in \mathbb{R}^{m \times d}$,

$$y_i \langle \nabla f(x_i; \mathbf{W}), \mathbf{V} \rangle \geq \begin{cases} \frac{199\gamma}{200} \|\mu\|, & i \in \mathcal{C} \\ -\frac{201}{200} \|\mu\|, & i \in \mathcal{N} \end{cases}, y_i \langle \nabla f(\tilde{x}_i; \mathbf{W}), \mathbf{V} \rangle \geq \begin{cases} \frac{199\gamma}{200} \|\mu\| - \gamma\alpha, & i \in \mathcal{C} \\ -\frac{201}{200} \|\mu\| - \alpha, & i \in \mathcal{N} \end{cases}$$

This allows for a lower bound on $\left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle$, since

$$\begin{aligned} &\left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle \\ &= \frac{1}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) y_i \langle \nabla f(\tilde{x}_i^t; \mathbf{W}^t), \mathbf{V} \rangle \end{aligned} \quad (12)$$

$$\begin{aligned}
&\geq \frac{1}{n} \sum_{i \in \mathcal{C}} \tilde{g}_i(\mathbf{W}^t) \left(\frac{199\gamma}{200} \|\mu\| - \gamma\alpha \right) - \frac{1}{n} \sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^t) \left(\frac{201}{200} \|\mu\| + \alpha \right) \\
&= \left(\frac{199\gamma \|\mu\|}{200} - \gamma\alpha \right) \left[\widehat{G}_{\text{rob}}(\mathbf{W}^t) - \left(1 + \frac{\frac{201}{200} \|\mu\| + \alpha}{\frac{199\gamma \|\mu\|}{200} - \gamma\alpha} \right) \frac{1}{n} \sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^t) \right] \\
&\geq \left(\frac{199\gamma \|\mu\|}{200} - \gamma\alpha \right) \left[\widehat{G}_{\text{rob}}(\mathbf{W}^t) - \left(1 + \frac{201 \|\mu\| + 200\alpha}{199\gamma \|\mu\| - 200\gamma\alpha} \right) \cdot C_r(\beta + \sqrt{\frac{2}{C}}) \widehat{G}_{\text{rob}}(\mathbf{W}^t) \right] \\
&\hspace{25em} \text{(Equation (8))} \\
&\geq \left(\frac{199\gamma \|\mu\|}{400} - \frac{\gamma\alpha}{2} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t), \tag{13}
\end{aligned}$$

where the last line holds by Assumption (A4), $\alpha \leq 0.99 \|\mu\|$ and C being sufficiently large so that

$$\begin{aligned}
&\left(1 + \frac{201 \|\mu\| + 200\alpha}{199\gamma \|\mu\| - 200\gamma\alpha} \right) \cdot C_r(\beta + \sqrt{\frac{2}{C}}) \\
&\leq \left(1 + \frac{201 + 198}{\gamma} \right) \cdot C_r \left(\frac{1}{C} + \sqrt{\frac{2}{C}} \right) \\
&\leq 0.5.
\end{aligned}$$

Thus we have

$$\widehat{G}_{\text{rob}}(\mathbf{W}^t) \leq \frac{400}{199\gamma \|\mu\| - 200\gamma\alpha} \left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle \leq \frac{400}{199\gamma \|\mu\| - 200\gamma\alpha} \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F. \tag{14}$$

We next give an upper bound on $\|\mathbf{W}^t\|_F^2$ as follows:

$$\begin{aligned}
&\|\mathbf{W}^{t+1}\|_F^2 \tag{15} \\
&= \left\| \mathbf{W}^t - \eta \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 \\
&= \|\mathbf{W}^t\|_F^2 + \eta^2 \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \\
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \tag{Equation (5)} \\
&= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \sum_{s=1}^m \frac{a_s}{\sqrt{m}} \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle \\
&= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \sum_{s=1}^m \frac{a_s}{\sqrt{m}} \phi(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \\
&\quad + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \sum_{s=1}^m \frac{a_s}{\sqrt{m}} (\phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle - \phi(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle)) \\
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \\
&\quad + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) \sum_{s=1}^m \frac{1}{\sqrt{m}} |\phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle - \phi(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle)| \\
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3} \eta + 2\eta \frac{1}{n} \sum_{i=1}^n \sum_{s=1}^m \frac{1}{\sqrt{m}} (c_1 + c_2 |\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle|^\zeta) \tag{(g(z)z \leq \frac{1}{3}, g(z) \leq 1)} \\
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3} \eta + 2\eta c_1 \sqrt{m} + \frac{2\eta c_2}{\sqrt{m}} \frac{1}{n} \sum_{i=1}^n \|\tilde{\mathbf{x}}_i^t\|_2^\zeta \sum_{s=1}^m \|\mathbf{w}_s^t\|_2^\zeta
\end{aligned}$$

$$\begin{aligned}
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + \frac{2\eta c_2}{\sqrt{m}} \frac{1}{n} \sum_{i=1}^n \|\tilde{x}_i^t\|_2^\zeta m^{1-\frac{\zeta}{2}} \|\mathbf{W}^t\|_F^\zeta \\
&\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \|\mathbf{W}^t\|_F^\zeta.
\end{aligned}$$

Then we have

$$\|\mathbf{W}^{t+1}\|_F^{2-\zeta} \leq (\|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \|\mathbf{W}^t\|_F^\zeta)^{\frac{2-\zeta}{2}}.$$

If $\|\mathbf{W}^t\|_F \leq 1$,

$$\begin{aligned}
\|\mathbf{W}^{t+1}\|_F^{2-\zeta} &\leq \left(1 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)^{\frac{2-\zeta}{2}} \\
&\leq \|\mathbf{W}^t\|_F^{2-\zeta} + \left(1 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)^{\frac{2-\zeta}{2}} \\
&\leq \|\mathbf{W}^t\|_F^{2-\zeta} + 1 + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right).
\end{aligned}$$

If $\|\mathbf{W}^t\|_F > 1$,

$$\begin{aligned}
\|\mathbf{W}^{t+1}\|_F^{2-\zeta} &\leq \|\mathbf{W}^t\|_F^{2-\zeta} \left(1 + \frac{\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \|\mathbf{W}^t\|_F^\zeta}{\|\mathbf{W}^t\|_F^2}\right)^{\frac{2-\zeta}{2}} \\
&\leq \|\mathbf{W}^t\|_F^{2-\zeta} \left(1 + \frac{2-\zeta}{2} \cdot \frac{\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \|\mathbf{W}^t\|_F^\zeta}{\|\mathbf{W}^t\|_F^2}\right) \\
&\leq \|\mathbf{W}^t\|_F^{2-\zeta} + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right).
\end{aligned}$$

Combining the two cases, we have

$$\|\mathbf{W}^{t+1}\|_F^{2-\zeta} \leq \|\mathbf{W}^t\|_F^{2-\zeta} + 1 + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right).$$

Summing up the above inequality, we get

$$\|\mathbf{W}^T\|_F^{2-\zeta} \leq \|\mathbf{W}^0\|_F^{2-\zeta} + T \cdot \left(1 + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)\right).$$

Thus,

$$\begin{aligned}
&\|\mathbf{W}^T\|_F \\
&\leq \left(\|\mathbf{W}^0\|_F^{2-\zeta} + T \cdot \left(1 + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)\right)\right)^{\frac{1}{2-\zeta}} \\
&\leq \|\mathbf{W}^0\|_F + T^{\frac{1}{2-\zeta}} \cdot \left(1 + \frac{2-\zeta}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)\right)^{\frac{1}{2-\zeta}} \\
&\leq \|\mathbf{W}^0\|_F + T^{\frac{1}{2-\zeta}} \cdot \left(1 + \frac{1}{2} \cdot \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + 2\eta c_1 \sqrt{m} + 2\eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right)\right) \\
&= \|\mathbf{W}^0\|_F + T^{\frac{1}{2-\zeta}} \cdot \left(1 + \eta^2 \frac{C_2^2 d}{2n} + \frac{1}{3}\eta + \eta c_1 \sqrt{m} + \eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta\right).
\end{aligned}$$

Consider the correlation between iterates weight and \mathbf{V} as follows:

$$\begin{aligned}
\langle \mathbf{W}^{t+1}, \mathbf{V} \rangle &= \langle \mathbf{W}^t - \eta \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle \\
&= \langle \mathbf{W}^t, \mathbf{V} \rangle - \eta \langle \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle \\
&= \dots \\
&= \langle \mathbf{W}^0, \mathbf{V} \rangle - \eta \sum_{s=0}^t \langle \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^s), \mathbf{V} \rangle.
\end{aligned} \tag{16}$$

Recall from Lemma B.7 that $\forall t \geq 0, y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^{t+1}) \geq y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^t), \forall \tilde{\mathbf{x}}_k \in \mathcal{B}_2(\mathbf{x}_k, \alpha), \forall k \in [n]$. Then we have $y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^T) \geq y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^t)$, and therefore $\ell(y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^T)) \leq \ell(y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^t)) \leq \ell(y_k f(\tilde{\mathbf{x}}_k^t; \mathbf{W}^t))$ by definition that $\tilde{\mathbf{x}}_i^t = \arg \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)), t \leq T$. As a result, we have

$$\begin{aligned}
&\widehat{L}_{\text{rob}}(\mathbf{W}^T) \\
&= \frac{1}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^T)) \\
&\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)) \\
&\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{2}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} -\ell'(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)) \quad (\ell(z) \leq -2\ell'(z) \text{ when } z \geq 0, \text{ Equation (4)}) \\
&= \frac{2}{T} \sum_{t=0}^{T-1} \widehat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma T} \sum_{t=0}^{T-1} \langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle \quad (\text{Equation (14)}) \\
&= \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} (\langle \mathbf{W}^T, \mathbf{V} \rangle - \langle \mathbf{W}^0, \mathbf{V} \rangle) \quad (\text{Equation (16)}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} (\|\mathbf{W}^T\|_F + \|\mathbf{W}^0\|_F) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} \left(2 \|\mathbf{W}^0\|_F + T^{\frac{1}{2-\zeta}} \cdot \left(1 + \eta^2 \frac{C_2^2 d}{2n} + \frac{1}{3} \eta + \eta c_1 \sqrt{m} + \eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \right) \right) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}} \left(2 \|\mathbf{W}^0\|_F + 1 + \eta^2 \frac{C_2^2 d}{2n} + \frac{1}{3} \eta + \eta c_1 \sqrt{m} + \eta c_2 m^{\frac{1-\zeta}{2}} (\sqrt{C_1 d} + \alpha)^\zeta \right) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}} \left(\omega_{\text{init}} \sqrt{6md} + \frac{4}{3} + \eta \left(\frac{C_2^2 d \eta}{2n} + (c_1 + c_2) \sqrt{m} (\sqrt{C_1 d} + \alpha) \right) \right) \\
&\hspace{15em} (\text{Lemma B.2}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}} \left(\sqrt{6} \eta + \frac{4}{3} + \eta \left(\frac{C_2^2 d \eta}{2n} + (c_1 + c_2) \sqrt{m} (\sqrt{C_1 d} + \alpha) \right) \right) \\
&\hspace{15em} (\text{Assumption (A6)}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}} \left(\sqrt{6} + \frac{4}{3} + \frac{1}{C d^2} \left(\frac{C_2^2 d}{2n C d^2} + (c_1 + c_2) \sqrt{m} (\sqrt{C_1 d} + \alpha) \right) \right) \\
&\hspace{15em} (\text{Assumption (A5)}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T^{\frac{1-\zeta}{2}}} \left(\sqrt{6} + \frac{4}{3} + \frac{1}{2} + \sqrt{\frac{m}{d^3}} \right) \\
&\hspace{10em} (\text{choose } C > \max\{1, C_1, C_2, 4(c_1 + c_2)^2\} \text{ be large enough})
\end{aligned}$$

$$\leq \frac{3500 + 800\sqrt{\frac{m}{d^3}}}{(199\|\mu\| - 200\alpha)\gamma\eta T^{\frac{1-\zeta}{2}}}.$$

$$\forall 0 < \varepsilon < \frac{1}{2n}, T \geq \left(\frac{3500 + 800\sqrt{\frac{m}{d^3}}}{(199\|\mu\| - 200\alpha)\gamma\eta\varepsilon} \right)^{\frac{2}{1-\zeta}} \text{ guarantees } \widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon. \quad \square$$

Corollary 3.2. For any γ -leaky H -smooth ReLU activation ϕ_{SLReLU} defined in Equation (1), and for all $\kappa \in (0, 1)$, $\lambda > 0$, given Assumption 1 holds, we have that with probability at least $1 - 2\delta$ over the random initialization and the draws of the samples, the robust training loss satisfies

$$\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \mathcal{O} \left(\frac{1 + \sqrt{(1-\gamma)/H} m^{1/4}}{(199\|\mu\| - 200\alpha)\gamma\sqrt{\eta}\sqrt{T}} \right).$$

Proof of Corollary 3.2. Here our activation function ϕ is ϕ_{SLReLU} . The definition of ϕ_{SLReLU} gives us that

$$\phi'(z)z = \begin{cases} z = \phi(z) + \frac{1-\gamma}{4H}, & z \geq \frac{1}{H} \\ \frac{1-\gamma}{2}Hz^2 + \frac{1+\gamma}{2}z = \phi(z) + \frac{1-\gamma}{4}Hz^2, & |z| \leq \frac{1}{H} \\ \gamma z = \phi(z) + \frac{1-\gamma}{4H}, & z \leq -\frac{1}{H} \end{cases}.$$

Therefore, $\phi'(z)z - \phi(z) \leq \frac{1-\gamma}{4H}$. Similar as Lemma B.11, we have $\widehat{G}_{\text{rob}}(\mathbf{W}^t) \leq \frac{400}{199\gamma\|\mu\| - 200\gamma\alpha} \left\| \widehat{\nabla L}_{\text{rob}}(\mathbf{W}^t) \right\|_F$. In terms of the upper bound on $\|\mathbf{W}^t\|_F^2$, we have

$$\begin{aligned} & \|\mathbf{W}^{t+1}\|_F^2 \\ &= \left\| \mathbf{W}^t - \eta \widehat{\nabla L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \left\| \widehat{\nabla L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \quad (\text{Equation (5)}) \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \frac{1}{\sqrt{m}} \sum_{s=1}^m a_s \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \frac{1}{\sqrt{m}} \sum_{s=1}^m a_s \phi(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \\ &\quad + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) \frac{1-\gamma}{4H} \sqrt{m} \quad (\phi'(z)z - \phi(z) \leq \frac{1-\gamma}{4H}) \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) + \eta \sqrt{m} \frac{1-\gamma}{2H} \\ &\quad (g(z) \leq 1) \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3}\eta + \eta \sqrt{m} \frac{1-\gamma}{2H} \quad (g(z)z \leq \frac{1}{3}) \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta \frac{C_2^2 d}{nC d^2} + \frac{2}{3}\eta + \eta \sqrt{m} \frac{1-\gamma}{2H} \quad (\text{Assumption (A5)}) \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta \left(\frac{5}{3} + \frac{1-\gamma}{2H} \sqrt{m} \right). \end{aligned}$$

Telescoping gives us that

$$\|\mathbf{W}^T\|_F^2 \leq \|\mathbf{W}_0\|_F^2 + \eta \left(\frac{5}{3} + \frac{1-\gamma}{2H} \sqrt{m} \right) T.$$

As a result, we have

$$\begin{aligned}
\widehat{L}_{\text{rob}}(\mathbf{W}^T) &= \frac{1}{n} \sum_{i=1}^n \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W}^T)) \\
&\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^n \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} \ell(y_i f(\tilde{x}_i; \mathbf{W}^t)) \\
&\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{2}{n} \sum_{i=1}^n \max_{\tilde{x}_i \in \mathcal{B}_2(x_i, \alpha)} -\ell'(y_i f(\tilde{x}_i; \mathbf{W}^t)) \\
&\hspace{15em} (\ell(z) \leq -2\ell'(z) \text{ when } z \geq 0, \text{ Equation (4)}) \\
&= \frac{2}{T} \sum_{t=0}^{T-1} \widehat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} \sum_{t=0}^{T-1} \langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle \hspace{5em} (\text{Equation (14)}) \\
&= \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} (\langle \mathbf{W}^T, \mathbf{V} \rangle - \langle \mathbf{W}^0, \mathbf{V} \rangle) \hspace{5em} (\text{Equation (16)}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} (\|\mathbf{W}^T\|_F + \|\mathbf{W}^0\|_F) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} \left(2\|\mathbf{W}^0\|_F + \sqrt{\eta \left(\frac{5}{3} + \frac{1-\gamma}{2H} \sqrt{m} \right) T} \right) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} \left(\omega_{\text{init}} \sqrt{6md} + \sqrt{\eta \left(\frac{5}{3} + \frac{1-\gamma}{2H} \sqrt{m} \right) T} \right) \hspace{2em} (\text{Lemma B.2}) \\
&\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma \eta T} \left(\sqrt{6\eta} + \sqrt{\eta \left(\frac{5}{3} + \frac{1-\gamma}{2H} \sqrt{m} \right) T} \right) \hspace{2em} (\text{Lemma (A6)}) \\
&\leq \frac{3000 + 600\sqrt{(1-\gamma)/H} m^{0.25}}{(199 \|\mu\| - 200\alpha) \gamma \sqrt{\eta} \sqrt{T}}.
\end{aligned}$$

□

B.2 MISSING PROOFS IN SECTION 3.2

Theorem 3.3. Let $0 < \varepsilon \leq \frac{1}{2n}$, $\delta \in (0, 1/2)$. Let ϕ be a non-smooth activation with $\gamma \in (0, 1]$. Let $\kappa \in (0, 1)$, $\lambda > 0$. Then, given that Assumption 2 holds with some constant $C > 0$, there exists a constant $c > 0$ such that after running Algorithm 1 for $T \geq \Omega\left(\left(\frac{1}{(199 \|\mu\| - 200\alpha) \gamma \sqrt{\eta} \varepsilon}\right)^{-2}\right)$ iterations, we have that with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample of size n , the following holds:

1. The robust training loss satisfies $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon$.
2. The clean test error satisfies $L^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(1 - \frac{\alpha}{\|\mu\|}\right)^2\right)$.
3. For $\frac{\alpha}{\|\mu\|} \leq \frac{\sqrt{n} \|\mu\|}{\sqrt{n} \|\mu\| + C\sqrt{d}}$, the robust test error satisfies
$$L_{\text{rob}}^{0/1}(\mathbf{W}^T) \leq \beta + 2 \exp\left(-c\lambda \|\mu\|^2 \left(\frac{\sqrt{n} \|\mu\|}{C\sqrt{d}} \left(1 - \frac{\alpha}{\|\mu\|}\right) - \frac{\alpha}{\|\mu\|}\right)^2\right).$$

Proof of Theorem 3.3. This proof is similar with the proof of Theorem 3.1. The robust training loss bound is proved in Lemma B.21. For the generalization guarantee, apply Lemma B.1 with Lemma B.26, with probability at least $1 - 2\delta$,

$$L_{\text{rob}}^{0/1}(\mathbf{W}^T) = \mathbb{P}_{(x,y) \sim \mathcal{D}}[\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y \neq \text{sign}(f(\tilde{x}; \mathbf{W}^T))]$$

$$\begin{aligned}
&\leq \beta + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}^T) | y_c = 1]}{\|\mathbf{W}^T\|_2} - \alpha \right)^2 \right) \\
&\quad + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}^T) | y_c = -1]}{\|\mathbf{W}^T\|_2} - \alpha \right)^2 \right) \\
&\leq \beta + 2 \exp \left(-c\lambda \left(\frac{\sqrt{n}}{C\sqrt{d}} (\|\mu\|^2 - \|\mu\| \alpha) - \alpha \right)^2 \right), \quad (\text{Choose } C \geq \frac{16C_2}{c_9})
\end{aligned}$$

where the last line holds for $\frac{\alpha}{\|\mu\|} \leq \frac{\|\mu\|\sqrt{n}}{\sqrt{n}\|\mu\| + C\sqrt{d}}$, so that $\frac{\sqrt{n}}{C\sqrt{d}} (\|\mu\|^2 - \|\mu\| \alpha) - \alpha \geq 0$.

Similarly, we have

$$\begin{aligned}
L^{0/1}(\mathbf{W}^T) &\leq \beta + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}^T) | y_c = 1]}{\|\mathbf{W}^T\|_2} \right)^2 \right) \\
&\quad + \exp \left(-c\lambda \left(\frac{\mathbb{E}_{(x,y_c) \sim \mathcal{D}_c} [y_c f(x; \mathbf{W}^T) | y_c = -1]}{\|\mathbf{W}^T\|_2} \right)^2 \right) \\
&\leq \beta + 2 \exp \left(-\frac{c\lambda n \|\mu\|^4}{C^2 d} \left(1 - \frac{\alpha}{\|\mu\|} \right)^2 \right).
\end{aligned}$$

□

The proof of Theorem 3.3 builds upon a sequence of Propositions and Lemmas, which we show below.

Proposition B.12 defines a set \mathcal{G} to characterize the index of noise that have large variance, and show the size of set \mathcal{G} is large.

Proposition B.12. Let $\xi = [\xi^1, \dots, \xi^d]^\top$ denote the random vector sampled from $\mathcal{D}_{\text{clust}}$. Define $\mathcal{G} = \{i \in [d] | \mathbb{E}(\xi^i)^2 \geq \frac{\kappa}{2}\}$. Then the number of elements $|\mathcal{G}| \geq \frac{\kappa}{2-\kappa}d$.

Proof of Proposition B.12. Since each ξ^i has subgaussian norm at most 1, we have $2 \geq \mathbb{E} \exp((\xi^i)^2) \geq 1 + \mathbb{E}(\xi^i)^2$, so $\mathbb{E}(\xi^i)^2 \leq 1$. Suppose $|\mathcal{G}| < \frac{\kappa}{2-\kappa}d$, then we have

$$\kappa d \leq \mathbb{E} \|\xi\|^2 = \sum_{i \in \mathcal{G}} \mathbb{E}(\xi^i)^2 + \sum_{i \notin \mathcal{G}} \mathbb{E}(\xi^i)^2 \leq |\mathcal{G}| \cdot 1 + (d - |\mathcal{G}|) \cdot \frac{\kappa}{2} < d \cdot \frac{\kappa}{2} + (1 - \frac{\kappa}{2}) \frac{\kappa}{2-\kappa} d = \kappa d,$$

which is a contradiction. □

Lemma B.13 provides properties of the initialized network weights, similar to Lemma B.2 except we have an additional result that the averaged initialized weights variance that belongs to set \mathcal{G} are not small. This additional result will be used in proving Lemma B.15.

Lemma B.13. There is a universal constant $C_0 > 1$ such that with probability at least $1 - 3\delta/4$ over the random initialization,

$$\begin{aligned}
\frac{1}{2} \omega_{\text{init}}^2 d &\leq \|\mathbf{w}_s^0\|_2^2 \leq \frac{3}{2} \omega_{\text{init}}^2 d, \forall s \in [m]; \|\mathbf{W}^0\|_2 \leq C_0 \omega_{\text{init}} (\sqrt{m} + \sqrt{d}); \\
\frac{\sum_{i \in \mathcal{G}} (\mathbf{w}_{s,i}^0)^2}{|\mathcal{G}|} &\geq \frac{1}{4} \frac{\sum_{i \in [d]} (\mathbf{w}_{s,i}^0)^2}{d}, \forall s \in [m].
\end{aligned}$$

Here \mathcal{G} is defined in Proposition B.12.

Proof of Lemma B.13. For any fixed s , note that $\|\mathbf{w}_s^0\|_2^2$ is a ω_{init}^2 -multiple of a chi-squared random variable with d degrees of freedom. By concentration of the χ^2 distribution, for any $t \in (0, 1]$,

$$\mathbb{P}\left(\left|\frac{1}{d\omega_{\text{init}}^2}\|\mathbf{w}_s^0\|_2^2 - 1\right| \geq t\right) \leq 2\exp(-dt^2/8).$$

In particular, if we choose $t = 1/2$, with probability at least $1 - 2\exp(-d/32)$, we have

$$\frac{1}{2}\omega_{\text{init}}^2 d \leq \|\mathbf{w}_s^0\|_2^2 \leq \frac{3}{2}\omega_{\text{init}}^2 d.$$

Applying a union bound, with probability at least $1 - 2m\exp(-d/32)$, we have

$$\frac{1}{2}\omega_{\text{init}}^2 d \leq \|\mathbf{w}_s^0\|_2^2 \leq \frac{3}{2}\omega_{\text{init}}^2 d, \forall s \in [m].$$

Note that

$$\begin{aligned} & 1 - 2m\exp(-d/32) \\ & \geq 1 - 2\delta\exp(n/C - d/32) && \text{(Assumption (B7))} \\ & \geq 1 - 2\delta\exp(-d/64) \\ & \quad (d \geq 64n \text{ from Assumption (B3), Assumption (B2) and } C \text{ sufficiently large}) \\ & \geq 1 - \delta/4. \quad (d \geq 192 \text{ from Assumption (B3), Assumption (B2) and } C \text{ sufficiently large}) \end{aligned}$$

Therefore $\frac{1}{2}\omega_{\text{init}}^2 d \leq \|\mathbf{w}_s^0\|_2^2 \leq \frac{3}{2}\omega_{\text{init}}^2 d, \forall s \in [m]$ holds with probability at least $1 - \delta/4$.

From the proof of Lemma B.2, $\|\mathbf{W}^0\|_2 \leq C_0\omega_{\text{init}}(\sqrt{m} + \sqrt{d})$ holds with probability at least $1 - \delta/4$.

For the last part, for any fixed $s \in [m]$, by concentration of the χ^2 distribution, we have

$$\begin{aligned} & \mathbb{P}\left(\left|\frac{1}{|\mathcal{G}|\omega_{\text{init}}^2}\sum_{i \in \mathcal{G}}(\mathbf{w}_{s,i}^0)^2 - 1\right| \geq \frac{1}{2}\right) \leq 2\exp(-|\mathcal{G}|/32), \\ & \mathbb{P}\left(\left|\frac{1}{d\omega_{\text{init}}^2}\sum_{i \in [d]}(\mathbf{w}_{s,i}^0)^2 - 1\right| \geq 1\right) \leq 2\exp(-d/8). \end{aligned}$$

Applying a union bound over all $s \in [m]$,

$$\frac{\sum_{i \in \mathcal{G}}(\mathbf{w}_{s,i}^0)^2}{|\mathcal{G}|} \geq \frac{1}{2}\omega_{\text{init}}^2 \geq \frac{1}{4}\frac{\sum_{i \in [d]}(\mathbf{w}_{s,i}^0)^2}{d}$$

holds with probability at least

$$\begin{aligned} & 1 - 2m\exp(-|\mathcal{G}|/32) - 2m\exp(-d/8) \\ & \geq 1 - 2m\exp\left(-\frac{\kappa}{32(2-\kappa)}d\right) - 2m\exp(-d/8) && \text{(Proposition B.12)} \\ & \geq 1 - 2\delta\exp\left(n/C - \frac{\kappa}{32(2-\kappa)}d\right) - 2\delta\exp(n/C - d/8) && \text{(Assumption (B7))} \\ & \geq 1 - 2\delta\exp\left(-\left(\frac{\kappa}{32(2-\kappa)}C\log(2) - 1/C\right)\right) - 2\delta\exp(-(C\log(2)/8 - 1/C)) \\ & \quad (d \geq Cn^2\log(2) \text{ from Assumption (B3); } C \text{ sufficiently large}) \\ & \geq 1 - \delta/4. && (C \text{ sufficiently large}) \end{aligned}$$

The proof is complete by taking a union bound over the above claims. \square

The following anti-concentration inequality is helpful in proving Lemma B.15.

Proposition B.14. [Anti-concentration of subgaussian random variables] Assume X is a c -subgaussian random variable with $\mathbb{E}[X] = 0, \mathbb{E}[X^2] = 1$, then

$$\mathbb{P}\left(X \geq \frac{1}{10c^2}\right) \geq \frac{0.04}{220^{\frac{2}{3}}c^4}.$$

Proof of Lemma B.14. Denote $a = \frac{1}{10c^2}$, $A = 220^{\frac{1}{3}}c^2$ and $B = \frac{0.04}{A^2}$. From $2 \geq \mathbb{E} \exp\left(\frac{X^2}{c^2}\right) \geq 1 + \mathbb{E} \frac{X^2}{c^2}$ we know $c \geq 1$. Consider a truncated version of X defined by $\tilde{X} = X \cdot \mathbb{1}_{a \leq |X| \leq A}$. We have

$$\mathbb{E} \tilde{X} = \mathbb{E} \tilde{X} - \mathbb{E} X = -\mathbb{E}(X \cdot \mathbb{1}_{|X| < a}) - \mathbb{E}(X \cdot \mathbb{1}_{|X| > A}).$$

It is trivial that $|\mathbb{E}(X \cdot \mathbb{1}_{|X| < a})| \leq a$. By subgaussian tail bound one may compute

$$\begin{aligned} & |\mathbb{E}(X \cdot \mathbb{1}_{|X| > A})| \\ & \leq \mathbb{E}(|X| \cdot \mathbb{1}_{|X| > A}) \\ & \leq \int_0^A \mathbb{P}(|X| > A) dt + \int_A^\infty \mathbb{P}(|X| > t) dt \\ & \leq 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \int_A^\infty 2 \exp\left(-\frac{t^2}{c^2}\right) dt \\ & \leq 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right). \end{aligned}$$

These imply

$$|\mathbb{E} \tilde{X}| \leq a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right).$$

Similarly,

$$\begin{aligned} 1 = \mathbb{E} X^2 &= \mathbb{E}(X^2 \cdot \mathbb{1}_{|X| < a}) + \mathbb{E}(X^2 \cdot \mathbb{1}_{|X| > A}) + \mathbb{E} \tilde{X}^2 \\ &\leq a^2 + 2A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) + 2c^2 \exp\left(-\frac{A^2}{c^2}\right) + \mathbb{E} \tilde{X}^2. \end{aligned}$$

Thus,

$$|\mathbb{E} \tilde{X}^2 - 1| \leq a^2 + 2A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) + 2c^2 \exp\left(-\frac{A^2}{c^2}\right).$$

Let $\tilde{X}_+ = \max(\tilde{X}, 0)$ and $\tilde{X}_- = \max(-\tilde{X}, 0)$, then \tilde{X}_+ and \tilde{X}_- are non-negative and $\tilde{X} = \tilde{X}_+ - \tilde{X}_-$, $\tilde{X}^2 = \tilde{X}_+^2 + \tilde{X}_-^2$. We thus have

$$\begin{aligned} |\mathbb{E} \tilde{X}_+ - \mathbb{E} \tilde{X}_-| &\leq a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right), \\ \mathbb{E} \tilde{X}_+^2 + \mathbb{E} \tilde{X}_-^2 &\geq 1 - a^2 - 2A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) - 2c^2 \exp\left(-\frac{A^2}{c^2}\right). \end{aligned}$$

Now assume to the contrary that $\mathbb{P}(X \geq a) < B$. We have

$$\begin{aligned} \mathbb{E} \tilde{X}_+ &\leq A \cdot \mathbb{P}(|X| \geq a) < A \cdot B, \\ \mathbb{E} \tilde{X}_+^2 &< A^2 \cdot B, \\ \mathbb{E} \tilde{X}_- &\leq \mathbb{E} \tilde{X}_+ + a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right) \\ &< A \cdot B + a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right), \\ \mathbb{E} \tilde{X}_-^2 &\leq A \cdot \mathbb{E} \tilde{X}_- < A \left(A \cdot B + a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right) \right). \end{aligned}$$

These together imply

$$\begin{aligned} \mathbb{E} \tilde{X}_+^2 + \mathbb{E} \tilde{X}_-^2 &< A^2 \cdot B + A \left(A \cdot B + a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right) \right) \\ &< 1 - a^2 - 2A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) - 2c^2 \exp\left(-\frac{A^2}{c^2}\right), \end{aligned}$$

which is a contradiction. The last inequality holds since

$$\begin{aligned}
RHS - LHS &= 1 - a^2 - 2A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) - 2c^2 \exp\left(-\frac{A^2}{c^2}\right) \\
&\quad - A^2 \cdot B - A \left(A \cdot B + a + 2A \cdot \exp\left(-\frac{A^2}{c^2}\right) + \frac{c^2}{A} \exp\left(-\frac{A^2}{c^2}\right) \right) \\
&= 1 - a^2 - 2A^2 \cdot B - A \cdot a - 4A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) - 3c^2 \exp\left(-\frac{A^2}{c^2}\right) \\
&= 1 - \frac{1}{100c^4} - 0.08 - A \cdot a - 4A^2 \cdot \exp\left(-\frac{A^2}{c^2}\right) - 3c^2 \exp\left(-\frac{A^2}{c^2}\right) \\
&\geq 1 - \frac{1}{100} - 0.08 - A \cdot a - 8A^2 / \left(\frac{A^2}{c^2}\right)^2 - 3c^2 / \frac{A^2}{c^2} \\
&= 1 - 0.09 - 1.5(22c^4 a^2)^{\frac{1}{3}} = 1 - 0.09 - 1.5(0.22)^{\frac{1}{3}} > 0.
\end{aligned}$$

This completes the proof. \square

By combining Proposition B.12, Lemma B.13, and Proposition B.14, we establish the existence of enough neurons with positive activation at the first step of adversarial training, as stated in Lemma B.15. This lemma plays a vital role throughout the entire proof of networks with non-smooth activation functions.

Lemma B.15. Suppose the events in Lemma B.3 and Lemma B.13 hold. Given Assumption 2, there exists a constant $c_0 > 0$ that only depends on κ such that with probability at least $1 - \delta/5$,

$$\begin{aligned}
\forall s \in [m], |\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| &\geq c_0 n; \\
\forall i \in [n], |\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| &\geq c_0 m.
\end{aligned}$$

Proof of Lemma B.15. Fix any given (x_i, y_i) . Recall that $\mathbf{x}_i = y_i^c \mu + \xi_i$, where y_i^c is the clean label, $\xi_i \sim \mathcal{D}_{\text{clust}}$ is the noise. Note that $\langle \mathbf{w}_s^0, \mathbf{x}_i \rangle = y_i^c \langle \mathbf{w}_s^0, \mu \rangle + \langle \mathbf{w}_s^0, \xi_i \rangle$. The first term is a centered Gaussian with variance $\omega_{\text{init}}^2 \|\mu\|^2$, therefore applying concentration argument gives us with probability at least $1 - \delta/20$, $\max_{s \in [m]} |\langle \mathbf{w}_s^0, \mu \rangle| \leq 4\omega_{\text{init}} \|\mu\| \sqrt{\log(m/\delta)}$. For the second term, condition on ξ_i , which is a centered Gaussian with variance $\omega_{\text{init}}^2 \|\xi_i\|^2$, $\forall s \in [m]$. Since $\mathbb{P}\left(\langle \mathbf{w}_s^0, \xi_i \rangle \geq \frac{\omega_{\text{init}} \|\xi_i\|}{10}\right) \geq \frac{1}{5}$, applying the Hoeffding's inequality gives us with probability at least $1 - \exp(-m/225)$, there exists a subset $J_i \in [m]$ with $|J_i| \geq m/15$ such that $\langle \mathbf{w}_s^0, \xi_i \rangle \geq \frac{\omega_{\text{init}} \|\xi_i\|}{10}$ and $a_s = y_i, \forall s \in J_i$. Conditioning on $\|\xi_i\|^2 \geq \frac{\kappa d}{2}$ obtained in Lemma B.3, we have that

$$\begin{aligned}
\langle \mathbf{w}_s^0, \mathbf{x}_i \rangle - \alpha \|\mathbf{w}_s^0\| &\geq -4\omega_{\text{init}} \|\mu\| \sqrt{\log(m/\delta)} + \omega_{\text{init}} \|\xi_i\| / 10 - 2\omega_{\text{init}} \sqrt{d} \alpha \quad (\text{Proposition B.13}) \\
&\geq \omega_{\text{init}} \left(\sqrt{\kappa d}/20 - 4 \|\mu\| \sqrt{\log(m/\delta)} - 2\sqrt{d} \alpha \right) \\
&\geq \omega_{\text{init}} \left(\sqrt{\kappa d}/20 - 4\sqrt{d}/C - 2\sqrt{d}/\sqrt{C} \right) > 0,
\end{aligned}$$

where the last line holds via Assumption (B3) and (B7) for large enough C as well as $\alpha \leq \sqrt{n/d} \|\mu\| \leq \frac{1}{\sqrt{C}}$. Combining the above arguments, we have with probability at least $1 - \delta/20 - n \exp(-m/225)$,

$$\forall i \in [n], |\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle > \alpha \|\mathbf{w}_s^0\|\}| \geq \frac{m}{15}.$$

Given $m \geq C \log(n/\delta)$, the above holds with probability at least $1 - \delta/10$.

For the other statement, we can condition on \mathbf{w}_s^0 similarly. Denote $X_i = \left\langle \frac{\mathbf{w}_s^0}{\|\mathbf{w}_s^0\|_2}, \xi_i \right\rangle, Y_i = \frac{X_i}{\sqrt{\mathbb{E}X_i^2}}$.

It is obvious that $\mathbb{E}Y_i = 0$ and $\mathbb{E}Y_i^2 = 1$.

$$\mathbb{E}X_i^2 = \frac{\sum_{j \in [d]} (\mathbf{w}_{s,j}^0)^2 \mathbb{E}(\xi_i^j)^2}{\|\mathbf{w}_s^0\|_2^2}$$

$$\begin{aligned}
&\geq \frac{\sum_{j \in \mathcal{G}} (\mathbf{w}_{s,j}^0)^2 \mathbb{E}(\xi_i^j)^2}{\sum_{j \in [d]} (\mathbf{w}_{s,j}^0)^2} && \text{(here } \mathcal{G} \text{ is defined in Proposition B.12)} \\
&\geq \frac{|\mathcal{G}| \cdot \kappa/2}{4d} \geq \frac{\kappa^2}{16 - 8\kappa}. && \text{(Lemma B.13 and Proposition B.12)}
\end{aligned}$$

From Proposition 2.6.1 in Vershynin (2018), there exists a universal constant c^1 such that $\|Y_i\|_{\psi_2} = \frac{\|X_i\|_{\psi_2}}{\sqrt{\mathbb{E}X_i^2}} \leq \frac{c^1 \sqrt{16-8\kappa}}{\kappa}$. Applying Proposition B.14 gives us that

$$\mathbb{P}\left(X_i \geq \frac{\kappa^3}{10(c^1)^2(16-8\kappa)^{\frac{3}{2}}}\right) \geq \mathbb{P}\left(Y_i \geq \frac{\kappa^2}{10(c^1)^2(16-8\kappa)}\right) \geq 0.04 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right).$$

Therefore, $\langle \mathbf{w}_s^0, \xi_i \rangle \geq \frac{\kappa^3}{10(c^1)^2(16-8\kappa)^{\frac{3}{2}}} \|\mathbf{w}_s^0\|$ holds with probability at least $0.04 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right)$. Consider y_i^c be the clean label that is uniformly distributed on $\{-1, +1\}$ and is independent of ξ_i , then we have

$$\begin{aligned}
&\mathbb{P}\left(\langle \mathbf{w}_s^0, \xi_i \rangle \geq \frac{\kappa^3}{10(c^1)^2(16-8\kappa)^{\frac{3}{2}}} \|\mathbf{w}_s^0\|, a_s = y_i^c | \mathbf{w}_s^0\right) \\
&= \frac{1}{2} \mathbb{P}\left(\langle \mathbf{w}_s^0, \xi_i \rangle \geq \frac{\kappa^3}{10(c^1)^2(16-8\kappa)^{\frac{3}{2}}} \|\mathbf{w}_s^0\| | \mathbf{w}_s^0\right) \\
&\geq 0.02 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right).
\end{aligned}$$

Similar with the previous part, since $\omega_{\text{init}}^2 3d/2 \geq \|\mathbf{w}_s^0\|^2 \geq \omega_{\text{init}}^2 d/2$ holds, applying Hoeffding's inequality, with probability at least $1 - \delta/20 - m \exp\left(-5 \times 10^{-5} \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right)^2 n\right)$,

$$\forall s \in [m], |\{i \in [n] : y_i^c = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle > \alpha \|\mathbf{w}_s^0\|\}| \geq 0.015 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right) n.$$

When $n \geq C \log(m/\delta)$ as assumed in Assumption (B7), the above holds with probability at least $1 - \delta/10$. Note that $|\{i \in [n] : y_i \neq y_i^c\}| \leq (1/C + \sqrt{2/C})n \leq 0.005 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right) n$ holds for a sufficient large C . Thus,

$$\forall s \in [m], |\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle > \alpha \|\mathbf{w}_s^0\|\}| \geq 0.01 \left(\frac{\kappa^4}{220^{\frac{2}{3}}(c^1)^4(16-8\kappa)^2}\right) n.$$

The proof is complete by taking a union bound over the above two claims. \square

Remark B.16. Given a designed network initialization to be $\mathbf{w}_s^0 = a_s \omega_{\text{init}} \sqrt{d} \frac{\hat{\mu}}{\|\hat{\mu}\|}$, where $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n y_i \mathbf{x}_i$, then we can relax Assumption (B8) to be $\alpha \leq \mathcal{O}(\|\mu\|)$. Under the modified Assumption 2 and suppose the events in Lemma B.3 hold, Lemma B.15 still holds; i.e., there exists a constant $c_0 > 0$ that only depends on κ such that with probability at least $1 - \delta/5$,

$$\begin{aligned}
&\forall s \in [m], |\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| \geq c_0 n; \\
&\forall i \in [n], |\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| \geq c_0 m.
\end{aligned}$$

Proof of Remark B.16. If $y_i = a_s$, we have,

$$\left\langle \frac{\mathbf{w}_s^0}{\|\mathbf{w}_s^0\|}, \mathbf{x}_i \right\rangle = \left\langle a_s \frac{\sum_{j=1}^n y_j \mathbf{x}_j}{\|\sum_{j=1}^n y_j \mathbf{x}_j\|}, \mathbf{x}_i \right\rangle$$

$$\begin{aligned}
&\geq \frac{\|\mathbf{x}_i\|^2}{\left\| \sum_{j=1}^n y_j \mathbf{x}_j \right\|} - \sum_{j \neq i} \frac{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}{\left\| \sum_{j=1}^n y_j \mathbf{x}_j \right\|} \\
&\geq \frac{d/C_1}{\left\| \sum_{j=1}^n y_j \mathbf{x}_j \right\|} - \sum_{j \neq i} \frac{C_1(\|\mu\|^2 + \sqrt{d \log(n/\delta)})}{\left\| \sum_{j=1}^n y_j \mathbf{x}_j \right\|} \\
&\geq \frac{d/(2C_1)}{3\sqrt{dn}} \\
&\geq \|\mu\| \geq \alpha.
\end{aligned}$$

Therefore, $|\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| = |\{s \in [m] : y_i = a_s\}| = 0.5m$.
 $|\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \mathbf{x}_i \rangle \geq \alpha \|\mathbf{w}_s^0\|\}| = |\{i \in [n] : y_i = a_s\}| \geq (0.4 - \beta)n \geq 0.3n$ holds with probability at least $1 - \delta/5$. \square

Definition. If the events in Lemma B.3, B.13 and B.15 occur, let us say that we have a good run.

A good run occurs with probability at least $1 - 2\delta$. In the following proof, we condition on a good run occurs.

The following proposition presents several properties of the distribution $\mathcal{D}_{\text{clust}}$, which are crucial for establishing the generalization guarantees.

Proposition B.17. Assume $\xi \sim \mathcal{D}_{\text{clust}}$. Then the following holds:

- (D1) For any fixed $\mathbf{v} \in \mathbb{R}^d$, for any $\bar{\delta} < 0.5$, with probability at least $1 - \bar{\delta}$ w.r.t ξ , $|\langle \mathbf{v}, \xi \rangle| \leq c_6 \|\mathbf{v}\| \sqrt{\log(1/\bar{\delta})}$, where c_6 is an absolute constant.
- (D2) For any $\bar{\delta} < 0.5$, with probability at least $1 - \bar{\delta}$ w.r.t ξ , $\|\xi\|^2 \leq 9d(\log(1/\bar{\delta}))$. In particular, denote event $\mathcal{E} = \{\|\xi\| \geq 6\sqrt{\log(1/\bar{\delta})}d\}$, we have $\mathbb{E}[\|\xi\| \mathbf{1}(\mathcal{E})] \leq 8\bar{\delta}^4 \cdot \sqrt{d \cdot \log(1/\bar{\delta})}$.

Proof of B.17. We first prove (D1). Note that the coordinates of ξ are independent variables with ψ_2 norm at most 1. From Hoeffding's inequality, there exists a universal constant c such that

$$\mathbb{P}\left(|\langle \mathbf{v}, \xi \rangle| > c_6 \|\mathbf{v}\| \sqrt{\log(1/\bar{\delta})}\right) \leq 2 \exp\left(-cc_6^2 \|\mathbf{v}\|^2 \log(1/\bar{\delta}) / \|\mathbf{v}\|^2\right) = 2\bar{\delta}^{cc_6^2}.$$

By selecting $c_6 = \sqrt{\frac{2}{c}}$, we get

$$\mathbb{P}\left(|\langle \mathbf{v}, \xi \rangle| > c_6 \|\mathbf{v}\| \sqrt{\log(1/\bar{\delta})}\right) \leq 2\bar{\delta}^2 \leq \bar{\delta}.$$

Next we prove (D2). Note that the coordinates of $\xi = [\xi^1, \xi^2, \dots, \xi^d]^\top$ are independent variables with ψ_2 norm at most 1. From Bernstein's inequality, there exists a universal constant c such that for every $t > 0$,

$$\mathbb{P}\left(\|\xi\|^2 - \mathbb{E}\|\xi\|^2 > t\right) \leq \exp\left(-c \min\left\{\frac{t^2}{d}, t\right\}\right).$$

Since $2 \geq \mathbb{E} \exp((\xi^i)^2) \geq 1 + \mathbb{E}(\xi^i)^2$, we have $\mathbb{E}\|\xi\|^2 = \sum_{i=1}^d \mathbb{E}(\xi^i)^2 \leq d$.

Select $t = \max \left\{ \sqrt{\frac{d \log(1/\bar{\delta})}{c}}, \frac{\log(1/\bar{\delta})}{c} \right\}$, then $\mathbf{P} \left(\|\xi\|^2 - \mathbb{E} \|\xi\|^2 > t \right) \leq \bar{\delta}$. Therefore, with probability at least $1 - \bar{\delta}$ w.r.t ξ ,

$$\begin{aligned} \|\xi\|^2 &\leq \mathbb{E} \|\xi\|^2 + \max \left\{ \sqrt{\frac{d \log(1/\bar{\delta})}{c}}, \frac{\log(1/\bar{\delta})}{c} \right\} \\ &\leq d + \sqrt{\frac{d \log(1/\bar{\delta})}{c}} + \frac{\log(1/\bar{\delta})}{c} \\ &\leq 9d \log(1/\bar{\delta}). \quad (d \text{ is sufficiently large since } C \text{ is sufficiently large; } \bar{\delta} < 0.5) \end{aligned}$$

As for the last statement,

$$\begin{aligned} &\mathbb{E} [\|\xi\| \mathbf{1}(\mathcal{E})] \\ &= \int_0^\infty \mathbf{P}(\|\xi\| \mathbf{1}(\mathcal{E}) \geq t) dt \\ &= \int_0^{6\sqrt{\log(1/\bar{\delta})d}} \mathbf{P}(\|\xi\| \mathbf{1}(\mathcal{E}) \geq t) dt + \int_{6\sqrt{\log(1/\bar{\delta})d}}^\infty \mathbf{P}(\|\xi\| \mathbf{1}(\mathcal{E}) \geq t) dt \\ &\leq 6\sqrt{\log(1/\bar{\delta})d} \cdot \mathbf{P} \left(\|\xi\| \mathbf{1}(\mathcal{E}) \geq 6\sqrt{\log(1/\bar{\delta})d} \right) + \int_{6\sqrt{\log(1/\bar{\delta})d}}^\infty \exp\left(-\frac{t^2}{9d}\right) dt \\ &\quad \left(\mathbf{P} \left(\|\xi\|^2 > 9d \log(1/\bar{\delta}) \right) \leq \bar{\delta} \right) \\ &\leq 6\sqrt{\log(1/\bar{\delta})d} \cdot \bar{\delta}^4 + \int_0^\infty \exp\left(-\frac{\left(t + 6\sqrt{\log(1/\bar{\delta})d}\right)^2}{9d}\right) dt \\ &\leq 6\sqrt{\log(1/\bar{\delta})d} \cdot \bar{\delta}^4 + \int_0^\infty \exp\left(-\frac{12t\sqrt{\log(1/\bar{\delta})d} + 36\log(1/\bar{\delta})d}{9d}\right) dt \\ &= 6\sqrt{\log(1/\bar{\delta})d} \cdot \bar{\delta}^4 + \bar{\delta}^4 \frac{9d}{12\sqrt{\log(1/\bar{\delta})d}} \\ &\leq 8\bar{\delta}^4 \cdot \sqrt{d \cdot \log(1/\bar{\delta})}. \quad (\bar{\delta} < 0.5) \end{aligned}$$

□

Before proceeding to the next Lemma, we define some important notations which will be used frequently later. We define

$$\lambda_i(\mathbf{x}; \mathbf{W}^t) = \frac{1}{m} \sum_{s=1}^m \frac{\phi \left(\left\langle \mathbf{w}_s^t - \eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle \right) - \phi(\langle \mathbf{w}_s^t, \mathbf{x} \rangle)}{\left\langle -\eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle} \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle),$$

so that the following holds:

$$\begin{aligned} &yf(\mathbf{x}; \mathbf{W}^{t+1}) - yf(\mathbf{x}; \mathbf{W}^t) \\ &= \frac{1}{\sqrt{m}} \sum_{s=1}^m a_s \frac{\phi \left(\left\langle \mathbf{w}_s^t - \eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle \right) - \phi(\langle \mathbf{w}_s^t, \mathbf{x} \rangle)}{\left\langle -\eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle} \left\langle -\eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), y\mathbf{x} \right\rangle \\ &= \frac{\eta}{mn} \sum_{i=1}^n \sum_{s=1}^m \frac{\phi \left(\left\langle \mathbf{w}_s^t - \eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle \right) - \phi(\langle \mathbf{w}_s^t, \mathbf{x} \rangle)}{\left\langle -\eta \nabla_{\mathbf{w}_s} \widehat{\mathcal{L}}_{\text{rob}}(\mathbf{W}^t), \mathbf{x} \right\rangle} \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle y_i \tilde{\mathbf{x}}_i^t, y\mathbf{x} \rangle \end{aligned}$$

$$= \frac{\eta}{n} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \lambda_i(\mathbf{x}; \mathbf{W}^t) \langle y_i \tilde{\mathbf{x}}_i^t, y \mathbf{x} \rangle.$$

We define $\mathcal{A}(t)$ as the set of pairs (i, s) such that the neurons s is active for the adversarial examples $\forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$; i.e.,

$$\mathcal{A}(t) := \{(i, s) \in [n] \times [m] : \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle > 0, \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)\}.$$

Denote its coordinate as

$$\begin{aligned} \mathcal{A}^i(t) &:= \{s \in [m] : (i, s) \in \mathcal{A}(t)\}, \\ \mathcal{A}_s(t) &:= \{i \in [n] : (i, s) \in \mathcal{A}(t)\}. \end{aligned}$$

Proposition B.18. For any pair $(i, s) \in \mathcal{A}(t)$, we have $\forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$, $\phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) \geq \gamma$. Moreover, $\lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \geq \gamma^2 |\mathcal{A}^i(t) \cap \mathcal{A}^i(t+1)| / m$.

Proof. The definition of $\mathcal{A}(t)$ implies that $\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle > 0, \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)$. By the definition of activation function we have $\phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) \geq \gamma$.

The definition of $\lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t)$ gives us the following:

$$\begin{aligned} \lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) &\geq \frac{\gamma}{m} \sum_{s \in \mathcal{A}^i(t) \cap \mathcal{A}^i(t+1)} \frac{\phi(\langle \mathbf{w}_s^{t+1}, \tilde{\mathbf{x}}_i \rangle) - \phi(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle)}{\langle \mathbf{w}_s^{t+1}, \tilde{\mathbf{x}}_i \rangle - \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle} \\ &\geq \frac{\gamma^2}{m} |\mathcal{A}^i(t) \cap \mathcal{A}^i(t+1)|. \end{aligned}$$

□

We further define $\mathcal{T} := \{(i, s) \in [n] \times [m] : y_i = a_s\}$ and similarly we denote

$$\begin{aligned} \mathcal{T}^i &:= \{s \in [m] : (i, s) \in \mathcal{T}\}, \\ \mathcal{T}_s &:= \{i \in [n] : (i, s) \in \mathcal{T}\} \end{aligned}$$

Lemma B.19. On a good run we have

$$\begin{aligned} |\mathcal{A}^i(0) \cap \mathcal{T}^i| &\geq c_0 m, \\ |\mathcal{A}_s(0) \cap \mathcal{T}_s| &\geq c_0 n. \end{aligned}$$

Proof of Lemma B.19. Lemma B.15 implies that $\forall i \in [n], s \in [m]$,

$$\begin{aligned} |\{i \in [n] : y_i = a_s, \langle \mathbf{w}_s^0, \tilde{\mathbf{x}}_i \rangle > 0, \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)\}| &\geq c_0 n, \\ |\{s \in [m] : y_i = a_s, \langle \mathbf{w}_s^0, \tilde{\mathbf{x}}_i \rangle > 0, \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)\}| &\geq c_0 m. \end{aligned}$$

Combine with the definition of \mathcal{A} and \mathcal{T} conclude the proof. □

In the following Lemma, we will 1) prove the number of neurons with positive activation increases as the training epochs increases; 2) provide both an upper bound and a lower bound on the increment in the un-normalized margin for arbitrary adversarial training examples; 3) show the loss g is at the same scale across all adversarial training examples. An analog of Lemma B.20 is Lemma B.7 for neural networks with smooth activation functions.

Lemma B.20. On a good run, there exists a constant $C_r > 0$ that only depends on κ, γ such that for any $t \geq 0$, we have

$$(E1) \quad \mathcal{A}(t) \cap \mathcal{T} \subset \mathcal{A}(t+1) \cap \mathcal{T}.$$

$$(E2) \quad \frac{\eta c_0 \gamma^2 \tilde{g}_i(\mathbf{W}^t)}{2n} \left(\sqrt{\frac{d}{C_1}} - \alpha \right)^2 \leq y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^{t+1}) - y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \leq \frac{3\eta}{n} (C_1 d + \alpha^2) \tilde{g}_i(\mathbf{W}^t), \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha), \forall i \in [n].$$

$$(E3) \quad \max_{i, j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq C_r.$$

Proof of Lemma B.20. We prove via induction. Recall $\tilde{\mathbf{x}}_i^t = \arg \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t))$. Similar as network with smooth activation functions, we have the following

$$\max_{i,j \in [n]} \frac{g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{g(y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq \max \left(2, 2 \cdot \max_{i,j \in [n]} \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \right).$$

Therefore we only need $\max_{i,j \in [n]} \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^t; \mathbf{W}^t))} \leq C_r/2$ to hold. Note that at initialization,

$$\max_{i,j \in [n]} \frac{\exp(-y_i f(\tilde{\mathbf{x}}_i^0; \mathbf{W}^0))}{\exp(-y_j f(\tilde{\mathbf{x}}_j^0; \mathbf{W}^0))} = 1 \leq C_r/2.$$

Without loss of generality, we choose $i = 1, j = 2$. Through induction, at iteration t , we have $\frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \leq C_r/2$. Now we are proving with the following order: (E1), (E2) and (E3) for $t + 1$. Given for any $(i, s) \in \mathcal{A}(t) \cap \mathcal{T}$, we have $y_i a_s = 1$ and $\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle > 0, \forall \tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha) > 0$ by definition, we have

$$\begin{aligned} & \langle \mathbf{w}_s^{t+1}, \tilde{\mathbf{x}}_i \rangle - \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle \\ &= \frac{\eta}{n\sqrt{m}} \sum_{k=1}^n y_k a_s \tilde{g}_k(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_k^t \rangle) \langle \tilde{\mathbf{x}}_k^t, \tilde{\mathbf{x}}_i \rangle \\ &= \frac{\eta}{n\sqrt{m}} \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \tilde{\mathbf{x}}_i^t, \tilde{\mathbf{x}}_i \rangle + \frac{\eta}{n\sqrt{m}} \sum_{k \neq i} y_k a_s \tilde{g}_k(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_k^t \rangle) \langle \tilde{\mathbf{x}}_k^t, \tilde{\mathbf{x}}_i \rangle \\ &\geq \frac{\eta}{n\sqrt{m}} \frac{\hat{G}_{\text{rob}}(\mathbf{W}^t)}{C_r} \gamma \left(\sqrt{d/C_1} - \alpha \right)^2 \\ &\quad - \frac{\eta}{\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \\ &\geq \frac{\eta}{n\sqrt{m}} \frac{\hat{G}_{\text{rob}}(\mathbf{W}^t)}{C_r} \gamma \left(\sqrt{d/C_1} - \alpha \right)^2 - \frac{\eta}{n\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) d \left(\frac{C_1 + 1}{C} + \frac{\sqrt{5} C_1}{\sqrt{C}} \right) \\ &\hspace{15em} \text{(Assumption (B3))} \\ &\geq \frac{\eta}{2n\sqrt{m}} \frac{\hat{G}_{\text{rob}}(\mathbf{W}^t)}{C_r} \gamma \left(\sqrt{d/C_1} - \alpha \right)^2 > 0. \\ &\hspace{15em} (\alpha \leq \sqrt{n/d} \|\mu\| \leq \|\mu\| \leq \frac{1}{2} \sqrt{d/C_1}; C \text{ sufficiently large}) \end{aligned}$$

This implies that $(i, s) \in \mathcal{A}(t+1) \cap \mathcal{T}$, therefore (E1) holds.

Next we consider the following

$$\begin{aligned} & y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^{t+1}) - y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \\ &= \frac{\eta}{n} \sum_{k=1}^n \tilde{g}_k(\mathbf{W}^t) \lambda_k(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle y_k \tilde{\mathbf{x}}_k^t, y_i \tilde{\mathbf{x}}_i \rangle \\ &= \frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle \tilde{\mathbf{x}}_i^t, \tilde{\mathbf{x}}_i \rangle + \frac{\eta}{n} \sum_{k \neq i} \tilde{g}_k(\mathbf{W}^t) \lambda_k(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle y_k \tilde{\mathbf{x}}_k^t, y_i \tilde{\mathbf{x}}_i \rangle. \end{aligned}$$

The first term gives us

$$\begin{aligned} \frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle \tilde{\mathbf{x}}_i^t, \tilde{\mathbf{x}}_i \rangle &\geq \frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \left(\sqrt{d/C_1} - \alpha \right)^2 \\ &\geq \frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \frac{\gamma^2}{m} |\mathcal{A}^i(t) \cap \mathcal{A}^i(t+1)| \left(\sqrt{d/C_1} - \alpha \right)^2 \\ &\hspace{15em} \text{(Proposition B.18)} \\ &\geq \frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \frac{\gamma^2}{m} |\mathcal{A}^i(0) \cap \mathcal{T}^i| \left(\sqrt{d/C_1} - \alpha \right)^2 \quad \text{(E1)} \\ &\geq \frac{\eta \gamma^2 c_0}{n} \left(\sqrt{d/C_1} - \alpha \right)^2 \tilde{g}_i(\mathbf{W}^t). \quad \text{(Lemma B.19)} \end{aligned}$$

On the other hand,

$$\frac{\eta}{n} \tilde{g}_i(\mathbf{W}^t) \lambda_i(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle \tilde{\mathbf{x}}_i^t, \tilde{\mathbf{x}}_i \rangle \leq \frac{\eta}{n} \left(\sqrt{C_1 d} + \alpha \right)^2 \tilde{g}_i(\mathbf{W}^t).$$

The second terms tells us that

$$\begin{aligned} & \frac{\eta}{n} \left| \sum_{k \neq i} \tilde{g}_k(\mathbf{W}^t) \lambda_k(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle y_k \tilde{\mathbf{x}}_k^t, y_i \tilde{\mathbf{x}}_i \rangle \right| \\ & \leq \eta \left(C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t) \\ & \leq \frac{\eta d \tilde{g}_i(\mathbf{W}^t)}{n}. \quad (\widehat{G}_{\text{rob}}(\mathbf{W}^t) \leq C_r \tilde{g}_i(\mathbf{W}^t); \text{assumption (B3) for large enough } C) \end{aligned}$$

Similarly,

$$\begin{aligned} & \frac{\eta}{n} \sum_{k \neq i} \tilde{g}_k(\mathbf{W}^t) \lambda_k(\tilde{\mathbf{x}}_i; \mathbf{W}^t) \langle y_k \tilde{\mathbf{x}}_k^t, y_i \tilde{\mathbf{x}}_i \rangle \\ & \geq -\eta \left(C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t) \\ & \geq -\eta C_r \left(C_1 \left(\|\mu\|^2 + \sqrt{d \log(n/\delta)} \right) + 2\alpha \sqrt{C_1 d} + \alpha^2 \right) \tilde{g}_i(\mathbf{W}^t) \quad (\widehat{G}_{\text{rob}}(\mathbf{W}^t) \leq C_r \tilde{g}_i(\mathbf{W}^t)) \\ & \geq -\frac{\eta \gamma^2 c_0}{2n} \left(0.5 \sqrt{d/C_1} \right)^2 \tilde{g}_i(\mathbf{W}^t) \quad (\text{Assumption (B3) for large enough } C) \\ & \geq -\frac{\eta \gamma^2 c_0}{2n} \left(\sqrt{d/C_1} - \alpha \right)^2 \tilde{g}_i(\mathbf{W}^t). \end{aligned}$$

Summing the two terms together we have (E2) holds. This implies that $y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \geq 0, \forall t \geq 0, i \in [n]$. Applying the above gives us the following,

$$\begin{aligned} & \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\ & = \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \cdot \frac{\exp(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t) - y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t) - y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\ & \leq \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \cdot \frac{\exp(y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^t) - y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t) - y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\ & \leq \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \cdot \exp \left(\frac{3\eta}{n} (C_1 d + \alpha^2) \tilde{g}_2(\mathbf{W}^t) - \frac{\eta c_0 \gamma^2}{2n} \left(\sqrt{\frac{d}{C_1}} - \alpha \right)^2 \tilde{g}_1(\mathbf{W}^t) \right) \\ & = \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \cdot \exp \left(\frac{\eta c_0 \gamma^2}{2n} \left(\sqrt{\frac{d}{C_1}} - \alpha \right)^2 \tilde{g}_2(\mathbf{W}^t) \left(\frac{3\eta}{n} (C_1 d + \alpha^2) - \frac{\tilde{g}_1(\mathbf{W}^t)}{\tilde{g}_2(\mathbf{W}^t)} \right) \right), \end{aligned}$$

where the first inequality holds since $\exp(y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^t)) \geq \exp(y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))$, $\exp(y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1})) \leq \exp(y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^{t+1}))$ by the definition of $\tilde{\mathbf{x}}_1^t, \tilde{\mathbf{x}}_2^{t+1}$.

If $\frac{\tilde{g}_1(\mathbf{W}^t)}{\tilde{g}_2(\mathbf{W}^t)} \geq \frac{\frac{3\eta}{n} (C_1 d + \alpha^2)}{\frac{\eta c_0 \gamma^2}{2n} \left(\sqrt{d/C_1} - \alpha \right)^2} = \frac{6(C_1 d + \alpha^2)}{c_0 \gamma^2 \left(\sqrt{d/C_1} - \alpha \right)^2}$, then $\frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \leq \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^t; \mathbf{W}^t))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^t; \mathbf{W}^t))} \leq C_r/2$. Otherwise we have

$$\begin{aligned} & \frac{\exp(-y_1 f(\tilde{\mathbf{x}}_1^{t+1}; \mathbf{W}^{t+1}))}{\exp(-y_2 f(\tilde{\mathbf{x}}_2^{t+1}; \mathbf{W}^{t+1}))} \\ & \leq \frac{2\tilde{g}_1(\mathbf{W}^t)}{\tilde{g}_2(\mathbf{W}^t)} \cdot \exp \left(\frac{3\eta}{n} (C_1 d + \alpha^2) \tilde{g}_2(\mathbf{W}^t) - \frac{\eta c_0 \gamma^2}{2n} \left(\sqrt{d/C_1} - \alpha \right)^2 \tilde{g}_1(\mathbf{W}^t) \right) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{12(C_1 d + \alpha^2)}{c_0 \gamma^2 (\sqrt{d/C_1} - \alpha)^2} \exp\left(\frac{3\eta}{n}(C_1 d + \alpha^2)\right) && (\tilde{g}_i(\mathbf{W}^t) \leq 1) \\
&\leq \frac{24(C_1 d + \alpha^2)}{c_0 \gamma^2 (\sqrt{d/C_1} - \alpha)^2} \leq C_r/2,
\end{aligned}$$

where the last line holds due to $\exp\left(\frac{3\eta}{n}(C_1 d + \alpha^2)\right) \leq 2$ for $\eta \leq 1/Cd^2$ with $C \geq \frac{6C_1}{\log(2)}$ given by Assumption (B6). Assumption (B8) and Assumption (B3) gives us that $\alpha \leq \sqrt{n/d} \|\mu\| \leq \|\mu\| \leq 0.5\sqrt{\frac{d}{C_1}}$. As a result, there exists a constant $C_r = \frac{192C_1(C_1+1)/(4C_1)}{c_0\gamma^2} \geq \frac{48(C_1 d + \alpha^2)}{c_0\gamma^2(\sqrt{d/C_1}-\alpha)^2}$ such that the Lemma statement holds, where C_1 comes from Lemma B.3. \square

Lemma B.21, an analog of Lemma B.11, aims at providing the convergence guarantees of robust training loss for networks with non-smooth activation functions.

Lemma B.21. For a non-smooth homogeneous activation function ϕ , provided $C > 1$ is sufficiently large, then on a good run, the robust training loss satisfies

$$\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \frac{3200}{(199\|\mu\| - 200\alpha)\gamma c_0 \sqrt{\eta T}}.$$

where $c_0 > 1$ is the constant in Lemma B.15.

Proof of Lemma B.21. The proof is similar as Lemma B.11. We first need to show a lower bound for $\|\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t)\|_F = \sup_{\mathbf{U}: \|\mathbf{U}\|_F=1} \langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{U} \rangle$, and it suffices to construct a matrix \mathbf{V} with Frobenius norm at most one such that $\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle$ is bounded from below by a positive constant. To this end, choose $\mathbf{V} \in \mathbb{R}^{m \times d}$ be the matrix with rows $\mathbf{v}_s = \frac{1}{\sqrt{m}} a_s \mu / \|\mu\|, \forall s \in [m]$. Then $\|\mathbf{V}\|_F = 1$ since $a_s = \pm 1$, and we have for any \mathbf{W}^t ,

$$\begin{aligned}
\langle \nabla f(\tilde{\mathbf{x}}_i; \mathbf{W}^t), \mathbf{V} \rangle &= \sum_{s=1}^m \frac{1}{\sqrt{m}} a_s \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) \langle \mathbf{v}_s, \tilde{\mathbf{x}}_i \rangle = \left\langle \frac{\mu}{\|\mu\|}, \tilde{\mathbf{x}}_i \right\rangle \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle). \\
1 &\geq \frac{1}{m} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) \\
&\geq \frac{1}{m} \sum_{s \in \mathcal{A}^i(\tilde{t}) \cap \mathcal{T}^i} \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) && \text{(Only count the neurons that satisfy } \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle > 0) \\
&\geq \frac{1}{m} \sum_{s \in \mathcal{A}^i(0) \cap \mathcal{T}^i} \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i \rangle) \geq c_0 \gamma. && \text{(Lemma B.20)}
\end{aligned}$$

By Lemma B.3, we have

$$\left\{ \begin{array}{l} y_i \langle \mu, \mathbf{x}_i \rangle \geq \frac{199}{200} \|\mu\|^2, \quad i \in \mathcal{C} \\ |\langle \mu, \mathbf{x}_i \rangle| \leq \frac{201}{200} \|\mu\|^2, \quad i \in \mathcal{N} \end{array} \right\}, \quad \left\{ \begin{array}{l} y_i \langle \mu, \tilde{\mathbf{x}}_i \rangle \geq \frac{199}{200} \|\mu\|^2 - \|\mu\| \alpha, \quad i \in \mathcal{C} \\ |\langle \mu, \tilde{\mathbf{x}}_i \rangle| \leq \frac{201}{200} \|\mu\|^2 + \|\mu\| \alpha, \quad i \in \mathcal{N} \end{array} \right.$$

And $\forall z > 0, \phi'(z) \geq \gamma > 0$, apply Lemma B.19, we have the following lower bound for any \mathbf{W}^t ,

$$y_i \langle \nabla f(\tilde{\mathbf{x}}_i; \mathbf{W}^t), \mathbf{V} \rangle \geq \begin{cases} \frac{199\gamma}{200} \|\mu\| \cdot c_0, & i \in \mathcal{C} \\ -\frac{201}{200} \|\mu\|, & i \in \mathcal{N} \end{cases}, \quad y_i \langle \nabla f(\tilde{\mathbf{x}}_i; \mathbf{W}^t), \mathbf{V} \rangle \geq \begin{cases} \left(\frac{199\gamma}{200} \|\mu\| - \gamma\alpha\right) \cdot c_0, & i \in \mathcal{C} \\ -\left(\frac{201}{200} \|\mu\| + \alpha\right), & i \in \mathcal{N} \end{cases}$$

Similar as Lemma B.11, we have

$$\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \rangle \geq \left(\frac{199\gamma \|\mu\|}{400} - \frac{\gamma\alpha}{2} \right) c_0 \widehat{G}_{\text{rob}}(\mathbf{W}^t).$$

Thus we have

$$\widehat{G}_{\text{rob}}(\mathbf{W}^t) \leq \frac{400}{(199\gamma \|\mu\| - 200\gamma\alpha) c_0} \left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle \leq \frac{400}{(199 \|\mu\| - 200\alpha) \gamma c_0} \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F.$$

We next give an upper bound on $\|\mathbf{W}^t\|_F^2$ as follows:

$$\begin{aligned} & \|\mathbf{W}^{t+1}\|_F^2 \\ &= \left\| \mathbf{W}^t - \eta \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \left\| \nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t) \right\|_F^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 - 2\eta \frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \langle \nabla f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t), \mathbf{W}^t \rangle \quad (\text{Equation (5)}) \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i \sum_{s=1}^m \frac{a_s}{\sqrt{m}} \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle \\ &= \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} \widehat{G}_{\text{rob}}(\mathbf{W}^t)^2 + 2\eta \frac{1}{n} \sum_{i=1}^n g(y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t)) y_i f(\tilde{\mathbf{x}}_i^t; \mathbf{W}^t) \\ &\leq \|\mathbf{W}^t\|_F^2 + \eta^2 \frac{C_2^2 d}{n} + \frac{2}{3} \eta. \end{aligned} \quad (g(z)z \leq \frac{1}{3}, g(z) \leq 1)$$

Telescoping gives us that

$$\|\mathbf{W}^t\|_F^2 \leq \|\mathbf{W}^0\|_F^2 + \left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3} \eta \right) t.$$

We apply the same argument as [B.11](#). Recall from [Lemma B.20 \(E2\)](#) that $\forall t \geq 0, y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^{t+1}) \geq y_k f(\tilde{\mathbf{x}}_k; \mathbf{W}^t), \forall \tilde{\mathbf{x}}_k \in \mathcal{B}_2(\mathbf{x}_k, \alpha), \forall k \in [n]$. Then we have $y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^T) \geq y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^t)$, and therefore $\ell(y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^T)) \leq \ell(y_k f(\tilde{\mathbf{x}}_k^T; \mathbf{W}^t)) \leq \ell(y_k f(\tilde{\mathbf{x}}_k^t; \mathbf{W}^t))$ by definition that $\tilde{\mathbf{x}}_i^t = \arg \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t))$, $t \leq T$. As a result, we have

$$\begin{aligned} \widehat{L}_{\text{rob}}(\mathbf{W}^T) &= \frac{1}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^T)) \\ &\leq \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)) \\ &= \frac{1}{T} \sum_{t=0}^{T-1} \frac{2}{n} \sum_{i=1}^n \max_{\tilde{\mathbf{x}}_i \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} -\ell'(y_i f(\tilde{\mathbf{x}}_i; \mathbf{W}^t)) \\ &\quad (\ell(z) \leq -2\ell'(z) \text{ when } z \geq 0, \text{ Lemma B.20 (E2)}) \\ &= \frac{2}{T} \sum_{t=0}^{T-1} \widehat{G}_{\text{rob}}(\mathbf{W}^t) \\ &\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma c_0 T} \sum_{t=0}^{T-1} \left\langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^t), \mathbf{V} \right\rangle \\ &= \frac{800}{(199 \|\mu\| - 200\alpha) \gamma c_0 \eta T} (\langle \mathbf{W}^T, \mathbf{V} \rangle - \langle \mathbf{W}^0, \mathbf{V} \rangle) \quad (\text{Equation (16)}) \\ &\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma c_0 \eta T} (\|\mathbf{W}^T\|_F + \|\mathbf{W}^0\|_F) \\ &\leq \frac{800}{(199 \|\mu\| - 200\alpha) \gamma c_0 \eta T} \left(2\|\mathbf{W}^0\|_F + \sqrt{\left(\eta^2 \frac{C_2^2 d}{n} + \frac{2}{3} \eta \right) T} \right) \end{aligned}$$

$$\begin{aligned}
&\leq \frac{800}{(199\|\mu\| - 200\alpha)\gamma c_0\eta T} \left(\omega_{\text{init}}\sqrt{6md} + \sqrt{\left(\frac{C_2^2 d}{nC d^2} + \frac{2}{3}\right)\eta T} \right) \\
&\hspace{15em} \text{(Lemma B.2, Assumption (B6))} \\
&\leq \frac{800}{(199\|\mu\| - 200\alpha)\gamma c_0\eta T} \left(\sqrt{6}\eta + \sqrt{\left(\frac{1}{nd} + \frac{2}{3}\right)\eta T} \right). \quad \text{(Choose } C \geq C_2^2\text{)} \\
&\leq \frac{3200}{(\|\mu\| - 2\alpha)\gamma c_0\sqrt{\eta T}}.
\end{aligned}$$

Thus $\forall 0 < \varepsilon < \frac{1}{2n}$, for $T \geq \left(\frac{3200}{(199\|\mu\| - 200\alpha)\gamma c_0\sqrt{\eta\varepsilon}}\right)^2$ guarantees $\widehat{L}_{\text{rob}}(\mathbf{W}^T) \leq \varepsilon$. \square

Now we switch to prove generalization guarantees. Lemma B.23 provides lower bound on both the local difference of $a_s y (\langle \mathbf{w}_s^{t+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^t, \mathbf{x} \rangle)$ and the global difference of $a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle)$, which serves a similar purpose as Lemma B.9 in networks with smooth activation functions.

Lemma B.22. Assume $(\mathbf{x}, y) \sim \mathcal{D}_c$, $\mathbf{x} = y\mu + \xi$. Fix some $t > 0$. On a good run, there exist constants $c_\tau, C' > 0, C'' > 0$ such that the following holds for all $s \in [m]$ and for all $\tau < t$ with probability at least $1 - 3(d/n)^{-11}$ w.r.t. ξ ,

$$\begin{aligned}
&a_s y (\langle \mathbf{w}_s^{\tau+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^\tau, \mathbf{x} \rangle) \\
&\geq \frac{c_\tau \eta}{\sqrt{m}} \left(\|\mu\|^2 - 2\alpha \|\mu\| - C' \max_{i \in [n]} |\langle \xi_i, \xi \rangle| - C' \alpha \sqrt{d \log \left(\frac{d}{n}\right)} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Concurrently,

$$\begin{aligned}
&a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle) \\
&\geq \frac{c_\tau \eta}{\sqrt{m}} \left(\|\mu\|^2 - 2\alpha \|\mu\| - C'' \sqrt{d \log(dm/n)/n} - C'' \alpha \sqrt{d \log \left(\frac{d}{n}\right)} \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Lemma B.23. Assume $(\mathbf{x}, y) \sim \mathcal{D}_c$, $\mathbf{x} = y\mu + \xi$. Fix some $t > 0$. On a good run, there exist constants $c_\tau, C' > 0, C'' > 0$ such that the following holds for all $s \in [m]$ and for all $\tau < t$ with probability at least $1 - 3(d/n)^{-11}$ w.r.t. ξ ,

$$\begin{aligned}
&a_s y (\langle \mathbf{w}_s^{\tau+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^\tau, \mathbf{x} \rangle) \\
&\geq \frac{c_\tau \eta}{\sqrt{m}} \left(\frac{199\|\mu\|^2}{100} - 2\alpha \|\mu\| - C' \max_{i \in [n]} |\langle \xi_i, \xi \rangle| - C' \alpha \sqrt{d \log \left(\frac{d}{n}\right)} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Concurrently,

$$\begin{aligned}
&a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle) \\
&\geq \frac{c_\tau \eta}{\sqrt{m}} \left(\frac{199\|\mu\|^2}{100} - 2\alpha \|\mu\| - C'' \sqrt{d \log(dm/n)/n} \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Proof of Lemma B.23. Consider $\mathbf{x}_i = y_i^c \mu + \xi_i$, $\mathbf{x} = y\mu + \xi$, where y_i^c and y are the clean label. Denote $\tilde{\mathbf{x}}_i^t = \arg \max_{\tilde{\mathbf{x}} \in \mathcal{B}_2(\mathbf{x}_i, \alpha)} \ell(y_i f(\tilde{\mathbf{x}}; \mathbf{W}^t))$, $\epsilon_i = \tilde{\mathbf{x}}_i^t - \mathbf{x}_i$. Then we have

$$\begin{aligned}
&a_s y (\langle \mathbf{w}_s^{t+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^t, \mathbf{x} \rangle) \\
&= \frac{\eta}{n\sqrt{m}} \sum_{i=1}^n y_i y \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle y_i^c \mu + \xi_i + \epsilon_i, y\mu + \xi \rangle \\
&= \frac{\eta}{n\sqrt{m}} \sum_{i=1}^n y_i y \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \left(y_i^c y \|\mu\|^2 + y \langle \mu, \xi_i + \epsilon_i \rangle + y_i^c \langle \mu, \xi \rangle + \langle \epsilon_i + \xi_i, \xi \rangle \right)
\end{aligned}$$

$$\begin{aligned}
&= \underbrace{\frac{\eta}{n\sqrt{m}} \sum_{i=1}^n \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \left(\|\mu\|^2 + y_i^c \langle \mu, \xi_i + \epsilon_i \rangle \right)}_{B_1(t)} \\
&\quad - \underbrace{\frac{2\eta}{n\sqrt{m}} \sum_{i \in \mathcal{N}} \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \left(\|\mu\|^2 + y_i^c \langle \mu, \xi_i + \epsilon_i \rangle \right)}_{B_2(t)} \\
&\quad + \underbrace{\frac{\eta}{n\sqrt{m}} \sum_{i=1}^n y_i y_i^c \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \mu, \xi \rangle}_{B_3(t)} + \underbrace{\frac{\eta}{n\sqrt{m}} \sum_{i=1}^n y_i \tilde{g}_i(\mathbf{W}^t) \phi'(\langle \mathbf{w}_s^t, \tilde{\mathbf{x}}_i^t \rangle) \langle \xi_i + \epsilon_i, \xi \rangle}_{B_4(t)}.
\end{aligned}$$

We now bound each of the term separately. Recall from Lemma B.3 that $|\langle \mu, \xi_i \rangle| \leq \frac{\|\mu\|^2}{200}$, also lemma B.19 gives us that $|\mathcal{A}_s(t)| \geq |\mathcal{A}_s(0) \cap \mathcal{T}_s| \geq c_0 n$, together with Lemma B.20 gives us that $\sum_{i \in \mathcal{A}_s(t)} g_i(\mathbf{W}^t) \geq \frac{n c_0 \hat{G}_{\text{rob}}(\mathbf{W}^t)}{C_r}$. Therefore, for $B_1(t)$, we have

$$\begin{aligned}
B_1(t) &\geq \frac{\eta\gamma}{n\sqrt{m}} \left(\frac{199\|\mu\|^2}{200} - \alpha\|\mu\| \right) \sum_{i \in \mathcal{A}_s(t)} \tilde{g}_i(\mathbf{W}^t) && \text{(Lemma B.20)} \\
&\geq \frac{\eta c_0 \gamma}{2C_r \sqrt{m}} \left(\frac{199\|\mu\|^2}{100} - 2\alpha\|\mu\| \right) \hat{G}_{\text{rob}}(\mathbf{W}^t)
\end{aligned}$$

For $B_2(t)$, we have

$$\begin{aligned}
|B_2(t)| &\leq \frac{\eta C_r}{\sqrt{m}} \left(\frac{1}{C} + \sqrt{\frac{2}{C}} \right) \left(\frac{201\|\mu\|^2}{100} + 2\|\mu\|\alpha \right) \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\leq \frac{\eta c_0 \gamma}{8C_r \sqrt{m}} \left(\frac{199\|\mu\|^2}{100} - 2\alpha\|\mu\| \right) \hat{G}_{\text{rob}}(\mathbf{W}^t),
\end{aligned}$$

where the last inequality holds for large enough C and $\alpha \leq \sqrt{n/d} \|\mu\| \leq 0.99\|\mu\|$.

For $B_3(t)$, define event $\mathcal{E}_1 = \left\{ |\langle \mu, \xi \rangle| \leq c_6 \sqrt{11 \log(d/n)} \|\mu\| \right\}$. Apply Proposition B.17 (D1) gives us that $\mathbb{P}(\mathcal{E}_1) \geq 1 - (d/n)^{-11}$. Therefore conditioning on \mathcal{E}_1 gives us that,

$$\begin{aligned}
|B_3(t)| &\leq \frac{\eta C_r}{\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) |\langle \mu, \xi \rangle| \\
&\leq \frac{c_6 \eta C_r}{\sqrt{m}} \sqrt{11 \log(d/n)} \|\mu\| \hat{G}_{\text{rob}}(\mathbf{W}^t) \\
&\leq \frac{c_6 \eta C_r}{\sqrt{m}} \sqrt{11 \log(\|\mu\|^4 / C^4)} \|\mu\| \hat{G}_{\text{rob}}(\mathbf{W}^t) && \text{(Assumption (B2))} \\
&\leq \frac{\eta c_0 \gamma}{8C_r \sqrt{m}} 0.01 \|\mu\|^2 \hat{G}_{\text{rob}}(\mathbf{W}^t) && \text{(Choose large enough } C) \\
&\leq \frac{\eta c_0 \gamma}{8C_r \sqrt{m}} \left(\frac{199\|\mu\|^2}{100} - 2\alpha\|\mu\| \right) \hat{G}_{\text{rob}}(\mathbf{W}^t).
\end{aligned}$$

For $B_4(t)$, we have

$$|B_4(t)| \leq \frac{\eta}{\sqrt{m}} \max_{i \in [n]} \tilde{g}_i(\mathbf{W}^t) \max_{i \in [n]} |\langle \xi_i + \epsilon_i, \xi \rangle| \leq \frac{\eta C_r}{\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(\max_{i \in [n]} |\langle \xi_i, \xi \rangle| + \alpha \|\xi\| \right).$$

Proposition B.17 (D2) gives us that with probability at least $1 - (d/n)^{-11}$ w.r.t. ξ , $\|\xi\| \leq 10\sqrt{d \log\left(\frac{d}{n}\right)}$. Conditioning on this event gives us that

$$|B_4(t)| \leq \frac{\eta C_r}{\sqrt{m}} \hat{G}_{\text{rob}}(\mathbf{W}^t) \left(\max_{i \in [n]} |\langle \xi_i, \xi \rangle| + 10\alpha \sqrt{d \log\left(\frac{d}{n}\right)} \right).$$

Combining $B_1(t), B_2(t), B_3(t), B_4(t)$ gives us the following:

$$\begin{aligned} a_s y (\langle \mathbf{w}_s^{t+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^t, \mathbf{x} \rangle) \\ \geq \frac{c_7 \eta}{\sqrt{m}} \left(\frac{199 \|\mu\|^2}{100} - 2\alpha \|\mu\| - C' \max_{i \in [n]} |\langle \xi_i, \xi \rangle| - C' \alpha \sqrt{d \log \left(\frac{d}{n} \right)} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t). \end{aligned}$$

For $\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle$, we need to consider the cumulative of the four terms. For $B_4(t)$ in specific, we have

$$\begin{aligned} \left\| \sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) \xi_i \right\|^2 &\leq 4d \sum_{i=1}^n \left(\sum_{\tau=0}^{t-1} \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) \right)^2 \\ &\leq 4C_r^2 d n \left(\sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \right)^2, \end{aligned}$$

where the first inequality comes from applying Lemma B.3 C3, and the second inequality uses the fact that $\phi'(z) \leq 1$ together with $\tilde{g}_i(\mathbf{W}^\tau) \leq C_r \widehat{G}_{\text{rob}}(\mathbf{W}^\tau)$. Thus,

$$\begin{aligned} \left\| \sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) (\xi_i + \epsilon_i^\tau) \right\| &\leq 2C_r \sqrt{dn} \left(\sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \right) + \alpha n \left(\sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \right) \\ &\leq 3C_r \sqrt{dn} \left(\sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \right). \end{aligned}$$

($\alpha \leq \sqrt{n/d} \|\mu\| \leq 0.99 \|\mu\|$, $d \geq Cn \|\mu\|^2$ from Assumption (B3) and C large enough)

Therefore,

$$\left\langle \sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) (\xi_i + \epsilon_i^\tau), \xi \right\rangle = \left\| \sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) (\xi_i + \epsilon_i^\tau) \right\| \left\langle \psi_s^t, \xi \right\rangle,$$

where for simplicity we define $\psi_s^t = \frac{\sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) (\xi_i + \epsilon_i^\tau)}{\left\| \sum_{i=1}^n \sum_{\tau=0}^{t-1} y_i y \tilde{g}_i(\mathbf{W}^\tau) \phi'(\langle \mathbf{w}_s^\tau, \tilde{\mathbf{x}}_i^\tau \rangle) (\xi_i + \epsilon_i^\tau) \right\|} \in \mathbb{R}^d$, which is independent of ξ . This gives us that

$$\left| \sum_{\tau=0}^{t-1} B_4(\tau) \right| \leq \frac{3C_r \eta \sqrt{d}}{\sqrt{mn}} \left(\sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \right) \max_{s \in [m]} |\langle \psi_s^t, \xi \rangle|.$$

And therefore we have

$$\begin{aligned} a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle) \\ = \sum_{\tau=0}^{t-1} B_1(\tau) + B_2(\tau) + B_3(\tau) + B_4(\tau) \\ \geq \frac{c_7 \eta}{\sqrt{m}} \left(\frac{199 \|\mu\|^2}{100} - 2\alpha \|\mu\| - 3C' \sqrt{d/n} \max_{s \in [m]} |\langle \psi_s^t, \xi \rangle| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau). \end{aligned}$$

Define another event $\mathcal{E}_2 = \left\{ \max_{s \in [m]} |\langle \psi_s^t, \xi \rangle| \leq c_6 \sqrt{11 \log(dm/n)} \right\}$. Applying Proposition B.17 (D1) gives us that $\mathbb{P}(\mathcal{E}_2) \geq 1 - (d/n)^{-11}$. Conditioning on the above events, we know

$$\begin{aligned} a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle) \\ \geq \frac{c_7 \eta}{\sqrt{m}} \left(\frac{199 \|\mu\|^2}{100} - 2\alpha \|\mu\| - C'' \sqrt{d \log(dm/n)/n} \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau). \end{aligned}$$

Applying a union bound, the above holds with probability at least $1 - 3(d/n)^{-11}$. \square

Corollary B.24. Assume $(x, y) \sim \mathcal{D}_c$. Fix some $t > 0$. On a good run, the following holds for all $s \in [m]$ with probability at least $1 - 4(d/n)^{-11}$,

$$a_s y \langle \mathbf{w}_s^t, \mathbf{x} \rangle \geq \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau),$$

where c_7 and C'' come from Lemma B.23.

Proof of Corollary B.24. With proper C , Assumption (B2) gives us that $\|\mu\|^2 \geq 200C'' \sqrt{d \log(md/n)/n}$. Therefore, Lemma B.23 tells us that with probability at least $1 - 3(d/n)^{-11}$,

$$a_s y (\langle \mathbf{w}_s^t, \mathbf{x} \rangle - \langle \mathbf{w}_s^0, \mathbf{x} \rangle) \geq \frac{c_7 \eta}{\sqrt{m}} \left(1.985 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).$$

Using Proposition B.17 as well as Lemma B.2, the following holds with probability at least $1 - (d/n)^{-11}$,

$$\begin{aligned} |\langle \mathbf{w}_s^0, \mathbf{x} \rangle| &= |\langle \mathbf{w}_s^0, y\mu + \xi \rangle| \\ &\leq \|\mathbf{w}_s^0\| \|\mu\| + c_6 \sqrt{11 \log(dm/n)} \|\mathbf{w}_s^0\| && \text{(Proposition B.17 (D1))} \\ &\leq \|\mathbf{w}_s^0\| \|\mu\| + c_6 \sqrt{\frac{11n \|\mu\|^4}{C^4 d}} \|\mathbf{w}_s^0\| && \text{(Assumption (B2))} \\ &\leq \|\mathbf{w}_s^0\| \|\mu\| + c_6 \sqrt{\frac{11n \|\mu\|^4}{C^5 n \|\mu\|^2}} \|\mathbf{w}_s^0\| && \text{(Assumption (B3))} \\ &\leq 2 \|\mu\| \|\mathbf{w}_s^0\| && \text{(Choose sufficiently large } C\text{)} \\ &\leq 4\omega_{\text{init}} \|\mu\| \sqrt{d} && \text{(Lemma B.13)} \\ &\leq \frac{4\eta \|\mu\|}{\sqrt{m}} && \text{(Assumption (B5))} \\ &= \frac{8\eta \|\mu\|}{\sqrt{m}} \widehat{G}_{\text{rob}}(\mathbf{W}^0) && (\tilde{g}_i(\mathbf{W}^0) = 0.5) \\ &\leq \frac{0.001c_7\eta}{\sqrt{m}} \|\mu\|^2 \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau), && (18) \end{aligned}$$

where the last line holds from $\|\mu\|^2 \geq C \log(2)$ by Assumption (B2) and C being large enough.

Therefore we have

$$a_s y \langle \mathbf{w}_s^t, \mathbf{x} \rangle \geq \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).$$

□

Lemma B.25. For any $(x, y) \sim \mathcal{D}_c$ with $\mathbf{x} = y\mu + \xi$, on a good run, there exists some constant $c_8 > 0$ such that

$$|\langle \mathbf{w}_s^t, \mathbf{x} \rangle| \leq \frac{c_8 \eta}{\sqrt{m}} \left(\|\mu\|^2 + \|\mu\| \|\xi\| + \sqrt{d/n} \|\xi\| + \|\mu\| \alpha + \|\xi\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau), \forall s \in [m].$$

Proof of Lemma B.25. Consider $\forall s \in [m]$,

$$\langle \mathbf{w}_s^t, \mathbf{x} \rangle = \langle \mathbf{w}_s^0, \mathbf{x} \rangle + \sum_{\tau=0}^{t-1} (\langle \mathbf{w}_s^{\tau+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^\tau, \mathbf{x} \rangle).$$

Decompose $\langle \mathbf{w}_s^{t+1}, \mathbf{x} \rangle - \langle \mathbf{w}_s^t, \mathbf{x} \rangle$ into $B_1(t), B_2(t), B_3(t), B_4(t)$ the same way as in Lemma B.23.

$$\begin{aligned}
|B_1(t)| &\leq \frac{\eta C_r}{\sqrt{m}} \left(\frac{201}{200} \|\mu\|^2 + \|\mu\| \alpha \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t), \\
|B_2(t)| &\leq \frac{\eta C_r}{\sqrt{m}} \left(\beta + \sqrt{\frac{2}{C}} \right) \left(\frac{201}{100} \|\mu\|^2 + 2 \|\mu\| \alpha \right) \widehat{G}_{\text{rob}}(\mathbf{W}^t), \\
|B_3(t)| &\leq \frac{\eta C_r}{\sqrt{m}} \widehat{G}_{\text{rob}}(\mathbf{W}^t) |\langle \mu, \xi \rangle| \leq \frac{\eta C_r}{\sqrt{m}} \|\mu\| \|\xi\| \widehat{G}_{\text{rob}}(\mathbf{W}^t), \\
\left| \sum_{\tau=0}^{t-1} B_4(\tau) \right| &\leq \frac{2C_r \eta \sqrt{d}}{\sqrt{mn}} \|\xi\| \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) + \frac{\eta C_r}{\sqrt{m}} \alpha \|\xi\| \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau), \\
|\langle \mathbf{w}_s^0, \mathbf{x} \rangle| &\leq \|\mathbf{w}_s^0\| \cdot \|\mathbf{x}\| \leq \|\mathbf{w}_s^0\| (\|\mu\| + \|\xi\|) \leq \frac{2\eta}{\sqrt{m}} (\|\mu\| + \|\xi\|) \\
&\quad \text{(Assumption (B5), Lemma B.13)} \\
&\leq \frac{2\eta}{\sqrt{C \log(2)} \sqrt{m}} (\|\mu\|^2 + \|\mu\| \|\xi\|). \quad \text{(Assumption (B2))}
\end{aligned}$$

Therefore,

$$\begin{aligned}
|\langle \mathbf{w}_s^t, \mathbf{x} \rangle| &\leq |\langle \mathbf{w}_s^0, \mathbf{x} \rangle| + \sum_{\tau=0}^{t-1} |B_1(\tau)| + \sum_{\tau=0}^{t-1} |B_2(\tau)| + \sum_{\tau=0}^{t-1} |B_3(\tau)| + \left| \sum_{\tau=0}^{t-1} B_4(\tau) \right| \\
&\leq \frac{c_8 \eta}{\sqrt{m}} \left(\|\mu\|^2 + \|\mu\| \|\xi\| + \sqrt{d/n} \|\xi\| + \|\mu\| \alpha + \|\xi\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

□

We finally demonstrate the lower bound on the normalized expected conditional margin, similar as Lemma B.10.

Lemma B.26. On a good run, there exists some constant $c_9 > 0$ such that

$$\begin{aligned}
\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [y f(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} &\geq \frac{c_9 \sqrt{n}}{16C_2 \sqrt{d}} \left(\|\mu\|^2 - \|\mu\| \alpha \right); \\
\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [y f(x; \mathbf{W}^t) | y = -1]}{\|\mathbf{W}^t\|_2} &\geq \frac{c_9 \sqrt{n}}{16C_2 \sqrt{d}} \left(\|\mu\|^2 - \|\mu\| \alpha \right).
\end{aligned}$$

Proof of Lemma B.26. Consider the following

$$y f(x; \mathbf{W}^t) = \frac{1}{\sqrt{m}} \sum_{a_s=y} \phi(\langle \mathbf{w}_s^t, \mathbf{x} \rangle) - \frac{1}{\sqrt{m}} \sum_{a_s=-y} \phi(\langle \mathbf{w}_s^t, \mathbf{x} \rangle).$$

We first consider $y = 1$. Denote event \mathcal{E} as the conclusion of Corollary B.24 holds, then $P(\mathcal{E}) \geq 1 - 4(d/n)^{-11}$. Corollary B.24 indicates that

$$\begin{aligned}
a_s y \langle \mathbf{w}_s^t, \mathbf{x} \rangle &\geq \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&\geq \frac{c_7 \eta}{\sqrt{m}} 0.004 \|\mu\|^2 \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) > 0. \quad (\alpha \leq \sqrt{n/d} \|\mu\| \leq 0.99 \|\mu\|)
\end{aligned}$$

Therefore, on event \mathcal{E} we have

$$\begin{aligned}
yf(\mathbf{x}; \mathbf{W}^t) &= \frac{1}{\sqrt{m}} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^t, \mathbf{x} \rangle) a_s y \langle \mathbf{w}_s^t, \mathbf{x} \rangle \\
&\geq \frac{1}{\sqrt{m}} \sum_{s=1}^m \phi'(\langle \mathbf{w}_s^t, \mathbf{x} \rangle) \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&\geq \frac{1}{\sqrt{m}} \sum_{a_s=y} \phi'(\langle \mathbf{w}_s^t, \mathbf{x} \rangle) \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&\geq \frac{1}{\sqrt{m}} \frac{m}{2} \gamma \frac{c_7 \eta}{\sqrt{m}} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&= \frac{\gamma c_7 \eta}{2} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau),
\end{aligned}$$

and thus

$$\begin{aligned}
&\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} [yf(\mathbf{x}; \mathbf{W}^t) \mathbf{1}(\mathcal{E}) \mid y = 1] \\
&\geq \mathbf{P}(\mathcal{E}) \frac{\gamma c_7 \eta}{2} \left(1.984 \|\mu\|^2 - 2\alpha \|\mu\| \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&\geq \mathbf{P}(\mathcal{E}) \frac{\gamma c_7 \eta}{2} 0.004 \|\mu\|^2 \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
&\geq \frac{c_9}{4} \eta \|\mu\|^2 \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau). \quad (\mathbf{P}(\mathcal{E}) \geq 1 - 4(d/n)^{-11} \geq \frac{3}{4}; \text{ choose sufficiently large } C)
\end{aligned}$$

We now consider on event \mathcal{E}^c . Using Lemma B.25 gives us that,

$$\begin{aligned}
&|\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} [yf(\mathbf{x}; \mathbf{W}^t) \mathbf{1}(\mathcal{E}^c) \mid y = 1]| \\
&\leq \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_c} \left[\frac{1}{\sqrt{m}} \sum_{s=1}^m |\langle \mathbf{w}_s^t, \mathbf{x} \rangle| \mathbf{1}(\mathcal{E}^c) \right] \\
&\leq c_8 \eta \left(\|\mu\|^2 + \|\mu\| \alpha \right) \mathbf{P}(\mathcal{E}^c) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) + c_8 \eta \left(\|\mu\| + \sqrt{\frac{d}{n}} + \alpha \right) \mathbb{E}(\|\xi\| \mathbf{1}(\mathcal{E}^c)) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Now we denote another event $\tilde{\mathcal{E}} = \left\{ \|\xi\| \leq 600 \sqrt{d \log \left(\frac{d}{n} \right)} \right\}$, then

$$\begin{aligned}
\mathbb{E} \left(\|\xi\| \mathbf{1}(\mathcal{E}^c) \mathbf{1}(\tilde{\mathcal{E}}) \right) &\leq \mathbb{E} \left(600 \sqrt{d \log \left(\frac{d}{n} \right)} \mathbf{1}(\mathcal{E}^c) \right) \leq 600 \sqrt{d \log \left(\frac{d}{n} \right)} 4 \left(\frac{d}{n} \right)^{-11} \leq 0.5 \left(\frac{d}{n} \right)^{-8}, \\
&\quad (d \geq \log(2) C n^2 \text{ from Assumption (B3) with sufficiently large } C) \\
\mathbb{E} \left(\|\xi\| \mathbf{1}(\mathcal{E}^c) \mathbf{1}(\tilde{\mathcal{E}}^c) \right) &\leq \mathbb{E} \left(\|\xi\| \mathbf{1}(\tilde{\mathcal{E}}^c) \right) \leq 8 \left(\frac{d}{n} \right)^{-40000} \sqrt{d} \cdot (100 \sqrt{\log \left(\frac{d}{n} \right)}) \leq 0.5 \left(\frac{d}{n} \right)^{-8}. \\
&\quad (\text{Proposition B.17 (D2); } d \geq \log(2) C n^2 \text{ with sufficiently large } C)
\end{aligned}$$

Putting them back gives us that

$$\begin{aligned}
& \left| \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) \mathbb{1}(\mathcal{E}^c) | y = 1] \right| \\
& \leq c_8 \eta \left(\|\mu\|^2 + \|\mu\| \alpha \right) \mathbb{P}(\mathcal{E}^c) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) + c_8 \eta \left(\|\mu\| + \sqrt{d/n} + \alpha \right) \mathbb{E}(\|\xi\| \mathbb{1}(\mathcal{E}^c)) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& \leq c_8 \eta \left(\|\mu\|^2 + \|\mu\| \alpha \right) 4(d/n)^{-11} \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) + c_8 \eta \left(\|\mu\| + \sqrt{d/n} + \alpha \right) (d/n)^{-8} \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& \leq c_8 \eta \left(\|\mu\|^2 + \|\mu\| \alpha + \|\mu\| + \sqrt{d/n} + \alpha \right) (d/n)^{-8} \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& \hspace{15em} \text{(Assumption (B3); choose sufficiently large } C) \\
& \leq \frac{c_9 \eta}{8} \left(\|\mu\|^2 + \|\mu\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau). \\
& \hspace{5em} (\|\mu\|^2 \geq C \log(2) \text{ from Assumption (B2); } d \geq \log(2) C n^2; \text{ choose sufficiently large } C)
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
& \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1] \\
& = \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) \mathbb{1}(\mathcal{E}) | y = 1] + \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) \mathbb{1}(\mathcal{E}^c) | y = 1] \\
& \geq \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) \mathbb{1}(\mathcal{E}) | y = 1] - \left| \mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) \mathbb{1}(\mathcal{E}^c) | y = 1] \right| \\
& \geq \frac{c_9 \eta}{8} \left(\|\mu\|^2 - \|\mu\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau).
\end{aligned}$$

Then similar as Lemma B.10, recall that $\|\mathbf{W}^0\|_F \leq 2\omega_{\text{init}}\sqrt{md} \leq 2\eta \leq \eta\sqrt{d/n}\widehat{G}_{\text{rob}}(\mathbf{W}^0)$. If $\|\mathbf{W}^t\|_F \leq 2\|\mathbf{W}^0\|_F$,

$$\begin{aligned}
\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} & \geq \frac{c_9 \eta}{16 \|\mathbf{W}^0\|_F} \left(\|\mu\|^2 - \|\mu\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& \geq \frac{c_9 \eta}{16 \eta \sqrt{d/n} \widehat{G}_{\text{rob}}(\mathbf{W}^0)} \left(\|\mu\|^2 - \|\mu\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& \geq \frac{c_9 \sqrt{n}}{16 \sqrt{d}} \left(\|\mu\|^2 - \|\mu\| \alpha \right). \quad (\sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s) \geq \widehat{G}_{\text{rob}}(\mathbf{W}^0))
\end{aligned}$$

If $\|\mathbf{W}^t\|_F > 2\|\mathbf{W}^0\|_F$, by Lemma B.8, we have

$$2\|\mathbf{W}^0\|_F \leq \|\mathbf{W}^t\|_F \leq \|\mathbf{W}^0\|_F + C_2 \eta \sqrt{d/n} \sum_{s=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^s).$$

Thus,

$$\begin{aligned}
\frac{\mathbb{E}_{(x,y) \sim \mathcal{D}_c} [yf(x; \mathbf{W}^t) | y = 1]}{\|\mathbf{W}^t\|_2} & \geq \frac{c_9}{16 C_2 \sqrt{d/n} \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau)} \left(\|\mu\|^2 - \|\mu\| \alpha \right) \sum_{\tau=0}^{t-1} \widehat{G}_{\text{rob}}(\mathbf{W}^\tau) \\
& = \frac{c_9 \sqrt{n}}{16 C_2 \sqrt{d}} \left(\|\mu\|^2 - \|\mu\| \alpha \right).
\end{aligned}$$

All the above proof holds for expected condition on $y = -1$.

□

B.3 MISSING PROOFS IN SECTION 3.3

Theorem 3.4. We consider independent label flip with probability β . Let $p(x)$ be the density function of $\mathcal{D}_{\text{clust}}$. For any given classifier $f(\cdot; \mathbf{W})$, when $\alpha < \|\mu\|$, we have $L_{\text{rob}}^{0/1}(\mathbf{W}) \geq \beta + \frac{1-2\beta}{4} \int_{\mathbb{R}^d} \min\{p(\xi), p(\xi + \mathbf{v})\} d\xi$, where $\mathbf{v} = 2(1 - \alpha/\|\mu\|)\mu$. When $\alpha \geq \|\mu\|$, the robust test error satisfies $L_{\text{rob}}^{0/1}(\mathbf{W}) \geq 0.5$.

Proof of Theorem 3.4.

$$\begin{aligned} L_{\text{rob}}^{0/1}(\mathbf{W}) &= \mathbf{P}_{(x,y) \sim \mathcal{D}} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } yf(\tilde{x}; \mathbf{W}) \leq 0] \\ &= (1 - \beta) \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y_c f(\tilde{x}; \mathbf{W}) \leq 0] \\ &\quad + \beta \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} [\exists \tilde{x} \in \mathcal{B}_2(x, \alpha) \text{ s.t. } y_c f(\tilde{x}; \mathbf{W}) \geq 0] \\ &= (1 - \beta) \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right] \\ &\quad + \beta \left(1 - \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[\max_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) < 0 \right] \right) \\ &\geq \beta + (1 - 2\beta) \mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left[\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right]. \end{aligned}$$

Recall that $\mathbf{x} = y_c \mu + \xi$.

Case 1: Consider the case that $\alpha \geq \|\mu\|$. We have

$$\begin{aligned} &\mathbf{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbf{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \\ &\geq \mathbf{1} (f(\xi; \mathbf{W}) \leq 0) + \mathbf{1} (f(\xi; \mathbf{W}) \geq 0) \\ &\geq 1, \end{aligned}$$

where the first inequality holds because $\xi \in \mathcal{B}_2(\mu + \xi, \alpha)$ and $\xi \in \mathcal{B}_2(-\mu + \xi, \alpha)$. Therefore we have

$$\begin{aligned} &\mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right) \\ &= \frac{1}{2} \mathbf{P}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \frac{1}{2} \mathbf{P}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \\ &= 0.5 \mathbb{E}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\mathbf{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbf{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) \\ &\geq 0.5. \end{aligned}$$

As a result,

$$L_{\text{rob}}^{0/1}(\mathbf{W}) \geq \beta + 0.5(1 - 2\beta) = 0.5.$$

Case 2: Consider the case that $\alpha < \|\mu\|$. Let $p(x)$ denote the density function of $\mathcal{D}_{\text{clust}}$. Define $\mathbf{v} = (2 - 2\frac{\alpha}{\|\mu\|})\mu$. We have

$$\begin{aligned} &\left(\mathbf{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbf{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) \\ &\quad + \left(\mathbf{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\xi + \mathbf{v} + \mu, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbf{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(\xi + \mathbf{v} - \mu, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) \\ &\geq \left(\mathbf{1} \left(f\left(\xi + \frac{\mathbf{v}}{2}; \mathbf{W}\right) \leq 0 \right) + 0 \right) + \left(0 + \mathbf{1} \left(f\left(\xi + \frac{\mathbf{v}}{2}; \mathbf{W}\right) \geq 0 \right) \right) \\ &\geq 1, \end{aligned}$$

where the first inequality holds because $\xi + \frac{\mathbf{v}}{2} \in \mathcal{B}_2(\mu + \xi, \alpha)$ and $\xi + \frac{\mathbf{v}}{2} \in \mathcal{B}_2(\xi + \mathbf{v} - \mu, \alpha)$. Therefore we have

$$\mathbf{P}_{(x,y_c) \sim \mathcal{D}_c} \left(\min_{\tilde{x} \in \mathcal{B}_2(x, \alpha)} y_c f(\tilde{x}; \mathbf{W}) \leq 0 \right)$$

$$\begin{aligned}
&= \left(\frac{1}{2} \mathbb{P}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \frac{1}{2} \mathbb{P}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) \\
&= \frac{1}{2} \mathbb{E}_{\xi \sim \mathcal{D}_{\text{clust}}} \left(\mathbb{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbb{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) \\
&= \frac{1}{4} \int_{\mathbb{R}^d} \left\{ \left(\mathbb{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbb{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(-\mu + \xi, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) p(\xi) \right. \\
&\quad \left. + \left(\mathbb{1} \left(\min_{\tilde{x} \in \mathcal{B}_2(\xi + \nu + \mu, \alpha)} f(\tilde{x}; \mathbf{W}) \leq 0 \right) + \mathbb{1} \left(\max_{\tilde{x} \in \mathcal{B}_2(\xi + \nu - \mu, \alpha)} f(\tilde{x}; \mathbf{W}) \geq 0 \right) \right) p(\xi + \nu) \right\} d\xi \\
&\geq \frac{1}{4} \int_{\mathbb{R}^d} \min\{p(\xi), p(\xi + \nu)\} d\xi.
\end{aligned}$$

As a result,

$$L_{\text{rob}}^{0/1}(\mathbf{W}) \geq \beta + \frac{1 - 2\beta}{4} \int_{\mathbb{R}^d} \min\{p(\xi), p(\xi + \nu)\} d\xi.$$

Consider a special instance where $\mathcal{D}_{\text{clust}}$ is a standard Gaussian distribution; i.e., $\mathcal{N}(0, \mathbf{I}_d)$. Then the result can be simplified as

$$L_{\text{rob}}^{0/1}(\mathbf{W}) \geq \beta + \frac{1 - 2\beta}{2} \Phi(-(\|\mu\| - \alpha)),$$

where $\Phi(x) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-t^2/2) dt$ is the normal cumulative distribution function. \square

Proposition 3.5. Consider the same setting as in Theorem 3.1. Then, for some absolute constant $C > 1$, with probability at least $1 - 2\delta$ over the random initialization and the draw of an i.i.d. sample, we have that $\frac{\|\mathbf{W}^1 - \mathbf{W}^0\|_F}{\|\mathbf{W}^0\|_F} \geq \frac{\gamma(199\|\mu\| - 200\alpha)}{1000}$.

Proof of Proposition 3.5. Consider $\mathbf{V} \in \mathbb{R}^{m \times d}$ be the matrix with rows $v_s = \frac{a_s \mu}{\|\mu\| \sqrt{m}}$, then we have

$$\begin{aligned}
\frac{\|\mathbf{W}^1 - \mathbf{W}^0\|_F}{\|\mathbf{W}^0\|_F} &\geq \frac{\langle \mathbf{W}^1 - \mathbf{W}^0, \mathbf{V} \rangle}{\|\mathbf{W}^0\|_F} \\
&= \frac{\eta \langle -\nabla \widehat{L}_{\text{rob}}(\mathbf{W}^0), \mathbf{V} \rangle}{\|\mathbf{W}^0\|_F} \\
&\geq \left(\frac{199\gamma \|\mu\| - 200\gamma\alpha}{400} \right) \frac{\eta \widehat{G}_{\text{rob}}(\mathbf{W}^0)}{\|\mathbf{W}^0\|_F} && \text{(Equation (13))} \\
&\geq \left(\frac{199\gamma \|\mu\| - 200\gamma\alpha}{400} \right) \frac{\eta \widehat{G}_{\text{rob}}(\mathbf{W}^0)}{\sqrt{3/2} m d \omega_{\text{init}}} && \text{(Lemma B.2)} \\
&\geq \left(\frac{199\gamma \|\mu\| - 200\gamma\alpha}{500} \right) \widehat{G}_{\text{rob}}(\mathbf{W}^0) && \text{(Assumption (A6))} \\
&= \frac{\gamma}{1000} (199 \|\mu\| - 200\alpha). && \text{(Equation (9))}
\end{aligned}$$

\square