Benign Overfitting in Two-layer Convolutional Neural Networks

Yuan Cao*
Department of Statistics & Actuarial Science
Department of Mathematics
The University of Hong Kong
yuancao@hku.hk

Zixiang Chen*
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095, USA
chenzx19@cs.ucla.edu

Mikhail Belkin
Halicioglu Data Science Institute
University of California San Diego
La Jolla, CA 92093, USA
mbelkin@ucsd.edu

Quanquan Gu
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095, USA
qgu@cs.ucla.edu

Abstract

Modern neural networks often have great expressive power and can be trained to overfit the training data, while still achieving a good test performance. This phenomenon is referred to as "benign overfitting". Recently, there emerges a line of works studying "benign overfitting" from the theoretical perspective. However, they are limited to linear models or kernel/random feature models, and there is still a lack of theoretical understanding about when and how benign overfitting occurs in neural networks. In this paper, we study the benign overfitting phenomenon in training a two-layer convolutional neural network (CNN). We show that when the signal-to-noise ratio satisfies a certain condition, a two-layer CNN trained by gradient descent can achieve arbitrarily small training and test loss. On the other hand, when this condition does not hold, overfitting becomes harmful and the obtained CNN can only achieve constant level test loss. These together demonstrate a sharp phase transition between benign overfitting and harmful overfitting, driven by the signal-to-noise ratio. To the best of our knowledge, this is the first work that precisely characterizes the conditions under which benign overfitting can occur in training convolutional neural networks.

1 Introduction

Modern deep learning models often consist of a huge number of model parameters, which is more than the number of training data points and therefore over-parameterized. These over-parameterized models can be trained to overfit the training data (achieving a close to 100% training accuracy), while still making accurate prediction on the unseen test data. This phenomenon has been observed in a number of prior works (Zhang et al., 2017; Neyshabur et al., 2019), and is often referred to as *benign overfitting* (Bartlett et al., 2020). It revolutionizes the classical understanding about the bias-variance trade-off in statistical learning theory, and has drawn great attention from the community (Belkin et al., 2018, 2019a,b; Hastie et al., 2019).

There exist a number of works towards understanding the benign overfitting phenomenon. While they offered important insights into the benign overfitting phenomenon, most of them are limited

*Equal contribution.

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

to the settings of linear models (Belkin et al., 2019b; Bartlett et al., 2020; Hastie et al., 2019; Wu and Xu, 2020; Chatterji and Long, 2020; Zou et al., 2021b; Cao et al., 2021) and kernel/random features models (Belkin et al., 2018; Liang and Rakhlin, 2020; Montanari and Zhong, 2020), and cannot be applied to neural network models that are of greater interest. The only notable exceptions are (Adlam and Pennington, 2020; Li et al., 2021), which attempted to understand benign overfitting in neural network models. However, they are still limited to the "neural tagent kernel regime" (Jacot et al., 2018) where the neural network learning problem is essentially equivalent to kernel regression. Thus, it remains a largely open problem to show how and when benign overfitting can occur in neural networks.

Clearly, understanding benign overfitting in neural networks is much more challenging than that in linear models, kernel methods or random feature models. The foremost challenge stems from nonconvexity: previous works on linear models and kernel methods/random features are all in the convex setting, while neural network training is a highly nonconvex optimization problem. Therefore, while most of the previous works can study the minimum norm interpolators/maximum margin classifiers according to the *implicit bias* (Soudry et al., 2018) results for the corresponding models, existing implicit bias results for neural networks (e.g., Lyu and Li (2019)) are not sufficient and a new analysis of the neural network learning process is in demand.

In this work, we provide one such algorithmic analysis for learning two-layer convolutional neural networks (CNNs) with the second layer parameters being fixed as $+1$'s and -1's and polynomial ReLU activation function: $\sigma(z) = \max\{0, z\}^q$, where $q > 2$ is a hyperparameter. We consider a setting where the input data consist of *label dependent signals* and *label independent noises*, and utilize a *signal-noise decomposition* of the CNN filters to precisely characterize the signal learning and noise memorization processes during neural network training. Our result not only demonstrates that benign overfitting can occur in learning two-layer neural networks, but also gives precise conditions under which the overfitted CNN trained by gradient descent can achieve small population loss. Our paper makes the following major contributions:

- We establish population loss bounds of overfitted CNN models trained by gradient descent, and theoretically demonstrate that benign overfitting can occur in learning over-parameterized neural networks. We show that under certain conditions on the signal-to-noise ratio, CNN models trained by gradient descent will prioritize learning the signal over memorizing the noise, and thus achieving both small training and test losses. To the best of our knowledge, this is the first result on the benign overfitting of neural networks that is beyond the neural tangent kernel regime.
- We also establish a negative result showing that when the conditions on the signal-to-noise ratio do not hold, then the overfitted CNN model will achieve at least a constant population loss. This result, together with our upper bound result, reveals an interesting phase transition between benign overfitting and harmful overfitting.
- Our analysis is based on a new proof technique namely *signal-noise decomposition*, which decomposes the convolutional filters into a linear combination of initial filters, the signal vectors and the noise vectors. We convert the neural network learning into a discrete dynamical system of the coefficients from the decomposition, and perform a two-stage analysis that decouples the complicated relation among the coefficients. This enables us to analyze the non-convex optimization problem, and bound the population loss of the CNN trained by gradient descent. We believe our proof technique is of independent interest and can potentially be applied to deep neural networks.

We note that a concurrent work (Frei et al., 2022) studies learning log-Concave mixture data with label flip noise using fully-connected two-layer neural networks with smoothed leaky ReLU activation. Notably, their risk bound matches the risk bound for linear models given in Cao et al. (2021) when the label flip noise is zero. However, their analysis only focuses on upper bounding the risk, and cannot demonstrate the phase transition between benign and harmful overfitting. Compared with (Frei et al., 2022), we focus on CNNs, and consider a different data model to better capture the nature of image classification problems. Moreover, we present both positive and negative results under different SNR regimes, and demonstrate a sharp phase transition between benign and harmful overfitting.

Notation. Given two sequences $\{x_n\}$ and $\{y_n\}$, we denote $x_n = O(y_n)$ if there exist some absolute constant $C_1 > 0$ and $N > 0$ such that $|x_n| \leq C_1|y_n|$ for all $n \geq N$. Similarly, we denote $x_n = \Omega(y_n)$ if there exist $C_2 > 0$ and $N > 0$ such that $|x_n| \geq C_2|y_n|$ for all $n > N$. We say

$x_n = \Theta(y_n)$ if $x_n = O(y_n)$ and $x_n = \Omega(y_n)$ both holds. We use $\widetilde{O}(\cdot)$, $\widetilde{\Omega}(\cdot)$, and $\widetilde{\Theta}(\cdot)$ to hide logarithmic factors in these notations respectively. Moreover, we denote $x_n = \mathrm{poly}(y_n)$ if $x_n = O(y_n^D)$ for some positive constant D, and $x_n = \mathrm{polylog}(y_n)$ if $x_n = \mathrm{poly}(\log(y_n))$. Finally, for two scalars a and b, we denote $a \vee b = \max\{a, b\}$.

2 Related Work

A line of recent works have attempted to understand why overfitted predictors can still achieve a good test performance. Belkin et al. (2019a) first empirically demonstrated that in many machine learning models such as random Fourier features, decision trees and ensemble methods , the population risk curve has a *double descent* shape with respect to the number of model parameters. Belkin et al. (2019b) further studied two specific data models, namely the Gaussian model and Fourier series model, and theoretically demonstrated the double descent risk curve in linear regression. Bartlett et al. (2020) studied over-parameterized linear regression to fit data produced by a linear model with additive noises, and established matching upper and lower bounds of the risk achieved by the minimum norm interpolator on the training dataset. It is shown that under certain conditions on the spectrum of the data covariance matrix, the population risk of the interpolator can be asymptotically optimal. Hastie et al. (2019); Wu and Xu (2020) studied linear regression in the setting where both the dimension and sample size grow together with a fixed ratio, and showed double descent of the risk with respect to this ratio. Chatterji and Long (2020) studied the population risk bounds of over-parameterized linear logistic regression on sub-Gaussian mixture models with label flipping noises, and showed how gradient descent can train over-parameterized linear models to achieve nearly optimal population risk. Cao et al. (2021) tightened the upper bound given by Chatterji and Long (2020) in the case without the label flipping noises, and established a matching lower bound of the risk achieved by over-parameterized maximum margin interpolators. Shamir (2022) proposed a generic data model for benign overfitting of linear predictors, and studied different problem settings under which benign overfitting can or cannot occur.

Besides the studies on linear models, several recent works also studied the benign overfitting and double descent phenomena in kernel methods or random feature models. Zhang et al. (2017) first pointed out that overfitting kernel predictors can sometimes still achieve good population risk. Liang and Rakhlin (2020) studied how interpolating kernel regression with radial basis function (RBF) kernels (and variants) can generalize and how the spectrum of the data covariance matrix affects the population risk of the interpolating kernel predictor. Li et al. (2021) studied the benign overfitting phenomenon of random feature models defined as two-layer neural networks whose first layer parameters are fixed at random initialization. Mei and Montanari (2019); Liao et al. (2020) demonstrated the double descent phenomenon for the population risk of interpolating random feature predictors with respect to the ratio between the dimensions of the random feature and the data input. Adlam and Pennington (2020) shows that neural tangent kernel (Jacot et al., 2018) based kernel regression has a triple descent risk curve with respect to the total number of trainable parameters. Montanari and Zhong (2020) further pointed out an interesting phase transition of the generalization error achieved by neural networks trained in the neural tangent kernel regime.

3 Problem Setup

In this section, we introduce the data generation model and the convolutional neural network we consider in this paper. We focus on binary classification, and present our data distribution $\mathcal{D}$ in the following definition.

Definition 3.1. *Let $\boldsymbol{\mu} \in \mathbb{R}^d$ be a fixed vector representing the signal contained in each data point. Then each data point $(\mathbf{x}, y)$ with $\mathbf{x} = [\mathbf{x}^{(1)\top}, \mathbf{x}^{(2)\top}]^\top \in \mathbb{R}^{2d}$ and $y \in \{-1, 1\}$ is generated from the following distribution $\mathcal{D}$:*

1. *The label y is generated as a Rademacher random variable.*
2. *A noise vector $\boldsymbol{\xi}$ is generated from the Gaussian distribution $N(\mathbf{0}, \sigma_p^2 \cdot (\mathbf{I} - \boldsymbol{\mu}\boldsymbol{\mu}^\top \cdot \|\boldsymbol{\mu}\|_2^{-2}))$.*
3. *One of $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}$ is randomly selected and then assigned as $y \cdot \boldsymbol{\mu}$, which represents the signal; the other is then given by $\boldsymbol{\xi}$, which represents noises.*

Our data generation model is inspired by image data, where the inputs consist of different patches, and only some of the patches are related to the class label of the image. In detail, the patch assigned as $y \cdot \boldsymbol{\mu}$ is the signal patch that is correlated to the label of the data, and the patch assigned as $\boldsymbol{\xi}$ is the noise patch that is independent of the label of the data and therefore is irrelevant for prediction. We assume that the noise patch is generated from the Gaussian distribution $N(\mathbf{0}, \sigma_p^2 \cdot (\mathbf{I} - \boldsymbol{\mu}\boldsymbol{\mu}^\top \cdot \|\boldsymbol{\mu}\|_2^{-2}))$ to ensure that the noise vector is orthogonal to the signal vector $\boldsymbol{\mu}$ for simplicity. Note that when the dimension d is large, $\|\boldsymbol{\xi}\|_2 \approx \sigma_p\sqrt{d}$ by standard concentration bounds. Therefore, we can treat $\|\boldsymbol{\mu}\|_2/(\sigma_p\sqrt{d}) \approx \|\boldsymbol{\mu}\|_2/\|\boldsymbol{\xi}\|_2$ as the signal-to-noise ratio (SNR). For the ease of discussion, we denote $\mathrm{SNR} = \|\boldsymbol{\mu}\|_2/(\sigma_p\sqrt{d})$. Note that the Bayes risk for learning our model is zero. We can also add label flip noise similar to Chatterji and Long (2020); Frei et al. (2022) to make the Bayes risk equal to the label flip noise and therefore nonzero, but this will not change the key message of our paper.

Intuitively, if a classifier learns the signal $\boldsymbol{\mu}$ and utilizes the signal patch of the data to make prediction, it can perfectly fit a given training data set $\{(\mathbf{x}_i, y_i) : i \in [n]\}$ and at the same time have a good performance on the test data. However, when the dimension d is large ($d > n$), a classifier that is a function of the noises $\boldsymbol{\xi}_i$, $i \in [n]$ can also perfectly fit the training data set, while the prediction will be totally random on the new test data. Therefore, the data generation model given in Definition 3.1 is a useful model to study the population loss of overfitted classifiers. Similar models have been studied in some recent works by Li et al. (2019); Allen-Zhu and Li (2020a,b); Zou et al. (2021a).

Two-layer CNNs. We consider a two-layer convolutional neural network whose filters are applied to the two patches $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ separately, and the second layer parameters of the network are fixed as $+1/m$ and $-1/m$ respectively. Then the network can be written as $f(\mathbf{W}, \mathbf{x}) = F_{+1}(\mathbf{W}_{+1}, \mathbf{x}) - F_{-1}(\mathbf{W}_{-1}, \mathbf{x})$, where $F_{+1}(\mathbf{W}_{+1}, \mathbf{x})$, $F_{-1}(\mathbf{W}_{-1}, \mathbf{x})$ are defined as:

$$F_j(\mathbf{W}_j, \mathbf{x}) = \frac{1}{m}\sum_{r=1}^{m}\big[\sigma(\langle \mathbf{w}_{j,r}, \mathbf{x}^{(1)}\rangle) + \sigma(\langle \mathbf{w}_{j,r}, \mathbf{x}^{(2)}\rangle)\big] = \frac{1}{m}\sum_{r=1}^{m}\big[\sigma(\langle \mathbf{w}_{j,r}, y \cdot \boldsymbol{\mu}\rangle) + \sigma(\langle \mathbf{w}_{j,r}\boldsymbol{\xi}\rangle)\big]$$

for $j \in \{+1, -1\}$. Here, m is the number of convolutional filters in F_{+1} and F_{-1}, $\sigma(z) = (\max\{0, z\})^q$ is the ReLU^q activation function where $q > 2$, $\mathbf{w}_{j,r} \in \mathbb{R}^d$ denotes the weight for the r-th filter (i.e., neuron), and $\mathbf{W}_j$ is the collection of model weights associated with F_j. We also use $\mathbf{W}$ to denote the collection of all model weights. We note that our CNN model can also be viewed as a CNN with average global pooling (Lin et al., 2013). We train the above CNN model by minimizing the empirical cross-entropy loss function

$$L_S(\boldsymbol{W}) = \frac{1}{n}\sum_{i=1}^{n}\ell[y_i \cdot f(\mathbf{W}, \mathbf{x}_i)],$$

where $\ell(z) = \log(1 + \exp(-z))$, and $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is the training data set. We further define the true loss (test loss) $L_{\mathcal{D}}(\mathbf{W}) := \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\ell[y \cdot f(\mathbf{W}, \mathbf{x})]$.

We consider gradient descent starting from Gaussian initialization, where each entry of $\mathbf{W}_{+1}$ and $\mathbf{W}_{-1}$ is sampled from a Gaussian distribution $N(0, \sigma_0^2)$, and σ_0^2 is the variance. The gradient descent update of the filters in the CNN can be written as

$$\begin{aligned}\mathbf{w}_{j,r}^{(t+1)} &= \mathbf{w}_{j,r}^{(t)} - \eta \cdot \nabla_{\mathbf{w}_{j,r}} L_S(\boldsymbol{W}^{(t)}) \\ &= \mathbf{w}_{j,r}^{(t)} - \frac{\eta}{nm}\sum_{i=1}^{n}\ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle) \cdot j y_i \boldsymbol{\xi}_i - \frac{\eta}{nm}\sum_{i=1}^{n}\ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle) \cdot j\boldsymbol{\mu} \quad (3.1)\end{aligned}$$

for $j \in \{\pm 1\}$ and $r \in [m]$, where we introduce a shorthand notation $\ell_i'^{(t)} = \ell'[y_i \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i)]$.

4 Main Results

In this section, we present our main theoretical results. At the core of our analyses and results is a *signal-noise decomposition* of the filters in the CNN trained by gradient descent. By the gradient descent update rule (3.1), it is clear that the gradient descent iterate $\mathbf{w}_{j,r}^{(t)}$ is a linear combination of its random initialization $\mathbf{w}_{j,r}^{(0)}$, the signal vector $\boldsymbol{\mu}$ and the noise vectors in the training data $\boldsymbol{\xi}_i$, $i \in [n]$. Motivated by this observation, we introduce the following definition.

Definition 4.1. *Let $\mathbf{w}_{j,r}^{(t)}$ for $j \in \{\pm 1\}$, $r \in [m]$ be the convolution filters of the CNN at the t-th iteration of gradient descent. Then there exist unique coefficients $\gamma_{j,r}^{(t)} \geq 0$ and $\rho_{j,r,i}^{(t)}$ such that*

$$\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + j \cdot \gamma_{j,r}^{(t)} \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu} + \sum_{i=1}^{n} \rho_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i.$$

We further denote $\overline{\rho}_{j,r,i}^{(t)} := \rho_{j,r,i}^{(t)} \mathbb{1}(\rho_{j,r,i}^{(t)} \geq 0)$, $\underline{\rho}_{j,r,i}^{(t)} := \rho_{j,r,i}^{(t)} \mathbb{1}(\rho_{j,r,i}^{(t)} \leq 0)$. Then we have that

$$\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + j \cdot \gamma_{j,r}^{(t)} \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu} + \sum_{i=1}^{n} \overline{\rho}_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i + \sum_{i=1}^{n} \underline{\rho}_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i. \tag{4.1}$$

We refer to (4.1) as the *signal-noise decomposition* of $\mathbf{w}_{j,r}^{(t)}$. We add normalization factors $\|\boldsymbol{\mu}\|_2^{-2}, \|\boldsymbol{\xi}_i\|_2^{-2}$ in the definition so that $\gamma_{j,r}^{(t)} \approx \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu} \rangle, \rho_{j,r,i}^{(t)} \approx \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$. In this decomposition, $\gamma_{j,r}^{(t)}$ characterizes the progress of learning the signal vector $\boldsymbol{\mu}$, and $\rho_{j,r,i}^{(t)}$ characterizes the degree of noise memorization by the filter. Evidently, based on this decomposition, for some iteration t, (i) If some of $\gamma_{j,r}^{(t)}$'s are large enough while $|\rho_{j,r,i}^{(t)}|$ are relatively small, then the CNN will have small training and test losses; (ii) If some $\overline{\rho}_{j,r,i}^{(t)}$'s are large and all $\gamma_{j,r}^{(t)}$'s are small, then the CNN will achieve a small training loss, but a large test loss. Thus, Definition 4.1 provides a handle for us to study the convergence of the training loss as well as the the population loss of the CNN trained by gradient descent.

Our results are based on the following conditions on the dimension d, sample size n, neural network width m, learning rate η, initialization scale σ_0.

Condition 4.2. *Suppose that*

1. *Dimension d is sufficiently large: $d = \widetilde{\Omega}(m^{2\vee[4/(q-2)]} n^{4\vee[(2q-2)/(q-2)]})$.*
2. *Training sample size n and neural network width m satisfy $n, m = \Omega(\mathrm{polylog}(d))$.*
3. *The learning rate η satisfies $\eta \leq \widetilde{O}(\min\{\|\boldsymbol{\mu}\|_2^{-2}, \sigma_p^{-2} d^{-1}\})$.*
4. *The standard deviation of Gaussian initialization σ_0 is appropriately chosen such that $\widetilde{O}(nd^{-1/2}) \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\} \leq \sigma_0 \leq \widetilde{O}(m^{-2/(q-2)} n^{-[1/(q-2)]\vee 1}) \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\}$.*

A few remarks on Condition 4.2 are in order. The condition on d is to ensure that the learning is in a sufficiently over-parameterized setting, and similar conditions have been made in the study of learning over-parameterized linear models (Chatterji and Long, 2020; Cao et al., 2021). For example, if we choose $q = 3$, then the condition on d becomes $d = \widetilde{\Omega}(m^4 n^4)$. Furthermore, we require the sample size and neural network width to be at least polylogarithmic in the dimension d to ensure some statistical properties of the training data and weight initialization to hold with probability at least $1 - d^{-1}$, which is a mild condition. Finally, the conditions on σ_0 and η are to ensure that gradient descent can effectively minimize the training loss, and they depend on the scale of the training data points. When $\sigma_p = O(d^{-1/2})$ and $\|\boldsymbol{\mu}\|_2 = O(1)$, the step size η can be chosen as large as $\widetilde{O}(1)$ and the initialization σ_0 can be as large as $\widetilde{O}(m^{-2/(q-2)} n^{-[1/(q-2)]\vee 1})$. In our paper, we only require $m, n = \Omega(\mathrm{polylog}(d))$, so our initialization and step-size can be chosen as an almost constant order. Based on these conditions, we give our main result on signal learning in the following theorem.

Theorem 4.3. *For any $\epsilon > 0$, let $T = \widetilde{\Theta}(\eta^{-1} m \sigma_0^{-(q-2)} \|\boldsymbol{\mu}\|_2^{-q} + \eta^{-1}\epsilon^{-1} m^3 \|\boldsymbol{\mu}\|_2^{-2})$. Under Condition 4.2, if $n \cdot \mathrm{SNR}^q = \widetilde{\Omega}(1)$[*], then with probability at least $1 - d^{-1}$, there exists $0 \leq t \leq T$ such that:*

1. *The CNN learns the signal: $\max_r \gamma_{j,r}^{(t)} = \Omega(1)$ for $j \in \{\pm 1\}$.*
2. *The CNN does not memorize the noises in the training data: $\max_{j,r,i} |\rho_{j,r,i}^{(T)}| = \widetilde{O}(\sigma_0 \sigma_p \sqrt{d})$.*

[*]Here the $\widetilde{\Omega}(\cdot)$ hides an $\mathrm{polylog}(\epsilon^{-1})$ factor. This applies to Theorem 4.4 as well.

3. *The training loss converges to ϵ, i.e., $L_S(\mathbf{W}^{(t)}) \leq \epsilon$.*
4. *The trained CNN achieves a small test loss: $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \leq 6\epsilon + \exp(-n^2)$.*

Theorem 4.3 characterizes the case of signal learning. It shows that, if $n \cdot \mathrm{SNR}^q = \widetilde{\Omega}(1)$, then at least one CNN filter can learn the signal by achieving $\gamma_{j,r_j^*}^{(t)} \geq \Omega(1)$, and as a result, the learned neural network can achieve small training and test losses. To demonstrate the sharpness of this condition, we also present the following theorem for the noise memorization by the CNN.

Theorem 4.4. *For any $\epsilon > 0$, let $T = \widetilde{\Theta}(\eta^{-1} m \cdot n(\sigma_p\sqrt{d})^{-q} \cdot \sigma_0^{-(q-2)} + \eta^{-1}\epsilon^{-1} n m^3 d^{-1} \sigma_p^{-2})$. Under Condition 4.2, if $n^{-1} \cdot \mathrm{SNR}^{-q} = \widetilde{\Omega}(1)$, then with probability at least $1 - d^{-1}$, there exists $0 \leq t \leq T$ such that:*

1. *The CNN memorizes noises in the training data: $\max_r \overline{\rho}_{y_i,r,i}^{(t)} = \Omega(1)$.*
2. *The CNN does not sufficiently learn the signal: $\max_{j,r} \gamma_{j,r}^{(t)} \leq \widetilde{O}(\sigma_0 \|\boldsymbol{\mu}\|_2)$.*
3. *The training loss converges to ϵ, i.e., $L_S(\mathbf{W}^{(t)}) \leq \epsilon$.*
4. *The trained CNN has a constant order test loss: $L_{\mathcal{D}}(\mathbf{W}^{(t)}) = \Theta(1)$.*

Theorem 4.4 holds under the condition that $n^{-1} \cdot \mathrm{SNR}^{-q} = \widetilde{\Omega}(1)$. Clearly, this is the opposite regime (up to some logarithmic factors) compared with Theorem 4.3. In this case, the CNN trained by gradient descent mainly memorizes noises in the training data and does not learn enough signal. This, together with the results in Theorem 4.3, reveals a clear phase transition between signal learning and noise memorization in CNN training:

- If $n \cdot \mathrm{SNR}^q = \widetilde{\Omega}(1)$, then the CNN learns the signal and achieves a $O(\epsilon + \exp(-n^2))$ test loss. This is the regime of benign overfitting.
- If $n^{-1} \cdot \mathrm{SNR}^{-q} = \widetilde{\Omega}(1)$ then the CNN can only memorize noises and will have a $\Theta(1)$ test loss. This is the regime of harmful overfitting.

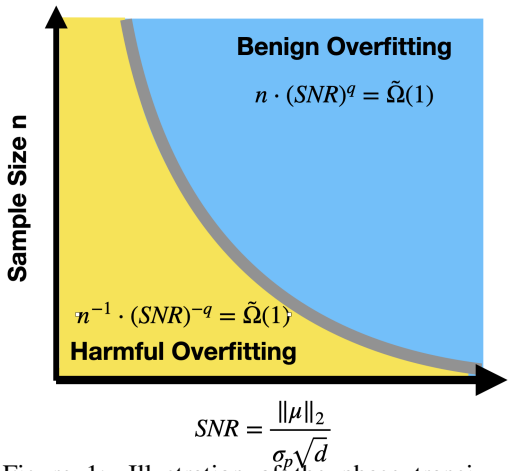

Figure 1: Illustration of the phase transition between benign and harmful overfitting. The blue region represents the setting under which the overfitted CNN trained by gradient descent is guaranteed to have small population loss, and the yellow region represents the setting under which the population loss is guaranteed to be of constant order. The slim gray band region is the setting where the population loss is not well characterized.

The phase transition is illustrated in Figure 1. Clearly, $n \cdot \mathrm{SNR}^q = \widetilde{\Omega}(1)$ is the precise condition under which benign overfitting occurs. Remarkably, in this case the population loss decreases *exponentially* with the sample size n. Under our condition that $n = \Omega(\mathrm{polylog}(d))$, this term can also be upper bounded by $1/\mathrm{poly}(d)$, which is small in the high-dimensional setting. Note that when $\|\boldsymbol{\mu}\|_2 = \Theta(1)$ and $\sigma_p = \Theta(d^{-1/2})$, applying standard uniform convergence based bounds (Bartlett et al., 2017; Neyshabur et al., 2018) or stability based bounds (Hardt et al., 2016; Mou et al., 2017; Chen et al., 2018) typically give $\widetilde{O}(n^{-1/2})$ bounds on the generalization gap, which are vacuous when $n = O(\mathrm{polylog}(d))$. Our bound under the same setting is $O(1/\mathrm{poly}(d))$, which is non-vacuous. This is attributed to our precise analysis of signal learning and noise memorization in Theorems 4.3 and 4.4.

Comparison with neural tangent kernel (NTK) results. We want to emphasize that our analysis is beyond the so-called neural tangent kernel regime. In the NTK regime, it has been shown that gradient descent can train an over-parameterized neural network to achieve good training and test accuracies (Jacot et al., 2018; Du et al., 2019b,a; Allen-Zhu et al., 2019b; Zou et al., 2019; Arora et al., 2019a; Cao and Gu, 2019a; Chen et al., 2019). However, it is widely believed in literature that the NTK analyses cannot fully explain the success of deep learning, as the neural networks in the NTK regime are almost "linearized" (Lee et al., 2019; Cao and Gu, 2019a). Our analysis and results are not in the NTK regime: In the NTK regime, the network parameters stay close to their initialization throughout training, i.e., $\|\mathbf{W}^{(t)} - \mathbf{W}^{(0)}\|_F = O(1)$, so that the NN model

can be approximated by its linearization (Allen-Zhu et al., 2019b; Cao and Gu, 2019a; Chen et al., 2019). In comparison, our analysis does not rely on linearizing the neural network function, and $\|\mathbf{W}^{(t)} - \mathbf{W}^{(0)}\|_F$ can be as large as $O(\text{poly}(m))$.

5 Overview of Proof Technique

In this section, we discuss the main challenges in the study of CNN training under our setting, and explain some key techniques we implement in our proofs to overcome these challenges. The complete proofs of all the results are given in the appendix.

Main challenges. Studying benign overfitting under our setting is a challenging task. The first challenge is the nonconvexity of the training objective function $L_S(\mathbf{W})$. Nonconvexity has introduced new challenges in the study of benign overfitting particularly because our goal is not only to show the convergence of the training loss, but also to study the population loss in the over-parameterized setting, which requires a precise algorithmic analysis of the learning problem.

5.1 Iterative Analysis of the Signal-Noise Decomposition

In order to study the learning process based on the nonconvex optimization problem, we propose a key technique which enables the iterative analysis of the coefficients in the signal-noise decomposition in Definition 4.1. This technique is given in the following lemma.

Lemma 5.1. *The coefficients $\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ in Definition 4.1 satisfy the following equations:*

$$\gamma_{j,r}^{(0)}, \overline{\rho}_{j,r,i}^{(0)}, \underline{\rho}_{j,r,i}^{(0)} = 0, \tag{5.1}$$

$$\gamma_{j,r}^{(t+1)} = \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \cdot \boldsymbol{\mu} \rangle) \cdot \|\boldsymbol{\mu}\|_2^2, \tag{5.2}$$

$$\overline{\rho}_{j,r,i}^{(t+1)} = \overline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = j), \tag{5.3}$$

$$\underline{\rho}_{j,r,i}^{(t+1)} = \underline{\rho}_{j,r,i}^{(t)} + \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = -j). \tag{5.4}$$

Remark 5.2. *With the decomposition* (4.1)*, the signal learning and noise memorization processes of a CNN can be formally studied by analyzing the dynamics of $\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ based on the dynamical system* (5.2)-(5.4)*. Note that prior to our work, several existing results have utilized the inner products $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu} \rangle$ during the neural network training process in order to establish generalization bounds (Brutzkus et al., 2018; Chatterji and Long, 2020; Frei et al., 2021). Similar inner product based arguments are also implemented in Allen-Zhu and Li (2020a,b); Zou et al. (2021a), which study different topics related to learning neural networks. Compared with the inner product based argument, our method has two major advantages: (i) Based on the definition* (5.2)-(5.4) *and the fact that $\ell_i'^{(t)} < 0$, it is clear that $\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}$ are monotonically increasing, while $\underline{\rho}_{j,r,i}^{(t)}$ is monotonically decreasing throughout the whole training process. In comparison, monotonicity does not hold in the inner product based argument, especially for $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$. (ii) Our signal-noise decomposition also enables a clean homogeneity-based proof for the convergence of the training loss to an arbitrarily small error rate $\epsilon > 0$, which will be presented in Subsection 5.2.*

With Lemma 5.1, we can reduce the study of the CNN learning process to the analysis of the discrete dynamical system given by (5.1)-(5.4). Our proof then focuses on a careful assessment of the values of the coefficients $\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ throughout training. To prepare for more detailed analyses, we first present the following bounds of the coefficients, which hold throughout training.

Proposition 5.3. *Under Condition 4.2, for any $T^* = \eta^{-1} \text{poly}(\epsilon^{-1}, \|\boldsymbol{\mu}\|_2^{-1}, d^{-1}\sigma_p^{-2}, \sigma_0^{-1}, n, m, d)$, the following bounds hold for $t \in [0, T^*]$:*

1. $0 \leq \gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)} \leq 4\log(T^*)$ *for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$.*
2. $0 \geq \underline{\rho}_{j,r,i}^{(t)} \geq -2\max_{i,j,r}\{|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle|\} - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}} \cdot 4\log(T^*)$ *for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$.*

We can then prove the following lemma, which demonstrates that the training objective function $L_S(\mathbf{W})$ can dominate the gradient norm $\|\nabla L_S(\mathbf{W}^{(t)})\|_F$ along the gradient descent path.

Lemma 5.4. *Under Condition 4.2, for any $T^* = \eta^{-1}poly(\epsilon^{-1}, \|\boldsymbol{\mu}\|_2^{-1}, d^{-1}\sigma_p^{-2}, \sigma_0^{-1}, n, m, d)$, the following result holds for $t \in [0, T^*]$:*

$$\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 = O\big(\max\{\|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d\}\big) \cdot L_S(\mathbf{W}^{(t)}).$$

Lemma 5.4 plays a key role in the convergence proof of training loss function. However, note that our study of benign overfitting requires carefully monitoring the changes of the coefficients in the signal-noise decomposition, which cannot be directly done by Lemma 5.4. This is quite a challenging task, due to the complicated interactions among $\gamma_{j,r}^{(t)}$, $\overline{\rho}_{j,r,i}^{(t)}$ and $\underline{\rho}_{j,r,i}^{(t)}$. Note that even $\gamma_{j,r}^{(t)}$, which has the simplest formula (5.2), depends on *all* the quantities $\gamma_{j',r'}^{(t)}$, $\overline{\rho}_{j',r',i}^{(t)}$ and $\underline{\rho}_{j',r',i}^{(t)}$ for $j' \in \{\pm 1\}$, $r' \in [m]$ and $i \in [n]$. This is because the cross-entropy loss derivative term $\ell_i'^{(t)} = \ell'[y_i \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i)]$ depends on all the neurons of the network. To overcome this challenge, we introduce in the next subsection a decoupling technique based on a two-stage analysis.

5.2 Decoupling with a Two-Stage Analysis.

We utilize a two-stage analysis to decouple the complicated relation among the coefficients $\gamma_{j,r}^{(t)}$, $\overline{\rho}_{j,r,i}^{(t)}$ and $\underline{\rho}_{j,r,i}^{(t)}$. Intuitively, the initial neural network weights are small enough so that the neural network at initialization has constant level cross-entropy loss derivatives on all the training data: $\ell_i'^{(0)} = \ell'[y_i \cdot f(\mathbf{W}^{(0)}, \mathbf{x}_i)] = \Theta(1)$ for all $i \in [n]$. This is guaranteed under Condition 4.2 and matches neural network training in practice. Motivated by this, we can consider the first stage of the training process where $\ell_i'^{(t)} = \Theta(1)$, in which case we can show significant scale differences among $\gamma_{j,r}^{(t)}$, $\overline{\rho}_{j,r,i}^{(t)}$ and $\underline{\rho}_{j,r,i}^{(t)}$. Based on the result in the first stage, we then proceed to the second stage of the training process where the loss derivatives are no longer at a constant level and show that the training loss can be optimized to be arbitrarily small and meanwhile, the scale differences shown in the first learning stage remain the same throughout the training process. In the following, we focus on explaining the key proof steps for Theorem 4.3. The proof idea for Theorem 4.4 is similar, so we defer the details to the appendix.

Stage 1. It can be shown that, until some of the coefficients $\gamma_{j,r}^{(t)}$, $\rho_{j,r,i}^{(t)}$ reach $\Theta(1)$, we have $\ell_i'^{(t)} = \ell'[y_i \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i)] = \Theta(1)$ for all $i \in [n]$. Therefore, we first focus on this first stage of the training process, where the dynamics of the coefficients in (5.2) - (5.4) can be greatly simplified by replacing the $\ell_i'^{(t)}$ factors by their constant upper and lower bounds. The following lemma summarizes our main conclusion at stage 1 for signal learning:

Lemma 5.5. *Under the same conditions as Theorem 4.3, there exists $T_1 = \widetilde{O}(\eta^{-1} m \sigma_0^{2-q} \|\boldsymbol{\mu}\|_2^{-q})$ such that*

1. $\max_r \gamma_{j,r}^{(T_1)} = \Omega(1)$ *for* $j \in \{\pm 1\}$.
2. $|\rho_{j,r,i}^{(t)}| = O(\sigma_0 \sigma_p \sqrt{d})$ *for all* $j \in \{\pm 1\}$, $r \in [m]$, $i \in [n]$ *and* $0 \le t \le T_1$.

Lemmas 5.5 takes advantage of the training period when the loss function derivatives remain a constant order to show that the CNN can capture the signal. At the end of stage 1 in signal learning, $\max_r \gamma_{j,r}$ reaches $\Theta(1)$, and is significantly larger than $\rho_{j,r,i}^{(t)}$. After this, it is no longer guaranteed that the loss derivatives $\ell_i'^{(t)}$ will remain constant order, and thus starts the training stage 2.

Stage 2. In this stage, we take into full consideration the exact definition $\ell_i'^{(t)} = \ell'[y_i \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i)]$ and show that the training loss function will converge to $L_S(\mathbf{W}^{(t)}) < \epsilon$. Thanks to the analysis in stage 1, we know that some $\gamma_{j,r}^{(t)}$ is significantly larger than all $\rho_{j,r,i}^{(t)}$'s at the end of stage 1. This scale difference is the key to our analysis in stage 2. Based on this scale difference and the monotonicity of $\gamma_{j,r}^{(t)}$, $\overline{\rho}_{j,r,i}^{(t)}$, $\underline{\rho}_{j,r,i}^{(t)}$ in the signal-noise decomposition, it can be shown that there exists $\mathbf{W}^*$ such that $y_i \cdot \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle \ge q \log(2q/\epsilon)$ throughout stage 2. Moreover, since the neural network

$f(\mathbf{W}, \mathbf{x})$ is q-homogeneous in $\mathbf{W}$, we have $\langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}), \mathbf{W}^{(t)} \rangle = q \cdot f(\mathbf{W}^{(t)}, \mathbf{x})$. Therefore,

$$
\begin{aligned}
\langle \nabla L_S(\mathbf{W}^{(t)}), \mathbf{W}^{(t)} - \mathbf{W}^* \rangle &= \frac{1}{n} \sum_{i=1}^{n} \ell_i'^{(t)} \cdot y_i \cdot \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^{(t)} - \mathbf{W}^* \rangle \\
&= \frac{1}{n} \sum_{i=1}^{n} \ell_i'^{(t)} \cdot [y_i \cdot q \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i) - y_i \cdot \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle] \\
&\geq \frac{1}{n} \sum_{i=1}^{n} \ell'[y_i \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i)] \cdot [y_i \cdot q \cdot f(\mathbf{W}^{(t)}, \mathbf{x}_i) - q \log(2q/\epsilon)] \\
&\geq q \cdot \frac{1}{n} \sum_{i=1}^{n} [\ell(f(\mathbf{W}^{(t)}, \mathbf{x}_i)) - \ell(\log(2q/\epsilon))] \\
&\geq q \cdot L_S(\mathbf{W}^{(t)}) - \epsilon/2,
\end{aligned}
$$

where the second inequality follows by the convexity of the cross-entropy loss function. With the above key technique, we can prove the following lemma.

Lemma 5.6. *Let T, T_1 be defined in Theorem 4.3 and Lemma 5.5 respectively. Then under the same conditions as Theorem 4.3, for any $t \in [T_1, T]$, it holds that $|\rho_{j,r,i}^{(t)}| \leq \sigma_0 \sigma_p \sqrt{d}$ for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$. Moreover, let $\mathbf{W}^*$ be the collection of CNN parameters with convolution filters $\mathbf{w}_{j,r}^* = \mathbf{w}_{j,r}^{(0)} + 2qm \log(2q/\epsilon) \cdot j \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu}$. Then the following bound holds*

$$
\frac{1}{t - T_1 + 1} \sum_{s=T_1}^{t} L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(t - T_1 + 1)} + \frac{\epsilon}{(2q-1)}
$$

for all $t \in [T_1, T]$, where we denote $\|\mathbf{W}\|_F = \sqrt{\|\mathbf{W}_{+1}\|_F^2 + \|\mathbf{W}_{-1}\|_F^2}$.

Lemma 5.6 states two main results on signal learning. First of all, during this training period, it is guaranteed that the coefficients of noise vectors $\rho_{j,r,i}^{(t)}$ in the signal-noise decomposition remain sufficiently small. Moreover, it also gives an optimization type result that the best iterate in $[T_1, T]$ is small as long as T is large enough. Clearly, the convergence of the training loss stated in Theorems 4.3 directly follows by choosing T to be sufficiently large in Lemmas 5.6. The lemma below further gives an upper bound on the test loss.

Lemma 5.7. *Let T be defined in Theorem 4.3. Under the same conditions as Theorem 4.3, for any $t \leq T$ with $L_S(\mathbf{W}^{(t)}) \leq 1$, it holds that $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \leq 6 \cdot L_S(\mathbf{W}^{(t)}) + \exp(-n^2)$.*

Below we finalize the proof of Theorem 4.3. The proofs of other results are in the appendix.

Proof of Theorem 4.3. The first part of Theorem 4.3 follows by Lemma 5.5 and the monotonicity of $\gamma_{j,r}^{(t)}$. The second part of Theorem 4.3 follows by Lemma 5.6. For the third part, let $\mathbf{W}^*$ be defined in Lemma 5.6. Then by the definition of $\mathbf{W}^*$, we have

$$
\begin{aligned}
\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F &\leq \|\mathbf{W}^{(T_1)} - \mathbf{W}^{(0)}\|_F + \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F \\
&\leq \sum_{j,r} \gamma_{j,r}^{(T_1)} \|\boldsymbol{\mu}\|_2^{-1} + \sum_{j,r,i} \frac{\overline{\rho}_{j,r,i}^{(T_1)}}{\|\boldsymbol{\xi}_i\|_2} + \sum_{j,r,i} \frac{\underline{\rho}_{j,r,i}^{(T_1)}}{\|\boldsymbol{\xi}_i\|_2} + \Theta(m^{3/2} \log(1/\epsilon)) \|\boldsymbol{\mu}\|_2^{-1} \\
&= \widetilde{O}(m^{3/2} \|\boldsymbol{\mu}\|_2^{-1}),
\end{aligned}
$$

where the first inequality is by triangle inequality, the second inequality is by the signal-noise decomposition of $\mathbf{W}^{(T_1)}$ and the definition of $\mathbf{W}^*$, and the last equality is by Proposition 5.3 and Lemma 5.5. Therefore, choosing $T = \widetilde{\Theta}(\eta^{-1} T_1 + \eta^{-1} \epsilon^{-1} m^3 \|\boldsymbol{\mu}\|_2^{-2}) = \widetilde{\Theta}(\eta^{-1} \sigma_0^{-(q-2)} \|\boldsymbol{\mu}\|_2^{-q} + \eta^{-1} \epsilon^{-1} m^3 \|\boldsymbol{\mu}\|_2^{-2})$ in Lemma 5.6 ensures that

$$
\frac{1}{T - T_1 + 1} \sum_{t=T_1}^{T} L_S(\mathbf{W}^{(t)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(T - T_1 + 1)} + \frac{\epsilon}{2q-1} \leq \frac{\widetilde{O}(m^3 \|\boldsymbol{\mu}\|_2^{-2})}{(2q-1)\eta(T - T_1 + 1)} + \frac{\epsilon}{2q-1} \leq \epsilon,
$$

and there exists $t \in [T_1, T]$ such that $L_S(\mathbf{W}^{(t)}) \leq \epsilon$. This completes the proof of the third part of Theorem 4.3. Finally, combining this bound with Lemma 5.7 gives

$$L_{\mathcal{D}}(\mathbf{W}^{(t)}) \leq 6 \cdot L_S(\mathbf{W}^{(t)}) + \exp(-n^2) \leq 6\epsilon + \exp(-n^2),$$

which proves the last part of Theorem 4.3. □

6 Conclusion and Future Work

This paper utilizes a signal-noise decomposition to study the signal learning and noise memorization process in the training of a two-layer CNN. We precisely give the conditions under which the CNN will mainly focus on learning signals or memorizing noises, and reveals a phase transition of the population loss with respect to the sample size, signal strength, noise level, and dimension. Our result theoretically demonstrates that benign overfitting can happen in neural network training. An important future work direction is to study the benign overfitting phenomenon of neural networks in learning other data models. Moreover, it is also important to generalize our analysis to deep convolutional neural networks.

Acknowledgments and Disclosure of Funding

We would like to thank Spencer Frei for valuable comment and discussion on the earlier version of this paper, and pointing out a related work. ZC and QG are supported in part by the National Science Foundation CAREER Award 1906169, IIS-2008981 and the Sloan Research Fellowship. MB is grateful for the support from the National Science Foundation (NSF) and the Simons Foundation for the Collaboration on the Theoretical Foundations of Deep Learning[†] through awards DMS-2031883 and #814639 as well as NSF IIS-1815697 and the TILOS institute (NSF CCF-2112665).

References

ADLAM, B. and PENNINGTON, J. (2020). The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization. In *International Conference on Machine Learning*. PMLR.

ALLEN-ZHU, Z. and LI, Y. (2020a). Feature purification: How adversarial training performs robust deep learning. *arXiv preprint arXiv:2005.10190* .

ALLEN-ZHU, Z. and LI, Y. (2020b). Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816* .

ALLEN-ZHU, Z., LI, Y. and LIANG, Y. (2019a). Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in Neural Information Processing Systems*.

ALLEN-ZHU, Z., LI, Y. and SONG, Z. (2019b). A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*.

ARORA, S., DU, S., HU, W., LI, Z. and WANG, R. (2019a). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*.

ARORA, S., DU, S. S., HU, W., LI, Z., SALAKHUTDINOV, R. and WANG, R. (2019b). On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*.

ARORA, S., GE, R., NEYSHABUR, B. and ZHANG, Y. (2018). Stronger generalization bounds for deep nets via a compression approach. In *International Conference on Machine Learning*.

BARTLETT, P. L., FOSTER, D. J. and TELGARSKY, M. J. (2017). Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*.

[†] `https://deepfoundations.ai/`

BARTLETT, P. L., LONG, P. M., LUGOSI, G. and TSIGLER, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences* .

BELKIN, M., HSU, D., MA, S. and MANDAL, S. (2019a). Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences* **116** 15849–15854.

BELKIN, M., HSU, D. and XU, J. (2019b). Two models of double descent for weak features. *arXiv preprint arXiv:1903.07571* .

BELKIN, M., MA, S. and MANDAL, S. (2018). To understand deep learning we need to understand kernel learning. In *International Conference on Machine Learning*.

BOUSQUET, O. and ELISSEEFF, A. (2002). Stability and generalization. *Journal of machine learning research* **2** 499–526.

BRUTZKUS, A., GLOBERSON, A., MALACH, E. and SHALEV-SHWARTZ, S. (2018). Sgd learns over-parameterized networks that provably generalize on linearly separable data. In *International Conference on Learning Representations*.

CAO, Y. and GU, Q. (2019a). Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*.

CAO, Y. and GU, Q. (2019b). Tight sample complexity of learning one-hidden-layer convolutional neural networks. In *Advances in Neural Information Processing Systems*.

CAO, Y., GU, Q. and BELKIN, M. (2021). Risk bounds for over-parameterized maximum margin classification on sub-gaussian mixtures. *Advances in Neural Information Processing Systems* **34**.

CHATTERJI, N. S. and LONG, P. M. (2020). Finite-sample analysis of interpolating linear classifiers in the overparameterized regime. *arXiv preprint arXiv:2004.12019* .

CHEN, Y., JIN, C. and YU, B. (2018). Stability and convergence trade-off of iterative optimization algorithms. *arXiv preprint arXiv:1804.01619* .

CHEN, Z., CAO, Y., ZOU, D. and GU, Q. (2019). How much over-parameterization is sufficient to learn deep relu networks? *arXiv preprint arXiv:1911.12360* .

DENG, L. (2012). The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* **29** 141–142.

DU, S., LEE, J., LI, H., WANG, L. and ZHAI, X. (2019a). Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*.

DU, S. S., LEE, J. D. and TIAN, Y. (2018a). When is a convolutional filter easy to learn? In *International Conference on Learning Representations*.

DU, S. S., LEE, J. D., TIAN, Y., SINGH, A. and POCZOS, B. (2018b). Gradient descent learns one-hidden-layer CNN: Dont be afraid of spurious local minima. In *International Conference on Machine Learning*.

DU, S. S., ZHAI, X., POCZOS, B. and SINGH, A. (2019b). Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*.

FREI, S., CAO, Y. and GU, Q. (2021). Provable generalization of sgd-trained neural networks of any width in the presence of adversarial label noise. In *International Conference on Machine Learning*. PMLR.

FREI, S., CHATTERJI, N. S. and BARTLETT, P. L. (2022). Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. *arXiv preprint arXiv:2202.05928* .

GOLOWICH, N., RAKHLIN, A. and SHAMIR, O. (2018). Size-independent sample complexity of neural networks. In *Conference On Learning Theory*.

HARDT, M., RECHT, B. and SINGER, Y. (2016). Train faster, generalize better: stability of stochastic gradient descent. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*. JMLR. org.

HASTIE, T., MONTANARI, A., ROSSET, S. and TIBSHIRANI, R. J. (2019). Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560* .

JACOT, A., GABRIEL, F. and HONGLER, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*.

JI, Z. and TELGARSKY, M. (2020). Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. In *International Conference on Learning Representations*.

LEE, J., XIAO, L., SCHOENHOLZ, S. S., BAHRI, Y., SOHL-DICKSTEIN, J. and PENNINGTON, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in Neural Information Processing Systems*.

LI, Y. and LIANG, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*.

LI, Y., WEI, C. and MA, T. (2019). Towards explaining the regularization effect of initial large learning rate in training neural networks. In *Advances in Neural Information Processing Systems*.

LI, Y. and YUAN, Y. (2017). Convergence analysis of two-layer neural networks with relu activation. In *Advances in Neural Information Processing Systems*.

LI, Z., ZHOU, Z.-H. and GRETTON, A. (2021). Towards an understanding of benign overfitting in neural networks. *arXiv preprint arXiv:2106.03212* .

LIANG, T. and RAKHLIN, A. (2020). Just interpolate: Kernel ridgeless regression can generalize. *The Annals of Statistics* **48** 1329–1347.

LIAO, Z., COUILLET, R. and MAHONEY, M. (2020). A random matrix analysis of random fourier features: beyond the gaussian kernel, a precise phase transition, and the corresponding double descent. In *34th Conference on Neural Information Processing Systems (NeurIPS 2020)*.

LIN, M., CHEN, Q. and YAN, S. (2013). Network in network. *arXiv preprint arXiv:1312.4400* .

LYU, K. and LI, J. (2019). Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890* .

MEI, S. and MONTANARI, A. (2019). The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355* .

MONTANARI, A. and ZHONG, Y. (2020). The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *arXiv preprint arXiv:2007.12826* .

MOU, W., WANG, L., ZHAI, X. and ZHENG, K. (2017). Generalization bounds of sgld for non-convex learning: Two theoretical viewpoints. *arXiv preprint arXiv:1707.05947* .

NEYSHABUR, B., BHOJANAPALLI, S., MCALLESTER, D. and SREBRO, N. (2018). A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representation*.

NEYSHABUR, B., LI, Z., BHOJANAPALLI, S., LECUN, Y. and SREBRO, N. (2019). Towards understanding the role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*.

NEYSHABUR, B., TOMIOKA, R. and SREBRO, N. (2015). Norm-based capacity control in neural networks. In *Conference on Learning Theory*.

SHAMIR, O. (2022). The implicit bias of benign overfitting. *arXiv preprint arXiv:2201.11489* .

SOLTANOLKOTABI, M. (2017). Learning ReLUs via gradient descent. In *Advances in Neural Information Processing Systems*.

SOUDRY, D., HOFFER, E., NACSON, M. S., GUNASEKAR, S. and SREBRO, N. (2018). The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research* **19** 2822–2878.

WU, D. and XU, J. (2020). On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems* **33**.

ZHANG, C., BENGIO, S., HARDT, M., RECHT, B. and VINYALS, O. (2017). Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*.

ZHANG, X., YU, Y., WANG, L. and GU, Q. (2019). Learning one-hidden-layer ReLU networks via gradient descent. In *The 22nd International Conference on Artificial Intelligence and Statistics*.

ZHONG, K., SONG, Z., JAIN, P., BARTLETT, P. L. and DHILLON, I. S. (2017). Recovery guarantees for one-hidden-layer neural networks. In *International Conference on Machine Learning*.

ZOU, D., CAO, Y., LI, Y. and GU, Q. (2021a). Understanding the generalization of adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371* .

ZOU, D., CAO, Y., ZHOU, D. and GU, Q. (2019). Gradient descent optimizes over-parameterized deep ReLU networks. *Machine Learning* .

ZOU, D., WU, J., BRAVERMAN, V., GU, Q. and KAKADE, S. (2021b). Benign overfitting of constant-stepsize sgd for linear regression. In *Conference on Learning Theory*. PMLR.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] This is the first work that precisely characterizes the conditions under which benign overfitting can occur in training convolutional neural networks
   (b) Did you describe the limitations of your work? [Yes] As we have discussed in the conclusion, we only analysis two-layer CNNs. It's important to generalize our analysis to deep convolutional neural networks.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A] This paper focuses on theoretical analyses of neural network training, and does not have any direct negative social impact.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] Please check the appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] The purpose of experiments in this paper is to back up our theoretical findings. The error bars are not required to verify our theory.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] The purpose of experiments in this paper is to back up our theoretical findings. Information about the computation resources are irrelevant to our purpose.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Additional Related Work

There has also been a large number of works studying the optimization and generalization of neural networks. A series of work (Li and Yuan, 2017; Soltanolkotabi, 2017; Du et al., 2018a,b; Zhong et al., 2017; Zhang et al., 2019; Cao and Gu, 2019b) studied the parameter recovery problem in two-layer neural networks, where the data are given by a teacher network and the task is to recover the parameters in the teacher network. These works either focus on the noiseless setting, or requires the number of training data points to be larger than the number of parameters in the network, and therefore does not cover the setting where the neural network can overfit the training data. Another line of works (Neyshabur et al., 2015; Bartlett et al., 2017; Neyshabur et al., 2018; Golowich et al., 2018; Arora et al., 2018) have studied the generalization gap between the training and test losses of neural networks with uniform convergence based arguments. However, these results are not algorithm-dependent and cannot explain benign overfitting. Some recent works studied the generalization gap based on stability based arguments (Bousquet and Elisseeff, 2002; Hardt et al., 2016; Mou et al., 2017; Chen et al., 2018). A more recent line of works studied the convergence (Jacot et al., 2018; Li and Liang, 2018; Du et al., 2019b; Allen-Zhu et al., 2019b; Du et al., 2019a; Zou et al., 2019) and test error bounds (Allen-Zhu et al., 2019a; Arora et al., 2019a,b; Cao and Gu, 2019a; Ji and Telgarsky, 2020; Chen et al., 2019) of over-parameterized networks in the neural tangent kernel regime. However, these works depend on the equivalence between neural network training and kernel methods, which cannot fully explain the success of deep learning. Compared with the works mentioned above, our work has a different focus which is to study the conditions for benign and harmful overfitting.

B Preliminary Lemmas

In this section, we present some pivotal lemmas that give some important properties of the data and the neural network parameters at their random initialization.

Lemma B.1. *Suppose that $\delta > 0$ and $n \geq 8\log(4/\delta)$. Then with probability at least $1-\delta$,*

$$|\{i \in [n] : y_i = 1\}|,\ |\{i \in [n] : y_i = -1\}| \geq n/4.$$

Proof of Lemma B.1. By Hoeffding's inequality, with probability at least $1 - \delta/2$, we have

$$\left|\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\{y_i = 1\} - \frac{1}{2}\right| \leq \sqrt{\frac{\log(4/\delta)}{2n}}.$$

Therefore, as long as $n \geq 8\log(4/\delta)$, we have

$$|\{i \in [n] : y_i = 1\}| = \sum_{i=1}^{n} \mathbb{1}\{y_i = 1\} \geq \frac{n}{2} - n \cdot \sqrt{\frac{\log(4/\delta)}{2n}} \geq \frac{n}{4}.$$

This proves the result for $|\{i \in [n] : y_i = 1\}|$. The proof for $|\{i \in [n] : y_i = -1\}|$ is exactly the same, and we can conclude the proof by applying a union bound. □

The following lemma estimates the norms of the noise vectors $\boldsymbol{\xi}_i$, $i \in [n]$, and gives an upper bound of their inner products with each other.

Lemma B.2. *Suppose that $\delta > 0$ and $d = \Omega(\log(4n/\delta))$. Then with probability at least $1-\delta$,*

$$\sigma_p^2 d/2 \leq \|\boldsymbol{\xi}_i\|_2^2 \leq 3\sigma_p^2 d/2,$$
$$|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_{i'} \rangle| \leq 2\sigma_p^2 \cdot \sqrt{d\log(4n^2/\delta)}$$

for all $i, i' \in [n]$.

Proof of Lemma B.2. By Bernstein's inequality, with probability at least $1 - \delta/(2n)$ we have

$$\big|\|\boldsymbol{\xi}_i\|_2^2 - \sigma_p^2 d\big| = O(\sigma_p^2 \cdot \sqrt{d\log(4n/\delta)}).$$

Therefore, as long as $d = \Omega(\log(4n/\delta))$, we have

$$\sigma_p^2 d/2 \le \|\boldsymbol{\xi}_i\|_2^2 \le 3\sigma_p^2 d/2.$$

Moreover, clearly $\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_{i'} \rangle$ has mean zero. For any i, i' with $i \neq i'$, by Bernstein's inequality, with probability at least $1 - \delta/(2n^2)$ we have

$$|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_{i'} \rangle| \le 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}.$$

Applying a union bound completes the proof. □

The following lemma studies the inner product between a randomly initialized CNN convolutional filter $\mathbf{w}_{j,r}^{(0)}$, $j \in \{+1, -1\}$ and $r \in [m]$ and the signal/noise vectors in the training data. The calculations characterize how the neural network at initialization randomly captures signal and noise information.

Lemma B.3. *Suppose that $d \ge \Omega(\log(mn/\delta))$, $m = \Omega(\log(1/\delta))$. Then with probability at least $1 - \delta$,*

$$|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle| \le \sqrt{2\log(8m/\delta)} \cdot \sigma_0 \|\boldsymbol{\mu}\|_2,$$
$$|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle| \le 2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}$$

for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$. Moreover,

$$\sigma_0 \|\boldsymbol{\mu}\|_2/2 \le \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle \le \sqrt{2\log(8m/\delta)} \cdot \sigma_0 \|\boldsymbol{\mu}\|_2,$$
$$\sigma_0 \sigma_p \sqrt{d}/4 \le \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle \le 2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}$$

for all $j \in \{\pm 1\}$ and $i \in [n]$.

Proof of Lemma B.3. It is clear that for each $r \in [m]$, $j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle$ is a Gaussian random variable with mean zero and variance $\sigma_0^2 \|\boldsymbol{\mu}\|_2^2$. Therefore, by Gaussian tail bound and union bound, with probability at least $1 - \delta/4$,

$$j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle \le |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle| \le \sqrt{2\log(8m/\delta)} \cdot \sigma_0 \|\boldsymbol{\mu}\|_2.$$

Moreover, $\mathbb{P}(\sigma_0 \|\boldsymbol{\mu}\|_2/2 > j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle)$ is an absolute constant, and therefore by the condition on m, we have

$$\begin{aligned}
\mathbb{P}\big(\sigma_0 \|\boldsymbol{\mu}\|_2/2 \le \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle\big) &= 1 - \mathbb{P}(\sigma_0 \|\boldsymbol{\mu}\|_2/2 > \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle\big) \\
&= 1 - \mathbb{P}\big(\sigma_0 \|\boldsymbol{\mu}\|_2/2 > j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle\big)^{2m} \\
&\ge 1 - \delta/4.
\end{aligned}$$

By Lemma B.2, with probability at least $1 - \delta/4$, $\sigma_p \sqrt{d}/\sqrt{2} \le \|\boldsymbol{\xi}_i\|_2 \le \sqrt{3/2} \cdot \sigma_p \sqrt{d}$ for all $i \in [n]$. Therefore, the result for $\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle$ follows the same proof as $j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle$. □

C Signal-noise Decomposition Analysis

In this section, we establish a series of results on the signal-noise decomposition. These results are based on the conclusions in Section B, which hold with high probability. Denote by $\mathcal{E}_{\mathrm{prelim}}$ the event that all the results in Section B hold. Then for simplicity and clarity, we state all the results in this and the following sections conditional on $\mathcal{E}_{\mathrm{prelim}}$.

Lemma C.1 (Restatement of Lemma 5.1)**.** *The coefficients $\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ defined in Definition 4.1 satisfy the following iterative equations:*

$$\gamma_{j,r}^{(0)}, \overline{\rho}_{j,r,i}^{(0)}, \underline{\rho}_{j,r,i}^{(0)} = 0,$$

$$\gamma_{j,r}^{(t+1)} = \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \cdot \boldsymbol{\mu} \rangle) \cdot \|\boldsymbol{\mu}\|_2^2,$$
$$\overline{\rho}_{j,r,i}^{(t+1)} = \overline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = j),$$
$$\underline{\rho}_{j,r,i}^{(t+1)} = \underline{\rho}_{j,r,i}^{(t)} + \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = -j)$$

for all $r \in [m]$, $j \in \{\pm 1\}$ *and* $i \in [n]$.

Proof of Lemma C.1. By our data model in Definition 3.1 and Gaussian initialization of the CNN weights, it is clear that with probability 1, the vectors are linearly independent. Therefore, the decomposition (4.1) is unique. Now consider $\widetilde{\gamma}_{j,r}^{(0)}, \widetilde{\rho}_{j,r,i}^{(0)} = 0$ and

$$\widetilde{\gamma}_{j,r}^{(t+1)} = \widetilde{\gamma}_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \cdot \boldsymbol{\mu} \rangle) \cdot \|\boldsymbol{\mu}\|_2^2,$$
$$\widetilde{\rho}_{j,r,i}^{(t+1)} = \widetilde{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot j y_i,$$

It is then easy to check by (3.1) that

$$\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + j \cdot \widetilde{\gamma}_{j,r}^{(t)} \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu} + \sum_{i=1}^{n} \widetilde{\rho}_{j,r,i}^{(t)} \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i.$$

Hence by the uniqueness of the decomposition we have $\gamma_{j,r}^{(t)} = \widetilde{\gamma}_{j,r}^{(t)}$ and $\rho_{j,r,i}^{(t)} = \widetilde{\rho}_{j,r,i}^{(t)}$. Then we have that

$$\rho_{j,r,i}^{(t)} = -\sum_{s=0}^{t-1} \frac{\eta}{nm} \cdot \ell_i'^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot j y_i$$

Moreover, note that $\ell_i'^{(t)} < 0$ by the definition of the cross-entropy loss. Therefore,

$$\overline{\rho}_{j,r,i}^{(t)} = -\sum_{s=0}^{t-1} \frac{\eta}{nm} \cdot \ell_i'^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = j), \tag{C.1}$$

$$\underline{\rho}_{j,r,i}^{(t)} = -\sum_{s=0}^{t-1} \frac{\eta}{nm} \cdot \ell_i'^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \boldsymbol{\xi}_i \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_i = -j). \tag{C.2}$$

Writing out the iterative versions of (C.1) and (C.2) completes the proof. □

We can futher plug the signal-noise decomposition (4.1) into the iterative formulas in Lemma C.1. By the second equation in Lemma C.1, we have

$$\gamma_{j,r}^{(t+1)} = \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(y_i \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle + y_i \cdot j \cdot \gamma_{j,r}^{(t)}) \cdot \|\boldsymbol{\mu}\|_2^2, \tag{C.3}$$

Moreover, by the third equation in Lemma C.1, we have

$$\overline{\rho}_{j,r,i}^{(t+1)} = \overline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \ell_i'^{(t)} \sigma' \bigg(\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \sum_{i'=1}^{n} \overline{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} + \sum_{i'=1}^{n} \underline{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} \bigg) \cdot \|\boldsymbol{\xi}_i\|_2^2 \tag{C.4}$$

if $j = y_i$, and $\overline{\rho}_{j,r,i}^{(t)} = 0$ for all $t \geq 0$ if $j = -y_i$. Similarly, by the last equation in Lemma C.1, we have

$$\underline{\rho}_{j,r,i}^{(t+1)} = \underline{\rho}_{j,r,i}^{(t)} + \frac{\eta}{nm} \cdot \ell_i'^{(t)} \sigma' \bigg(\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \sum_{i'=1}^{n} \overline{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} + \sum_{i'=1}^{n} \underline{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} \bigg) \cdot \|\boldsymbol{\xi}_i\|_2^2 \tag{C.5}$$

if $j = -y_i$, and $\underline{\rho}_{j,r,i}^{(t)} = 0$ for all $t \geq 0$ if $j = y_i$.

We will now show that the parameter of the signal-noise decomposition will stay a reasonable scale during a long time of training. Let us consider the learning period $0 \leq t \leq T^*$, where $T^* = \eta^{-1}\text{poly}(\epsilon^{-1}, \|\boldsymbol{\mu}\|_2^{-1}, d^{-1}\sigma_p^{-2}, \sigma_0^{-1}, n, m, d)$ is the maximum admissible iterations. Note that we can consider any polynomial training time T^*. Denote $\alpha = 4\log(T^*)$. Here we list the exact conditions for η, σ_0, d required by the proofs in this section, which are part of Condition 4.2:

$$\eta = O\Big(\min\{nm/(q\sigma_p^2 d), nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2 d), nm/(q2^{q+2}\alpha^{q-2}\|\boldsymbol{\mu}\|_2^2)\}\Big), \tag{C.6}$$

$$\sigma_0 \leq [16\sqrt{\log(8mn/\delta)}]^{-1}\min\{\|\boldsymbol{\mu}\|_2^{-1}, (\sigma_p\sqrt{d})^{-1}\}, \tag{C.7}$$

$$d \geq 1024\log(4n^2/\delta)\alpha^2 n^2. \tag{C.8}$$

Denote $\beta = 2\max_{i,j,r}\{|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle|\}$. By Lemma B.3, with probability at least $1-\delta$, we can upper bound β by $4\sqrt{\log(8mn/\delta)}\cdot\sigma_0\cdot\max\{\|\boldsymbol{\mu}\|_2, \sigma_p\sqrt{d}\}$. Then, by (C.7) and (C.8), it is straightforward to verify the following inequality:

$$4\max\left\{\beta, 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\right\} \leq 1. \tag{C.9}$$

Suppose the conditions listed in (C.6), (C.7) and (C.8) hold, we claim that for $0 \leq t \leq T^*$ the following property holds.

Proposition C.2 (Restatement of Proposition 5.3)**.** *Under Condition 4.2, for $0 \leq t \leq T^*$, we have that*

$$0 \leq \gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)} \leq \alpha, \tag{C.10}$$

$$0 \geq \underline{\rho}_{j,r,i}^{(t)} \geq -\beta - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \geq -\alpha. \tag{C.11}$$

for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$.

We will use induction to prove Proposition C.2. We first introduce several technical lemmas that will be used for the proof of Proposition C.2.

Lemma C.3. *For any $t \geq 0$, it holds that $\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle = j\cdot\gamma_{j,r}^{(t)}$ for all $r \in [m]$, $j \in \{\pm 1\}$.*

Proof of Lemma C.3. For any time $t \geq 0$, we have that

$$\begin{aligned}\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle &= j\cdot\gamma_{j,r}^{(t)} + \sum_{i'=1}^{n}\overline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'}, \boldsymbol{\mu}\rangle + \sum_{i'=1}^{n}\underline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'}, \boldsymbol{\mu}\rangle \\ &= j\cdot\gamma_{j,r}^{(t)},\end{aligned}$$

where the equation is by our orthogonal assumption. □

Lemma C.4. *Under Condition 4.2, suppose* (C.10) *and* (C.11) *hold at iteration t. Then*

$$\underline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq \langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle \leq \underline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,\ j \neq y_i,$$

$$\overline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq \langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle \leq \overline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,\ j = y_i$$

for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$.

Proof of Lemma C.4. For $j \neq y_i$, we have that $\overline{\rho}_{j,r,i}^{(t)} = 0$ and

$$\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle = \sum_{i'=1}^{n}\overline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle + \sum_{i'=1}^{n}\underline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle$$

$$\leq 4\sqrt{\frac{\log(4n^2/\delta)}{d}}\sum_{i'\neq i}|\overline{\rho}_{j,r,i'}^{(t)}| + 4\sqrt{\frac{\log(4n^2/\delta)}{d}}\sum_{i'\neq i}|\underline{\rho}_{j,r,i'}^{(t)}| + \underline{\rho}_{j,r,i}^{(t)}$$
$$\leq \underline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,$$

where the second inequality is by Lemma B.2 and the last inequality is by $|\overline{\rho}_{j,r,i'}^{(t)}|, |\underline{\rho}_{j,r,i'}^{(t)}| \leq \alpha$ in (C.10) Similarly, for $y_i = j$, we have that $\underline{\rho}_{j,r,i}^{(t)} = 0$ and

$$\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle = \sum_{i'=1}^{n}\overline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle + \sum_{i'=1}^{n}\underline{\rho}_{j,r,i'}^{(t)}\|\boldsymbol{\xi}_{i'}\|_2^{-2}\cdot\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle$$
$$\leq \overline{\rho}_{j,r,i}^{(t)} + 4\sqrt{\frac{\log(4n^2/\delta)}{d}}\sum_{i'\neq i}|\overline{\rho}_{j,r,i'}^{(t)}| + 4\sqrt{\frac{\log(4n^2/\delta)}{d}}\sum_{i'\neq i}|\underline{\rho}_{j,r,i'}^{(t)}|$$
$$\leq \overline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,$$

where the first inequality is by Lemma B.1 and the second inequality is by $|\overline{\rho}_{j,r,i'}^{(t)}|, |\underline{\rho}_{j,r,i'}^{(t)}| \leq \alpha$ in (C.10). Similarly, we can show that $\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle \geq \underline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\log(4n^2/\delta)/d}\cdot\alpha$ and $\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle \geq \overline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\log(4n^2/\delta)/d}\cdot\alpha$, which completes the proof. □

Lemma C.5. *Under Condition 4.2, suppose* (C.10) *and* (C.11) *hold at iteration t. Then*

$$\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle,$$
$$\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,$$
$$F_j(\mathbf{W}_j^{(t)}, \mathbf{x}_i) \leq 1$$

for all $r \in [m]$ and $j \neq y_i$.

Proof of Lemma C.5. For $j \neq y_i$, we have that

$$\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle = \langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle + y_i\cdot j\cdot\gamma_{j,r}^{(t)} \leq \langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle, \tag{C.12}$$

where the inequality is by $\gamma_{j,r}^{(t)} \geq 0$. In addition, we have

$$\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + \underline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha, \tag{C.13}$$

where the first inequality is by Lemma C.4 and the second inequality is due to $\underline{\rho}_{j,r,i}^{(t)} \leq 0$. Then we can get that

$$F_j(\mathbf{W}_j^{(t)}, \mathbf{x}_i) = \frac{1}{m}\sum_{r=1}^{m}[\sigma(\langle \mathbf{w}_{j,r}^{(t)}, -j\cdot\boldsymbol{\mu}\rangle) + \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)]$$
$$\leq 2^{q+1}\max_{j,r,i}\left\{|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle|, 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\right\}^q$$
$$\leq 1,$$

where the first inequality is by (C.12), (C.13) and the second inequality is by (C.9). □

Lemma C.6. *Under Condition 4.2, suppose* (C.10) *and* (C.11) *hold at iteration t. Then*

$$\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle = \langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle + \gamma_{j,r}^{(t)},$$
$$\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + \overline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$$

for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$. If $\max\{\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}\} = O(1)$, we further have that $F_j(\mathbf{W}_j^{(t)}, \mathbf{x}_i) = O(1)$.

Proof of Lemma C.6. For $j = y_i$, we have that

$$\langle \mathbf{w}_{j,r}^{(t)}, y_i \boldsymbol{\mu} \rangle = \langle \mathbf{w}_{j,r}^{(0)}, y_i \boldsymbol{\mu} \rangle + \gamma_{j,r}^{(t)}, \tag{C.14}$$

where the equation is by Lemma C.3. We also have that

$$\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \overline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha, \tag{C.15}$$

where the inequality is by Lemma C.4. If $\max\{\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}\} = O(1)$, we have following bound

$$\begin{aligned}
F_j(\mathbf{W}_j^{(t)}, \mathbf{x}_i) &= \frac{1}{m}\sum_{r=1}^{m}[\sigma(\langle \mathbf{w}_{j,r}^{(t)}, -j \cdot \boldsymbol{\mu}\rangle + \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)] \\
&\leq 2 \cdot 3^q \max_{j,r,i}\left\{\gamma_{j,r}^{(t)}, \overline{\rho}_{j,r,i}^{(t)}, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle|, 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\right\}^q \\
&= O(1),
\end{aligned}$$

where the first inequality is by (C.14), (C.15) and the second inequality is by (C.9) where $\beta = 2\max_{i,j,r}\{|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle|\}$. □

Now we are ready to prove Proposition C.2.

Proof of Proposition C.2. Our proof is based on induction. The results are obvious at $t = 0$ as all the coefficients are zero. Suppose that there exists $\widetilde{T} \leq T^*$ such that the results in Proposition C.2 hold for all time $0 \leq t \leq \widetilde{T} - 1$. We aim to prove that they also hold for $t = \widetilde{T}$.

We first prove that (C.11) holds for $t = \widetilde{T}$, i.e., $\underline{\rho}_{j,r,i}^{(t)} \geq -\beta - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$ for $t = \widetilde{T}$, $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$. Notice that $\underline{\rho}_{j,r,i}^{(t)} = 0, \forall j = y_i$. Therefore, we only need to consider the case that $j \neq y_i$. When $\underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} \leq -0.5\beta - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$, by Lemma C.4 we have that

$$\langle \mathbf{w}_{j,r}^{(\widetilde{T}-1)}, \boldsymbol{\xi}_i\rangle \leq \underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} + \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0,$$

and thus

$$\begin{aligned}
\underline{\rho}_{j,r,i}^{(\widetilde{T})} &= \underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} + \frac{\eta}{nm} \cdot \ell_i'^{(\widetilde{T}-1)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(\widetilde{T}-1)}, \boldsymbol{\xi}_i\rangle) \cdot \mathbb{1}(y_i = -j)\|\boldsymbol{\xi}_i\|_2^2 \\
&= \underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} \\
&\geq -\beta - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,
\end{aligned}$$

where the last inequality is by induction hypothesis. When $\underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} \geq -0.5\beta - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$, we have that

$$\begin{aligned}
\underline{\rho}_{j,r,i}^{(\widetilde{T})} &= \underline{\rho}_{j,r,i}^{(\widetilde{T}-1)} + \frac{\eta}{nm} \cdot \ell_i'^{(\widetilde{T}-1)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(T-1)}, \boldsymbol{\xi}_i\rangle) \cdot \mathbb{1}(y_i = -j)\|\boldsymbol{\xi}_i\|_2^2 \\
&\geq -0.5\beta - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha - O\bigg(\frac{\eta\sigma_p^2 d}{nm}\bigg)\sigma'\bigg(0.5\beta + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\bigg) \\
&\geq -0.5\beta - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha - O\bigg(\frac{\eta q\sigma_p^2 d}{nm}\bigg)\bigg(0.5\beta + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\bigg) \\
&\geq -\beta - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,
\end{aligned}$$

where we use $-\ell_i'^{(\widetilde{T}-1)} \leq 1$ and $\|\boldsymbol{\xi}_i\|_2 = O(\sigma_p^2 d)$ in the first inequality, the second inequality is by $0.5\beta + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 1$, and the last inequality is by $\eta = O\big(nm/(q\sigma_p^2 d)\big)$ in (C.6).

Next we prove (C.10) holds for $t = \widetilde{T}$. We have

$$
\begin{aligned}
|\ell_i'^{(t)}| &= \frac{1}{1+\exp\{y_i \cdot [F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i) - F_{-1}(\mathbf{W}_{-1}^{(t)}, \mathbf{x}_i)]\}} \\
&\leq \exp\{-y_i \cdot [F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i) - F_{-1}(\mathbf{W}_{-1}^{(t)}, \mathbf{x}_i)]\} \\
&\leq \exp\{-F_{y_i}(\mathbf{W}_{y_i}^{(t)}, \mathbf{x}_i) + 1\}. \qquad \text{(C.16)}
\end{aligned}
$$

where the last inequality is due to Lemma C.5. Moreover, recall the update rule of $\gamma_{j,r}^{(t)}$ and $\overline{\rho}_{j,r,i}^{(t)}$,

$$
\begin{aligned}
\gamma_{j,r}^{(t+1)} &= \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \cdot \boldsymbol{\mu} \rangle) \|\boldsymbol{\mu}\|_2^2, \\
\overline{\rho}_{j,r,i}^{(t+1)} &= \overline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \mathbb{1}(y_i = j) \|\boldsymbol{\xi}_i\|_2^2.
\end{aligned}
$$

Let $t_{j,r,i}$ be the last time $t < T^*$ that $\overline{\rho}_{j,r,i}^{(t)} \leq 0.5\alpha$. Then we have that

$$
\begin{aligned}
\overline{\rho}_{j,r,i}^{(\widetilde{T})} = \overline{\rho}_{j,r,i}^{(t_{j,r,i})} &\underbrace{- \frac{\eta}{nm} \cdot \ell_i'^{(t_{j,r,i})} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t_{j,r,i})}, \boldsymbol{\xi}_i \rangle) \cdot \mathbb{1}(y_i = j) \|\boldsymbol{\xi}_i\|_2^2}_{I_1} \\
&\underbrace{- \sum_{t_{j,r,i} < t < T} \frac{\eta}{nm} \cdot \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \mathbb{1}(y_i = j) \|\boldsymbol{\xi}_i\|_2^2}_{I_2}. \qquad \text{(C.17)}
\end{aligned}
$$

We first bound I_1 as follows,

$$
|I_1| \leq 2qn^{-1}m^{-1}\eta \bigg(\overline{\rho}_{j,r,i}^{(t_{j,r,i})} + 0.5\beta + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \bigg)^{q-1} \sigma_p^2 d \leq q2^q n^{-1}m^{-1}\eta\alpha^{q-1}\sigma_p^2 d \leq 0.25\alpha,
$$

where the first inequality is by Lemmas C.4 and B.2, the second inequality is by $\beta \leq 0.1\alpha$ and $8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0.1\alpha$, the last inequality is by $\eta \leq nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2 d)$.

Second, we bound I_2. For $t_{j,r,i} < t < \widetilde{T}$ and $y_i = j$, we can lower bound $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$ as follows,

$$
\begin{aligned}
\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle &\geq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \overline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\geq -0.5\beta + 0.5\alpha - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\geq 0.25\alpha,
\end{aligned}
$$

where the first inequality is by Lemma C.4, the second inequality is by $\overline{\rho}_{j,r,i}^{(t)} > 0.5\alpha$ and $\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle \geq -0.5\beta$ due to the definition of $t_{j,r,i}$ and β, the last inequality is by $\beta \leq 0.1\alpha$ and $8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0.1\alpha$. Similarly, for $t_{j,r,i} < t < \widetilde{T}$ and $y_i = j$, we can also upper bound $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$ as follows,

$$
\begin{aligned}
\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle &\leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \overline{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\leq 0.5\beta + \alpha + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\leq 2\alpha,
\end{aligned}
$$

where the first inequality is by Lemma C.4, the second inequality is by induction hypothesis $\overline{\rho}_{j,r,i}^{(t)} \leq \alpha$, the last inequality is by $\beta \leq 0.1\alpha$ and $8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0.1\alpha$. Thus, plugging the upper and lower bounds of $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$ into I_2 gives

$$
|I_2| \leq \sum_{t_{j,r,i} < t < \widetilde{T}} \frac{\eta}{nm} \cdot \exp(-\sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) + 1) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot \mathbb{1}(y_i = j) \|\boldsymbol{\xi}_i\|_2^2
$$

$$
\begin{aligned}
&\leq \frac{eq2^q\eta T^*}{nm}\exp(-\alpha^q/4^q)\alpha^{q-1}\sigma_p^2 d \\
&\leq 0.25T^*\exp(-\alpha^q/4^q)\alpha \\
&\leq 0.25T^*\exp(-\log(T^*)^q)\alpha \\
&\leq 0.25\alpha,
\end{aligned}
$$

where the first inequality is by (C.16), the second inequality is by Lemma B.2, the third inequality is by $\eta = O\big(nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2 d)\big)$ in (C.6), the fourth inequality is by our choice of $\alpha = 4\log(T^*)$ and the last inequality is due to the fact that $\log(T^*)^q \geq \log(T^*)$. Plugging the bound of I_1, I_2 into (C.17) completes the proof for $\overline{\rho}$. Similarly, we can prove that $\gamma_{j,r}^{(\widetilde{T})} \leq \alpha$ using $\eta = O\big(nm/(q2^{q+2}\alpha^{q-2}\|\boldsymbol{\mu}\|_2^2)\big)$ in (C.6). Therefore Proposition C.2 holds for $t = \widetilde{T}$, which completes the induction. □

Based on Proposition C.2, we introduce some important properties of the training loss function for $0 \leq t \leq T^*$.

Lemma C.7 (Restatement of Lemma 5.4)**.** *Under Condition 4.2, for $0 \leq t \leq T^*$, the following result holds.*

$$
\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \leq O(\max\{\|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d\})L_S(\mathbf{W}^{(t)}).
$$

Proof of Lemma C.7. We first prove that

$$
-\ell'\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) \cdot \|\nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i)\|_F^2 = O(\max\{\|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d\}). \tag{C.18}
$$

Without loss of generality, we suppose that $y_i = 1$ and $\mathbf{x}_i = [\boldsymbol{\mu}^\top, \boldsymbol{\xi}_i]$. Then we have that

$$
\begin{aligned}
\|\nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i)\|_F &\leq \frac{1}{m}\sum_{j,r}\Big\|\big[\sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu}\rangle)\boldsymbol{\mu} + \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\boldsymbol{\xi}_i\big]\Big\|_2 \\
&\leq \frac{1}{m}\sum_{j,r}\sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu}\rangle)\|\boldsymbol{\mu}\|_2 + \frac{1}{m}\sum_{j,r}\sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\|\boldsymbol{\xi}_i\|_2 \\
&\leq 2q\Big[F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i)\Big]^{(q-1)/q}\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\} \\
&\quad + 2q\Big[F_{-1}(\mathbf{W}_{-1}^{(t)}, \mathbf{x}_i)\Big]^{(q-1)/q}\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\} \\
&\leq 2q\Big[F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i)\Big]^{(q-1)/q}\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\} + 2q\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\},
\end{aligned}
$$

where the first and second inequalities are by triangle inequality, the third inequality is by Jensen's inequality and Lemma B.2, and the last inequality is due to Lemma C.5. Denote $A = F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i)$. Then we have that $A \geq 0$, and besides, $F_{-1}(\mathbf{W}_{-1}^{(t)}, \mathbf{x}_i) \leq 1$ by Lemma C.5. Then we have that

$$
\begin{aligned}
&-\ell'\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) \cdot \|\nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i)\|_F^2 \\
&\quad\leq -\ell'(A-1)\Big(2q\cdot A^{(q-1)/q}\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\} + 2q\cdot\max\{\|\boldsymbol{\mu}\|_2, 2\sigma_p\sqrt{d}\}\Big)^2 \\
&\quad= -4q^2\ell'(A-1)(A^{(q-1)/q}+1)^2\cdot\max\{\|\boldsymbol{\mu}\|_2^2, 4\sigma_p^2 d\} \\
&\quad\leq \big(\max_{z>0} -4q^2\ell'(z-1)(z^{(q-1)/q}+1)^2\big)\max\{\|\boldsymbol{\mu}\|_2^2, 4\sigma_p^2 d\} \\
&\quad\overset{(i)}{=} O(\max\{\|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d\}),
\end{aligned}
$$

where (i) is by $\max_{z\geq 0} -4q^2\ell'(z-1)(z^{(q-1)/q}+1)^2 < \infty$ because ℓ' has an exponentially decaying tail. Now we can upper bound the gradient norm $\|\nabla L_S(\mathbf{W}^{(t)})\|_F$ as follows,

$$
\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \leq \bigg[\frac{1}{n}\sum_{i=1}^n \ell'\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big)\|\nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i)\|_F\bigg]^2
$$

$$\leq \bigg[\frac{1}{n}\sum_{i=1}^{n}\sqrt{-O(\max\{\|\boldsymbol{\mu}\|_2^2,\sigma_p^2 d\})\ell'\big(y_i f(\mathbf{W}^{(t)},\mathbf{x}_i)\big)}\bigg]^2$$
$$\leq O(\max\{\|\boldsymbol{\mu}\|_2^2,\sigma_p^2 d\})\cdot\frac{1}{n}\sum_{i=1}^{n}-\ell'\big(y_i f(\mathbf{W}^{(t)},\mathbf{x}_i)\big)$$
$$\leq O(\max\{\|\boldsymbol{\mu}\|_2^2,\sigma_p^2 d\})L_S(\mathbf{W}^{(t)}),$$

where the first inequality is by triangle inequality, the second inequality is by (C.18), the third inequality is by Cauchy-Schwartz inequality and the last inequality is due to the property of the cross entropy loss $-\ell' \leq \ell$. □

D Signal Learning

In this section, we consider the signal learning case under the condition that $n\|\boldsymbol{\mu}\|_2^q \geq \widetilde{\Omega}(\sigma_p^q(\sqrt{d})^q)$. We remind the readers that the proofs in this section are based on the results in Section B, which hold with high probability.

D.1 First stage

Lemma D.1 (Restatement of Lemma 5.5)**.** *Under the same conditions as Theorem 4.3, in particular if we choose*

$$n\cdot\mathrm{SNR}^q \geq C\log(6/\sigma_0\|\boldsymbol{\mu}\|_2)2^{2q+6}[4\log(8mn/\delta)]^{(q-1)/2}, \tag{D.1}$$

where $C = O(1)$ is a positive constant, there exists time

$$T_1 = \frac{C\log(6/\sigma_0\|\boldsymbol{\mu}\|_2)2^{q+1}m}{\eta\sigma_0^{q-2}\|\boldsymbol{\mu}\|_2^q}$$

such that

- $\max_r \gamma_{j,r}^{(T_1)} \geq 2$ *for* $j\in\{\pm 1\}$.
- $|\rho_{j,r,i}^{(t)}| \leq \sigma_0\sigma_p\sqrt{d}/2$ *for all* $j\in\{\pm1\}, r\in[m]$, $i\in[n]$ *and* $0\leq t\leq T_1$.

Proof of Lemma D.1. Let

$$T_1^+ = \frac{nm\eta^{-1}\sigma_0^{2-q}\sigma_p^{-q}d^{-q/2}}{2^{q+4}q[4\log(8mn/\delta)]^{(q-1)/2}}. \tag{D.2}$$

We first prove the second bullet. Define $\Psi^{(t)} = \max_{j,r,i}|\rho_{j,r,i}^{(t)}| = \max_{j,r,i}\{\overline{\rho}_{j,r,i}^{(t)}, -\underline{\rho}_{j,r,i}^{(t)}\}$. We use induction to show that

$$\Psi^{(t)} \leq \sigma_0\sigma_p\sqrt{d}/2 \tag{D.3}$$

for all $0\leq t\leq T_1^+$. By definition, clearly we have $\Psi^{(0)}=0$. Now suppose that there exists some $\widetilde{T}\leq T_1^+$ such that (D.3) holds for $0<t\leq\widetilde{T}-1$. Then by (C.4) and (C.5) we have

$$\begin{aligned}
\Psi^{(t+1)} &\leq \Psi^{(t)} + \max_{j,r,i}\bigg\{\frac{\eta}{nm}\cdot|\ell_i'^{(t)}|\cdot\sigma'\bigg(\langle\mathbf{w}_{j,r}^{(0)},\boldsymbol{\xi}_i\rangle + \sum_{i'=1}^{n}\Psi^{(t)}\cdot\frac{|\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2} + \sum_{i'=1}^{n}\Psi^{(t)}\cdot\frac{|\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2}\bigg)\cdot\|\boldsymbol{\xi}_i\|_2^2\bigg\}\\
&\leq \Psi^{(t)} + \max_{j,r,i}\bigg\{\frac{\eta}{nm}\cdot\sigma'\bigg(\langle\mathbf{w}_{j,r}^{(0)},\boldsymbol{\xi}_i\rangle + 2\cdot\sum_{i'=1}^{n}\Psi^{(t)}\cdot\frac{|\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2}\bigg)\cdot\|\boldsymbol{\xi}_i\|_2^2\bigg\}\\
&= \Psi^{(t)} + \max_{j,r,i}\bigg\{\frac{\eta}{nm}\cdot\sigma'\bigg(\langle\mathbf{w}_{j,r}^{(0)},\boldsymbol{\xi}_i\rangle + 2\Psi^{(t)} + 2\cdot\sum_{i'\neq i}^{n}\Psi^{(t)}\cdot\frac{|\langle\boldsymbol{\xi}_{i'},\boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2}\bigg)\cdot\|\boldsymbol{\xi}_i\|_2^2\bigg\}\\
&\leq \Psi^{(t)} + \frac{\eta q}{nm}\cdot\bigg[2\cdot\sqrt{\log(8mn/\delta)}\cdot\sigma_0\sigma_p\sqrt{d} + \bigg(2+\frac{4n\sigma_p^2\cdot\sqrt{d\log(4n^2/\delta)}}{\sigma_p^2 d/2}\bigg)\cdot\Psi^{(t)}\bigg]^{q-1}\cdot 2\sigma_p^2 d
\end{aligned}$$

$$\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \big(2 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d} + 4\Psi^{(t)}\big)^{q-1} \cdot 2\sigma_p^2 d$$
$$\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \big(4 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}\big)^{q-1} \cdot 2\sigma_p^2 d,$$

where the second inequality is by $|\ell_i'^{(t)}| \leq 1$, the third inequality is due to Lemmas B.2 and B.3, the fourth inequality follows by the condition that $d \geq 16n^2 \log(4n^2/\delta)$, and the last inequality follows by the induction hypothesis (D.3). Taking a telescoping sum over $t = 0, 1, \ldots, \widetilde{T} - 1$ then gives

$$\begin{aligned}
\Psi^{(\widetilde{T})} &\leq \widetilde{T} \cdot \frac{\eta q}{nm} \cdot \big(4 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}\big)^{q-1} \cdot 2\sigma_p^2 d \\
&\leq T_1^+ \cdot \frac{\eta q}{nm} \cdot \big(4 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}\big)^{q-1} \cdot 2\sigma_p^2 d \\
&\leq \frac{\sigma_0 \sigma_p \sqrt{d}}{2},
\end{aligned}$$

where the second inequality follows by $\widetilde{T} \leq T_1^+$ in our induction hypothesis. Therefore, by induction, we have $\Psi^{(t)} \leq \sigma_0 \sigma_p \sqrt{d}/2$ for all $t \leq T_1^+$.

Now, without loss of generality, let us consider $j = 1$ first. Denote by $T_{1,1}$ the last time for t in the period $[0, T_1^+]$ satisfying that $\max_r \gamma_{1,r}^{(t)} \leq 2$. Then for $t \leq T_{1,1}$, $\max_{j,r,i}\{|\rho_{j,r,i}^{(t)}|\} = O(\sigma_0 \sigma_p \sqrt{d}) = O(1)$ and $\max_r \gamma_{1,r}^{(t)} \leq 2$. Therefore, by Lemmas C.5 and C.6, we know that $F_{-1}(\mathbf{W}_{-1}^{(t)}, \mathbf{x}_i), F_{+1}(\mathbf{W}_{+1}^{(t)}, \mathbf{x}_i) = O(1)$ for all i with $y_i = 1$. Thus there exists a positive constant C_1 such that $-\ell_i'^{(t)} \geq C_1$ for all i with $y_i = 1$.

By (C.3), for $t \leq T_{1,1}$ we have

$$\begin{aligned}
\gamma_{1,r}^{(t+1)} &= \gamma_{1,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(y_i \cdot \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle + y_i \cdot \gamma_{1,r}^{(t)}) \cdot \|\boldsymbol{\mu}\|_2^2 \\
&\geq \gamma_{1,r}^{(t)} + \frac{C_1 \eta}{nm} \cdot \sum_{y_i = 1} \sigma'(\langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle + \gamma_{1,r}^{(t)}) \cdot \|\boldsymbol{\mu}\|_2^2.
\end{aligned}$$

Denote $\widehat{\gamma}_{1,r}^{(t)} = \gamma_{1,r}^{(t)} + \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle$ and let $A^{(t)} = \max_r \widehat{\gamma}_{1,r}^{(t)}$. Then we have

$$\begin{aligned}
A^{(t+1)} &\geq A^{(t)} + \frac{C_1 \eta}{nm} \cdot \sum_{y_i = 1} \sigma'(A^{(t)}) \cdot \|\boldsymbol{\mu}\|_2^2 \\
&\geq A^{(t)} + \frac{C_1 \eta q \|\boldsymbol{\mu}\|_2^2}{4m} \big[A^{(t)}\big]^{q-1} \\
&\geq \bigg(1 + \frac{C_1 \eta q \|\boldsymbol{\mu}\|_2^2}{4m} \big[A^{(0)}\big]^{q-2}\bigg) A^{(t)} \\
&\geq \bigg(1 + \frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^q m}\bigg) A^{(t)},
\end{aligned}$$

where the second inequality is by the lower bound on the number of positive data in Lemma B.1, the third inequality is due to the fact that $A^{(t)}$ is an increasing sequence, and the last inequality follows by $A^{(0)} = \max_r \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle \geq \sigma_0 \|\boldsymbol{\mu}\|_2/2$ proved in Lemma B.3. Therefore, the sequence $A^{(t)}$ will exponentially grow and we have that

$$A^{(t)} \geq \bigg(1 + \frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^q m}\bigg)^t A^{(0)} \geq \exp\bigg(\frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^{q+1} m} t\bigg) A^{(0)} \geq \exp\bigg(\frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^{q+1} m} t\bigg) \frac{\sigma_0 \|\boldsymbol{\mu}\|_2}{2},$$

where the second inequality is due to the fact that $1 + z \geq \exp(z/2)$ for $z \leq 2$ and our condition of η and σ_0 listed in Condition 4.2, and the last inequality follows by Lemma B.3 and $A^{(0)} = \max_r \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle$. Therefore, $A^{(t)}$ will reach 3 within $T_1 = \frac{\log(6/\sigma_0 \|\boldsymbol{\mu}\|_2) 2^{q+1} m}{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}$ iterations. Since $\max_r \gamma_{1,r}^{(t)} \geq A^{(t)} - \max_r |\langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle| \geq A^{(t)} - 1$, $\max_r \gamma_{1,r}^{(t)}$ will reach 2 within T_1 iterations. We can next verify that

$$T_1 = \frac{\log(6/\sigma_0 \|\boldsymbol{\mu}\|_2) 2^{q+1} m}{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q} \leq \frac{nm\eta^{-1} \sigma_0^{2-q} \sigma_p^{-q} d^{-q/2}}{2^{q+5} q [4 \log(8mn/\delta)]^{(q-1)/2}} = T_1^+/2,$$

where the inequality holds due to our SNR condition in (D.1). Therefore, by the definition of $T_{1,1}$, we have $T_{1,1} \leq T_1 \leq T_1^+/2$, where we use the non-decreasing property of γ. The proof for $j = -1$ is similar, and we can prove that $\max_r \gamma_{-1,r}^{(T_{1,-1})} \geq 2$ while $T_{1,-1} \leq T_1 \leq T_1^+/2$, which completes the proof. □

D.2 Second Stage

By the results we get in the first stage we know that

$$\mathbf{w}_{j,r}^{(T_1)} = \mathbf{w}_{j,r}^{(0)} + j \cdot \gamma_{j,r}^{(T_1)} \cdot \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2} + \sum_{i=1}^{n} \overline{\rho}_{j,r,i}^{(T_1)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} + \sum_{i=1}^{n} \underline{\rho}_{j,r,i}^{(T_1)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2}.$$

And at the beginning of the second stage, we have following property holds:

- $\max_r \gamma_{j,r}^{(T_1)} \geq 2, \forall j \in \{\pm 1\}$.
- $\max_{j,r,i} |\rho_{j,r,i}^{(T_1)}| \leq \widehat{\beta}$ where $\widehat{\beta} = \sigma_0 \sigma_p \sqrt{d}/2$.

Lemma 5.1 implies that the learned feature $\gamma_{j,r}^{(t)}$ will not get worse, i.e., for $t \geq T_1$, we have that $\gamma_{j,r}^{(t+1)} \geq \gamma_{j,r}^{(t)}$, and therefore $\max_r \gamma_{j,r}^{(t)} \geq 2$. Now we choose $\mathbf{W}^*$ as follows:

$$\mathbf{w}_{j,r}^* = \mathbf{w}_{j,r}^{(0)} + 2qm \log(2q/\epsilon) \cdot j \cdot \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2}.$$

Based on the above definition of $\mathbf{W}^*$, we have the following lemma.

Lemma D.2. *Under the same conditions as Theorem 4.4, we have that* $\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F \leq \widetilde{O}(m^{3/2}\|\boldsymbol{\mu}\|_2^{-1})$.

Proof of Lemma D.2. We have

$$\begin{aligned}
\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F &\leq \|\mathbf{W}^{(T_1)} - \mathbf{W}^{(0)}\|_F + \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F \\
&\leq \sum_{j,r} \frac{\gamma_{j,r}^{(T_1)}}{\|\boldsymbol{\mu}\|_2} + \sum_{j,r,i} \frac{|\overline{\rho}_{j,r,i}^{(T_1)}|}{\|\boldsymbol{\xi}_i\|_2} + \sum_{j,r,i} \frac{|\underline{\rho}_{j,r,i}^{(T_1)}|}{\|\boldsymbol{\xi}_i\|_2} + O(m^{3/2}\log(1/\epsilon))\|\boldsymbol{\mu}\|_2^{-1} \\
&\leq \widetilde{O}(m\|\boldsymbol{\mu}\|^{-1}) + O(nm\sigma_0) + O(m^{3/2}\log(1/\epsilon))\|\boldsymbol{\mu}\|_2^{-1} \\
&\leq \widetilde{O}(m^{3/2}\|\boldsymbol{\mu}\|_2^{-1}),
\end{aligned}$$

where the first inequality is by triangle inequality, the second inequality is by our decomposition of $\mathbf{W}^{(T_1)}$ and the definition of $\mathbf{W}^*$, the third inequality is by Proposition C.2 and Lemma D.1, and the last inequality is by our condition of σ_0 in Condition 4.2. □

Lemma D.3. *Under the same conditions as Theorem 4.3, we have that* $y_i \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle \geq q\log(2q/\epsilon)$ *for all* $i \in [n]$ *and* $T_1 \leq t \leq T^*$.

Proof of Lemma D.3. Recall that $f(\mathbf{W}^{(t)}, \mathbf{x}_i) = (1/m)\sum_{j,r} j \cdot \big[\sigma(\langle \mathbf{w}_{j,r}, y_i \cdot \boldsymbol{\mu}\rangle) + \sigma(\langle \mathbf{w}_{j,r}, \boldsymbol{\xi}_i\rangle)\big]$, so we have

$$\begin{aligned}
y_i \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle &= \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\langle \boldsymbol{\mu}, j\mathbf{w}_{j,r}^*\rangle + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\langle y_i\boldsymbol{\xi}_i, j\mathbf{w}_{j,r}^*\rangle \\
&= \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle) 2qm\log(2q/\epsilon) + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\langle \boldsymbol{\mu}, j\mathbf{w}_{j,r}^{(0)}\rangle \\
&\quad + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\langle y_i\boldsymbol{\xi}_i, j\mathbf{w}_{j,r}^{(0)}\rangle \\
&\geq \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle) 2qm\log(2q/\epsilon) - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2)
\end{aligned}$$

$$-\frac{1}{m}\sum_{j,r}\sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\widetilde{O}(\sigma_0\sigma_p\sqrt{d}), \tag{D.4}$$

where the inequality is by Lemma B.3. Next we will bound the inner-product terms in (D.4) respectively. By Lemma C.6, we have that for $j = y_i$

$$\max_r\{\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle\} = \max_r\{\gamma_{j,r}^{(t)} + \langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle\} \geq 2 - \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2) \geq 1. \tag{D.5}$$

We can also get the upper bound of the inner products between the parameter and the signal (noise) as follows,

$$\begin{aligned}
&|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu}\rangle| \overset{(i)}{\leq} |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle| + |\gamma_{j,r}^{(t)}| \overset{(ii)}{\leq} \widetilde{O}(1)\\
&|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle| \overset{(iii)}{\leq} |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle| + |\underline{\rho}_{j,r,i}^{(t)}| + |\overline{\rho}_{j,r,i}^{(t)}| + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \overset{(iv)}{\leq} \widetilde{O}(1),
\end{aligned} \tag{D.6}$$

where (i) is by Lemma C.3, (iii) is by Lemma C.4, (ii) and (iv) are due to Proposition C.2. Plugging (D.5) and (D.6) into (D.4) gives,

$$y_i\langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^*\rangle \geq 2q\log(2q/\epsilon) - \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2) - \widetilde{O}(\sigma_0\sigma_p\sqrt{d}) \geq q\log(2q/\epsilon),$$

where the last inequality is by $\sigma_0 \leq \widetilde{O}(m^{-2/(q-2)}n^{-1})\cdot\min\{(\sigma_p\sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\}$ in Condition 4.2. This completes the proof. □

Lemma D.4. *Under the same conditions as Theorem 4.3, we have that*

$$\|\mathbf{W}^{(t)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(t+1)} - \mathbf{W}^*\|_F^2 \geq (2q-1)\eta L_S(\mathbf{W}^{(t)}) - \eta\epsilon$$

for all $T_1 \leq t \leq T^*$.

Proof of Lemma D.4. We first apply a proof technique similar to Lemma 2.6 in Ji and Telgarsky (2020). The difference between our analysis and Ji and Telgarsky (2020) is that here the neural network is q homogeneous rather than 1 homogeneous.

$$\begin{aligned}
&\|\mathbf{W}^{(t)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(t+1)} - \mathbf{W}^*\|_F^2\\
&= 2\eta\langle \nabla L_S(\mathbf{W}^{(t)}), \mathbf{W}^{(t)} - \mathbf{W}^*\rangle - \eta^2\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2\\
&= \frac{2\eta}{n}\sum_{i=1}^n \ell_i'^{(t)}[qy_i f(\mathbf{W}^{(t)}, \mathbf{x}_i) - \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^*\rangle] - \eta^2\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2\\
&\geq \frac{2\eta}{n}\sum_{i=1}^n \ell_i'^{(t)}[qy_i f(\mathbf{W}^{(t)}, \mathbf{x}_i) - q\log(2q/\epsilon)] - \eta^2\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2\\
&\geq \frac{2q\eta}{n}\sum_{i=1}^n [\ell\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) - \epsilon/(2q)] - \eta^2\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2\\
&\geq (2q-1)\eta L_S(\mathbf{W}^{(t)}) - \eta\epsilon,
\end{aligned}$$

where the first inequality is by Lemma D.3, the second inequality is due to the convexity of the cross entropy function, and the last inequality is due to Lemma C.7. □

Lemma D.5 (Restatement of Lemma 5.6)**.** *Under the same conditions as Theorem 4.3, let $T = T_1 + \Big\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \Big\rfloor = T_1 + \widetilde{O}(m^3\eta^{-1}\epsilon^{-1}\|\boldsymbol{\mu}\|_2^{-2})$. Then we have $\max_{j,r,i}|\rho_{j,r,i}^{(t)}| \leq 2\widehat{\beta} = \sigma_0\sigma_p\sqrt{d}$ for all $T_1 \leq t \leq T$. Besides,*

$$\frac{1}{t - T_1 + 1}\sum_{s=T_1}^t L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(t - T_1 + 1)} + \frac{\epsilon}{2q-1}$$

for all $T_1 \leq t \leq T$, and we can find an iterate with training loss smaller than ϵ within T iterations.

Proof of Lemma D.5. By Lemma D.4, for any $t \in [T_1, T]$, we have that

$$\|\mathbf{W}^{(s)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(s+1)} - \mathbf{W}^*\|_F^2 \geq (2q-1)\eta L_S(\mathbf{W}^{(s)}) - \eta\epsilon$$

holds for $s \leq t$. Taking a summation, we obtain that

$$\sum_{s=T_1}^{t} L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2 + \eta\epsilon(t - T_1 + 1)}{(2q-1)\eta} \tag{D.7}$$

for all $T_1 \leq t \leq T$. Dividing $(t - T_1 + 1)$ on both side of (D.7) gives that

$$\frac{1}{t - T_1 + 1}\sum_{s=T_1}^{t} L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(t - T_1 + 1)} + \frac{\epsilon}{2q-1}.$$

Then we can take $t = T$ and have that

$$\frac{1}{T - T_1 + 1}\sum_{s=T_1}^{T} L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(T - T_1 + 1)} + \frac{\epsilon}{2q-1} \leq \frac{3\epsilon}{2q-1} < \epsilon,$$

where we use the fact that $q > 2$ and our choice that $T = T_1 + \left\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \right\rfloor$. Because the mean is smaller than ϵ, we can conclude that there exist $T_1 \leq t \leq T$ such that $L_S(\mathbf{W}^{(t)}) < \epsilon$.

Finally, we will prove that $\max_{j,r,i} |\rho_{j,r,i}^{(t)}| \leq 2\widehat{\beta}$ for all $t \in [T_1, T]$. Plugging $T = T_1 + \left\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \right\rfloor$ into (D.7) gives that

$$\sum_{s=T_1}^{T} L_S(\mathbf{W}^{(s)}) \leq \frac{2\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta} = \widetilde{O}(\eta^{-1}m^3\|\boldsymbol{\mu}\|_2^2), \tag{D.8}$$

where the inequality is due to $\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F \leq \widetilde{O}(m^{3/2}\|\boldsymbol{\mu}\|_2^{-1})$ in Lemma D.2. Define $\Psi^{(t)} = \max_{j,r,i} |\rho_{j,r,i}^{(t)}|$. We will use induction to prove $\Psi^{(t)} \leq 2\widehat{\beta}$ for all $t \in [T_1, T]$. At $t = T_1$, by the definition of $\widehat{\beta}$, clearly we have $\Psi^{(T_1)} \leq \widehat{\beta} \leq 2\widehat{\beta}$. Now suppose that there exists $\widetilde{T} \in [T_1, T]$ such that $\Psi^{(t)} \leq 2\widehat{\beta}$ for all $t \in [T_1, \widetilde{T} - 1]$. Then for $t \in [T_1, \widetilde{T} - 1]$, by (C.4) and (C.5) we have

$$\begin{aligned}
\Psi^{(t+1)} &\leq \Psi^{(t)} + \max_{j,r,i}\left\{\frac{\eta}{nm} \cdot |\ell_i'^{(t)}| \cdot \sigma'\left(\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + 2\sum_{i'=1}^{n}\Psi^{(t)} \cdot \frac{|\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2}\right) \cdot \|\boldsymbol{\xi}_{i'}\|_2^2\right\} \\
&= \Psi^{(t)} + \max_{j,r,i}\left\{\frac{\eta}{nm} \cdot |\ell_i'^{(t)}| \cdot \sigma'\left(\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + 2\Psi^{(t)} + 2\sum_{i'\neq i}^{n}\Psi^{(t)} \cdot \frac{|\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2^2}\right) \cdot \|\boldsymbol{\xi}_{i'}\|_2^2\right\} \\
&\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \max_i |\ell_i'^{(t)}| \cdot \Bigg[2 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0\sigma_p\sqrt{d} \\
&\quad + \left(2 + \frac{4n\sigma_p^2 \cdot \sqrt{d\log(4n^2/\delta)}}{\sigma_p^2 d/2}\right) \cdot \Psi^{(t)}\Bigg]^{q-1} \cdot 2\sigma_p^2 d \\
&\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \max_i |\ell_i'^{(t)}| \cdot \big(2 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0\sigma_p\sqrt{d} + 4 \cdot \Psi^{(t)}\big)^{q-1} \cdot 2\sigma_p^2 d,
\end{aligned}$$

where the second inequality is due to Lemmas B.2 and B.3, and the last inequality follows by the assumption that $d \geq 16n^2\log(4n^2/\delta)$. Taking a telescoping sum over $t = 0, 1, \ldots, \widetilde{T} - 1$, we have that

$$\begin{aligned}
\Psi^{(T)} &\overset{(i)}{\leq} \Psi^{(T_1)} + \frac{\eta q}{nm}\sum_{s=T_1}^{\widetilde{T}-1} \max_i |\ell_i'^{(s)}| \widetilde{O}(\sigma_p^2 d)\widehat{\beta}^{q-1} \\
&\overset{(ii)}{\leq} \Psi^{(T_1)} + \frac{\eta q}{nm}\widetilde{O}(\sigma_p^2 d)\widehat{\beta}^{q-1}\sum_{s=T_1}^{\widetilde{T}-1} \max_i \ell_i^{(s)}
\end{aligned}$$

$$
\begin{aligned}
&\overset{(iii)}{\leq} \Psi^{(T_1)} + \widetilde{O}(\eta m^{-1}\sigma_p^2 d)\widehat{\beta}^{q-1}\sum_{s=T_1}^{\widetilde{T}-1} L_S(\mathbf{W}^{(s)}) \\
&\overset{(iv)}{\leq} \Psi^{(T_1)} + \widetilde{O}(m^2\mathrm{SNR}^{-2})\widehat{\beta}^{q-1} \\
&\overset{(v)}{\leq} \widehat{\beta} + \widetilde{O}(m^2 n^{2/q}\widehat{\beta}^{q-2})\widehat{\beta} \\
&\overset{(vi)}{\leq} 2\widehat{\beta},
\end{aligned}
$$

where (i) is by out induction hypothesis that $\Psi^{(t)} \leq 2\widehat{\beta}$, (ii) is by $|\ell'| \leq \ell$, (iii) is by $\max_i \ell_i^{(s)} \leq \sum_i \ell_i^{(s)} = nL_S(\mathbf{W}^{(s)})$, (iv) is due to $\sum_{s=T_1}^{\widetilde{T}-1} L_S(\mathbf{W}^{(s)}) \leq \sum_{s=T_1}^{T} L_S(\mathbf{W}^{(s)}) = \widetilde{O}(\eta^{-1}m^3\|\boldsymbol{\mu}\|_2^2)$ in (D.8), (v) is by $n\mathrm{SNR}^q \geq \widetilde{\Omega}(1)$, and (vi) is by the definition that $\widehat{\beta} = \sigma_0\sigma_p\sqrt{d}/2$ and $\widetilde{O}(m^2 n^{2/q}\widehat{\beta}^{q-2}) = \widetilde{O}(m^2 n^{2/q}(\sigma_0\sigma_p\sqrt{d})^{q-2}) \leq 1$ by Condition 4.2. Therefore, $\Psi^{(\widetilde{T})} \leq 2\widehat{\beta}$, which completes the induction. □

D.3 Population Loss

Consider a new data point $(\mathbf{x}, y)$ drawn from the distribution defined in Definition 3.1. Without loss of generality, we suppose that the first patch is the signal patch and the second patch is the noise patch, i.e., $\mathbf{x} = [y\boldsymbol{\mu}, \boldsymbol{\xi}]$. Moreover, by the signal-noise decomposition, the learned neural network has parameter

$$
\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + j\cdot\gamma_{j,r}^{(t)}\cdot\frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2} + \sum_{i=1}^{n}\overline{\rho}_{j,r,i}^{(t)}\cdot\frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} + \sum_{i=1}^{n}\underline{\rho}_{j,r,i}^{(t)}\cdot\frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2}
$$

for $j \in \{\pm 1\}$ and $r \in [m]$.

Lemma D.6. *Under the same conditions as Theorem 4.3, we have that* $\max_{j,r}|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle| \leq 1/2$ *for all* $0 \leq t \leq T$.

Proof. We can get the upper bound of the inner products between the parameter and the noise as follows:

$$
\begin{aligned}
|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle| &\overset{(i)}{\leq} |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle| + |\underline{\rho}_{j,r,i}^{(t)}| + |\overline{\rho}_{j,r,i}^{(t)}| + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\overset{(ii)}{\leq} 2\sqrt{\log(8mn/\delta)}\cdot\sigma_0\sigma_p\sqrt{d} + \sigma_0\sigma_p\sqrt{d} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\
&\overset{(iii)}{\leq} 1/2
\end{aligned}
$$

for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$, where (i) is by Lemma C.3, (ii) is due to $|\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle| \leq 2\sqrt{\log(8mn/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}$ in Lemma B.3 and $\max_{j,r,i}|\rho_{j,r,i}^{(t)}| \leq \sigma_0\sigma_p\sqrt{d}$ in Lemma D.5, and (iii) is due to our condition of $\sigma_0 \leq \widetilde{O}(m^{-2/(q-2)}n^{-1})\cdot(\sigma_p\sqrt{d})^{-1}$ and $d \geq \widetilde{\Omega}(m^2n^4)$ in Condition 4.2. □

Lemma D.7. *Under the same conditions as Theorem 4.3, with probability at least* $1 - 4mT\cdot\exp(-C_2^{-1}\sigma_0^{-2}\sigma_p^{-2}d^{-1})$, *we have that* $\max_{j,r}|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle| \leq 1/2$ *for all* $0 \leq t \leq T$, *where* $C_2 = \widetilde{O}(1)$.

Proof of Lemma D.7. Let $\widetilde{\mathbf{w}}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(t)} - j\cdot\gamma_{j,r}^{(t)}\cdot\frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2}$, then we have that $\langle\widetilde{\mathbf{w}}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle = \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle$ and

$$
\|\widetilde{\mathbf{w}}_{j,r}^{(t)}\|_2 \leq \widetilde{O}(\sigma_0\sqrt{d} + n\sigma_0) = \widetilde{O}(\sigma_0\sqrt{d}), \tag{D.9}
$$

where the equality is due to $d \geq \widetilde{\Omega}(m^2n^4)$ by Condition 4.2.

By (D.9), $\max_{j,r} \|\widetilde{\mathbf{w}}_{j,r}^{(t)}\|_2 \leq C_1\sigma_0\sqrt{d}$, where $C_1 = \widetilde{O}(1)$. Clearly $\langle \widetilde{\mathbf{w}}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle$ is a Gaussian distribution with mean zero and standard deviation smaller than $C_1\sigma_0\sigma_p\sqrt{d}$. Therefore, the probability is bounded by

$$\mathbb{P}\big(|\langle \widetilde{\mathbf{w}}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle| \geq 1/2\big) \leq 2\exp\bigg(-\frac{1}{8C_1^2\sigma_0^2\sigma_p^2 d}\bigg).$$

Applying a union bound over j, r, t completes the proof. □

Lemma D.8 (Restatement of Lemma 5.7)**.** *Let T be defined in Lemma 5.5 respectively. Under the same conditions as Theorem 4.3, for any $0 \leq t \leq T$ with $L_S(\mathbf{W}^{(t)}) \leq 1$, it holds that $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \leq 6 \cdot L_S(\mathbf{W}^{(t)}) + \exp(-n^2)$.*

Proof of Lemma D.8. Let event $\mathcal{E}$ be the event that Lemma D.7 holds. Then we can divide $L_{\mathcal{D}}(\mathbf{W}^{(t)})$ into two parts:

$$\mathbb{E}\big[\ell\big(yf(\mathbf{W}^{(t)}, \mathbf{x})\big)\big] = \underbrace{\mathbb{E}[\mathbb{1}(\mathcal{E})\ell\big(yf(\mathbf{W}^{(t)}, \mathbf{x})\big)]}_{I_1} + \underbrace{\mathbb{E}[\mathbb{1}(\mathcal{E}^c)\ell\big(yf(\mathbf{W}^{(t)}, \mathbf{x})\big)]}_{I_2}. \tag{D.10}$$

In the following, we bound I_1 and I_2 respectively.

Bounding I_1: Since $L_S(\mathbf{W}^{(t)}) \leq 1$, there must exist one $(\mathbf{x}_i, y_i)$ such that $\ell\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) \leq L_S(\mathbf{W}^{(t)}) \leq 1$, which implies that $y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i) \geq 0$. Therefore, we have that

$$\exp(-y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)) \overset{(i)}{\leq} 2\log\big(1 + \exp(-y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i))\big) = 2\ell\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) \leq 2L_S(\mathbf{W}^{(t)}), \tag{D.11}$$

where (i) is by $z \leq 2\log(1+z), \forall z \leq 1$. If event $\mathcal{E}$ holds, we have that

$$\begin{aligned}
|yf(\mathbf{W}^{(t)}, \mathbf{x}) - y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)| &\leq \frac{1}{m}\sum_{j,r}\sigma(\langle \mathbf{w}_{j,r}, \boldsymbol{\xi}_i\rangle) + \frac{1}{m}\sum_{j,r}\sigma(\langle \mathbf{w}_{j,r}, \boldsymbol{\xi}\rangle) \\
&\leq \frac{1}{m}\sum_{j,r}\sigma(1/2) + \frac{1}{m}\sum_{j,r}\sigma(1/2) \\
&\leq 1,
\end{aligned} \tag{D.12}$$

where the second inequality is by $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle| \leq 1/2$ in Lemma D.7 and $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle| \leq 1/2$ in Lemma D.6. Thus we have that

$$\begin{aligned}
I_1 &\leq \mathbb{E}[\mathbb{1}(\mathcal{E})\exp(-yf(\mathbf{W}^{(t)}, \mathbf{x}))] \\
&\leq e \cdot \mathbb{E}[\mathbb{1}(\mathcal{E})\exp(-y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i))] \\
&\leq 2e \cdot \mathbb{E}[\mathbb{1}(\mathcal{E})L_S(\mathbf{W}^{(t)})],
\end{aligned}$$

where the first inequality is by the property of cross-entropy loss that $\ell(z) \leq \exp(-z)$ for all z, the second inequality is by (D.12), and the third inequality is by (D.11). Dropping the event in the expectation gives $I_1 \leq 6L_S(\mathbf{W}^{(t)})$.

Bounding I_2: Next we bound the second term I_2. We choose an arbitrary training data $(\mathbf{x}_{i'}, y_{i'})$ such that $y_{i'} = y$. Then we have

$$\begin{aligned}
\ell\big(yf(\mathbf{W}^{(t)}, \mathbf{x})\big) &\leq \log(1 + \exp(F_{-y}(\mathbf{W}^{(t)}, \mathbf{x}))) \\
&\leq 1 + F_{-y}(\mathbf{W}^{(t)}, \mathbf{x}) \\
&= 1 + \frac{1}{m}\sum_{j=-y, r\in[m]}\sigma(\langle \mathbf{w}_{j,r}^{(t)}, y\boldsymbol{\mu}\rangle) + \frac{1}{m}\sum_{j=-y, r\in[m]}\sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle) \\
&\leq 1 + F_{-y_{i'}}(\mathbf{W}_{-y_{i'}}, \mathbf{x}_{i'}) + \frac{1}{m}\sum_{j=-y, r\in[m]}\sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle)
\end{aligned}$$

$$\leq 2 + \frac{1}{m} \sum_{j=-y, r\in[m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi} \rangle)$$
$$\leq 2 + \widetilde{O}((\sigma_0\sqrt{d})^q)\|\boldsymbol{\xi}\|^q, \tag{D.13}$$

where the first inequality is due to $F_y(\mathbf{W}^{(t)}, \mathbf{x}) \geq 0$, the second inequality is by the property of cross-entropy loss, i.e., $\log(1+\exp(z)) \leq 1 + z$ for all $z \geq 0$, the third inequality is by $\frac{1}{m}\sum_{j=-y,r\in[m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, y\boldsymbol{\mu}\rangle) \leq F_{-y}(\mathbf{W}_{-y}, \mathbf{x}_{i'}) = F_{-y_{i'}}(\mathbf{W}_{-y_{i'}}, \mathbf{x}_{i'})$, the fourth inequality is by $F_{-y_{i'}}(\mathbf{W}_{-y_{i'}}, \mathbf{x}_{i'}) \leq 1$ in Lemma C.5, and the last inequality is due to $\langle \widetilde{\mathbf{w}}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle = \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle \leq \|\widetilde{\mathbf{w}}_{j,r}^{(t)}\|_2\|\boldsymbol{\xi}\|_2 \leq \widetilde{O}(\sigma_0\sqrt{d})\|\boldsymbol{\xi}\|_2$ in (D.9). Then we further have that

$$I_2 \leq \sqrt{\mathbb{E}[\mathbb{1}(\mathcal{E}^c)]} \cdot \sqrt{\mathbb{E}\Big[\ell\big(yf(\mathbf{W}^{(t)}, \mathbf{x})\big)^2\Big]}$$
$$\leq \sqrt{\mathbb{P}(\mathcal{E}^c)} \cdot \sqrt{4 + \widetilde{O}((\sigma_0\sqrt{d})^{2q})\mathbb{E}[\|\boldsymbol{\xi}\|_2^{2q}]}$$
$$\leq \exp[-\widetilde{\Omega}(\sigma_0^{-2}\sigma_p^{-2}d^{-1}) + \mathrm{polylog}(d)]$$
$$\leq \exp(-n^2),$$

where the first inequality is by Cauchy-Schwartz inequality, the second inequality is by (D.13), the third inequality is by Lemma D.7 and the fact that $\sqrt{4 + \widetilde{O}((\sigma_0\sqrt{d})^{2q})\mathbb{E}[\|\boldsymbol{\xi}\|_2^{2q}]} = O(\mathrm{poly}(d))$, and the last inequality is by our condition $\sigma_0 \leq \widetilde{O}(m^{-2/(q-2)}n^{-1}) \cdot (\sigma_p\sqrt{d})^{-1}$ in Condition 4.2. Plugging the bounds of I_1, I_2 into (D.10) completes the proof. □

E Noise Memorization

In this section, we will consider the noise memorization case under the condition that $\sigma_p^q(\sqrt{d})^q \geq \widetilde{\Omega}(n\|\boldsymbol{\mu}\|_2^q)$. We remind the readers that the proofs in this section are based on the results in Section B, which hold with high probability.

We also remind readers that $\alpha = 4\log(T^*)$ is defined in Appendix C. Denote $\bar{\beta} = \min_i \max_r \langle \mathbf{w}_{y_i,r}^{(0)}, \boldsymbol{\xi}_i\rangle$. The following lemma provides a lower bound of $\bar{\beta}$.

Lemma E.1. *Under the same conditions as Theorem 4.4, if in particular*

$$\sigma_0 \geq 80n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\}, \tag{E.1}$$

then we have that $\bar{\beta} \geq \sigma_0\sigma_p\sqrt{d}/4 \geq 20n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$.

Proof of Lemma E.1. Because $\sigma_p^q(\sqrt{d})^q \geq \widetilde{\Omega}(n\|\boldsymbol{\mu}\|_2^q)$, we have that $\sigma_p\sqrt{d} \geq \|\boldsymbol{\mu}\|_2$. Therefore we have that

$$\bar{\beta} \geq \sigma_0\sigma_p\sqrt{d}/4$$
$$= \sigma_0/4 \cdot \max\{\sigma_p\sqrt{d}, \|\boldsymbol{\mu}\|_2\}$$
$$\geq 20n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,$$

where the first inequality is by Lemma B.3 and the last inequality is by our lower bound condition of σ_0 in (E.1). □

E.1 First Stage

Lemma E.2. *Under the same conditions as Theorem 4.4, in particular if we choose*

$$n^{-1}\mathrm{SNR}^{-q} \geq \frac{C2^{q+2}\log\big(20/(\sigma_0\sigma_p\sqrt{d})\big)(\sqrt{2\log(8m/\delta)})^{q-2}}{0.15^{q-2}}, \tag{E.2}$$

where $C = O(1)$ is a positive constant, then there exist

$$T_1 = \frac{C\log\big(20/(\sigma_0\sigma_p\sqrt{d})\big)4mn}{0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}$$

such that

- $\max_{j,r}\overline{\rho}_{j,r,i}^{(T_1)} \geq 2$ *for all* $i \in [n]$.
- $\max_{j,r}\gamma_{j,r}^{(t)} = \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2)$ *for all* $0 \leq t \leq T_1$.
- $\max_{j,r,i}|\underline{\rho}_{j,r,i}^{(t)}| = \widetilde{O}(\sigma_0\sigma_p\sqrt{d})$ *for all* $0 \leq t \leq T_1$.

Proof of Lemma E.2. Let

$$T_1^+ = \frac{m}{\eta q 2^{q-1}(\sqrt{2\log(8m/\delta)})^{q-2}\sigma_0^{q-2}\|\boldsymbol{\mu}\|_2^q}. \tag{E.3}$$

By Proposition C.2, we have that $\underline{\rho}_{j,r,i}^{(t)} \geq -\beta - 16n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \geq -\beta - \bar{\beta}$ for all $j \in \{\pm 1\}$, $r \in [m]$, $i \in [n]$ and $0 \leq t \leq T^*$. Since $\underline{\rho}_{j,r,i}^{(t)} \leq 0$ and $\bar{\beta} \leq \beta = \widetilde{O}(\sigma_0\sigma_p\sqrt{d})$, we have that $\max_{j,r,i}|\underline{\rho}_{j,r,i}^{(t)}| = \widetilde{O}(\sigma_0\sigma_p\sqrt{d})$. Next, we will carefully compute the growth of the $\gamma_{j,r}^{(t)}$.

$$\begin{aligned}
\gamma_{j,r}^{(t+1)} &= \gamma_{j,r}^{(t)} - \frac{\eta}{nm}\cdot\sum_{i=1}^n \ell_i'^{(t)}\cdot\sigma'(\langle\mathbf{w}_{j,r}^{(t)}, y_i\cdot\boldsymbol{\mu}\rangle)\|\boldsymbol{\mu}\|_2^2 \\
&\leq \gamma_{j,r}^{(t)} + \frac{\eta}{nm}\cdot\sum_{i=1}^n \sigma'(|\langle\mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle| + \gamma_{j,r}^{(t)})\|\boldsymbol{\mu}\|_2^2,
\end{aligned}$$

where the inequality is by $|\ell'| \leq 1$. Let $A^{(t)} = \max_{j,r}\{\gamma_{j,r}^{(t)} + |\langle\mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|\}$, then we have that

$$A^{(t+1)} \leq A^{(t)} + \frac{\eta q\|\boldsymbol{\mu}\|_2^2}{m}[A^{(t)}]^{q-1}. \tag{E.4}$$

We will use induction to prove that $A^{(t)} \leq 2A^{(0)}$ for $t \leq T_1^+$. By definition, clearly we have that $A^{(0)} \leq 2A^{(0)}$. Now suppose that there exists some $\widetilde{T} \leq T_1^+$ such that $A^{(t)} \leq 2A^{(0)}$ holds for $0 \leq t \leq \widetilde{T} - 1$. Taking a telescoping sum of (E.4) gives that

$$\begin{aligned}
A^{(\widetilde{T})} &\leq A^{(0)} + \sum_{s=0}^{\widetilde{T}}\frac{\eta q\|\boldsymbol{\mu}\|_2^2}{m}[A^{(s)}]^{q-1} \\
&\leq A^{(0)} + \frac{\eta q\|\boldsymbol{\mu}\|_2^2 T_1^+ 2^{q-1}}{m}[A^{(0)}]^{q-1} \\
&\leq A^{(0)} + \frac{\eta q\|\boldsymbol{\mu}\|_2^2 T_1^+ 2^{q-1}}{m}[\sqrt{2\log(8m/\delta)}\cdot\sigma_0\|\boldsymbol{\mu}\|_2]^{q-2}A^{(0)} \\
&\leq 2A^{(0)},
\end{aligned}$$

where the second inequality is by our induction hypothesis, the third inequality is by $A_0 \leq \sqrt{2\log(8m/\delta)}\cdot\sigma_0\|\boldsymbol{\mu}\|_2$ in Lemma B.3, and the last inequality is by (E.3). Thus we have that $A^{(t)} \leq 2A^{(0)}$ for all $t \leq T_1^+$. Therefore, $\max_{j,r}\gamma_{j,r}^{(t)} \leq A^{(t)} + \max_{j,r}\{|\langle\mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu}\rangle|\} \leq 3A^{(0)}$ for all $0 \leq t \leq T_1^+$. Recall that

$$\overline{\rho}_{j,r,i}^{(t+1)} = \overline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm}\cdot\ell_i'^{(t)}\cdot\sigma'(\langle\mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\cdot\mathbb{1}(y_i = j)\|\boldsymbol{\xi}_i\|_2^2.$$

For $y_i = j$, Lemma C.4 implies that

$$\langle\mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle \geq \langle\mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i\rangle + \overline{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$$

$$\geq \overline{\rho}_{j,r,i}^{(t)} + \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle - 0.4\bar{\beta},$$

where the last inequality is by $\bar{\beta} \geq 20n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha$. Now let $B_i^{(t)} = \max_{j=y_i,r}\{\overline{\rho}_{j,r,i}^{(t)} + \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle - 0.4\bar{\beta}\}$. For each i, denote by $T_1^{(i)}$ the last time in the period $[0, T_1^+]$ satisfying that $\overline{\rho}_{j,r,i}^{(t)} \leq 2$. Then for $t \leq T_1^{(i)}$, $\max_{j,r}\{|\overline{\rho}_{j,r,i}^{(t)}|, |\underline{\rho}_{j,r,i}^{(t)}|\} = O(1)$ and $\max_{j,r}\gamma_{j,r}^{(t)} \leq 3A^{(0)} = O(1)$. Therefore, by Lemmas C.5 and C.6, we know that $F_{-1}(\mathbf{W}^{(t)}, \mathbf{x}_i), F_{+1}(\mathbf{W}^{(t)}, \mathbf{x}_i) = O(1)$. Thus there exists a positive constant C_1 such that $-\ell_i'^{(t)} \geq C_1$ for all $0 \leq t \leq T_1^{(i)}$. It is also easy to check that $B_i^{(0)} \geq 0.6\bar{\beta} \geq 0.15\sigma_0\sigma_p\sqrt{d}$. Then we can carefully compute the growth of $B_i^{(t)}$,

$$\begin{aligned}
B_i^{(t+1)} &\geq B_i^{(t)} + \frac{C_1\eta q\sigma_p^2 d}{2nm}[B_i^{(t)}]^{q-1} \\
&\geq B_i^{(t)} + \frac{C_1\eta q\sigma_p^2 d}{2nm}[B_i^{(0)}]^{q-2}B_i^{(t)} \\
&\geq \bigg(1 + \frac{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}{2nm}\bigg)B_i^{(t)},
\end{aligned}$$

where the second inequality is by the non-decreasing property of $B_i^{(t)}$. Therefore, $B_i^{(t)}$ is an exponentially increasing sequence and we have that

$$\begin{aligned}
B_i^{(t)} &\geq \bigg(1 + \frac{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}{2nm}\bigg)^t B_i^{(0)} \\
&\geq \exp\bigg(\frac{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}{4nm}t\bigg)B_i^{(0)} \\
&\geq \exp\bigg(\frac{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}{4nm}t\bigg)\cdot 0.15\sigma_0\sigma_p\sqrt{d},
\end{aligned}$$

where the second inequality is due to the fact that $1 + z \geq \exp(z/2)$ for $z \leq 2$ and our conditions of η and σ_0 listed in Condition 4.2, and the last inequality is due to $B_i^{(0)} \geq 0.15\sigma_0\sigma_p\sqrt{d}$. Therefore, $B_i^{(t)}$ will reach 3 within $T_1 = \frac{\log\big(20/(\sigma_0\sigma_p\sqrt{d})\big)4mn}{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q}$ iterations. Since $\max_{j=y_i,r}\overline{\rho}_{j,r,i}^{(t)} \geq B_i^{(t)} - \max_{j=y_i,r}|\langle \mathbf{w}_{j,r}^{(0)}, \xi_i\rangle| + 0.4\bar{\beta} \geq B_i^{(t)} - 1$, $\max_{j=y_i,r}\overline{\rho}_{j,r,i}^{(t)}$ will reach 2 within T_1 iterations. We can next verify that

$$T_1 = \frac{\log\big(20/(\sigma_0\sigma_p\sqrt{d})\big)4mn}{C_1 0.15^{q-2}\eta q\sigma_0^{q-2}(\sigma_p^2\sqrt{d})^q} \leq \frac{m}{\eta q 2^q(\sqrt{2\log(8m/\delta)})^{q-2}\sigma_0^{q-2}\|\boldsymbol{\mu}\|_2^q} = T_1^+/2,$$

where the inequality holds due to our SNR condition in (E.2). Therefore, by the definition of $T_1^{(i)}$, we have $T_1^{(i)} \leq T_1 \leq T_1^+/2$, where we use the non-decreasing property of $\overline{\rho}_{j,r,i}$. This completes the proof. □

E.2 Second Stage

By the signal-noise decompositon, at the end of the first stage, we have

$$\mathbf{w}_{j,r}^{(T_1)} = \mathbf{w}_{j,r}^{(0)} + j\cdot\gamma_{j,r}^{(T_1)}\cdot\frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2} + \sum_{i=1}^n \overline{\rho}_{j,r,i}^{(T_1)}\cdot\frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} + \sum_{i=1}^n \underline{\rho}_{j,r,i}^{(T_1)}\cdot\frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2}$$

for $j \in \{\pm 1\}$ and $r \in [m]$. By the results we get in the first stage, we know that at the beginning of this stage, we have following property holds:

- $\max_r \overline{\rho}_{y_i,r,i}^{(T_1)} \geq 2$ for all $i \in [n]$.
- $\max_{j,r,i}|\underline{\rho}_{j,r,i}^{(T_1)}| = \widetilde{O}(\sigma_0\sigma_p\sqrt{d})$.
- $\max_{j,r}\gamma_{j,r}^{(T_1)} \leq \widehat{\beta}'$, where $\widehat{\beta}' = \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2)$.

Note that Lemma 5.1 implies that the learned noise $\overline{\rho}_{j,r,i}^{(t)}$ will not decrease, i.e., $\overline{\rho}_{j,r,i}^{(t+1)} \geq \overline{\rho}_{j,r,i}^{(t)}$. Therefore, for all data index i, we have $\max_r \overline{\rho}_{y_i,r,i}^{(t)} \geq 2$ for all $t \geq T_1$. Now we choose $\mathbf{W}^*$ as follows

$$\mathbf{w}_{j,r}^* = \mathbf{w}_{j,r}^{(0)} + 2qm\log(2q/\epsilon))\bigg[\sum_{i=1}^n \mathbb{1}(j = y_i) \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2}\bigg].$$

Based on the definition of $\mathbf{W}^*$, we have the following lemma.

Lemma E.3. *Under the same conditions as Theorem 4.4, we have that* $\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F \leq \widetilde{O}(m^2 n^{1/2}\sigma_p^{-1}d^{-1/2})$.

Proof of Lemma E.3. We have

$$\begin{aligned}
\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F &\leq \|\mathbf{W}^{(T_1)} - \mathbf{W}^{(0)}\|_F + \|\mathbf{W}^{(0)} - \mathbf{W}^*\|_F \\
&\leq \sum_{j,r} \gamma_{j,r}^{(T_1)} \|\boldsymbol{\mu}\|_2^{-1} + O(\sqrt{m}) \max_{j,r} \bigg\| \sum_{i=1}^n \overline{\rho}_{j,r,i}^{(T_1)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} + \sum_{i=1}^n \underline{\rho}_{j,r,i}^{(T_1)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} \bigg\|_2 \\
&\quad + O(m^{3/2}n^{1/2}\log(1/\epsilon)\sigma_p^{-1}d^{-1/2}) \\
&\leq \widetilde{O}(m^{3/2}n^{1/2}\sigma_p^{-1}d^{-1/2}),
\end{aligned}$$

where the first inequality is by triangle inequality, the second inequality is by our decomposition of $\mathbf{W}^{(T_1)}, \mathbf{W}^*$ and Lemma B.2 (notice that different noises are almost orthogonal), and the last inequality is by Proposition C.2 and Lemma E.2. This completes the proof. □

Lemma E.4. *Under the same conditions as Theorem 4.4, we have that*

$$y_i\langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^*\rangle \geq q\log(2q/\epsilon)$$

for all $T_1 \leq t \leq T^*$.

Proof of Lemma E.4. Recall that $f(\mathbf{W}^{(t)}, \mathbf{x}_i) = (1/m)\sum_{j,r} j \cdot \big[\sigma(\langle \mathbf{w}_{j,r}, y_i \cdot \boldsymbol{\mu}\rangle) + \sigma(\langle \mathbf{w}_{j,r}, \boldsymbol{\xi}_i\rangle)\big]$, so we have

$$\begin{aligned}
&y_i\langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^*\rangle \\
&= \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\langle \boldsymbol{\mu}, j\mathbf{w}_{j,r}^*\rangle + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\langle y_i\boldsymbol{\xi}_i, j\mathbf{w}_{j,r}^*\rangle \\
&= \frac{1}{m}\sum_{j,r}\sum_{i'=1}^n \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle) 2qm\log(2q/\epsilon)\, \mathbb{1}(j = y_{i'}) \cdot \frac{\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle}{\|\boldsymbol{\xi}_{i'}\|_2} \\
&\quad + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\langle \boldsymbol{\mu}, j\mathbf{w}_{j,r}^{(0)}\rangle + \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\langle y_i\boldsymbol{\xi}_i, j\mathbf{w}_{j,r}^{(0)}\rangle \\
&\geq \frac{1}{m}\sum_{j=y_i,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle) 2qm\log(2q/\epsilon) - \frac{1}{m}\sum_{j,r}\sum_{i'\neq i} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle) 2qm\log(2q/\epsilon)\frac{|\langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_i\rangle|}{\|\boldsymbol{\xi}_{i'}\|_2} \\
&\quad - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2) - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\widetilde{O}(\sigma_0\sigma_p\sqrt{d}) \\
&\geq \frac{1}{m}\sum_{j=y_i,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle) 2qm\log(2q/\epsilon) - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\widetilde{O}(mnd^{-1/2}) \\
&\quad - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle)\widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2) - \frac{1}{m}\sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i\rangle)\widetilde{O}(\sigma_0\sigma_p\sqrt{d}),
\end{aligned} \tag{E.5}$$

where the first inequality is by Lemma B.3 and the last inequality is by Lemma B.2. Next we will bound the inner-product terms in (D.4) respectively. By Lemma C.6, we have that

$$|\langle \mathbf{w}_{j,r}^{(t)}, y_i\boldsymbol{\mu}\rangle| \leq |\langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle| + \gamma_{j,r}^{(t)} \leq \widetilde{O}(1), \tag{E.6}$$

where the last inequality is by Proposition C.2.

For $j = y_i$, we can bound the inner product between the parameter and the noise as follows

$$\max_{j,r} \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle \geq \max_{j,r} \big[\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + \overline{\rho}_{j,r,i}^{(t)} \big] - 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha \geq 1, \tag{E.7}$$

where the first inequality is by Lemma C.4, the second inequality is by Lemma E.2.

For $j = -y_i$, we can bound the inner product between the parameter and the noise as follows

$$\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha \leq 1, \tag{E.8}$$

where the first inequality is by Lemma C.5 and the last inequality is by Lemma B.3 and the conditions of σ_0 and d in Condition 4.2. Therefore, plugging (E.6), (E.7), (E.8) into (E.5) gives

$$\begin{aligned} y_i \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle &\geq 2q \log(2q/\epsilon)) - \widetilde{O}(mnd^{-1/2}) - \widetilde{O}(\sigma_0 \|\boldsymbol{\mu}\|_2) - \widetilde{O}(\sigma_0 \sigma_p \sqrt{d}) \\ &\geq q \log(2q/\epsilon), \end{aligned}$$

where the last inequality is by $d \geq \widetilde{\Omega}(m^2 n^4)$ and $\sigma_0 \leq \widetilde{O}(m^{-2/(q-2)} n^{-1}) \cdot \min\{(\sigma_p \sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\}$ in Condition 4.2. □

Lemma E.5. *Under the same conditions as Theorem 4.4, we have that*

$$\|\mathbf{W}^{(t)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(t+1)} - \mathbf{W}^*\|_F^2 \geq (2q-1)\eta L_S(\mathbf{W}^{(t)}) - \eta\epsilon$$

for all $T_1 \leq t \leq T^*$.

Proof of Lemma E.5. The proof is exactly same as the proof of Lemma D.4.

$$\begin{aligned} &\|\mathbf{W}^{(t)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(t+1)} - \mathbf{W}^*\|_F^2 \\ &\quad = 2\eta \langle \nabla L_S(\mathbf{W}^{(t)}), \mathbf{W}^{(t)} - \mathbf{W}^* \rangle - \eta^2 \|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \\ &\quad = \frac{2\eta}{n} \sum_{i=1}^n \ell_i'^{(t)} [q y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i) - \langle \nabla f(\mathbf{W}^{(t)}, \mathbf{x}_i), \mathbf{W}^* \rangle] - \eta^2 \|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \\ &\quad \geq \frac{2\eta}{n} \sum_{i=1}^n \ell_i'^{(t)} [q y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i) - q \log(6/\epsilon)] - \eta^2 \|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \\ &\quad \geq \frac{2q\eta}{n} \sum_{i=1}^n [\ell\big(y_i f(\mathbf{W}^{(t)}, \mathbf{x}_i)\big) - \epsilon/(2q)] - \eta^2 \|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \\ &\quad \geq (2q-1)\eta L_S(\mathbf{W}^{(t)}) - \eta\epsilon, \end{aligned}$$

where the first inequality is by Lemma E.4, the second inequality is due to the convexity of the cross entropy function and the last inequality is due to Lemma C.7. □

Lemma E.6. *Under the same conditions as Theorem 4.4, let $T = T_1 + \Big\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \Big\rfloor = T_1 + \widetilde{O}(\eta^{-1}\epsilon^{-1} m^3 n d^{-1} \sigma_p^{-2})$. Then we have $\max_{j,r} \gamma_{j,r}^{(t)} \leq 2\widehat{\beta}'$, $\max_{j,r,i} |\underline{\rho}_{j,r,i}^{(t)}| = \widetilde{O}(\sigma_0 \sigma_p \sqrt{d})$ for all $T_1 \leq t \leq T$. Besides,*

$$\frac{1}{t - T_1 + 1} \sum_{s=T_1}^t L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(t - T_1 + 1)} + \frac{\epsilon}{(2q-1)}$$

for all $T_1 \leq t \leq T$, and we can find an iterate with training loss smaller than ϵ within T iterations.

Proof of Lemma E.6. By Lemma E.5, for any $T_1 \leq t \leq T$, we obtain that

$$\|\mathbf{W}^{(s)} - \mathbf{W}^*\|_F^2 - \|\mathbf{W}^{(s+1)} - \mathbf{W}^*\|_F^2 \geq (2q-1)\eta L_S(\mathbf{W}^{(s)}) - \eta\epsilon \tag{E.9}$$

holds for $T_1 \leq s \leq t$. Taking a summation, we have that

$$\sum_{s=T_1}^{t} L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2 + \eta\epsilon(t - T_1 + 1)}{(2q-1)\eta}$$
$$\overset{(i)}{\leq} \frac{2\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta}$$
$$\overset{(ii)}{=} \widetilde{O}(\eta^{-1}m^3nd^{-1}\sigma_p^{-2}), \tag{E.10}$$

where (i) is by $t \leq T_2$ and (ii) is by Lemma E.3 Then we can use induction to prove that $\max_{j,r} \gamma_{j,r}^{(t)} \leq 2\widehat{\beta}'$ for all $t \in [T_1, T]$. Clearly, by the definition of $\widehat{\beta}'$, we have $\max_{j,r} \gamma_{j,r}^{(T_1)} \leq \widehat{\beta}' \leq 2\widehat{\beta}'$. Now suppose that there exists $\widetilde{T} \in [T_1, T]$ such that $\max_{j,r} \gamma_{j,r}^{(t)} \leq 2\widehat{\beta}'$ for all $t \in [T_1, \widetilde{T} - 1]$. Then by (C.3), we have

$$\gamma_{j,r}^{(\widetilde{T})} = \gamma_{j,r}^{(T_1)} - \frac{\eta}{nm} \sum_{s=T_1}^{\widetilde{T}-1} \sum_{i=1}^{n} \ell_i'^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \cdot \boldsymbol{\mu} \rangle) \|\boldsymbol{\mu}\|_2^2,$$
$$\overset{(i)}{\leq} \gamma_{j,r}^{(T_1)} + \frac{q3^{q-1}\eta}{nm} \|\boldsymbol{\mu}\|_2^2 \widehat{\beta}'^{q-1} \sum_{s=T_1}^{\widetilde{T}-1} \sum_{i=1}^{n} |\ell_i'^{(t)}|$$
$$\overset{(ii)}{\leq} \gamma_{j,r}^{(T_1)} + q3^{q-1}\eta m^{-1} \|\boldsymbol{\mu}\|_2^2 \widehat{\beta}'^{q-1} \sum_{s=T_1}^{\widetilde{T}-1} L_S(\mathbf{W}^{(s)})$$
$$\overset{(iii)}{\leq} \gamma_{j,r}^{(T_1)} + \widehat{\beta}'^{q-1} \widetilde{O}(m^2 n \mathrm{SNR}^2)$$
$$\overset{(iv)}{\leq} \gamma_{j,r}^{(T_1)} + \widehat{\beta}'^{q-1} \widetilde{O}(m^2 n^{1-2/q})$$
$$\overset{(v)}{\leq} 2\widehat{\beta}'$$

for all $j \in \{\pm 1\}$ and $r \in [m]$, where (i) is by induction hypothesis $\max_{j,r} \gamma_{j,r}^{(t)} \leq 2\widehat{\beta}'$, (ii) is by $|\ell'| \leq \ell$, (iii) is by (E.10), (iv) is by $n^{-1}\mathrm{SNR}^{-q} \geq \widetilde{\Omega}(1)$, and (v) is by $\widehat{\beta}' = \widetilde{O}(\sigma_0 \|\boldsymbol{\mu}\|_2)$ and $\widehat{\beta}'^{q-2} \widetilde{O}(m^2 n^{1-2/q}) = \widetilde{O}(m^2 n^{1-2/q} (\sigma_0 \|\boldsymbol{\mu}\|_2)^{q-2}) \leq 1$ by Condition 4.2. Therefore, we have $\max_{j,r} \gamma_{j,r}^{(\widetilde{T})} \leq 2\widehat{\beta}'$, which completes the induction. □

E.3 Population Loss

Lemma E.7 (4th statement of Theorem 4.4)**.** *Under the same conditions as Theorem 4.4, within* $\widetilde{O}(\eta^{-1} n \sigma_0^{2-q} \sigma_p^{-q} d^{-q/2} + \eta^{-1}\epsilon^{-1} m^3 n \sigma_p^{-2} d^{-1})$ *iterations, we can find* $\mathbf{W}^{(T)}$ *such that* $L_S(\mathbf{W}^{(T)}) \leq \epsilon$*. Besides, for any* $0 \leq t \leq T$ *we have that* $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \geq 0.1$.

Proof of Lemma E.7. Given a new example (x, y), we have that

$$\|\mathbf{w}_{j,r}^{(t)}\|_2 = \bigg\| \mathbf{w}_{j,r}^{(0)} + j \cdot \gamma_{j,r}^{(t)} \cdot \frac{\boldsymbol{\mu}}{\|\boldsymbol{\mu}\|_2^2} + \sum_{i=1}^{n} \overline{\rho}_{j,r,i}^{(t)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} + \sum_{i=1}^{n} \underline{\rho}_{j,r,i}^{(t)} \cdot \frac{\boldsymbol{\xi}_i}{\|\boldsymbol{\xi}_i\|_2^2} \bigg\|_2$$
$$\overset{(i)}{\leq} \|\mathbf{w}_{j,r}^{(0)}\|_2 + \frac{\gamma_{j,r}^{(t)}}{\|\boldsymbol{\mu}\|_2} + \sum_{i=1}^{n} \frac{\overline{\rho}_{j,r,i}^{(t)}}{\|\boldsymbol{\xi}_i\|_2} + \sum_{i=1}^{n} \frac{|\underline{\rho}_{j,r,i}^{(t)}|}{\|\boldsymbol{\xi}_i\|_2}$$
$$\overset{(ii)}{=} O(\sigma_0\sqrt{d}) + \widetilde{O}(n\sigma_p^{-1} d^{-1/2}),$$

where (i) is by triangle inequality and (ii) is by $\max_{j,r} \gamma_{j,r}^{(t)} = \widetilde{O}(\sigma_0 \|\boldsymbol{\mu}\|_2)$ in Lemma E.6 and $\max_{i,j,r} |\rho_{j,r,i}| \leq 4\log(T^*)$ in Proposition 5.3.

Therefore, we have that $\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi} \rangle \sim \mathcal{N}(0, \sigma_p^2 \|\mathbf{w}_{j,r}^{(t)}\|_2^2)$. So with probability $1 - 1/(4m)$,

$$|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi} \rangle| \leq \widetilde{O}(\sigma_0 \sigma_p \sqrt{d} + nd^{-1/2}).$$

Since the signal vector $\boldsymbol{\mu}$ is orthogonal to noises, by $\max_{j,r} \gamma_{j,r}^{(t)} \leq 2\widehat{\beta}' = \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2)$ in Lemma E.6, we also have that $|\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu}\rangle| \leq |\langle \mathbf{w}_{j,r}^{(0)}, y_i\boldsymbol{\mu}\rangle| + \gamma_{j,r}^{(t)} = \widetilde{O}(\sigma_0\|\boldsymbol{\mu}\|_2)$. Now by union bound, with probability at least $1 - 1/2$, we have that

$$
\begin{aligned}
F_j(\mathbf{W}_j^{(t)}, \mathbf{x}) &= \frac{1}{m}\sum_{r=1}^{m} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, y\boldsymbol{\mu}\rangle) + \frac{1}{m}\sum_{r=1}^{m} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle) \\
&\leq \max_r |\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu}\rangle|^q + \max_r |\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}\rangle|^q \\
&\leq \widetilde{O}(\sigma_0^q \sigma_p^q d^{q/2} + n^q d^{-q/2} + \sigma_0^q \|\boldsymbol{\mu}\|_2^2) \\
&\leq 1,
\end{aligned}
$$

where the last inequality is by $\sigma_0 \leq \widetilde{O}(n^{-1}m^{-2/(q-2)} \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\boldsymbol{\mu}\|_2^{-1}\})$ and $d \geq \widetilde{\Omega}(m^2 n^4)$ in Condition 4.2. Therefore, with probability at least $1 - 1/2$, we have that

$$
\ell\big(y \cdot f(\mathbf{W}^{(t)}, \mathbf{x})\big) \geq \log(1 + e^{-1}).
$$

Thus $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \geq \log(1 + e^{-1}) \cdot 0.5 \geq 0.1$. This completes the proof. □

F Experiments

We present simulations on synthetic data and experiments on real-world data to back up our theoretical analysis. The code and data for our experiments can be found on Github [‡].

Synthetic-data experiments. Here we generate synthetic data exactly following Definition 3.1. Specifically, we set dimension $d = 400$. Since the learning problem is rotation-invariant, without loss of generality, we set $\boldsymbol{\mu} /\!\!/ [1, 0, \ldots, 0]^\top$. We then generate the noise vector ξ in Definition 3.1 so that its first entry is zero, while the rest of the entries are standard Gaussian random vectors. In this way, to perform experiments under different SNRs, it suffices to change $\|\mu\|_2$, or equivalently, change the first entry of $\boldsymbol{\mu}$.

We train a two-layer CNN model defined in Section 3 with RELU^3 activation function. The number of filters is set as $m = 10$. We use the default initialization method in PyTorch to initialize the CNN parameters, and train the CNN with full-batch gradient descent with a learning rate of 0.01 for 50 epochs. We consider different training data sizes n ranging from 1 to 100, and different SNRs ranging from 0 to 1.0. Note that in all these training sample size and SNR settings, our training setup can guarantee a training loss smaller than 0.05. After training, we also calculate the test losses for each case using 100 test data points. The results are given in Figure 2.

It is clear that the results shown in Figure 2 match our theoretical results very well: the test loss values depend on both sample size and SNR, and a larger sample size or a higher SNR can both lead to smaller test losses. Moreover, Figure 2 shows a clear phase transition between benign and harmful overfitting, and is consistent with our illustration in Figure 1.

Real-data experiments. We further conduct real-world experiments on the MNIST data set (Deng, 2012), which consists of gray-scale hand-written digits of size 28×28. Note that for a given data set, we cannot accurately define the SNR, as it is not clear which part of the image is signal or noise. Therefore, in order to verify our theory, for each image, we first multiply each pixel in the image with a factor which we call "scaled SNR", and then add standard Gaussian random noises to the outer regions with a width of 5. In this way, we can roughly use the scaled SNR to represent the signal-to-noise ratio in the data. Two examples of the modified images with scaled SNRs 1 and 0.2 are given in Figure 3.

We train a simple ReLU CNN model which has two convolutional layers each followed by a max-pooling layer, and a fully-connected layer that gives the final output of the network. The first convolutional layer has 32 output channels (i.e., 32 different filters), with filter size 5, stride 1 and padding 0. The second convolutional layer has 64 output channels, with filter size 3, stride 1 and padding 0. Both max-pooling layers are of size 2 and stride 2.

[‡] https://github.com/uclaml/Benign-Overfitting-CNN

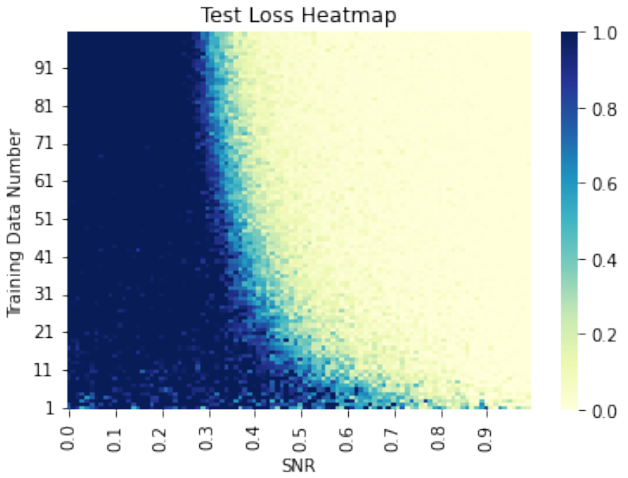

Figure 2: Heatmap of test losses on synthetic data under different training data sizes and SNRs. High test losses are marked in blue and low test losses are marked in yellow.

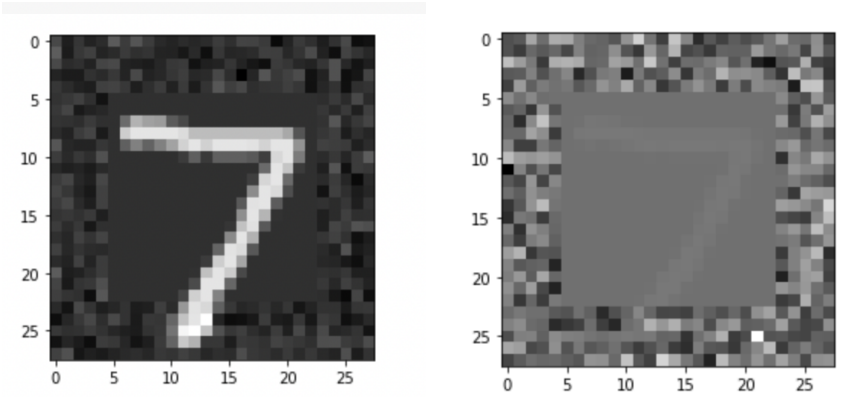

(a) Scaled SNR = 1 (b) Scaled SNR = 0.2

Figure 3: Illustration of modified MNIST images. (a) shows an example of a modified digit 7 with Scaled SNR = 1, and (b) shows the same image with Scaled SNR = 0.2.

We train the network with mini-batch stochastic gradient descent. We set the mini-batch size to be 128, and the learning rate to be 0.1. The network is trained for 2000 epochs in each setting. We consider different training data sizes n ranging from 3000 to 7500, and different scaled SNRs ranging from 0.01 to 0.2. Again, in all these training sample size and SNR settings, our training setup can guarantee a training loss smaller than 0.05. After training, we calculate the test losses for each case using the 10000 test data (modified in the same way as training data). The results are given in Figure 4.

Clearly, the results in Figure 4 also match our theoretical results very well, and show a phase transition between benign and harmful overfitting. Note that for this set of experiments, the setup does not satisfy our assumptions in many aspects:

- The CNN has two convolution layers and two max-pooling layers.
- The convolutions are with stride 1 and therefore the patches in the data have overlaps. This also implies that the noise patches in an image are not independent.
- The activation function is ReLU instead of ReLU^q with $q > 2$.
- The noise patches and signal patches are not orthogonal to each other, and are sometimes mixed together, as is shown in Figure 3.

Nevertheless, the experiment results still corroborate our theory to a certain extent. This demonstrates that our study of benign and harmful overfitting in CNNs captures the nature of real-world image classification problems.

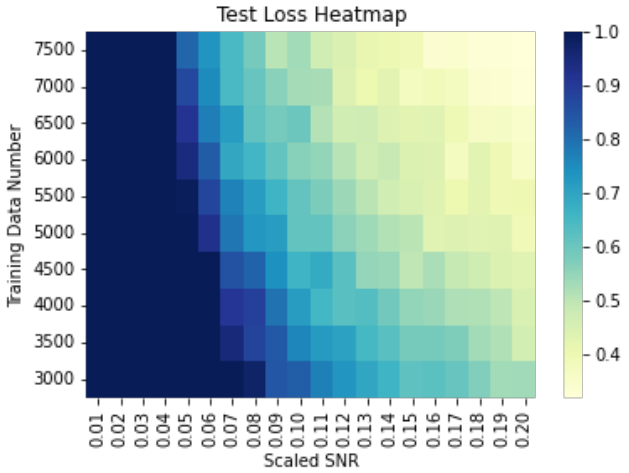

Figure 4: Heatmap of test losses on modified MNIST images under different training data sizes and scaled SNRs. High test losses are marked in blue and low test losses are marked in yellow.