
Graph Neural Networks Benefit from Structural Information Provably: A Feature Learning Perspective

Anonymous Author(s)

Affiliation

Address

email

Abstract

Graph neural networks (GNNs) have pioneered advancements in graph representation learning, exhibiting superior feature learning and performance over multilayer perceptrons (MLPs) when handling graph inputs. However, understanding the feature learning aspect of GNNs is still in its initial stage. This study aims to bridge this gap by investigating the role of graph convolution within the context of feature learning theory in neural networks using gradient descent training. We provide a distinct characterization of signal learning and noise memorization in two-layer graph convolutional networks (GCNs), contrasting them with two-layer convolutional neural networks (CNNs). Our findings reveal that graph convolution significantly augments the benign overfitting regime over the counterpart CNNs, where signal learning surpasses noise memorization, by approximately factor \sqrt{D}^{q-2} , with D denoting a node's expected degree and q being the power of the ReLU activation function where $q > 2$. These findings highlight a substantial discrepancy between GNNs and MLPs in terms of feature learning and generalization capacity after gradient descent training, a conclusion further substantiated by our empirical simulations.

1 Introduction

Graph neural networks (GNNs) have recently demonstrated remarkable capability in learning node or graph representations, yielding superior results across various downstream tasks, such as node classifications [1, 2, 3], graph classifications [4, 5, 6, 7] and link predictions [8, 9, 10], etc. However, the theoretical understanding of why GNNs can achieve such success is still in its infancy. Compared to multilayer perceptron (MLPs), GNNs enhance representation learning with an added message passing operation [11]. Take graph convolutional network (GCN) [1] as an example, it aggregates a node's attributes with those of its neighbors through a *graph convolution* operation. This operation, which leverages the structural information (adjacency matrix) of graph data, forms the core distinction between GNNs and MLPs. Empirical evidence from three node classification tasks, as shown in Figure 1, suggests GCNs outperform MLPs. Motivated by the superior performance of GNNs, we pose a critical question about graph convolution:

- *What role does graph convolution play during gradient descent training, and what mechanism enables a GCN to exhibit better generalization after training?*

Several recent studies have embarked on a theoretical exploration of graph convolution's role in GNNs. For instance, Baranwal et al. (2021) [12] considered a setting of linear classification of data generated from a contextual stochastic block model [13]. Their findings indicate that graph convolution extends the regime where data is linearly separable by a factor of approximately $1/\sqrt{D}$ compared to MLPs, with D denoting a node's expected degree. Baranwal et al. (2023) [14] further

investigated the impact of graph convolutions in multi-layer networks, showcasing improved linear separability. However, these examples, while insightful, operate within a linear neural network setting and do not account for non-linear activation, which significantly constrains the network’s capabilities. Additionally, these studies assume the maximum margin solution of GNNs, thereby losing a nuanced characterization of the GNNs’ optimization process. Consequently, there exists a notable gap between existing theoretical explorations and the detailed examination of GNNs incorporating non-linear activation, comprehensive characterization of optimization, and generalization ability.

To respond to the growing demand for a comprehensive theoretical understanding of graph convolution, we delve into the feature learning process [15, 16, 17] of graph neural networks. In our study, we introduce a data generation model—termed SNM-SBM—that combines a signal-noise model [15, 18] for feature creation and a stochastic block model [19] for graph construction. Our analysis is centered on the convergence and generalization attributes of two-layer graph convolution networks (GCNs) when trained via gradient descent, compared with the established outcomes for two-layer convolutional neural networks (CNNs) as presented by Cao et al. (2022) [15]. While both GCNs and CNNs demonstrate the potential to achieve near-zero training error, our study effectively sheds light on the discrepancies in their generalization abilities. We emphasize the crucial contribution of graph convolution to the enhanced performance of GNNs. Our study’s key contributions are as follows:




Figure 1: Performance (test accuracy) comparison between GCN and MLP on node classification tasks.

- We establish global convergence guarantees for graph neural networks training on data drawn from SNM-SBM model. We demonstrate that, despite the nonconvex optimization landscape, GCNs can achieve zero training error after a polynomial number of iterations.
- We further establish population loss bounds of overfitted GNN models trained by gradient descent. We show that under certain conditions on the signal-to-noise ratio, GNNs trained by gradient descent will prioritize learning the signal over memorizing the noise, and thus achieves small test losses.
- We delineate a marked contrast in the generalization capabilities of GCNs and CNNs following gradient descent training. We identify a specific regime where GCNs can attain nearly zero test error, whereas the performance of the model discovered by CNNs does not exceed random guessing. This conclusion is further substantiated by empirical verification.

2 Related Work

Role of Graph Convolution in GNNs. Enormous empirical studies of various GNNs models with graph convolution [20, 21, 22, 23, 24] have been demonstrating that graph convolutions can enhance the performance of traditional classification methods, such as a multi-layer perceptron (MLP). Towards theoretically understanding the role of graph convolution, Xu et al. (2020) [25] identify conditions under which MLPs and GNNs extrapolate, thereby highlighting the superiority of GNNs for extrapolation problems. Their theoretical analysis leveraged the concept of the over-parameterized networks and the neural tangent kernel [26]. Huang et al. (2021) [27] employed a similar approach to examine the role of graph convolution in deep GNNs within a node classification setting. They discovered that excessive graph convolution layers can hamper the optimization and generalization of GNNs, corroborating the well-known over-smoothing issue in deep GNNs [28]. Another pertinent work by Hou et al. (2022) [29] proposed two smoothness metrics to measure the quantity and quality of information derived from graph data, along with a novel attention-based framework. Some recent works [12, 14, 21] have demonstrated that graph convolution broadens the regime in which a multi-layer network can classify nodes, compared to methods that do not utilize the graph structure, especially when the graph is dense and exhibits homophily. Yang et al. (2022) [30] attributed the major performance gains of GNNs to their inherent generalization capability through graph neural tangent kernel (GNTK) and extrapolation analysis. As for neural network theory, these works either gleaned insights from GNTK [31, 27, 32] or studied the role of graph convolution within a linear neural network setting. Unlike them, our work extends beyond NTK and investigates a more realistic setting concerning the convergence and generalization of neural networks in terms of feature learning.

Feature learning in neural networks. This work builds upon a growing body of research on how neural networks learn features. Allen-Zhu et al. (2020) [18] formulated a theory illustrating that when data possess a “multi-view” feature, ensembles of independently trained neural networks can demonstrably improve test accuracy. Further, Allen-Zhu et al. (2022) [16] demonstrated that adversarial training can purge certain small dense mixtures from the hidden weights during the training process of a neural network, thus refining the hidden weights. Ba et al. (2022) [33] established that the initial gradient update contains a rank-1 ‘spike’, which leads to an alignment between the first-layer weights and the linear component feature of the teacher model. Cao et al. (2022) [15] investigated the benign overfitting phenomenon in training a two-layer convolutional neural network (CNN), illustrating that under certain conditions related to the signal-to-noise ratio, a two-layer CNN trained by gradient descent can achieve exceedingly low test loss through feature learning. Alongside related works [34, 35, 17, 36, 37, 38, 39, 40], all these studies have underscored the existence of feature learning in neural networks during gradient descent training, forming a critical line of inquiry that this work continues to explore. However, the neural tangent kernel (NTK) theory [41, 42, 43, 44], also known as “lazy training” [45], where the neural network function is approximately linear in its parameters, cannot demonstrate feature learning. Thus, the optimization and generalization analysis in our study extends beyond the NTK regime.

3 Problem Setup and Preliminary

3.1 Notations

We use lower bold-faced letters for vectors, upper bold-faced letters for matrices, and non-bold-faced letters for scalars. For a vector $\mathbf{v} = (v_1, v_2, \dots, v_d)^\top$, its ℓ_2 -norm is denoted as $\|\mathbf{v}\|_2 \triangleq \sqrt{\sum_{i=1}^d v_i^2}$. For a matrix $\mathbf{A} = (a_{ij}) \in \mathbb{R}^{m \times n}$, we use $\|\mathbf{A}\|_2$ to denote its spectral norm and $\|\mathbf{A}\|_F$ for its Frobenius norm. When comparing two sequences $\{a_n\}$ and $\{b_n\}$, we employ standard asymptotic notations such as $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$ to describe their limiting behavior. Specifically, we write $a_n = O(b_n)$ if there exists a positive real number C_1 and a positive integer N such that $|a_n| \leq C_1|b_n|$ for all $n \geq N$. Similarly, we write $a_n = \Omega(b_n)$ if there exists $C_2 > 0$ and $N > 0$ such that $|a_n| > C_2|b_n|$ for all $n \geq N$. We say $a_n = \Theta(b_n)$ if $a_n = O(b_n)$ and $a_n = \Omega(b_n)$. Besides, if $\lim_{n \rightarrow \infty} |a_n/b_n| = 0$, we express this as $a_n = o(b_n)$. We use $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$, and $\tilde{\Theta}(\cdot)$ to hide logarithmic factors in these notations respectively. Moreover, we denote $a_n = \text{poly}(b_n)$ if $a_n = O((b_n)^p)$ for some positive constant p and $a_n = \text{polylog}(b_n)$ if $a_n = \text{poly}(\log(b_n))$. Lastly, sequences of integers are denoted as $[n] = \{1, 2, \dots, n\}$ and $[m] = \{1, 2, \dots, m\}$.

3.2 Data model

In our approach, we utilize a signal-noise model for feature generation, combined with a stochastic block model for graph structure generation. Specifically, we define the feature matrix as $\mathbf{X} \in \mathbb{R}^{n \times 2d}$, with n representing the number of samples and $2d$ being the feature dimensionality. Each feature associated with a data point is generated from a *signal-noise model*, conditional on the Rademacher random variable $y \in -1, 1$, and a latent vector $\boldsymbol{\mu} \in \mathbb{R}^d$:

$$\mathbf{x} = [\mathbf{x}^{(1)}, \mathbf{x}^{(2)}] = [y \cdot \boldsymbol{\mu}, \boldsymbol{\xi}], \quad (1)$$

where $\mathbf{x}^{(1)}, \mathbf{x}^{(2)} \in \mathbb{R}^d$, and $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \sigma_p^2 \cdot (\mathbf{I} - \boldsymbol{\mu}\boldsymbol{\mu}^\top \cdot \|\boldsymbol{\mu}\|_2^{-2}))$ consists of independent standard normal entries with σ_p^2 as the variance. The term $\mathbf{I} - \boldsymbol{\mu}\boldsymbol{\mu}^\top \cdot \|\boldsymbol{\mu}\|_2^{-2}$ is employed to guarantee that the noise vector is orthogonal to the signal vector $\boldsymbol{\mu}$. It’s worth mentioning that a series of recent works [18, 15, 35, 46] have explored similar signal-noise models to illustrate the feature learning process and benign overfitting of neural networks.

Following this, we implement a stochastic block model with inter-class edge probability p and intra-class edge probability s , devoid of self-loops. Specifically, the adjacency matrix $\mathbf{A} = (a_{ij})_{n \times n}$ is Bernoulli distributed, with $a_{ij} \sim \text{Ber}(p)$ when $y_i = y_j$, and $a_{ij} \sim \text{Ber}(s)$ when $y_i = -y_j$. The combination of a stochastic block model with the signal-noise model (1) is represented as $\text{SNM} - \text{SBM}(n, p, s, \boldsymbol{\mu}, \sigma_p, d)$. Consequently, the raw feature and graph structure are generated as $(\mathbf{A}, \mathbf{X}, \mathbf{y}) \sim \text{SNM} - \text{SBM}(n, p, s, \boldsymbol{\mu}, \sigma_p, d)$, allowing the data model (1) used in MLP to be considered as a special case where $p = s = 0$.

3.3 Neural network model and training method

In this section, we present two distinct types of neural network models: a two-layer convolutional neural network (CNN), which falls under the category of a multilayer perceptron (MLP), and a Graph Convolutional Neural Network (GCN) [1].

CNN. We introduce a two-layer CNN model, denoted as f , which utilizes a non-linear activation function, $\sigma(\cdot)$. Specifically, we employ a polynomial ReLU activation function defined as $\sigma(z) = \max\{0, z\}^q$, where $q > 2$ is a hyperparameter. Note that the use of a polynomial ReLU activation function aligns with related studies [18, 16, 15, 35, 47] that investigate neural network feature learning. Mathematically, given the input data \mathbf{x} , the CNN’s output is represented as $f(\mathbf{W}, \mathbf{x}) = F_{+1}(\mathbf{W}_{+1}, \mathbf{x}) - F_{-1}(\mathbf{W}_{-1}, \mathbf{x})$, where $F_{+1}(\mathbf{W}_{+1}, \mathbf{x})$ and $F_{-1}(\mathbf{W}_{-1}, \mathbf{x})$ are defined as follows:

$$F_j(\mathbf{W}_j, \mathbf{x}) = \frac{1}{m} \sum_{r=1}^m \left[\sigma(\mathbf{w}_{j,r}^\top \mathbf{x}^{(1)}) + \sigma(\mathbf{w}_{j,r}^\top \mathbf{x}^{(2)}) \right], \quad (2)$$

where the second layer parameters are fixed as either $+1$ or -1 . We assume a poly-logarithmic network width in relation to the training sample size, i.e., $m = \text{polylog}(n)$, where m signifies the network’s width, and $\mathbf{w}_{j,r} \in \mathbb{R}^d$ refers to the weight of the first layer’s r -th neuron connected to the second layer’s j class. The symbol \mathbf{W} collectively represents the model’s weights. Moreover, each weight in the first layer is initialized from a random draw of a Gaussian random variable, $\mathbf{w}_{j,r} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \cdot \mathbf{I}_{d \times d})$ for all $r \in [m]$ and $j \in \{-1, 1\}$, with σ_0 regulating the initialization magnitude for the first layer’s weight.

Upon receiving training data $\mathcal{S} \triangleq \{\mathbf{x}_i, y_i\}_{i=1}^n$ drawn from SNM – SBM($n, p = 0, s = 0, \boldsymbol{\mu}, \sigma_p, d$), we aim to learn the network’s parameter \mathbf{W} by minimizing the empirical cross-entropy loss function:

$$L_S^{\text{CNN}}(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \ell(y_i \cdot f(\mathbf{W}, \mathbf{x}_i)), \quad (3)$$

where $\ell(y \cdot f(\mathbf{W}, \mathbf{x})) = \log(1 + \exp(-f(\mathbf{W}, \mathbf{x}) \cdot y))$. The update rule for the gradient descent used in the CNN is then given as:

$$\begin{aligned} \mathbf{w}_{j,r}^{(t+1)} &= \mathbf{w}_{j,r}^{(t)} - \eta \cdot \nabla_{\mathbf{w}_{j,r}} L_S^{\text{CNN}}(\mathbf{W}^{(t)}) \\ &= \mathbf{w}_{j,r}^{(t)} - \frac{\eta}{nm} \sum_{i=1}^n \ell'_i{}^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle) \cdot j y_i \boldsymbol{\xi}_i - \frac{\eta}{nm} \sum_{i=1}^n \ell'_i{}^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, y_i \boldsymbol{\mu} \rangle) \cdot j \boldsymbol{\mu}, \end{aligned} \quad (4)$$

where we define the loss derivative as $\ell'_i \triangleq \ell'(y_i \cdot f_i) = -\frac{\exp(-y_i \cdot f_i)}{1 + \exp(-y_i \cdot f_i)}$. It’s important to clarify that the model we use for the MLP part is a two-layer CNN network. We categorize it as an MLP for comparison purposes with the graph neural network.

GCN. Graph neural network (GNNs) fuse graph structure information and node features to learn representation of nodes. Consider a two-layer GCN f with graph convolution operation on the first layer. The output of the GCN is given by $f(\mathbf{W}, \tilde{\mathbf{x}}) = F_{+1}(\mathbf{W}_{+1}, \tilde{\mathbf{x}}) - F_{-1}(\mathbf{W}_{-1}, \tilde{\mathbf{x}})$, where $F_{+1}(\mathbf{W}_{+1}, \tilde{\mathbf{x}})$ and $F_{-1}(\mathbf{W}_{-1}, \tilde{\mathbf{x}})$ are defined as follows:

$$F_j(\mathbf{W}_j, \tilde{\mathbf{x}}) = \frac{1}{m} \sum_{r=1}^m \left[\sigma(\mathbf{w}_{j,r}^\top \tilde{\mathbf{x}}^{(1)}) + \sigma(\mathbf{w}_{j,r}^\top \tilde{\mathbf{x}}^{(2)}) \right]. \quad (5)$$

Here, $\tilde{\mathbf{X}} \triangleq [\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_n]^\top = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{X} \in \mathbb{R}^{n \times 2d}$ with $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}_n$ representing the adjacency matrix with self-loop, and $\tilde{\mathbf{D}}$ is a diagonal matrix that records the degree of each node, namely, $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$. For simplicity we denote $\tilde{D}_i \triangleq \tilde{D}_{ii}$. Therefore, in contrast to the CNN model (2), the GCNs (5) incorporate the normalized adjacency matrix $\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, also termed as graph convolution, which serves as a pivotal component.

With the training data $\mathcal{S} \triangleq \{\mathbf{x}_i, y_i\}_{i=1}^n$ and $\mathbf{A} \in \mathbb{R}^{n \times n}$ drawn from SNM – SBM($n, p, s, \boldsymbol{\mu}, \sigma_p, d$), we consider to learn the network’s parameter \mathbf{W} by optimizing the empirical cross-entropy loss function:

$$L_S^{\text{GCN}}(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \ell(y_i \cdot f(\mathbf{W}, \tilde{\mathbf{x}}_i)). \quad (6)$$

176 The gradient descent update for the first layer weight \mathbf{W} in GCN can be expressed as:

$$\begin{aligned}\mathbf{w}_{j,r}^{(t+1)} &= \mathbf{w}_{j,r}^{(t)} - \eta \cdot \nabla_{\mathbf{w}_{j,r}} L_S^{\text{GCN}}(\mathbf{W}^{(t)}) \\ &= \mathbf{w}_{j,r}^{(t)} - \frac{\eta}{nm} \sum_{i=1}^n \ell_i'^{(t)} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle) \cdot j y_i \tilde{\boldsymbol{\xi}}_i - \frac{\eta}{nm} \sum_{i=1}^n \ell_i'^{(t)} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) \cdot j \tilde{y}_i \boldsymbol{\mu},\end{aligned}\quad (7)$$

177 where we define “aggregated label” $\tilde{y}_i = D_i^{-1} \sum_{k \in \mathcal{N}(i)} y_k$ and “aggregated noise vector” $\tilde{\boldsymbol{\xi}}_i =$
178 $D_i^{-1} \sum_{k \in \mathcal{N}(i)} \boldsymbol{\xi}_k$, with $\mathcal{N}(i)$ being a set that contains all the neighbor of node i . In this study, our
179 primary objective is to demonstrate the enhanced feature learning capabilities of GNNs in comparison
180 to CNNs. This is achieved by examining the generalization ability of the GNN model through the
181 lens of population loss, which can be formulated as follows:

$$L_D^{\text{GCN}}(\mathbf{W}) = \mathbb{E}_{\mathbf{x}, y \sim \mathcal{D}=\text{SNM-SBM}} \ell(y \cdot f(\mathbf{W}, \tilde{\mathbf{x}})). \quad (8)$$

182 4 Theoretical Results

183 In this section, we introduce our key theoretical findings that elucidate the optimization and general-
184 ization processes of feature learning in GCNs. Through the application of the gradient descent rule
185 outlined in Equation (7), we observe that the gradient descent iterate $\mathbf{w}_{j,r}^{(t)}$ is a linear combination of
186 its random initialization $\mathbf{w}_r(0)$, the signal vector $\boldsymbol{\mu}$ and the noise vectors in the training data $\boldsymbol{\xi}_i$ ¹ for
187 $i \in [n]$. Consequently, for $r \in [m]$, the decomposition of weight vector iteration can be expressed:

$$\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + \gamma_{j,r}^{(t)} \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu} + \sum_{i=1}^n \rho_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i. \quad (9)$$

188 where $\gamma_{j,r}^{(t)}$ and $\rho_{j,r,i}^{(t)}$ serve as coefficients. We refer to Equation (9) as the signal-noise decomposition
189 of $\mathbf{w}_{j,r}^{(t)}$. The normalization factors $\|\boldsymbol{\mu}\|_2^{-2}$ and $\|\boldsymbol{\xi}_i\|_2^{-2}$ are introduced to ensure that $\gamma_{j,r}^{(t)} \approx \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\mu} \rangle$,
190 and $\rho_{j,r,i}^{(t)} \approx \langle \mathbf{w}_{j,r}^{(t)}, \boldsymbol{\xi}_i \rangle$. We employ $\gamma_{j,r}^{(t)}$ to characterize the process of signal learning and $\rho_{j,r,i}^{(t)}$ to
191 characterize the noisy represent. From an intuitive standpoint, if, for some iteration certain $\gamma_{j,r}^{(t)}$ values
192 are sufficiently large while all $|\rho_{j,r,i}^{(t)}|$ are relatively small, this indicates that the neural network is
193 primarily learning the label through feature learning. This scenario can lead to *benign overfitting*,
194 characterized by both minimal training and test errors. Conversely, if some $|\rho_{j,r,i}^{(t)}|$ values are relatively
195 large while all $\gamma_{j,r}^{(t)}$ are small, the neural network will achieve a low training loss but a high test loss.
196 This occurs as the neural network attempts to generalize by memorizing noise, resulting in a *harmful*
197 *overfitting* regime.

198 To facilitate a fine-grained analysis for the evolution of coefficients, we introduce the notations
199 $\bar{\rho}_{j,r,i}^{(t)} \triangleq \rho_{j,r,i}^{(t)} \mathbb{1}(\rho_{j,r,i}^{(t)} \geq 0)$, $\underline{\rho}_{j,r,i}^{(t)} \triangleq \rho_{j,r,i}^{(t)} \mathbb{1}(\rho_{j,r,i}^{(t)} \leq 0)$. Consequently, we further express the vector
200 weight decomposition (9) as:

$$\mathbf{w}_{j,r}^{(t)} = \mathbf{w}_{j,r}^{(0)} + j \cdot \gamma_{j,r}^{(t)} \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu} + \sum_{i=1}^n \bar{\rho}_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i + \sum_{i=1}^n \underline{\rho}_{j,r,i}^{(t)} \cdot \|\boldsymbol{\xi}_i\|_2^{-2} \cdot \boldsymbol{\xi}_i. \quad (10)$$

201 Our analysis will be made under the following set of assumptions:

202 **Assumption 4.1.** Suppose that

- 203 1. The dimension d is sufficiently large: $d = \tilde{\Omega}(m^{2 \vee [4/(q-2)]} n^{4 \vee [(2q-2)/(q-2)]})$.
- 204 2. The size of training sample n and width of GCNs m adhere to $n, m = \Omega(\text{polylog}(d))$.
- 205 3. The learning rate η satisfies $\eta \leq \tilde{O}(\min\{\|\boldsymbol{\mu}\|_2^{-2}, \sigma_p^{-2} d^{-1}\})$.

¹By referring to Equation (7), we assert that the gradient descent update moves in the direction of $\tilde{\boldsymbol{\xi}}_i$ for each $i \in [n]$. Then we can apply the definition of $\tilde{\boldsymbol{\xi}}_i = D_i^{-1} \sum_{k \in \mathcal{N}(i)} \boldsymbol{\xi}_k$.

- 206 4. The edge probability $p, s = \Omega(\sqrt{\log(nd)/n})$ and $\Xi \triangleq \frac{p-s}{p+s}$ is a positive constant.
- 207 5. The standard deviation of Gaussian initialization σ_0 is chosen such that $\sigma_0 \leq$
 208 $\tilde{O}(m^{-2/(q-2)}n^{-[1/(q-2)]\vee 1} \cdot \min\{(\sigma_p\sqrt{d/(n(p+s)))}^{-1}, \Xi^{-1}\|\boldsymbol{\mu}\|_2^{-1}\})$.

209 *Remark 4.2.* (1) The requirement for the dimension d ensures that the learning process operates in
 210 a suitably over-parameterized environment [48, 15] when the second layer remains fixed. (2) It's
 211 necessary for the sample size and neural network width to be at least polylogarithmic in the dimension
 212 d . This condition ensures certain statistical properties of the training data and weight initialization
 213 hold with a probability of at least $1 - d^{-1}$. (3) The condition on η is to ensure that gradient descent
 214 can effectively minimize the training loss. (4) The assumption regarding edge probability guarantees
 215 a sufficient level of concentration in the degree and an adequate display of homophily of graph data.
 216 (5) Lastly, the conditions imposed on initialization strength σ_0 are intended to guarantee that the
 217 training loss can effectively converge to a sufficiently small value and to discern the differential
 218 learning speed between signal and noise.

219 Given the above assumptions, we present our main result on feature learning of GCNs in the following
 220 theorem.

221 **Theorem 4.3.** Suppose $\epsilon > 0$, and let $T = \tilde{\Theta}(\eta^{-1}m\sigma_0^{-(q-2)}\Xi^{-q}\|\boldsymbol{\mu}\|_2^{-q} + \eta^{-1}\epsilon^{-1}m^3\|\boldsymbol{\mu}\|_2^{-2})$.
 222 Under Assumption 4.1, if $n \cdot \text{SNR}^q \cdot \sqrt{n(p+s)}^{q-2} = \tilde{\Omega}(1)$, where $\text{SNR} \triangleq \|\boldsymbol{\mu}\|_2/(\sigma_p\sqrt{d})$ is the
 223 signal-to-noise ratio, then with probability at least $1 - d^{-1}$, there exists a $0 \leq t \leq T$ such that:

- 224 • The GCN learns the signal: $\max_r \gamma_{j,r}^{(t)} = \Omega(1)$ for $j \in \{\pm 1\}$.
- 225 • The GCN does not memorize the noises in the training data: $\max_{j,r,i} |\rho_{j,r,i}^{(T)}| =$
 226 $\tilde{O}(\sigma_0\sigma_p\sqrt{d/n(p+s)})$.
- 227 • The training loss converges to ϵ , i.e., $L_S^{\text{GCN}}(\mathbf{W}^{(t)}) \leq \epsilon$.
- 228 • The trained GCN achieves a small test loss: $L_D^{\text{GCN}}(\mathbf{W}^{(t)}) \leq c_1\epsilon + \exp(-c_2n^2)$.

229 where c_1 and c_2 are positive constants.

230 Theorem 4.3 outlines the scenario of *benign overfitting* for GCNs. It reveals that, provided $n \cdot$
 231 $\text{SNR}^q \cdot \sqrt{n(p+s)}^{q-2} = \tilde{\Omega}(1)$, the GCN can learn the signal by achieving $\max_r \gamma_{j,r}^{(t)} = \Omega(1)$
 232 for $j \in \{\pm 1\}$, and on the other hand, the noise memorization during gradient descent training
 233 is suppressed by $\max_{j,r,i} |\rho_{j,r,i}^{(T)}| = \tilde{O}(\sigma_0\sigma_p\sqrt{d/n(p+s)})$, given that $\sigma_0\sigma_p\sqrt{d/n(p+s)} \ll 1$
 234 according to assumption 4.1. Because the signal learned by the network is large enough and much
 235 stronger than the noise memory, it can perfectly predict the label in the test sample according to the
 236 learned signal when it generalizes. Consequently, the learned neural network can attain small training
 237 and test losses.

238 To illustrate the pronounced divergence between GNN and CNN in terms of generalization capa-
 239 bility post-gradient descent training, we show that, under identical conditions, a GCN engages in
 240 signal learning while a CNN emphasizes noise memorization, and thus diverges in the ability of
 241 generalization:

242 **Corollary 4.4 (Informal).** Under assumption 4.1, if $n \cdot \text{SNR}^q \cdot \sqrt{n(p+s)}^{q-2} = \tilde{\Omega}(1)$ and $n^{-1} \cdot$
 243 $\text{SNR}^{-q} = \tilde{\Omega}(1)$, then with probability at least $1 - d^{-1}$, then there exists a t such that:

- 244 • The trained GNN achieves a small test loss: $L_D^{\text{GCN}}(\mathbf{W}^{(t)}) \leq c_1\epsilon + \exp(-c_2n^2)$.
- 245 • The trained CNN has a constant order test loss: $L_D^{\text{CNN}}(\mathbf{W}^{(t)}) = \Theta(1)$.

246 Corollary 4.4 clearly provides a condition that GNNs learn the signal and achieves a small test loss
 247 while the CNNs can only memorize noises and will have a $\Theta(1)$ test loss. The CNN results are derived
 248 from the work of Cao et al. (2022) [15]. The improvement in *benign overfitting* regime is facilitated
 249 by graph convolution, a process that will be elaborated on in the subsequent section. Through the
 250 precise characterization of neural network feature learning from optimization to generalization, we
 251 have successfully demonstrated that the graph neural network can gain superiority with the help of
 252 graph convolution.

5 Proof Sketches

In this section, we discuss the primary challenges encountered during the study of GNN training, and illustrate the key techniques we employed in our proofs to overcome these challenges:

- Graph convolution aggregates information from neighboring nodes to the central node, which often leads to the loss of statistical stability for the aggregated noise vectors and labels. To overcome this challenge, we utilize a dense graph input, achieved by setting the edge probability as per Assumption 4.1.
- Graph convolution can potentially cause erratic iterative dynamics of coefficients during the feature learning process. To mitigate this issue, we introduce the concept of homophily into the graph input, which helps in stabilizing the coefficient iterations.
- Lastly, for the generalization analysis, depicting the generalization ability of graph neural networks poses a significant challenge. To address this issue, we introduce an expectation over the distribution for a single data point and develop an algorithm-dependent test error analysis.

These main techniques are further elaborated upon in the following sections, and detailed proofs for all the results can be found in the appendix.

5.1 Iterative analysis of the signal-noise decomposition under graph convolution

To analyze the feature learning process of graph neural networks during gradient descent training, we introduce an iterative methodology, based on the signal-noise decomposition in decomposition (10) and gradient descent update (7). The following lemma offers us a means to monitor the iteration of the signal learning and noise memorization under graph convolution:

Lemma 5.1. *The coefficients $\gamma_{j,r}^{(t)}$, $\bar{\rho}_{j,r,i}^{(t)}$, $\underline{\rho}_{j,r,i}^{(t)}$ in decomposition (10) adhere to the following equations:*

$$\gamma_{j,r}^{(0)}, \bar{\rho}_{j,r,i}^{(0)}, \underline{\rho}_{j,r,i}^{(0)} = 0, \quad (11)$$

$$\gamma_{j,r}^{(t+1)} = \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^n \ell'_i{}^{(t)} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu}_i \rangle) y_i \tilde{y}_i \|\boldsymbol{\mu}\|_2^2, \quad (12)$$

$$\bar{\rho}_{j,r,i}^{(t+1)} = \bar{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell'_k{}^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = j), \quad (13)$$

$$\underline{\rho}_{j,r,i}^{(t+1)} = \underline{\rho}_{j,r,i}^{(t)} + \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell'_k{}^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = -j). \quad (14)$$

Lemma 5.1 simplifies the analysis of the feature learning in GCNs by reducing it to the examination of the discrete dynamical system expressed by equations (11)-(14). Our proof strategy emphasizes an in-depth evaluation of the coefficient values $\gamma_{j,r}^{(t)}$, $\bar{\rho}_{j,r,i}^{(t)}$, $\underline{\rho}_{j,r,i}^{(t)}$ throughout the training. We present the following bounds of the coefficients and loss derivative, which persist throughout the training period:

Proposition 5.2. *Under Assumption 4.1, 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 $t \in [0, T^*]$:*

$$0 \leq \gamma_{j,r}^{(t)}, \bar{\rho}_{j,r,i}^{(t)} \leq 4 \log(T^*) \text{ for all } j \in \{\pm 1\}, r \in [m] \text{ and } i \in [n].$$

$$0 \geq \underline{\rho}_{j,r,i}^{(t)} \geq -4 \log(T^*) \text{ for all } j \in \{\pm 1\}, r \in [m] \text{ and } i \in [n].$$

$$\|\nabla L_S^{\text{GCN}}(\mathbf{W}^{(t)})\|_F^2 \leq O\left(\max\{\Xi^2 \|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d / (n(p+s))\}\right) L_S^{\text{GCN}}(\mathbf{W}^{(t)}).$$

The proof of Proposition 5.2 is provided in Appendix B.2. As suggested by Proposition 5.2, both the coefficients related to signal learning and noise memorization can reach a logarithmic order relative to training time T^* . Furthermore, the training objective function $L_S^{\text{GCN}}(\mathbf{W})$ maintains dominance over the gradient norm $\|\nabla L_S^{\text{GCN}}(\mathbf{W}^{(t)})\|_F$ along the gradient descent trajectory. This observation sets the preliminary for our convergence analysis. We then propose a two-stage dynamics analysis to elucidate the behavior of these coefficients. Subsequently, we can depict the generalization ability of GCN with the learned weight.

5.2 A two-phase dynamics analysis

Stage 1. 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: $\ell'_i{}^{(0)} = \ell'[y_i \cdot f(\mathbf{W}^{(0)}, \tilde{\mathbf{x}}_i)] = \Theta(1)$ for all $i \in [n]$. This is guaranteed under Condition 4.1 on σ_0 . Motivated by this, the dynamics of the coefficients in (12) - (14) 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 5.3. *Under the same conditions as Theorem 4.3, there exists $T_1 = \tilde{O}(\eta^{-1} m \sigma_0^{2-q} \Xi^{-q} \|\boldsymbol{\mu}\|_2^{-q})$ such that*

- $\max_r \gamma_{j,r}^{(T_1)} = \Omega(1)$ for $j \in \{\pm 1\}$.
- $|\rho_{j,r,i}^{(t)}| = O\left(\sigma_0 \sigma_p \sqrt{d} / \sqrt{n(p+s)}\right)$ for all $j \in \{\pm 1\}$, $r \in [m]$, $i \in [n]$ and $0 \leq t \leq T_1$.

The proof can be found in Appendix C.1. Lemma 5.3 leverages the period of training when the derivatives of the loss function are of a constant order. It's important to note that graph convolution plays a significant role in diverging the learning speed between signal learning and noise memorization in this first stage. Originally, the learning speeds are roughly determined by $\|\boldsymbol{\mu}\|_2$ and $\|\boldsymbol{\xi}\|_2$ respectively. However, after applying graph convolution, the learning speeds are approximately determined by $|\tilde{y}| \cdot \|\boldsymbol{\mu}\|_2$ and $\|\boldsymbol{\xi}\|_2$ respectively. Here, $|\tilde{y}| \cdot \|\boldsymbol{\mu}\|_2$ is close to $\|\boldsymbol{\mu}\|_2$, but $\|\boldsymbol{\xi}\|_2$ is much smaller than $\|\boldsymbol{\xi}\|_2$. This means that graph convolution can slow down the learning speed of noise memorization, thus enabling GNNs to focus more on signal learning in the initial training stage.

Stage 2. Building on the results from the first stage, we then move to the second stage of the training process. In this stage, the loss derivatives are no longer constant, and we demonstrate that the training loss can be minimized to an arbitrarily small amount. Importantly, the scale differences established during the first stage of learning continue to be maintained throughout the training process:

Lemma 5.4. *Let T, T_1 be defined in Theorem 4.3 and Lemma 5.3 respectively and \mathbf{W}^* be the collection of GCN parameters $\mathbf{w}_{j,r}^* = \mathbf{w}_{j,r}^{(0)} + 2qm \log(2q/\epsilon) \cdot j \cdot \|\boldsymbol{\mu}\|_2^{-2} \cdot \boldsymbol{\mu}$. Then under the same conditions as Theorem 4.3, for any $t \in [T_1, T]$, it holds that:*

- $\max_r \gamma_{j,r}^{(T_1)} \geq 2, \forall j \in \{\pm 1\}$ and $|\rho_{j,r,i}^{(t)}| \leq \sigma_0 \sigma_p \sqrt{d/(n(p+s))}$ for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$.
- $\frac{1}{t-T_1+1} \sum_{s=T_1}^t L_S^{\text{GCN}}(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta(t-T_1+1)} + \frac{\epsilon}{(2q-1)}.$

Here we denote $\|\mathbf{W}\|_F \triangleq \sqrt{\|\mathbf{W}_{+1}\|_F^2 + \|\mathbf{W}_{-1}\|_F^2}$.

Lemma 5.4 presents two primary outcomes related to feature learning. Firstly, throughout this training phase, it ensures that the coefficients of noise vectors, denoted as $\rho_{j,r,i}^{(t)}$, retain a significantly small value while coefficients of feature vector, denoted as $\gamma_{j,r}^{(t)}$ can achieve large value. Furthermore, it offers an optimization-oriented outcome, indicating that the optimal iterate within the interval $[T_1, T]$. In this process, graph convolution and gradient descent will continue to maintain the speed gap between signal learning and noise memory, and when the time is large enough, the training loss will tend to receive an arbitrarily small value.

5.3 Test error analysis

Finally, we consider a new data point (\mathbf{x}, y) drawn from the distribution SNM-SBM. The lemma below further gives an upper bound on the test loss of GNNs post-training:

Lemma 5.5. *Let T be defined in Theorem 4.3. Under the same conditions as Theorem 4.3, for any $t \leq T$ with $L_S^{\text{GCN}}(\mathbf{W}^{(t)}) \leq 1$, it holds that $L_D^{\text{GCN}}(\mathbf{W}^{(t)}) \leq c_1 \cdot L_S^{\text{GCN}}(\mathbf{W}^{(t)}) + \exp(-c_2 n^2)$.*

The proof is presented in the appendix. Lemma 5.5 demonstrates that GNNs achieve *benign overfitting* and completes the last step of feature learning theory.

6 Experiments

In this section, we validate our theoretical findings through numerical simulations using synthetic data, specifically generated according to the SNM-SBM model. We set the signal vector, μ , to drawn from a standard normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. The noise vector, ξ , is drawn from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma_p^2 \mathbf{I})$. We train a two-layer CNN defined as per equation (2) and a two-layer GNN as per equation (5) with polynomial ReLU $q = 3$.

Feature learning dynamics. Firstly, we display the training loss, test loss, training accuracy, and test accuracy for both the CNN and GNN in Figure 2. In this case, we further set the training data size to $n = 250$, input dimension to $d = 500$, noise strength to $\sigma_p = 20$, and edge probability to $p = 0.5$, $s = 0.08$. We observe that both the GNN and CNN can achieve zero training error. However, while the GNN obtains nearly zero test error, the CNN fails to generalize effectively to the test set. This simulation result serves to validate our theoretical results in Theorem 4.3 and Corollary 4.4.




Figure 2: Training loss, testing loss, training accuracy, and testing accuracy for both CNN and GNN over a span of 100 training epochs.

Phase diagram. We then explore a range of Signal-to-Noise Ratios (SNRs) from 0.045 to 0.98, and a variety of sample sizes, n , ranging from 200 to 7200. Based on our results, we train the neural network for 200 steps for each combination of SNR and sample size n . After training, we calculate the test accuracy for each run. The results are presented as a heatmap in Figure 3. Compared to CNNs, GCNs demonstrate a perfect accuracy score of 1 across a more extensive range in the SNR and n plane, indicating that GNNs have a broader *benign overfitting* regime. This further validates our theoretical findings.




Figure 3: Test accuracy heatmap for CNNs and GCNs after training.

7 Conclusion and Limitations

This paper utilizes a signal-noise decomposition to study the signal learning and noise memorization process in training a two-layer GCN. We provide specific conditions under which a GNN will primarily concentrate on signal learning, thereby achieving low training and testing errors. Our results theoretically demonstrate that GCNs, by leveraging structural information, outperform CNNs in terms of generalization ability across a broader benign overfitting regime. As a pioneering work that studies feature learning of GNNs, our theoretical framework is constrained to examining the role of graph convolution within a specific two-layer GCN and a certain data generalization model. In fact, the feature learning of a neural network can be influenced by a myriad of other factors, such as activation function, optimization algorithm, and data model [47, 35, 37]. Future work can extend our framework to consider the influence of a wider array of factors on feature learning within GCNs.

References

- [1] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [2] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [3] Will Hamilton, Zhitaoy Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017.
- [4] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- [5] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- [6] Junhyun Lee, Inyeop Lee, and Jaewoo Kang. Self-attention graph pooling. *arXiv preprint arXiv:1904.08082*, 2019.
- [7] Hao Yuan and S. Ji. Structpool: Structured graph pooling via conditional random fields. In *ICLR*, 2020.
- [8] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- [9] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. *Advances in neural information processing systems*, 31, 2018.
- [10] Ajay Kumar, Shashank Sheshar Singh, Kuldeep Singh, and Bhaskar Biswas. Link prediction techniques, applications, and performance: A survey. *Physica A: Statistical Mechanics and its Applications*, 553:124289, 2020.
- [11] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI open*, 1:57–81, 2020.
- [12] Aseem Baranwal, Kimon Fountoulakis, and Aukosh Jagannath. Graph convolution for semi-supervised classification: Improved linear separability and out-of-distribution generalization. *arXiv preprint arXiv:2102.06966*, 2021.
- [13] Chen Lu and Subhabrata Sen. Contextual stochastic block model: Sharp thresholds and contiguity. *arXiv preprint arXiv:2011.09841*, 2020.
- [14] Aseem Baranwal, Kimon Fountoulakis, and Aukosh Jagannath. Effects of graph convolutions in multi-layer networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- [15] Yuan Cao, Zixiang Chen, Mikhail Belkin, and Quanquan Gu. Benign overfitting in two-layer convolutional neural networks. *arXiv preprint arXiv:2202.06526*, 2022.
- [16] Zeyuan Allen-Zhu and Yuanzhi Li. Feature purification: How adversarial training performs robust deep learning. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 977–988. IEEE, 2022.
- [17] Zixin Wen and Yuanzhi Li. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pages 11112–11122. PMLR, 2021.
- [18] Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*, 2020.
- [19] Emmanuel Abbe, Afonso S Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Transactions on information theory*, 62(1):471–487, 2015.

- [20] Zhengdao Chen, Xiang Li, and Joan Bruna. Supervised community detection with line graph neural networks. *arXiv preprint arXiv:1705.08415*, 2017.
- [21] Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. Is homophily a necessity for graph neural networks? *arXiv preprint arXiv:2106.06134*, 2021.
- [22] Si Zhang, Hanghang Tong, Jiejun Xu, and Ross Maciejewski. Graph convolutional networks: a comprehensive review. *Computational Social Networks*, 6(1):1–23, 2019.
- [23] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 639–648, 2020.
- [24] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019.
- [25] Keyulu Xu, Mozhi Zhang, Jingling Li, Simon S Du, Ken-ichi Kawarabayashi, and Stefanie Jegelka. How neural networks extrapolate: From feedforward to graph neural networks. *arXiv preprint arXiv:2009.11848*, 2020.
- [26] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [27] Wei Huang, Yayong Li, Weitao Du, Richard Yi Da Xu, Jie Yin, Ling Chen, and Miao Zhang. Towards deepening graph neural networks: A gntk-based optimization perspective. *arXiv preprint arXiv:2103.03113*, 2021.
- [28] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI conference on artificial intelligence*, 2018.
- [29] Yifan Hou, Jian Zhang, James Cheng, Kaili Ma, Richard TB Ma, Hongzhi Chen, and Ming-Chang Yang. Measuring and improving the use of graph information in graph neural networks. *arXiv preprint arXiv:2206.13170*, 2022.
- [30] Chenxiao Yang, Qitian Wu, Jiahua Wang, and Junchi Yan. Graph neural networks are inherently good generalizers: Insights by bridging gnns and mlps. *arXiv preprint arXiv:2212.09034*, 2022.
- [31] Simon S Du, Kangcheng Hou, Russ R Salakhutdinov, Barnabas Poczos, Ruosong Wang, and Keyulu Xu. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *Advances in neural information processing systems*, 32, 2019.
- [32] Mahalakshmi Sabanayagam, Pascal Esser, and Debarghya Ghoshdastidar. Representation power of graph convolutions: Neural tangent kernel analysis. *arXiv preprint arXiv:2210.09809*, 2022.
- [33] Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. *arXiv preprint arXiv:2205.01445*, 2022.
- [34] Greg Yang and Edward J Hu. Feature learning in infinite-width neural networks. *arXiv preprint arXiv:2011.14522*, 2020.
- [35] Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. Understanding the generalization of adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371*, 2021.
- [36] Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.
- [37] Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. The benefits of mixup for feature learning. *arXiv preprint arXiv:2303.08433*, 2023.

- 457 [38] Yongqiang Chen, Wei Huang, Kaiwen Zhou, Yatao Bian, Bo Han, and James Cheng. To-
 458 wards understanding feature learning in out-of-distribution generalization. *arXiv preprint*
 459 *arXiv:2304.11327*, 2023.
- 460 [39] Xuran Meng, Yuan Cao, and Difan Zou. Per-example gradient regularization improves learning
 461 signals from noisy data. *arXiv preprint arXiv:2303.17940*, 2023.
- 462 [40] Samy Jelassi, Michael Sander, and Yuanzhi Li. Vision transformers provably learn spatial
 463 structure. *Advances in Neural Information Processing Systems*, 35:37822–37836, 2022.
- 464 [41] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via
 465 over-parameterization. In *International Conference on Machine Learning*, pages 242–252.
 466 PMLR, 2019.
- 467 [42] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global
 468 minima of deep neural networks. In *International Conference on Machine Learning*, pages
 469 1675–1685. PMLR, 2019.
- 470 [43] Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and
 471 deep neural networks. *Advances in Neural Information Processing Systems*, 32:10836–10846,
 472 2019.
- 473 [44] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of op-
 474 timization and generalization for overparameterized two-layer neural networks. In *International*
 475 *Conference on Machine Learning*, pages 322–332. PMLR, 2019.
- 476 [45] Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable program-
 477 ming. *Advances in Neural Information Processing Systems*, 32, 2019.
- 478 [46] Ruoqi Shen, Sebastien Bubeck, and Suriya Gunasekar. Data augmentation as feature ma-
 479 nipulation. In *International Conference on Machine Learning*, pages 19773–19808. PMLR,
 480 2022.
- 481 [47] Yiwen Kou, Zixiang Chen, Yuanzhou Chen, and Quanquan Gu. Benign overfitting for two-layer
 482 relu networks. *arXiv preprint arXiv:2303.04145*, 2023.
- 483 [48] Niladri S Chatterji and Philip M Long. Finite-sample analysis of interpolating linear classifiers
 484 in the overparameterized regime. *J. Mach. Learn. Res.*, 22:129–1, 2021.

Appendix

Contents

487	A Preliminary Lemmas	14
488	A.1 Preliminary Lemmas on Sample Properties	14
489	A.2 Preliminary Lemmas on Noise Vector Properties	14
490	A.3 Preliminary Lemmas on Graph Properties	15
491	A.4 Preliminary Lemmas on Initialization Properties	16
492	B General Lemmas for Iterative Coefficient Analysis	17
493	B.1 Coefficient update rule	17
494	B.2 Scale of training dynamics	18
495	C Two Stage Dynamics Analysis	24
496	C.1 First stage: feature learning versus noise memorization	24
497	C.2 Second stage: convergence analysis	26
498	C.3 Population loss	30
499	D Additional Experimental Procedures and Results	32
500	D.1 Dataset in Node Classification	32
501	D.2 Phase transition in GCN	33

502 A Preliminary Lemmas

503 In this section, we present preliminary lemmas which form the foundation for the proofs to be detailed
504 in the subsequent sections. The proof will be developed after the lemmas presented.

505 A.1 Preliminary Lemmas on Sample Properties

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

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

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

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

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

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

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

511 A.2 Preliminary Lemmas on Noise Vector Properties

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

$$\begin{aligned} \sigma_p^2 d/2 &\leq \|\xi_i\|_2^2 \leq 3\sigma_p^2 d/2, \\ |\langle \xi_i, \xi_{i'} \rangle| &\leq 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}, \end{aligned}$$

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

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

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

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

$$\sigma_p^2 d/2 \leq \|\xi_i\|_2^2 \leq 3\sigma_p^2 d/2.$$

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

$$|\langle \xi_i, \xi_{i'} \rangle| \leq 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}.$$

518 Applying a union bound completes the proof. \square

519 **Lemma A.3.** Suppose that $\delta > 0$ and $d = \Omega(n^2(p+s)^2 \log(4n^2/\delta))$. Then with probability at least
520 $1 - \delta$,

$$\sigma_p^2 d/(4n(p+s)) \leq \|\tilde{\xi}_i\|_2^2 \leq 3\sigma_p^2 d/(4n(p+s)),$$

521 for all $i \in [n]$.

522 *Proof of Lemma A.3.* It is known that:

$$\|\tilde{\xi}_i\|_2^2 = \frac{1}{D_i^2} \sum_{j=1}^d \left(\sum_{k=1}^{D_i} \xi_{jk} \right)^2 = \frac{1}{D_i^2} \sum_{j=1}^d \sum_{k=1}^{D_i} \xi_{jk}^2 + \frac{1}{D_i^2} \sum_{j=1}^d \sum_{k \neq k'}^{D_i} \xi_{jk} \xi_{jk'}.$$

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

$$\left| \sum_{j=1}^d \sum_{k=1}^{D_i} \xi_{jk}^2 - \sigma_p^2 d D_i \right| = O(\sigma_p^2 \cdot \sqrt{d D_i \log(4n/\delta)}).$$

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

$$3\sigma_p^2 d D_i / 4 \leq \sum_{j=1}^d \sum_{k=1}^{D_i} \xi_{jk}^2 \leq 5\sigma_p^2 d D_i / 4.$$

By Lemma A.4, we have,

$$2\sigma_p^2 d / (4n(p+s)) \leq \frac{1}{D_i^2} \sum_{j=1}^d \sum_{k=1}^{D_i} \xi_{jk}^2 \leq 6\sigma_p^2 d / (4n(p+s)).$$

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

$$|\langle \xi_k, \xi_{k'} \rangle| \leq 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}.$$

Applying a union bound we have that with probability at least $1 - \delta$,

$$|\langle \xi_k, \xi_{k'} \rangle| \leq 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}.$$

Therefore, as long as $d = \Omega(n^2(p+s)^2 \log(4n^2/\delta))$, we have

$$\sigma_p^2 d / (4n(p+s)) \leq \|\tilde{\xi}_i\|_2^2 \leq 3\sigma_p^2 d / (4n(p+s)).$$

□

A.3 Preliminary Lemmas on Graph Properties

Lemma A.4 (Degree concentration). *Let $p, s = \Omega\left(\sqrt{\frac{\log(n/\delta)}{n}}\right)$ and $\delta > 0$, then with probability at least $1 - \delta$, we have*

$$n(p+s)/4 \leq D_i \leq 3n(p+s)/4.$$

Proof. It is known that the degrees are sums of Bernoulli random variables.

$$D_i = 1 + \sum_{j \neq i}^n a_{ij},$$

where $a_{ij} = [\mathbf{A}]_{ij}$. Hence, by the Hoeffding's inequality, with probability at least $1 - \delta/n$

$$|D_i - \mathbb{E}[D_i]| < \sqrt{\log(n/\delta)(n-1)}.$$

Note that $a_{ii} = 1$ is a fixed value, which means that it is not a random variable, thus the denominator in the exponential part is $n - 1$ instead of n . Now we calculate the expectation of degree:

$$\mathbb{E}[D_{ii}] = 1 + \frac{n}{2}s + \left(\frac{n}{2} - 1\right)p = n(p+s)/2 + 1 - p,$$

then we have

$$|D_i - n(p+s)/2 + 1 - p| \leq \sqrt{n \log(n/\delta)}.$$

Because that $p, s = \Omega\left(\sqrt{\frac{\log(n/\delta)}{n}}\right)$, we further have,

$$n(p+s)/4 \leq D_i \leq 3n(p+s)/4.$$

Applying a union bound over $i \in [n]$ conclude the proof. □

541 **Lemma A.5.** Suppose that $\delta > 0$ and $n \geq 8 \frac{p+s}{(p-s)^2} \log(4/\delta)$. Then with probability at least $1 - \delta$,

$$\frac{1}{2} \frac{p-s}{p+s} |y_i| \leq |\tilde{y}_i| \leq \frac{3}{2} \frac{p-s}{p+s} |y_i|.$$

542 *Proof of Lemma A.5.* By Hoeffding's inequality, with probability at least $1 - \delta/2$, we have

$$\left| \frac{1}{D_i} \sum_{k \in \mathcal{N}(i)} y_k - \frac{p-s}{p+s} y_i \right| \leq \sqrt{\frac{\log(4/\delta)}{2n(p+s)}}.$$

543 Therefore, as long as $n \geq 8 \frac{p+s}{(p-s)^2} \log(4/\delta)$, we have:

$$\frac{1}{2} \frac{p-s}{p+s} |y_i| \leq |\tilde{y}_i| \leq \frac{3}{2} \frac{p-s}{p+s} |y_i|.$$

544 This proves the result for the stability of sign of graph convoluted label. \square

545 A.4 Preliminary Lemmas on Initialization Properties

546 **Lemma A.6.** Suppose that $d = \Omega(n(p+s) \log(nm/\delta))$, $m = \Omega(\log(1/\delta))$. Then with probability
547 at least $1 - \delta$,

$$\begin{aligned} |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle| &\leq \sqrt{2 \log(8m/\delta)} \cdot \sigma_0 \|\boldsymbol{\mu}\|_2, \\ |\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle| &\leq 2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}, \\ |\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle| &\leq 4\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))}, \end{aligned}$$

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

$$\begin{aligned} \sigma_0 \|\boldsymbol{\mu}\|_2/2 &\leq \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle \leq \sqrt{2 \log(8m/\delta)} \cdot \sigma_0 \|\boldsymbol{\mu}\|_2, \\ \sigma_0 \sigma_p \sqrt{d}/4 &\leq \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle \leq 2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d}, \\ \sigma_0 \sigma_p \sqrt{d/(n(p+s))}/4 &\leq \max_{r \in [m]} j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle \leq 2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))}, \end{aligned}$$

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

550 *Proof of Lemma A.6.* 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
551 with mean zero and variance $\sigma_0^2 \|\boldsymbol{\mu}\|_2^2$. Therefore, by Gaussian tail bound and union bound, with
552 probability at least $1 - \delta/4$,

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

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

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

555 By Lemma A.2, with probability at least $1 - \delta/4$, $\sigma_p \sqrt{d}/\sqrt{2} \leq \|\boldsymbol{\xi}_i\|_2 \leq \sqrt{3/2} \cdot \sigma_p \sqrt{d}$ for all $i \in [n]$.

556 Therefore, the result for $\langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\xi}_i \rangle$ follows the same proof as $j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle$.

557 By Lemma A.3, with probability at least $1 - \delta/4$, $\sigma_p \sqrt{d/(n(p+s))}/\sqrt{2} \leq \|\tilde{\boldsymbol{\xi}}_i\|_2 \leq \sqrt{3/2} \cdot$

558 $\sigma_p \sqrt{d/(n(p+s))}$ for all $i \in [n]$. Therefore, the result for $\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle$ follows the same proof as

559 $j \cdot \langle \mathbf{w}_{j,r}^{(0)}, \boldsymbol{\mu} \rangle$. \square

560

B General Lemmas for Iterative Coefficient Analysis

In this section, we deliver lemmas that delineate the iterative behavior of coefficients under gradient descent. We commence with proving the coefficient update rules as stated in Lemma 5.1 in Section B.1. Subsequently, we establish the scale of training dynamics as declared in Proposition 5.2, in Section B.2.

B.1 Coefficient update rule

Lemma B.1 (Restatement of Lemma 5.1). *The coefficients $\gamma_{j,r}^{(t)}, \bar{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ defined in Eq. (10) satisfy the following iterative equations:*

$$\begin{aligned}\gamma_{j,r}^{(0)}, \bar{\rho}_{j,r,i}^{(0)}, \underline{\rho}_{j,r,i}^{(0)} &= 0, \\ \gamma_{j,r}^{(t+1)} &= \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^n \ell_i^{(t)} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu}_i \rangle) y_i \tilde{y}_i \|\boldsymbol{\mu}\|_2^2, \\ \bar{\rho}_{j,r,i}^{(t+1)} &= \bar{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell_k^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = j), \\ \underline{\rho}_{j,r,i}^{(t+1)} &= \underline{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell_k^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = -j),\end{aligned}$$

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

Proof of Lemma B.1. Considering our data model and the Gaussian initialization of the GCN weights, it becomes clear that the vectors are linearly independent with a probability of 1. Consequently, the decomposition expressed in (10) is guaranteed to be unique. Now, let's consider $\hat{\gamma}_{j,r}^{(0)}, \hat{\rho}_{j,r,i}^{(0)} = 0$ and

$$\begin{aligned}\hat{\gamma}_{j,r}^{(t+1)} &= \hat{\gamma}_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^n \ell_i^{(t)} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu}_i \rangle) y_i \tilde{y}_i \|\boldsymbol{\mu}\|_2^2, \\ \hat{\rho}_{j,r,i}^{(t+1)} &= \hat{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell_k^{(t)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot y_k,\end{aligned}$$

It is then easy to check by (7) that

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

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

$$\rho_{j,r,i}^{(t)} = - \sum_{s=0}^{t-1} \sum_{k \in \mathcal{N}(i)} D_k^{-1} \frac{\eta}{nm} \cdot \ell_k^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot j y_k.$$

Moreover, note that $\ell_i^{(t)} < 0$ by the definition of the cross-entropy loss. Therefore,

$$\bar{\rho}_{j,r,i}^{(t)} = - \sum_{s=0}^{t-1} \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell_k^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = j), \quad (15)$$

$$\underline{\rho}_{j,r,i}^{(t)} = - \sum_{s=0}^{t-1} \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell_k^{(s)} \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(s)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \|\boldsymbol{\xi}_i\|_2^2 \cdot \mathbb{1}(y_k = -j). \quad (16)$$

Writing out the iterative versions of (15) and (16) completes the proof. \square

578 B.2 Scale of training dynamics

579 Our proof hinges on a meticulous evaluation of the coefficient values $\gamma_{j,r}^{(t)}, \bar{\rho}_{j,r,i}^{(t)}, \underline{\rho}_{j,r,i}^{(t)}$ throughout the
 580 entire training process. In order to facilitate a more thorough analysis, we first establish the following
 581 bounds for these coefficients, which are maintained consistently throughout the training period.

582 We will now show that the parameter of the signal-noise decomposition will stay a reasonable
 583 scale during a long time of training. Let us consider the learning period $0 \leq t \leq T^*$, where
 584 $T^* = \eta^{-1} \text{poly}(\epsilon^{-1}, \|\boldsymbol{\mu}\|_2^{-1}, d^{-1} \sigma_p^{-2}, \sigma_0^{-1}, n, m, d)$ is the maximum admissible iterations. Note that
 585 we can consider any polynomial training time T^* .

586 Denote $\alpha = 4 \log(T^*)$. Here we list the exact conditions for η, σ_0, d required by the proofs in this
 587 section, which are part of Condition 4.1:

$$\eta = O\left(\min\{nm/(q\sigma_p^2 d), nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2 d), nm/(q2^{q+2}\alpha^{q-2}\|\boldsymbol{\mu}\|_2^2)\}\right), \quad (17)$$

$$\sigma_0 \leq [16\sqrt{\log(8mn/\delta)}]^{-1} \min\left\{\Xi^{-1}\|\boldsymbol{\mu}\|_2^{-1}, (\sigma_p\sqrt{d/(n(p+s))})^{-1}\right\}, \quad (18)$$

$$d \geq 1024 \log(4n^2/\delta) \alpha^2 n^2. \quad (19)$$

588 Denote $\beta = 2 \max_{i,j,r} \{|\langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \cdot \boldsymbol{\mu} \rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\xi}_i \rangle|\}$. By Lemma A.6 and Lemma A.5 with probabil-
 589 ity at least $1 - \delta$, we can upper bound β by $4\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \cdot \max\{\Xi\|\boldsymbol{\mu}\|_2, \sigma_p\sqrt{d/(n(p+s))}\}$.
 590 Then, by (18) and (19), it is straightforward to verify the following inequality:

$$4 \max\left\{\beta, 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha\right\} \leq 1. \quad (20)$$

591 Assuming that the conditions detailed in (17), (18), and (19) are satisfied, we assert that the following
 592 property is maintained for $0 \leq t \leq T^*$.

593 **Proposition B.2.** *Under Condition 4.1, for $0 \leq t \leq T^*$, where $T^* =$
 594 $\eta^{-1} \text{poly}(\epsilon^{-1}, \|\boldsymbol{\mu}\|_2^{-1}, d^{-1} \sigma_p^{-2}, \sigma_0^{-1}, n, m, d)$, we have that*

$$0 \leq \gamma_{j,r}^{(t)}, \bar{\rho}_{j,r,i}^{(t)} \leq \alpha, \quad (21)$$

$$0 \geq \underline{\rho}_{j,r,i}^{(t)} \geq -\alpha, \quad (22)$$

595 for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$, where $\alpha = 4 \log(T^*)$.

596 To establish Proposition B.2, we will employ an inductive approach. Before proceeding with the
 597 proof, we need to introduce several technical lemmas that are fundamental to our argument.

598 **Lemma B.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 $r \in [m]$, $j \in \{\pm 1\}$.*

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

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

600 where the equation is by our orthogonal assumption between feature vector and noise vector. \square

601 **Lemma B.4.** *Under Condition 4.1, suppose (21) and (22) hold at iteration t . Then*

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

602 where $\hat{\rho}_{j,r,i}^{(t)} \triangleq \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \rho_{j,r,i'}^{(t)}$, for all $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$.

603 *Proof of Lemma B.4.* It is known that,

$$\begin{aligned}
\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle &= \sum_{i'=1}^n \bar{\rho}_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_i \rangle + \sum_{i'=1}^n \rho_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_i \rangle \\
&= \sum_{i'=1}^n \sum_{k \in \mathcal{N}(i)} D_i^{-1} \bar{\rho}_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_k \rangle \\
&\quad + \sum_{i'=1}^n \sum_{k \in \mathcal{N}(i)} D_i^{-1} \rho_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_k \rangle \\
&\leq 4\sqrt{\frac{\log(4n^2/\delta)}{d}} \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} |\bar{\rho}_{j,r,i'}^{(t)}| \\
&\quad + 4\sqrt{\frac{\log(4n^2/\delta)}{d}} \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} |\rho_{j,r,i'}^{(t)}| \\
&\quad + \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \bar{\rho}_{j,r,i'}^{(t)} + \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \rho_{j,r,i'}^{(t)} \\
&\leq \hat{\rho}_{j,r,i}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,
\end{aligned}$$

604 where we define $\hat{\rho}_{j,r,i}^{(t)} \triangleq \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \rho_{j,r,i'}^{(t)}$ the second inequality is by Lemma A.2 and the
605 last inequality is by $|\bar{\rho}_{j,r,i'}^{(t)}|, |\rho_{j,r,i'}^{(t)}| \leq \alpha$ in (21).

606 Similarly, we can show that:

$$\begin{aligned}
\langle \mathbf{w}_{j,r}^{(t)} - \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle &= \sum_{i'=1}^n \bar{\rho}_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_i \rangle + \sum_{i'=1}^n \rho_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_i \rangle \\
&= \sum_{i'=1}^n \sum_{k \in \mathcal{N}(i)} D_i^{-1} \bar{\rho}_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_k \rangle \\
&\quad + \sum_{i'=1}^n \sum_{k \in \mathcal{N}(i)} D_i^{-1} \rho_{j,r,i'}^{(t)} \|\boldsymbol{\xi}_{i'}\|_2^{-2} \cdot \langle \boldsymbol{\xi}_{i'}, \boldsymbol{\xi}_k \rangle \\
&\geq -4\sqrt{\frac{\log(4n^2/\delta)}{d}} \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} |\bar{\rho}_{j,r,i'}^{(t)}| \\
&\quad - 4\sqrt{\frac{\log(4n^2/\delta)}{d}} \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} |\rho_{j,r,i'}^{(t)}| \\
&\quad + \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \bar{\rho}_{j,r,i'}^{(t)} + \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i' \neq k} \rho_{j,r,i'}^{(t)} \\
&\geq \hat{\rho}_{j,r,i}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha,
\end{aligned}$$

607 where the first inequality is by Lemma A.1 and the second inequality is by $|\bar{\rho}_{j,r,i'}^{(t)}|, |\rho_{j,r,i'}^{(t)}| \leq \alpha$ in
608 (21), which completes the proof. \square

609 **Lemma B.5.** Under Condition 4.1, suppose (21) and (22) hold at iteration t . Then

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

610 for all $r \in [m]$ and $j \neq y_i$. If $\max\{\gamma_{j,r}^{(t)}, \rho_{j,r,i}^{(t)}\} = O(1)$, we further have that $F_j(\mathbf{W}_j^{(t)}, \tilde{\mathbf{x}}_i) = O(1)$.

611 *Proof of Lemma B.5.* For $j \neq y_i$, we have that

$$\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle = \langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \boldsymbol{\mu} \rangle + \tilde{y}_i \cdot j \cdot \gamma_{j,r}^{(t)} \leq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \boldsymbol{\mu} \rangle, \quad (23)$$

612 where the inequality is by $\gamma_{j,r}^{(t)} \geq 0$ and Lemma A.5 stating that $\text{sign}(y_i) = \text{sign}(\tilde{y}_i)$ with a high
613 probability. In addition, we have

$$\begin{aligned} \langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle &= \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle + \sum_{k \in \mathcal{N}(i)} D_i^{-1} \sum_{i'=1}^n \rho_{j,r,i'} \langle \boldsymbol{\xi}_k, \boldsymbol{\xi}_{i'} \rangle \|\boldsymbol{\xi}_{i'}\|_2^{-2} \\ &\leq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle + D_i^{-1} \left(\sum_{y_k \neq j} \rho_{j,r,i}^{(t)} + \sum_{y_k = j} \bar{\rho}_{j,r,i}^{(t)} \right) + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha \\ &\leq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha, \end{aligned} \quad (24)$$

614 where the first inequality is by Lemma B.4 and the second inequality is due to $\hat{\rho}_{j,r,i}^{(t)} \leq 0$ based on
615 Lemma A.5. Then we can get that

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

616 where the first inequality is by (23), (24) and the second inequality is by (20) and $\max\{\gamma_{j,r}^{(t)}, \rho_{j,r,i}^{(t)}\} =$
617 $O(1)$. \square

618 **Lemma B.6.** Under Condition 4.1, suppose (21) and (22) hold at iteration t . Then

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

619 for all $r \in [m]$, $j = y_i$ and $i \in [n]$. If $\max\{\gamma_{j,r}^{(t)}, \rho_{j,r,i}^{(t)}\} = O(1)$, we further have that
620 $F_j(\mathbf{W}_j^{(t)}, \tilde{\mathbf{x}}_i) = O(1)$.

621 *Proof of Lemma B.6.* For $j = y_i$, we have that

$$\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle = \langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \boldsymbol{\mu} \rangle + \gamma_{j,r}^{(t)}, \quad (25)$$

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

$$\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle \leq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle + \hat{\rho}_{j,r,i}^{(t)} + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha, \quad (26)$$

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

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

where $\hat{\rho}_{j,r,i}^{(t)} = \frac{1}{D_i} \sum_{k \in \mathcal{N}(i)} \bar{\rho}_{j,r,k}^{(t)} \mathbb{1}(y_k = j) + \bar{\rho}_{j,r,k}^{(t)} \mathbb{1}(y_k \neq j)$, the first inequality is by (25), (26). Then the second inequality is by (20) where $\beta = 2 \max_{i,j,r} \{|\langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \cdot \boldsymbol{\mu} \rangle|, |\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle|\} \leq 1$ and condition that $\max\{\gamma_{j,r}^{(t)}, \rho_{j,r,i}^{(t)}\} = O(1)$. \square

Now we are ready to prove Proposition B.2.

Proof of Proposition B.2. The proof of Proposition B.2 relies on induction. At $t = 0$, the results are straightforward, given that all coefficients are zero. We assume that there is a time $\tilde{T} \leq T^*$ for which the Proposition B.2 is valid for all moments $0 \leq t \leq \tilde{T} - 1$. Our goal is to demonstrate that the proposition also stands true for $t = \tilde{T}$.

We first prove that (22) holds for $t = \tilde{T}$, i.e., $\rho_{j,r,i}^{(t)} \geq -\beta - 16n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha$ for $t = \tilde{T}$, $r \in [m]$, $j \in \{\pm 1\}$ and $i \in [n]$. Notice that $\rho_{j,r,i}^{(t)} = 0, \forall j = y_i$. Therefore, we only need to consider the case that $j \neq y_i$. When $\rho_{j,r,i}^{(\tilde{T}-1)} \leq -0.5\beta - 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha$, by Lemma B.4 we have that

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

and thus

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

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

$$\begin{aligned} \rho_{j,r,i}^{(\tilde{T})} &= \rho_{j,r,i}^{(\tilde{T}-1)} + \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \ell'_k(\tilde{T}-1) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(\tilde{T}-1)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \mathbb{1}(y_k = -j) \|\boldsymbol{\xi}_i\|_2^2 \\ &\geq -0.5\beta - 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha - O\left(\frac{\eta \sigma_p^2 d}{nm}\right) \sigma' \left(0.5\beta + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha\right) \\ &\geq -0.5\beta - 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha - O\left(\frac{\eta q \sigma_p^2 d}{nm}\right) \left(0.5\beta + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha\right) \\ &\geq -\beta - 16n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha, \end{aligned}$$

where we use $\ell'_i(\tilde{T}-1) \leq 1$ 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 $\eta = O(nm/(q\sigma_p^2 d))$ in (17).

Next we prove (21) holds for $t = \tilde{T}$. We have

$$\begin{aligned} |\ell'_i(t)| &= \frac{1}{1 + \exp\{y_i \cdot [F_{+1}(\mathbf{W}_{+1}^{(t)}, \tilde{\mathbf{x}}_i) - F_{-1}(\mathbf{W}_{-1}^{(t)}, \tilde{\mathbf{x}}_i)]\}} \\ &\leq \exp\{-y_i \cdot [F_{+1}(\mathbf{W}_{+1}^{(t)}, \tilde{\mathbf{x}}_i) - F_{-1}(\mathbf{W}_{-1}^{(t)}, \tilde{\mathbf{x}}_i)]\} \\ &\leq \exp\{-F_{y_i}(\mathbf{W}_{y_i}^{(t)}, \tilde{\mathbf{x}}_i) + 1\}. \end{aligned} \tag{27}$$

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

$$\begin{aligned}\gamma_{j,r}^{(t+1)} &= \gamma_{j,r}^{(t)} - \frac{\eta}{nm} \cdot \sum_{i=1}^n \ell'_i(t) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \cdot \boldsymbol{\mu} \rangle) y_i \tilde{y}_i \|\boldsymbol{\mu}\|_2^2, \\ \bar{\rho}_{j,r,i}^{(t+1)} &= \bar{\rho}_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \ell'_k(t) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \mathbb{1}(y_k = j) \|\boldsymbol{\xi}_i\|_2^2.\end{aligned}$$

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

$$\begin{aligned}\bar{\rho}_{j,r,i}^{(\tilde{T})} &= \bar{\rho}_{j,r,i}^{(t_{j,r,i})} - \underbrace{\frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell'_k(t_{j,r,i}) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t_{j,r,i})}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \mathbb{1}(y_k = j) \|\boldsymbol{\xi}_i\|_2^2}_{I_1} \\ &\quad - \underbrace{\sum_{t_{j,r,i} < t < T} \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \ell'_k(t) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle) \cdot \mathbb{1}(y_k = j) \|\boldsymbol{\xi}_i\|_2^2}_{I_2}.\end{aligned}\quad (28)$$

643 We first bound I_1 as follows,

$$\begin{aligned}|I_1| &\leq 2qn^{-1}m^{-1}\eta \left(\max_k \bar{\rho}_{j,r,k}^{(t_{j,r,i})} + 0.5\beta + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \right)^{q-1} \sigma_p^2 d \\ &\leq q2^q n^{-1}m^{-1}\eta \alpha^{q-1} \sigma_p^2 d \leq 0.25\alpha,\end{aligned}$$

644 where the first inequality is by Lemmas B.4 and A.2, the second inequality is by $\beta \leq 0.1\alpha$ and
645 $8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0.1\alpha$, the last inequality is by $\eta \leq nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2 d)$.

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

$$\begin{aligned}\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle &\geq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_k \rangle + \hat{\rho}_{j,r,k}^{(t)} - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\ &\geq -0.5\beta + \frac{1}{4} \frac{p-s}{p+s} \alpha - 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\ &\geq 0.25\alpha,\end{aligned}$$

647 where the first inequality is by Lemma B.4, the second inequality is by $\hat{\rho}_{j,r,i}^{(t)} > \frac{1}{4} \frac{p-s}{p+s} \alpha$ and
648 $\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle \geq -0.5\beta$ due to the definition of $t_{j,r,i}$ and β , the last inequality is by $\beta \leq 0.1\alpha$ and
649 $8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \leq 0.1\alpha$. Similarly, for $t_{j,r,i} < t < \tilde{T}$ and $y_k = j$, we can also upper bound
650 $\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle$ as follows,

$$\begin{aligned}\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_k \rangle &\leq \langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_k \rangle + \hat{\rho}_{j,r,k}^{(t)} + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\ &\leq 0.5\beta + \frac{3}{4} \frac{p-s}{p+s} \alpha + 8n\sqrt{\frac{\log(4n^2/\delta)}{d}}\alpha \\ &\leq 2\alpha,\end{aligned}$$

651 where the first inequality is by Lemma B.4, the second inequality is by induction hypothesis $\hat{\rho}_{j,r,i}^{(t)} \leq \alpha$,
652 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

653 bounds of $\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\xi}_k \rangle$ into I_2 gives

$$\begin{aligned}
|I_2| &\leq \sum_{t_{j,r,i} < t < \tilde{T}} \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \exp(-\sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\xi}_k \rangle) + 1) \cdot \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\xi}_k \rangle) \cdot \mathbb{1}(y_k = j) \|\xi_i\|_2^2 \\
&\leq \frac{eq2^q \eta T^*}{n} \exp(-\alpha^q/4^q) \alpha^{q-1} \sigma_p^2 d \\
&\leq 0.25 T^* \exp(-\alpha^q/4^q) \alpha \\
&\leq 0.25 T^* \exp(-\log(T^*)^q) \alpha \\
&\leq 0.25 \alpha,
\end{aligned}$$

654 where the first inequality is by (27), the second inequality is by Lemma A.2, the third inequality is by
655 $\eta = O(nm/(q2^{q+2}\alpha^{q-2}\sigma_p^2d))$ in (17), the fourth inequality is by our choice of $\alpha = 4\log(T^*)$ and
656 the last inequality is due to the fact that $\log(T^*)^q \geq \log(T^*)$. Plugging the bound of I_1, I_2 into (28)
657 completes the proof for $\bar{\rho}$.

658 Similarly, we can prove that $\gamma_{j,r}^{(\tilde{T})} \leq \alpha$ using $\eta = O(nm/(q2^{q+2}\alpha^{q-2}\|\mu\|_2^2))$ in (17). Therefore
659 Proposition B.2 holds for $t = \tilde{T}$, which completes the induction. \square

660 Building upon Proposition B.2, we introduce a key property of the training loss function within the
661 range of $0 \leq t \leq T^*$. This property will be instrumental in the forthcoming proof of convergence.

662 **Lemma B.7.** *Under Condition 4.1, for $0 \leq t \leq T^*$, the following result holds.*

$$\|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 \leq O(\max\{\Xi^2\|\mu\|_2^2, \sigma_p^2 d/(n(p+s))\}) L_S(\mathbf{W}^{(t)}).$$

663 *Proof of Lemma B.7.* We first prove that

$$-\ell'(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \cdot \|\nabla f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)\|_F^2 = O(\max\{\Xi^2\|\mu\|_2^2, \sigma_p^2 d/(n(p+s))\}). \quad (29)$$

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

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

665 where the first and second inequalities are by triangle inequality, the third inequality is by Jensen's
666 inequality and Lemma A.2, and the last inequality is due to Lemma B.5. Denote $A = F_{+1}(\mathbf{W}_{+1}^{(t)}, \tilde{\mathbf{x}}_i)$.

667 Then we have that $A \geq 0$, and besides, $F_{-1}(\mathbf{W}_{-1}^{(t)}, \tilde{\mathbf{x}}_i) \leq 1$ by Lemma B.5. Then we have that

$$\begin{aligned}
&-\ell'(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \cdot \|\nabla f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)\|_F^2 \\
&\leq -\ell'(A-1) \left((2q \cdot A^{(q-1)/q} + 1) \max\{\Xi\|\mu\|_2, 2\sigma_p \sqrt{d/(n(p+s))}\} \right)^2 \\
&= -4q^2 \ell'(A-1) (A^{(q-1)/q} + 1)^2 \cdot \max\{\Xi^2\|\mu\|_2^2, 4\sigma_p^2 d/(n(p+s))\} \\
&\leq \left(\max_{z>0} -4q^2 \ell'(z-1) (z^{(q-1)/q} + 1)^2 \right) \max\{\Xi^2\|\mu\|_2^2, 4\sigma_p^2 d/(n(p+s))\} \\
&\stackrel{(i)}{=} O(\max\{\Xi^2\|\mu\|_2^2, \sigma_p^2 d/(n(p+s))\}),
\end{aligned}$$

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,

$$\begin{aligned} \|\nabla L_S(\mathbf{W}^{(t)})\|_F^2 &\leq \left[\frac{1}{n} \sum_{i=1}^n \ell'(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \|\nabla f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)\|_F \right]^2 \\ &\leq \left[\frac{1}{n} \sum_{i=1}^n \sqrt{-O(\max\{\Xi^2 \|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d/(n(p+s))\}) \ell'(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i))} \right]^2 \\ &\leq O(\max\{\Xi^2 \|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d/(n(p+s))\}) \cdot \frac{1}{n} \sum_{i=1}^n -\ell'(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \\ &\leq O(\max\{\Xi^2 \|\boldsymbol{\mu}\|_2^2, \sigma_p^2 d/(n(p+s))\}) L_S(\mathbf{W}^{(t)}), \end{aligned}$$

where the first inequality is by triangle inequality, the second inequality is by (29), the third inequality is by Cauchy-Schwartz inequality and the last inequality is due to the property of the cross entropy loss $-\ell' \leq \ell$. \square

C Two Stage Dynamics Analysis

In this section, we employ a two-stage dynamics analysis to investigate the behavior of coefficient iterations. During the first stage, the derivative of the loss function remains almost constant due to the small weight initialization. In the second stage, the derivative of the loss function ceases to be constant, necessitating an analysis that meticulously takes this into account. Upon completion of the convergence analysis, we employ its results to facilitate the evaluation of the population loss.

C.1 First stage: feature learning versus noise memorization

Lemma C.1 (Restatement of Lemma 5.3). *Under the same conditions as Theorem 4.3, in particular if we choose*

$$n \cdot \text{SNR}^q \cdot (n(p+s))^{q/2-1} \geq C \log(6/\sigma_0 \|\boldsymbol{\mu}\|_2) 2^{2q+6} [4 \log(8mn/\delta)]^{(q-1)/2}, \quad (30)$$

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

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

such that

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

Proof of Lemma C.1. Let

$$T_1^+ = \frac{nm\eta^{-1} \sigma_0^{2-q} \sigma_p^{-q} d^{-q/2} (n(p+s))^{(q-2)/2}}{2^{q+4} q [4 \log(8mn/\delta)]^{(q-2)/2}}. \quad (31)$$

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

$$\Psi^{(t)} \leq \sigma_0 \sigma_p \sqrt{d/(n(p+s))} \quad (32)$$

for all $0 \leq t \leq T_1^+$. By definition, clearly we have $\Psi^{(0)} = 0$. Now suppose that there exists some $\tilde{T} \leq T_1^+$ such that (32) holds for $0 < t \leq \tilde{T} - 1$. Then by (13) and (14) we have

$$\begin{aligned}
\Psi^{(t+1)} &\leq \Psi^{(t)} + \frac{\eta}{nm} \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot |\ell_k^{(t)}|. \\
&\sigma' \left(\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\xi}_k \rangle + \sum_{i'=1}^n \Psi^{(t)} \cdot \frac{|\langle \xi_{i'}, \tilde{\xi}_k \rangle|}{\|\xi_{i'}\|_2^2} + \sum_{i'=1}^n \Psi^{(t)} \cdot \frac{|\langle \xi_{i'}, \tilde{\xi}_k \rangle|}{\|\xi_{i'}\|_2^2} \right) \cdot \|\xi_i\|_2^2 \\
&\leq \Psi^{(t)} + \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \sigma' \left(\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\xi}_k \rangle + 2 \cdot \sum_{i'=1}^n \Psi^{(t)} \cdot \frac{|\langle \xi_{i'}, \tilde{\xi}_k \rangle|}{\|\xi_{i'}\|_2^2} \right) \cdot \|\xi_i\|_2^2 \\
&= \Psi^{(t)} + \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \cdot \\
&\sigma' \left(\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\xi}_k \rangle + 2\Psi^{(t)} + 2 \cdot \sum_{i' \neq k'}^n \Psi^{(t)} \cdot D_k^{-1} \sum_{k' \in \mathcal{N}(k)} \frac{|\langle \xi_{i'}, \xi_{k'} \rangle|}{\|\xi_{i'}\|_2^2} \right) \cdot \|\xi_i\|_2^2 \\
&\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \left[2 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))} \right. \\
&\quad \left. + \left(2 + \frac{4n\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}}{\sigma_p^2 d} \right) \cdot \Psi^{(t)} \right]^{q-1} \cdot 2\sigma_p^2 d \\
&\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \left(2 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))} + 4\Psi^{(t)} \right)^{q-1} \cdot 2\sigma_p^2 d \\
&\leq \Psi^{(t)} + \frac{\eta q}{nm} \cdot \left(4 \cdot \sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))} \right)^{q-1} \cdot 2\sigma_p^2 d,
\end{aligned}$$

where the second inequality is by $|\ell_i^{(t)}| \leq 1$, the third inequality is due to Lemmas A.2 and A.6, the fourth inequality follows by the condition that $d \geq 16Dn^2 \log(4n^2/\delta)$, and the last inequality follows by the induction hypothesis (32). Taking a telescoping sum over $t = 0, 1, \dots, \tilde{T} - 1$ then gives

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

where the second inequality follows by $\tilde{T} \leq T_1^+$ in our induction hypothesis. Therefore, by induction, we have $\Psi^{(t)} \leq \sigma_0 \sigma_p \sqrt{d/(n(p+s))}/2$ 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 $\max_r \gamma_{1,r}^{(t)} \leq 2$, given that $\sigma_0 \leq \sqrt{n(p+s)}/d/\sigma_p$. Then for $t \leq T_{1,1}$, $\max_{j,r,i} \{|\rho_{j,r,i}^{(t)}|\} = O(\sigma_0 \sigma_p \sqrt{d/(n(p+s))}) = O(1)$ and $\max_r \gamma_{1,r}^{(t)} \leq 2$. Therefore, by Lemmas B.5 and B.6, we know that $F_{-1}(\mathbf{W}_{-1}^{(t)}, \tilde{\mathbf{x}}_i), F_{+1}(\mathbf{W}_{+1}^{(t)}, \tilde{\mathbf{x}}_i) = O(1)$ 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 (12), for $t \leq T_{1,1}$ we have

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

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

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

704 where the second inequality is by the lower bound on the number of positive data in Lemma A.1, the
705 third inequality is due to the fact that $A^{(t)}$ is an increasing sequence, and the last inequality follows
706 by $A^{(0)} = \max_r \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle \geq \sigma_0 \|\boldsymbol{\mu}\|_2 / 2$ proved in Lemma A.6. Therefore, the sequence $A^{(t)}$ will
707 exponentially grow and we have that

$$\begin{aligned} A^{(t)} &\geq \left(1 + \frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^q m} \Xi^q \right)^t A^{(0)} \\ &\geq \exp \left(\frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^{q+1} m} \Xi^q t \right) A^{(0)} \\ &\geq \exp \left(\frac{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q}{2^{q+1} m} \Xi^q t \right) \frac{\sigma_0 \|\boldsymbol{\mu}\|_2}{2}, \end{aligned}$$

708 where the second inequality is due to the fact that $1 + z \geq \exp(z/2)$ for $z \leq 2$ and our condition
709 of η and σ_0 listed in Condition 4.1, and the last inequality follows by Lemma A.6 and $A^{(0)} =$
710 $\max_r \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle$. Therefore, $A^{(t)}$ will reach 3 within $T_1 = \frac{\log(6/\sigma_0 \|\boldsymbol{\mu}\|_2) 2^{q+1} m}{C_1 \eta q \sigma_0^{q-2} \|\boldsymbol{\mu}\|_2^q \Xi^q}$ iterations. Since
711 $\max_r \gamma_{1,r}^{(t)} \geq A^{(t)} - \max_r \langle \mathbf{w}_{1,r}^{(0)}, \boldsymbol{\mu} \rangle \geq A^{(t)} - 1$, $\max_r \gamma_{1,r}^{(t)}$ will reach 2 within T_1 iterations. We
712 can next verify that

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

713 where the inequality holds due to our SNR condition in (30). Therefore, by the definition of $T_{1,1}$, we
714 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
715 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
716 proof.

717

□

718 C.2 Second stage: convergence analysis

719 After the first stage and at time step T_1 we know that:

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

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

- 721 • $\max_r \gamma_{j,r}^{(T_1)} \geq 2, \forall j \in \{\pm 1\}$.
- 722 • $\max_{j,r,i} |\rho_{j,r,i}^{(T_1)}| \leq \hat{\beta}$ where $\hat{\beta} = \sigma_0 \sigma_p \sqrt{d/(n(p+s))} / 2$.

723 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

724 $\gamma_{j,r}^{(t+1)} \geq \gamma_{j,r}^{(t)}$, and therefore $\max_r \gamma_{j,r}^{(t)} \geq 2$. Now we choose \mathbf{W}^* as follows:

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

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

726 **Lemma C.2.** Under the same conditions as Theorem 4.3, we have that $\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F \leq$
727 $\tilde{O}(m^{3/2}\|\boldsymbol{\mu}\|_2^{-1})$.

728 *Proof of Lemma C.2.* We have

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

729 where the first inequality is by triangle inequality, the second inequality is by our decomposition of
730 $\mathbf{W}^{(T_1)}$ and the definition of \mathbf{W}^* , the third inequality is by Proposition B.2 and Lemma C.1, and the
731 last inequality is by our condition of σ_0 in Condition 4.1. \square

732 **Lemma C.3.** Under the same conditions as Theorem 4.3, we have that $y_i \langle \nabla f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i), \mathbf{W}^* \rangle \geq$
733 $q \log(2q/\epsilon)$ for all $i \in [n]$ and $T_1 \leq t \leq T^*$.

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

$$\begin{aligned} &y_i \langle \nabla f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i), \mathbf{W}^* \rangle \tag{33} \\ &= \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) \langle \tilde{y}_i y_i \boldsymbol{\mu}, j \mathbf{w}_{j,r}^* \rangle + \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle) \langle y_i \tilde{\boldsymbol{\xi}}_i, j \mathbf{w}_{j,r}^* \rangle \\ &= \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) 2qm \tilde{y}_i y_i \log(2q/\epsilon) + \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) \tilde{y}_i y_i \langle \boldsymbol{\mu}, j \mathbf{w}_{j,r}^{(0)} \rangle \\ &\quad + \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle) \langle y_i \tilde{\boldsymbol{\xi}}_i, j \mathbf{w}_{j,r}^{(0)} \rangle \\ &\geq \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) 2qm \Xi \log(2q/\epsilon) - \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle) \tilde{O}(\sigma_0 \Xi \|\boldsymbol{\mu}\|_2) \\ &\quad - \frac{1}{m} \sum_{j,r} \sigma'(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle) \tilde{O}(\sigma_0 \sigma_p \sqrt{d/(n(p+s))}), \tag{34} \end{aligned}$$

736 where the inequality is by Lemma A.6. Next we will bound the inner-product terms in (34) respec-
737 tively. By Lemma B.6, we have that for $j = y_i$

$$\max_r \{\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \boldsymbol{\mu} \rangle\} = \max_r \{\gamma_{j,r}^{(t)} + \langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \boldsymbol{\mu} \rangle\} \geq 2 - \tilde{O}(\sigma_0 \Xi \|\boldsymbol{\mu}\|_2) \geq 1. \tag{35}$$

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

$$\begin{aligned} |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{y}_i \cdot \boldsymbol{\mu} \rangle| &\stackrel{(i)}{\leq} |\langle \mathbf{w}_{j,r}^{(0)}, \tilde{y}_i \cdot \boldsymbol{\mu} \rangle| + |\gamma_{j,r}^{(t)}| \stackrel{(ii)}{\leq} \tilde{O}(1) \\ |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle| &\stackrel{(iii)}{\leq} |\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle| + |\hat{\rho}_{j,r,i}^{(t)}| + 8n \sqrt{\frac{\log(4n^2/\delta)}{d}} \alpha \stackrel{(iv)}{\leq} \tilde{O}(1), \tag{36} \end{aligned}$$

740 where (i) is by Lemma B.3, (iii) is by Lemma B.4, (ii) and (iv) are due to Proposition B.2. Plugging
741 (35) and (36) into (34) gives,

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

742 where the last inequality is by $\sigma_0 \leq \tilde{O}(m^{-2/(q-2)} n^{-1}) \cdot \min\{(\sigma_p \sqrt{d/(n(p+s))})^{-1}, \Xi^{-1} \|\boldsymbol{\mu}\|_2^{-1}\}$
743 in Condition 4.1. This completes the proof. \square

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

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

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

746 *Proof of Lemma C.4.* It is known that:

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

747 where the first inequality is by Lemma C.3, the second inequality is due to the convexity of the cross
748 entropy function, and the last inequality is due to Lemma B.7. \square

749 **Lemma C.5** (Restatement of Lemma 5.4). *Under the same conditions as Theorem 4.3, let $T =$
750 $T_1 + \left\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \right\rfloor = T_1 + \tilde{O}(m^3\eta^{-1}\epsilon^{-1}\|\boldsymbol{\mu}\|_2^{-2})$. Then we have $\max_{j,r,i} |\rho_{j,r,i}^{(t)}| \leq 2\hat{\beta} =$
751 $\sigma_0\sigma_p\sqrt{d/(n(p+s))}$ for all $T_1 \leq t \leq T$. Besides,*

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

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

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

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

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

$$\sum_{s=T_1}^t L_S(\mathbf{W}^{(s)}) \leq \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2 + \eta\epsilon(t - T_1 + 1)}{(2q-1)\eta} \quad (37)$$

755 for all $T_1 \leq t \leq T$. Dividing $(t - T_1 + 1)$ on both side of (37) gives that

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

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

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

757 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
758 is smaller than ϵ , we can conclude that there exist $T_1 \leq t \leq T$ such that $L_S(\mathbf{W}^{(t)}) < \epsilon$.

759 Finally, we will prove that $\max_{j,r,i} |\rho_{j,r,i}^{(t)}| \leq 2\hat{\beta}$ for all $t \in [T_1, T]$. Plugging $T = T_1 +$
760 $\left\lfloor \frac{\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{2\eta\epsilon} \right\rfloor$ into (37) gives that

$$\sum_{s=T_1}^T L_S(\mathbf{W}^{(s)}) \leq \frac{2\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F^2}{(2q-1)\eta} = \tilde{O}(\eta^{-1}m^3\|\boldsymbol{\mu}\|_2^2), \quad (38)$$

761 where the inequality is due to $\|\mathbf{W}^{(T_1)} - \mathbf{W}^*\|_F \leq \tilde{O}(m^{3/2}\|\boldsymbol{\mu}\|_2^{-1})$ in Lemma C.2. Define $\Psi^{(t)} =$
762 $\max_{j,r,i} |\rho_{j,r,i}^{(t)}|$. We will use induction to prove $\Psi^{(t)} \leq 2\hat{\beta}$ for all $t \in [T_1, T]$. At $t = T_1$, by the
763 definition of $\hat{\beta}$, clearly we have $\Psi^{(T_1)} \leq \hat{\beta} \leq 2\hat{\beta}$. Now suppose that there exists $\tilde{T} \in [T_1, T]$ such
764 that $\Psi^{(t)} \leq 2\hat{\beta}$ for all $t \in [T_1, \tilde{T} - 1]$. Then for $t \in [T_1, \tilde{T} - 1]$, by the following expression:

$$\begin{aligned} \rho_{j,r,i}^{(t+1)} &= \rho_{j,r,i}^{(t)} - \frac{\eta}{nm} \cdot \sum_{k \in \mathcal{N}(i)} D_k^{-1} \ell_k^{(t)} \\ &\quad \sigma' \left(\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_k \rangle + \sum_{i'=1}^n \bar{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_k \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} + \sum_{i'=1}^n \underline{\rho}_{j,r,i'}^{(t)} \frac{\langle \boldsymbol{\xi}_{i'}, \tilde{\boldsymbol{\xi}}_k \rangle}{\|\boldsymbol{\xi}_{i'}\|_2^2} \right) \cdot \|\boldsymbol{\xi}_i\|_2^2 \end{aligned} \quad (39)$$

765 we have

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

766 where the second inequality is due to Lemmas A.2 and A.6, and the last inequality follows by the
767 assumption that $d \geq 16n^2 \log(4n^2/\delta)$. Taking a telescoping sum over $t = 0, 1, \dots, \tilde{T} - 1$, we have
768 that

$$\begin{aligned} \Psi^{(T)} &\stackrel{(i)}{\leq} \Psi^{(T_1)} + \frac{\eta q}{nm} \sum_{s=T_1}^{\tilde{T}-1} \max_i \sum_{k \in \mathcal{N}(i)} D_k^{-1} |\ell_k^{(s)}| \tilde{O}(\sigma_p^2 d) \hat{\beta}^{q-1} \\ &\stackrel{(ii)}{\leq} \Psi^{(T_1)} + \frac{\eta q}{nm} \tilde{O}(\sigma_p^2 d) \hat{\beta}^{q-1} \sum_{s=T_1}^{\tilde{T}-1} \max_i \sum_{k \in \mathcal{N}(i)} D_k^{-1} \ell_k^{(s)} \\ &\stackrel{(iii)}{\leq} \Psi^{(T_1)} + \tilde{O}(\eta m^{-1} \sigma_p^2 d) \hat{\beta}^{q-1} \sum_{s=T_1}^{\tilde{T}-1} L_{\mathcal{S}}(\mathbf{W}^{(s)}) \\ &\stackrel{(iv)}{\leq} \Psi^{(T_1)} + \tilde{O}(m^2 \text{SNR}^{-2}) \hat{\beta}^{q-1} \\ &\stackrel{(v)}{\leq} \hat{\beta} + \tilde{O}(m^2 n^{2/q} (n(p+s))^{1-2/q} \hat{\beta}^{q-2}) \hat{\beta} \\ &\stackrel{(vi)}{\leq} 2\hat{\beta}, \end{aligned}$$

769 where (i) is by out induction hypothesis that $\Psi^{(t)} \leq 2\hat{\beta}$, (ii) is by $|\ell'| \leq \ell$, (iii) is
770 by $\max_i \sum_{k \in \mathcal{N}(i)} D_k^{-1} \leq \sum_i \ell_i^{(s)} = nL_{\mathcal{S}}(\mathbf{W}^{(s)})$, (iv) is due to $\sum_{s=T_1}^{\tilde{T}-1} L_{\mathcal{S}}(\mathbf{W}^{(s)}) \leq$
771 $\sum_{s=T_1}^T L_{\mathcal{S}}(\mathbf{W}^{(s)}) = \tilde{O}(\eta^{-1} m^3 \|\boldsymbol{\mu}\|_2^2)$ in (38), (v) is by $n \text{SNR}^q \cdot (n(p+s))^{q/2-1} \geq \tilde{\Omega}(1)$, and
772 (vi) is by the definition that $\hat{\beta} = \sigma_0 \sigma_p \sqrt{d/(n(p+s))}/2$ and $\tilde{O}(m^2 n^{2/q} (n(p+s))^{1-2/q} \hat{\beta}^{q-2}) =$

773 $\tilde{O}(m^2 n^{2/q} (n(p+s))^{1-2/q} (\sigma_0 \sigma_p \sqrt{d/(n(p+s))})^{q-2}) \leq 1$ by Condition 4.1. Therefore, $\Psi^{(\tilde{T})} \leq 2\hat{\beta}$,
 774 which completes the induction. \square

775 C.3 Population loss

776 Consider a new data point (\mathbf{x}, y) drawn from the SNM-SBM distribution. Without loss of generality,
 777 we suppose that the first patch is the signal patch and the second patch is the noise patch, i.e.,
 778 $\mathbf{x} = [y\boldsymbol{\mu}, \boldsymbol{\xi}]$. Moreover, by the signal-noise decomposition, the learned neural network has parameter

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

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

780 **Lemma C.6.** *Under the same conditions as Theorem 4.3, we have that $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle| \leq 1/2$
 781 for all $0 \leq t \leq T$, and $i \in [n]$.*

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

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

784 for all $j \in \{\pm 1\}$, $r \in [m]$ and $i \in [n]$, where (i) is by Lemma B.3, (ii) is due to $|\langle \mathbf{w}_{j,r}^{(0)}, \tilde{\boldsymbol{\xi}}_i \rangle| \leq$
 785 $2\sqrt{\log(8mn/\delta)} \cdot \sigma_0 \sigma_p \sqrt{d/(n(p+s))}$ in Lemma A.6 and $\max_{j,r,i} |\rho_{j,r,i}^{(t)}| \leq \sigma_0 \sigma_p \sqrt{d/(n(p+s))}$
 786 in Lemma C.5, and (iii) is due to our condition of $\sigma_0 \leq \tilde{O}(m^{-2/(q-2)} n^{-1}) \cdot (\sigma_p \sqrt{d/(n(p+s))})^{-1}$
 787 and $d \geq \tilde{\Omega}(m^2 n^4)$ in Condition 4.1. \square

788 **Lemma C.7.** *Under the same conditions as Theorem 4.3, with probability at least $1 - 4mT \cdot$
 789 $\exp(-C_2^{-1} \sigma_0^{-2} \sigma_p^{-2} d^{-1} n(p+s))$, we have that $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle| \leq 1/2$ for all $0 \leq t \leq T$, where
 790 $C_2 = \tilde{O}(1)$.*

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

$$\|\bar{\mathbf{w}}_{j,r}^{(t)}\|_2 \leq \tilde{O}(\sigma_0 \sqrt{d} + n\sigma_0) = \tilde{O}(\sigma_0 \sqrt{d}), \quad (40)$$

792 where the equality is due to $d \geq \tilde{\Omega}(m^2 n^4)$ by Condition 4.1.

793 By (40), $\max_{j,r} \|\bar{\mathbf{w}}_{j,r}^{(t)}\|_2 \leq C_1 \sigma_0 \sqrt{d}$, where $C_1 = \tilde{O}(1)$. Clearly $\langle \bar{\mathbf{w}}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle$ is a Gaussian distribu-
 794 tion with mean zero and standard deviation smaller than $C_1 \sigma_0 \sigma_p \sqrt{d/(n(p+s))}$. Therefore, the
 795 probability is bounded by

$$\mathbb{P}(|\langle \bar{\mathbf{w}}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle| \geq 1/2) \leq 2 \exp\left(-\frac{n(p+s)}{8C_1^2 \sigma_0^2 \sigma_p^2 d}\right).$$

796 Applying a union bound over j, r, t completes the proof. \square

797 **Lemma C.8** (Restatement of Lemma 5.5). *Let T be defined in Lemma 5.3 respectively. Under
 798 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
 799 $L_{\mathcal{D}}(\mathbf{W}^{(t)}) \leq c_1 \cdot L_S(\mathbf{W}^{(t)}) + \exp(-c_2 n^2)$.*

800 *Proof of Lemma C.8.* Let event \mathcal{E} to be the event that Lemma C.7 holds. Then we can divide
 801 $L_{\mathcal{D}}(\mathbf{W}^{(t)})$ into two parts:

$$\mathbb{E}[\ell(yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}))] = \underbrace{\mathbb{E}[\mathbb{1}(\mathcal{E})\ell(yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}))]}_{I_1} + \underbrace{\mathbb{E}[\mathbb{1}(\mathcal{E}^c)\ell(yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}))]}_{I_2}. \quad (41)$$

802 In the following, we bound I_1 and I_2 respectively.

803 **Bounding I_1 :** Since $L_S(\mathbf{W}^{(t)}) \leq 1$, there must exist one $(\tilde{\mathbf{x}}_i, y_i)$ such that $\ell(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \leq$
804 $L_S(\mathbf{W}^{(t)}) \leq 1$, which implies that $y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i) \geq 0$. Therefore, we have that

$$\exp(-y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \stackrel{(i)}{\leq} 2 \log(1 + \exp(-y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i))) = 2\ell(y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i)) \leq 2L_S(\mathbf{W}^{(t)}), \quad (42)$$

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

$$\begin{aligned} |yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}^{(2)}) - y_i f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}_i^{(2)})| &\leq \frac{1}{m} \sum_{j,r} \sigma(\langle \mathbf{w}_{j,r}, \tilde{\boldsymbol{\xi}}_i \rangle) + \frac{1}{m} \sum_{j,r} \sigma(\langle \mathbf{w}_{j,r}, \tilde{\boldsymbol{\xi}} \rangle) \\ &\leq \frac{1}{m} \sum_{j,r} \sigma(1/2) + \frac{1}{m} \sum_{j,r} \sigma(1/2) \\ &\leq 1. \end{aligned} \quad (43)$$

806 By writing $f(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}^{(2)})$, we mean that the input is $\tilde{\mathbf{x}} = [0, \tilde{\mathbf{x}}^{(2)}]$. The second inequality is by
807 $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle| \leq 1/2$ in Lemma C.7 and $\max_{j,r} |\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}}_i \rangle| \leq 1/2$ in Lemma C.6. Thus we
808 have that

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

809 where the first inequality is by the property of cross-entropy loss that $\ell(z) \leq \exp(-z)$ for all z , the
810 second inequality is by (43) and Lemma A.5, and the third inequality is by (42). Dropping the event
811 in the expectation gives $I_1 \leq c_1 L_S(\mathbf{W}^{(t)})$.

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

$$\begin{aligned} \ell(yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}})) &\leq \log(1 + \exp(F_{-y}(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}))) \\ &\leq 1 + F_{-y}(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}) \\ &= 1 + \frac{1}{m} \sum_{j=-y, r \in [m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\mathbf{y}}\boldsymbol{\mu} \rangle) + \frac{1}{m} \sum_{j=-y, r \in [m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle) \\ &\leq 1 + F_{-y_i}(\mathbf{W}_{-y_{i'}}, \tilde{\mathbf{x}}_{i'}) + \frac{1}{m} \sum_{j=-y, r \in [m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle) \\ &\leq 2 + \frac{1}{m} \sum_{j=-y, r \in [m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle) \\ &\leq 2 + \tilde{O}((\sigma_0 \sqrt{d})^q) \|\tilde{\boldsymbol{\xi}}\|^q, \end{aligned} \quad (44)$$

814 where the first inequality is due to $F_y(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}) \geq 0$, the second inequality is by the property
815 of cross-entropy loss, i.e., $\log(1 + \exp(z)) \leq 1 + z$ for all $z \geq 0$, the third inequality is by
816 $\frac{1}{m} \sum_{j=-y, r \in [m]} \sigma(\langle \mathbf{w}_{j,r}^{(t)}, \tilde{\mathbf{y}}\boldsymbol{\mu} \rangle) \leq F_{-y}(\mathbf{W}_{-y}, \tilde{\mathbf{x}}_{i'}) = F_{-y_{i'}}(\mathbf{W}_{-y_{i'}}, \tilde{\mathbf{x}}_{i'})$, the fourth inequality is by
817 $F_{-y_{i'}}(\mathbf{W}_{-y_{i'}}, \tilde{\mathbf{x}}_{i'}) \leq 1$ in Lemma B.5, and the last inequality is due to $\langle \tilde{\mathbf{w}}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle = \langle \mathbf{w}_{j,r}^{(t)}, \tilde{\boldsymbol{\xi}} \rangle \leq$




Figure 4: Test accuracy heatmap for GCNs after training.

818 $\|\bar{\mathbf{w}}_{j,r}^{(t)}\|_2 \|\tilde{\xi}\|_2 \leq \tilde{O}(\sigma_0 \sqrt{d}) \|\tilde{\xi}\|_2$ in (40). Then we further have that

$$\begin{aligned}
I_2 &\leq \sqrt{\mathbb{E}[\mathbf{1}(\mathcal{E}^c)]} \cdot \sqrt{\mathbb{E}[\ell(yf(\mathbf{W}^{(t)}, \tilde{\mathbf{x}}))^2]} \\
&\leq \sqrt{\mathbb{P}(\mathcal{E}^c)} \cdot \sqrt{4 + \tilde{O}((\sigma_0 \sqrt{d})^{2q}) \mathbb{E}[\|\tilde{\xi}\|_2^{2q}]} \\
&\leq \exp[-\tilde{\Omega}(\sigma_0^{-2} \sigma_p^{-2} d^{-1} n(p+s)) + \text{polylog}(d)] \\
&\leq \exp(-c_1 n^2),
\end{aligned}$$

819 where c_1 is a constant, the first inequality is by Cauchy-Schwartz inequality, the second inequality is
820 by (44), the third inequality is by Lemma C.7 and the fact that $\sqrt{4 + \tilde{O}((\sigma_0 \sqrt{d})^{2q}) \mathbb{E}[\|\tilde{\xi}\|_2^{2q}]} =$
821 $O(\text{poly}(d))$, and the last inequality is by our condition $\sigma_0 \leq \tilde{O}(m^{-2/(q-2)} n^{-1})$.
822 $(\sigma_p \sqrt{d/(n(p+s))})^{-1}$ in Condition 4.1. Plugging the bounds of I_1, I_2 into (41) completes the
823 proof. \square

824 D Additional Experimental Procedures and Results

825 D.1 Dataset in Node Classification

826 In Figure 1, we execute node classification experiments on three frequently used citation networks:
827 Cora, Citeseer, and Pubmed [1]. Detailed information about these datasets is provided below and
828 summarized in Table 1.

Table 1: Details of Datasets

Dataset	Nodes	Edges	Classes	Features	Train/Val/Test
Cora	2,708	5,429	7	1,433	0.05/0.18/0.37
Citeseer	3,327	4,732	6	3,703	0.04/0.15/0.30
Pubmed	19,717	44,338	3	500	0.003/0.03/0.05

- 829 • The Cora dataset includes 2,708 scientific publications, each categorized into one of seven
830 classes, connected by 5,429 links. Each publication is represented by a binary word vector,
831 which denotes the presence or absence of a corresponding word from a dictionary of 1,433
832 unique words.

- 833 • The Citeseer dataset comprises 3,312 scientific publications, each classified into one of six
834 classes, connected by 4,732 links. Each publication is represented by a binary word vector,
835 indicating the presence or absence of a corresponding word from a dictionary that includes
836 3,703 unique words.
- 837 • The Pubmed Diabetes dataset includes 19,717 scientific publications related to diabetes,
838 drawn from the PubMed database and classified into one of three classes. The citation
839 network is made up of 44,338 links. Each publication is represented by a TF-IDF weighted
840 word vector from a dictionary consisting of 500 unique words.

841 **D.2 Phase transition in GCN**

842 In Figure 3, we illustrated the variance in test accuracy between CNN and GCN within a chosen range
843 of SNR and sample numbers, where GCN was shown to achieve near-perfect test accuracy. Here,
844 we broaden the SNR range towards the smaller end and display the corresponding phase diagram
845 of GCN in Figure 4. When the SNR is exceedingly small, we observe that GCNs return lower test
846 accuracy, suggesting the possibility of a phase transition in the test accuracy of GCNs.