

# Beyond Unconstrained Features: Neural Collapse for Shallow Neural Networks with General Data

**Wanli Hong** <sup>\*†</sup>

*Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning  
New York University Shanghai  
Shanghai, 200124, China*

WH992@NYU.EDU

*Center for Data Science  
New York University  
New York, NY 10011, USA*

**Shuyang Ling**<sup>\*</sup>

*Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning  
New York University Shanghai  
Shanghai, 200124, China*

SL3635@NYU.EDU

**Editor:** Weijie Su

## Abstract

Neural collapse ( $\mathcal{NC}$ ) is a phenomenon that emerges at the terminal phase of the training (TPT) of deep neural networks (DNNs). The features of the data in the same class collapse to their respective sample means and the sample means exhibit a simplex equiangular tight frame (ETF). In the past few years, there has been a surge of works that focus on explaining why the  $\mathcal{NC}$  occurs and how it affects generalization. Since the DNNs are notoriously difficult to analyze, most works mainly focus on the unconstrained feature model (UFM). While the UFM explains the  $\mathcal{NC}$  to some extent, it fails to provide a complete picture of how the network architecture and the dataset affect  $\mathcal{NC}$ . In this work, we focus on shallow ReLU neural networks and try to understand how the width, depth, data dimension, and statistical property of the training dataset influence the neural collapse. We provide a complete characterization of when the  $\mathcal{NC}$  occurs for two or three-layer neural networks. For two-layer ReLU neural networks, a sufficient condition on when the global minimizer of the regularized empirical risk function exhibits the  $\mathcal{NC}$  configuration depends on the data dimension, sample size, and the signal-to-noise ratio in the data instead of the network width. For three-layer neural networks, we show that the  $\mathcal{NC}$  occurs as long as the first layer is sufficiently wide. Regarding the connection between  $\mathcal{NC}$  and generalization, we

---

\*. Shanghai Frontiers Science Center of Artificial Intelligence and Deep Learning, New York University Shanghai, China. S.L. and W.H. are (partially) financially supported by the National Key R&D Program of China, Project Number 2021YFA1002800, National Natural Science Foundation of China (NSFC) Grant No.12571105 and No.12001372, Science and Technology Commission of Shanghai Grant No.24QA2706200 and No.24ZR1455300, Shanghai Shanghai Municipal Education Commission (SMEC) via SMEC AI Initiative Program and Grant 0920000112, and NYU Shanghai Boost Fund. W.H. is also supported by NYU Shanghai Ph.D. fellowship and acknowledges the NSF/NRT support.

†. Center for Data Science, New York University.

Correspondence to: Shuyang Ling (sl3635@nyu.edu)

show the generalization heavily depends on the SNR (signal-to-noise ratio) in the data: even if the  $\mathcal{NC}$  occurs, the generalization can still be bad provided that the SNR in the data is too low. Our results significantly extend the state-of-the-art theoretical analysis of the  $\mathcal{NC}$  under the UFM by characterizing the emergence of the  $\mathcal{NC}$  under shallow nonlinear networks and showing how it depends on data properties and network architecture.

**Keywords:** Neural Collapse, Constrained Feature Model, Gaussian Mixture Model, Generalization

## 1. Introduction

Deep neural networks have achieved tremendous success in the past few years in a variety of applications He et al. (2016); LeCun et al. (2015). However, the mystery behind deep neural networks and deep learning remains far behind its significant applications in practice. In this work, we will focus on a phenomenon called neural collapse ( $\mathcal{NC}$ ) that was observed in Pappayan et al. (2020) that some particular structures emerge in the feature representation layer and the classification layer of DNNs in the TPT regime for classification tasks when the training dataset is balanced.

To make our discussion more precise, we consider an  $L$ -layer feedforward neural network in the form of

$$\mathbf{h}_{\ell+1}(\mathbf{x}) = \sigma(\mathbf{W}_\ell^\top \mathbf{h}_\ell(\mathbf{x}) + \mathbf{b}_\ell), \quad 1 \leq \ell \leq L-1, \quad (1.1)$$

where  $\mathbf{W}_\ell$  and  $\mathbf{b}_\ell$  are the weight and bias on the  $\ell$ -th layer, and the output after  $L-1$  layers goes through the last linear layer gives  $\mathbf{W}_L^\top \mathbf{h}_L(\mathbf{x})$  which is used for classification. Here  $\sigma(\cdot)$  is a nonlinear activation function such as ReLU activation and sigmoid function.

From now on, we always set  $\mathbf{W} := \mathbf{W}_L$ , i.e., the linear classifier on the last layer, and let  $\mathbf{h}_\theta(\mathbf{x}) := \mathbf{h}_L(\mathbf{x})$  be the feature map of the data point  $\mathbf{x}$  where  $\theta$  represents all the parameters  $\{(\mathbf{W}_\ell, \mathbf{b}_\ell)\}_{\ell=1}^{L-1}$ . Given the training data  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ , the model training reduces to the empirical risk minimization:

$$R_N(\mathbf{W}, \theta) := \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{W}^\top \mathbf{h}_\theta(\mathbf{x}_i), \mathbf{y}_i)$$

where  $\mathbf{W} \in \mathbb{R}^{D \times K}$ ,  $\mathbf{h}_\theta(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^D$  is a feature map and  $\ell(\cdot, \cdot)$  is a loss function such  $\ell_2$ -loss and cross-entropy (CE) loss.

Suppose the training dataset consists of  $K$  classes with equal size  $\{\mathbf{x}_{ki}\}_{1 \leq i \leq n, 1 \leq k \leq K}$ , i.e., each class contains exactly  $n$  points, and the label  $\mathbf{y}_k = \mathbf{e}_k$  is a one-hot vector. Then at the nearly final stage of training, the linear classifier and features exhibit the following structures Pappayan et al. (2020)  $\mathcal{NC}_1$ - $\mathcal{NC}_3$ :

- $\mathcal{NC}_1$  variability collapse: the features of the samples from the same class converge to their mean feature vector, i.e.,

$$\mathbf{h}_{ki} := \mathbf{h}_\theta(\mathbf{x}_{ki}) \longrightarrow \bar{\mathbf{h}}_k := \frac{1}{n} \sum_{i=1}^n \mathbf{h}_\theta(\mathbf{x}_{ki}), \quad 1 \leq i \leq n, 1 \leq k \leq K,$$

as the training evolves where  $\mathbf{h}_{ki}$  is the feature of the  $i$ -th data point in the  $k$ -th class.

- $\mathcal{NC}_2$  convergence to the simplex ETF: these feature vectors form an equiangular tight frame (ETF), i.e., they share the same pairwise angles and length;

$$[\langle \bar{\mathbf{h}}_k, \bar{\mathbf{h}}_{k'} \rangle]_{k,k'} \propto \mathbf{I}_K - \mathbf{J}_K/K$$

where  $\mathbf{I}_K$  and  $\mathbf{J}_K$  are the  $K \times K$  identity and constant “1” matrices respectively.

- $\mathcal{NC}_3$  convergence to self-duality: the weight of the linear classifier converges to the corresponding feature mean (up to a scalar product):

$$\mathbf{W}^\top \mathbf{W} \propto \bar{\mathbf{H}}^\top \bar{\mathbf{H}}$$

where  $\bar{\mathbf{H}} \in \mathbb{R}^{D \times K}$  consists of the mean feature vectors  $\{\bar{\mathbf{h}}_k\}_{k=1}^K$ .

These empirical findings have sparked a series of theoretical works that try to explain the emergence of the  $\mathcal{NC}$  in the training of DNNs. The key question is to understand when neural collapse occurs and how the neural collapse is related to the generalization. Starting from E and Wojtowytsch (2022); Fang et al. (2021); Lu and Steinerberger (2022); Mixon et al. (2022), a thread of works considers the unconstrained feature model (UFM) or the layer peeled model to explain the emergence of neural collapse. The rationale behind UFM relies on the universal approximation theorem Cybenko (1989); Hornik et al. (1989): over-parameterize deep neural networks have an exceptional power to approximate many common function classes. Therefore, the expressiveness of the feature map  $\mathbf{h}_\theta(\mathbf{x})$  is extremely powerful so that it can be replaced by arbitrary vectors. Various versions of UFMs with different loss functions and regularizations are proposed in these works E and Wojtowytsch (2022); Mixon et al. (2022); Zhu et al. (2021); Fang et al. (2021); Dang et al. (2023a); Lu and Steinerberger (2022); Tirer and Bruna (2022); Mixon et al. (2022); Yaras et al. (2022); Zhou et al. (2022) including the  $\mathcal{NC}$  on imbalanced datasets Fang et al. (2021); Hong and Ling (2024); Dang et al. (2023a, 2024); Thrampoulidis et al. (2022). All the aforementioned works manage to find that the global minimizers of the empirical risk function under the UFMs and the global minimizers indeed match the characterization of  $\mathcal{NC}$  proposed in Pappayan et al. (2020) such as variability collapse  $\mathcal{NC}_1$  and the convergence to simplex ETF  $\mathcal{NC}_{2-3}$ . There are also some recent works trying to extend the UFMs to deep linear neural networks Dang et al. (2023b); Garrod and Keating (2024) and also nonlinear ReLU network Dang et al. (2024); Nguyen et al. (2023); Sůkeník et al. (2024); Tirer and Bruna (2022). In particular, Tirer and Bruna (2022) shows that for two-layer ReLU networks with nonlinear unconstrained features, the global minimizer to the empirical risk function with  $\ell_2$ -loss matches the  $\mathcal{NC}$  configuration. In Sůkeník et al. (2024), the authors prove that for a binary classification task with deep unconstrained feature networks, the global minimum exhibits all the typical properties of the  $\mathcal{NC}$ . Another streamline of works focuses on analyzing the dynamical side of neural collapse: whether the optimization algorithm such as gradient descent converges to the  $\mathcal{NC}$  Poggio and Liao (2021); Han et al. (2021); Ji et al. (2022); Mixon et al. (2022); Tirer et al. (2023); Rangamani and Banburski-Fahey (2022). For a review on the topic of neural collapse, the readers might refer to Kothapalli et al. (2022). More recently, the combination of the dynamical  $\mathcal{NC}$  property and general datasets (without the UFM) has also been considered. Wang et al. (2023) proves the convergence of deep linear networks to the  $\mathcal{NC}$  architecture under the assumptions that the data are

linearly separable and nearly orthogonal and that the network weights are balanced. Jacot et al. (2024) proves the emergence of  $\mathcal{NC}$  for deep networks with a pyramidal structure in the NTK regime. The convergence to the  $\mathcal{NC}$  configuration in three-layer networks is discussed in the mean-field regime, with the network evolving along the gradient flow in Wu and Mondelli (2025). The authors also show the collapsed model will generalize well if the SNR of the training dataset is high. Also, the  $\mathcal{NC}$  has been studied in the domain of transfer learning Li et al. (2022); Galanti et al. (2021) and the relation between  $\mathcal{NC}$  and neural tangent kernel is discussed in Seleznova et al. (2024).

As the  $\mathcal{NC}$  is rather challenging to understand for deep neural networks, the UFM’s have been a great surrogate to the complicated DNNs and make the theoretical analysis much less difficult. However, the unconstrained feature model also has its obvious disadvantage: by replacing  $\mathbf{h}(\mathbf{x})$  with any vector, the resulting feature map actually has nothing to do with the alignment of training data and labels. Consequently, the study of the  $\mathcal{NC}$  under UFM’s is more or less equal to an analysis of optimization phenomenon. In addition, the relation between input data and neural collapse remains unexplored yet but is the key to understanding this intriguing phenomenon of  $\mathcal{NC}$  Papayan et al. (2020); Yang et al. (2023). Therefore, the UFM’s do not seem to be suitable models to study the connection between the generalization of DNNs and  $\mathcal{NC}$ , as discussed in Hui et al. (2022).

To overcome the obvious limitations of UFM’s, we take one step further to study the  $\mathcal{NC}$  under a more realistic setting: we focus on the shallow ReLU neural networks (with the number of layers equal to  $L = 2$  or  $3$ ) with a given dataset  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ . In other words, we will study the  $\mathcal{NC}$  under a nonlinear feature model and more importantly, in the presence of a general dataset. More precisely, our work aims to address a few important questions regarding the emergence of NC and its connection to generalization:

*Does the neural collapse occur for a sufficiently wide shallow network?* (Question 1)

It is an interesting question since we believe that the family of two-layer neural networks already has the property of universal approximation. Therefore, it is crucial to see if the universal approximation of a two-layer neural network implies neural collapse, i.e., whether the global minimizer to the regularized empirical risk minimization equals the ETF configuration.

*Is the  $\mathcal{NC}$  more likely to occur on a dataset with cluster structures?* (Question 2)

The motivation behind this question arises from a simple observation: if the datasets are highly separated, i.e., each data point is very close to its mean and far away from the data in other classes, then intuitively the  $\mathcal{NC}$  is more likely to happen than on purely random noise. Therefore, it is natural to ask how the separation (which can be quantified by the signal-to-noise ratio) determines the emergence of  $\mathcal{NC}$ .

*Does the emergence of the  $\mathcal{NC}$  necessarily imply excellent generalization?* (Question 3)

Intuitively,  $\mathcal{NC}$  may help the generalization as the  $\mathcal{NC}$  maps data in the same class to a single point, and the classifier is simplified to a nearest class-mean decision rule, which is related to the max-margin classifier Papayan et al. (2020) and implicit bias Behnia et al.

(2023); Soudry et al. (2018). Thus, it is important to address the connection between  $\mathcal{NC}$  and generalization.

**Motivation of focusing on shallow networks:** We focus our analysis on shallow networks in order to identify the simplest architectures capable of inducing  $\mathcal{NC}$  and to assess how far the assumptions of the UFM model can be relaxed. Shallow networks also provide an analytically tractable framework, which enables us to derive explicit statistical relationships between  $\mathcal{NC}$ , network width, and dataset structure. This connection, to the best of our knowledge, has not been established in prior work. Although deep, fully-trained networks are more likely to exhibit  $\mathcal{NC}$  due to their greater expressivity, their complexity makes them much harder to analyze. Extending our theoretical framework to these richer architectures remains an important direction for future research.

We summarize our main contribution as follows: we provide a new proof for the  $\mathcal{NC}$  under positive unconstrained ReLU feature model via convex optimization that is very versatile and can potentially apply to many other unconstrained feature settings. More importantly, we derive sufficient conditions for the  $\mathcal{NC}$  under both general datasets and Gaussian mixture model (GMM) for cross-entropy loss and  $\ell_2$ -loss: (a) we provide a complete characterization of when the  $\mathcal{NC}$  occurs for two-layer neural networks, i.e., when the global minimizer exhibits the  $\mathcal{NC}$  configuration. For a two-layer neural network, even if it is sufficiently wide, the  $\mathcal{NC}$  may not occur if the dimension of input data is not sufficiently large. (b) If the data dimension is moderate, whether the  $\mathcal{NC}$  occurs mainly depends on the SNR in the data instead of the network width, which is fully elucidated under the GMM as a benchmark example. For three-layer neural networks, we show that the  $\mathcal{NC}$  occurs as long as the first layer is sufficiently wide. (c) In addition, we show that under the Gaussian mixture model, the generalization mainly depends on the SNR in the dataset even if the  $\mathcal{NC}$  occurs. Our results significantly extend the state-of-the-art works that mainly focus on the unconstrained feature models.

## 1.1 Notation

We let boldface letter  $\mathbf{X}$  and  $\mathbf{x}$  be a matrix and a vector respectively;  $\mathbf{X}^\top$  and  $\mathbf{x}^\top$  are the transpose of  $\mathbf{X}$  and  $\mathbf{x}$  respectively. The matrices  $\mathbf{I}_n$ ,  $\mathbf{J}_n$ , and  $\mathbf{e}_k$  are the  $n \times n$  identity matrix, a constant matrix with all entries equal to 1, and the one-hot vector with  $k$ -th entry equal to 1. Let

$$\mathbf{C}_K := \mathbf{I}_K - \mathbf{J}_K/K \tag{1.2}$$

be the  $K \times K$  centering matrix. For any vector  $\mathbf{x}$ ,  $\text{diag}(\mathbf{x})$  denotes the diagonal matrix whose diagonal entries equal  $\mathbf{x}$ . For any matrix  $\mathbf{X}$ , we let  $\|\mathbf{X}\|$ ,  $\|\mathbf{X}\|_F$ , and  $\|\mathbf{X}\|_*$  be the operator norm, Frobenius form, and nuclear norm.

## 1.2 Organization

The following sections are organized in the following way. In Section 2, we start with the model setting and then present our main theoretical results. In Section 3, we provide numerical experiments to support our theoretical findings, and Section 4 justifies all our theorems.

## 2. Preliminaries and main theorems

In this work, we focus on the empirical risk minimization with activation regularization:

$$R_N(\mathbf{W}, \boldsymbol{\theta}) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{W}^\top \mathbf{h}_\theta(\mathbf{x}_i), \mathbf{y}_i) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}_\theta(\mathbf{X})\|_F^2 \quad (2.1)$$

where  $\mathbf{H}_\theta(\mathbf{X}) = [\mathbf{h}(\mathbf{x}_1), \dots, \mathbf{h}(\mathbf{x}_N)] \in \mathbb{R}^{D \times N}$  is the aggregation of all features.

Throughout our discussion, we assume that the dataset is balanced, i.e.,

$$\mathbf{Y} := \mathbf{I}_K \otimes \mathbf{1}_n^\top = \underbrace{[\mathbf{e}_1, \dots, \mathbf{e}_1]}_{n \text{ times}}, \dots, \underbrace{[\mathbf{e}_K, \dots, \mathbf{e}_K]}_{n \text{ times}} \in \mathbb{R}^{K \times Kn} \quad (2.2)$$

is the label matrix where each column represents the label of the corresponding data point, and each class contains  $n$  samples. In particular, the mean feature matrix  $\bar{\mathbf{H}}$  is given by

$$\bar{\mathbf{H}} := \frac{1}{n} \mathbf{H} \mathbf{Y}^\top \in \mathbb{R}^{D \times K}, \quad \bar{\mathbf{h}}_k := \frac{1}{n} \sum_{i=1}^n \mathbf{h}_{ki} \quad (2.3)$$

where  $\bar{\mathbf{h}}_k$  is the  $k$ -th column of  $\bar{\mathbf{H}}$ . The regularization is crucial as it has been shown empirically that without the regularization, the  $\mathcal{NC}$  will not happen. In particular, Rangamani and Banburski-Fahey (2022) has shown that without regularization, the DNN will interpolate the data but the features will not exhibit  $\mathcal{NC}$  phenomenon. The reason for choosing activation regularization instead of weight decay is due to the simplicities of analyzing the activation regularization. We will leave the study of the role of weight decay to our future research agenda. Before proceeding, we will need to formalize a concept about the occurrence of neural collapse.

**Definition 1 (Neural collapse occurs).** *We say the neural collapse occurs for the neural network  $f_{\mathbf{W}, \boldsymbol{\theta}}(\mathbf{x}) = \mathbf{W}^\top \mathbf{h}_\theta(\mathbf{x})$  with data  $\{(\mathbf{x}_{ki}, \mathbf{y}_k)\}_{1 \leq i \leq n, 1 \leq k \leq K}$  if there exists  $(\mathbf{W}, \boldsymbol{\theta})$  such that it is equal to the global minimizer of (2.1) and also it holds*

$$\mathbf{H}_\theta(\mathbf{X}) = \bar{\mathbf{H}}_\theta(\mathbf{X}) \otimes \mathbf{1}_n^\top, \quad \bar{\mathbf{H}}_\theta(\mathbf{X})^\top \bar{\mathbf{H}}_\theta(\mathbf{X}) \propto \mathbf{I}_K, \quad \mathbf{W}^\top \mathbf{W} \propto \mathbf{I}_K - \mathbf{J}_K / K$$

*i.e., the feature matrix  $\mathbf{H}_\theta(\mathbf{X})$  satisfies the within-class variability collapse and the mean feature  $\bar{\mathbf{H}}_\theta(\mathbf{X})$  in (2.3) converges to an orthogonal frame.*

Here are two important remarks about the definition. (a) The reason why we use the convergence of the mean features to an orthogonal frame instead of a simplex ETF is because of the nonnegativity of the ReLU feature discussed in this work, also see Dang et al. (2024); Nguyen et al. (2023); Tirer and Bruna (2022); (b) This definition only concerns the global minimizer to (1.1) and its  $\mathcal{NC}$  properties, but we do not discuss whether an algorithm (gradient descent or SGD) converges to a global minimizer, which is beyond the scope of this work and will be investigated in the future. As (2.1) is quite challenging to analyze in general, we will focus on the following special yet non-trivial models.

**Two-layer and three-layer ReLU network** We consider the two-layer bias-free neural network with ReLU activation function  $\sigma_{\text{ReLU}}(x) = \max\{x, 0\}$ , i.e.,  $L = 2$  or  $3$  in (1.1). For two-layer neural networks, we have the empirical risk function as follows:

$$R_N(\mathbf{W}, \mathbf{W}_1) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{W}^\top \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{x}_i), \mathbf{y}_i) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})\|_F^2 \quad (2.4)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times D}$ ,  $\mathbf{W} \in \mathbb{R}^{D \times K}$ , and  $\mathbf{X} \in \mathbb{R}^{d \times N}$ . In addition, we also consider the three-layer counterpart whose empirical risk function is given by

$$\begin{aligned} R_N(\mathbf{W}, \mathbf{W}_1, \mathbf{W}_2) &= \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{W}^\top \sigma_{\text{ReLU}}(\mathbf{W}_2^\top \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{x}_i)), \mathbf{y}_i) \\ &\quad + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\sigma_{\text{ReLU}}(\mathbf{W}_2^\top \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}))\|_F^2 \end{aligned} \quad (2.5)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times d_1}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{d_1 \times D}$ ,  $\mathbf{W} \in \mathbb{R}^{D \times K}$ , and  $\mathbf{X} \in \mathbb{R}^{d \times N}$ . The two-layer neural network (two-layer NN) has become one important model to study from a theoretical perspective as it is possibly the simplest nonlinear neural network. The convergence of gradient descent and SGD in training two-layer neural networks has been discussed in Li and Yuan (2017); Du et al. (2018) as well as in Mei et al. (2018); Rotskoff and Vanden-Eijnden (2022) by using the mean-field analysis of shallow networks. Despite the shallowness, two-layer neural networks have powerful approximation properties Cybenko (1989); Hornik et al. (1989): the universal approximation theorem, which motivates the study of the unconstrained feature model or the layer-peeled model Fang et al. (2021); Zhou et al. (2022). This inspires us to understand the  $\mathcal{NC}$  phenomenon in shallow networks, especially when the global minimizers to (2.5) include the  $\mathcal{NC}$  configuration.

**Unconstrained positive feature model** One obvious difficulty to determine when the  $\mathcal{NC}$  occurs comes from the nonlinearity of (1.1) even if  $L = 2$  or  $3$ . Therefore, in most state-of-the-art literature, a significant simplification of the nonlinear model (1.1) is to assume the feature matrix  $\mathbf{H} = \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  is unconstrained, i.e.,  $\mathbf{H}$  is any  $D \times N$  nonnegative matrix. We will also start with analyzing whether the  $\mathcal{NC}$  occurs under the unconstrained positive feature model (UPFM). Under  $\ell_{CE}$ -loss function and unconstrained positive feature model, we have

$$R_N(\mathbf{W}, \mathbf{H}) = \frac{1}{N} \ell_{CE}(\mathbf{W}^\top \mathbf{H}, \mathbf{Y}) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 \quad (2.6)$$

subject to  $\mathbf{H} \geq 0$  where  $\mathbf{Y}$  is defined in (2.2),

$$\ell_{CE}(\mathbf{z}, \mathbf{e}_k) = \log \sum_{j=1}^n e^{z_j} - z_k$$

and  $\mathbf{e}_k$  is a one-hot vector. The counterpart under  $\ell_2$ -loss is

$$R_N(\mathbf{W}, \mathbf{H}) = \frac{1}{2N} \|\mathbf{W}^\top \mathbf{H} - \mathbf{Y}\|_F^2 + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 \quad (2.7)$$

subject to  $\mathbf{H} \geq 0$ . The study of the  $\mathcal{NC}$  on (2.6) and (2.7) is the first step towards to understanding the  $\mathcal{NC}$  for two/three-layer neural networks (2.4) and (2.5). In fact, analyzing (2.7) is quite straightforward as it is directly related to the singular value thresholding. For (2.4), it is slightly more complicated but we will provide proof of the  $\mathcal{NC}$  under UPFM via convex relaxation.

## 2.1 $\mathcal{NC}$ under unconstrained feature model with general data

Now we present our first theorem that characterizes the  $\mathcal{NC}$  under the unconstrained positive feature model for both  $\ell_2$ - and cross-entropy loss.

**Theorem 2 (Neural collapse under unconstrained positive feature models).** *Under the unconstrained positive feature model with a balanced dataset, the following holds.*

(a) *The global minimizer to (2.6) is given by*

$$\begin{aligned} \mathbf{H}^\top \mathbf{H} &= b \mathbf{Y}^\top \mathbf{Y} = b \mathbf{I}_K \otimes \mathbf{J}_n, & \mathbf{W}^\top \mathbf{W} &= \frac{a^2}{b} (\mathbf{I}_K - \mathbf{J}_K/K), \\ \mathbf{W}^\top \mathbf{H} &= a (\mathbf{I}_K - \mathbf{J}_K/K) \mathbf{Y}, \end{aligned} \quad (2.8)$$

*i.e., the mean feature  $\bar{\mathbf{H}}$  is nonnegative and satisfies*

$$\bar{\mathbf{H}}^\top \bar{\mathbf{H}} = b \mathbf{I}_K, \quad \mathbf{W}^\top \bar{\mathbf{H}} = a (\mathbf{I}_K - \mathbf{J}_K/K), \quad \mathbf{W} = \frac{a}{b} \bar{\mathbf{H}} (\mathbf{I}_K - \mathbf{J}_K/K) \quad (2.9)$$

where  $\mathbf{Y}$  is defined in (2.2),

$$b = \sqrt{\frac{K-1}{nK} \cdot \frac{\lambda_W}{\lambda_H}} a, \quad a = \max \left\{ \log \left( (K-1) \left[ \sqrt{\frac{1}{nK(K-1)} \cdot \frac{1}{\lambda_H \lambda_W}} - 1 \right] \right), 0 \right\}$$

*In particular,  $a > 0$  if*

$$\sqrt{\frac{n(K-1)}{K}} > Kn\lambda$$

where  $\lambda = \sqrt{\lambda_W \lambda_H}$ . Here (2.8) and (2.9) correspond to  $\mathcal{NC}_1$  and  $\mathcal{NC}_{2-3}$  of the UPFM under cross-entropy loss.

(b) *The global minimizer to (2.7) is given by*

$$\begin{aligned} \mathbf{H}^\top \mathbf{H} &= \sqrt{\frac{\lambda_W}{\lambda_H}} \frac{(1 - \sqrt{nK\lambda})_+}{\sqrt{n}} \mathbf{Y}^\top \mathbf{Y} = \sqrt{\frac{\lambda_W}{\lambda_H}} \frac{(1 - \sqrt{nK\lambda})_+}{\sqrt{n}} \mathbf{I}_K \otimes \mathbf{J}_n, \\ \mathbf{W}^\top \mathbf{W} &= \sqrt{\frac{\lambda_H}{\lambda_W}} n^{1/2} (1 - \sqrt{nK\lambda})_+ \mathbf{I}_K, \quad \mathbf{W}^\top \mathbf{H} = (1 - \sqrt{nK\lambda})_+ \mathbf{Y} \end{aligned} \quad (2.10)$$

*i.e., the mean feature  $\bar{\mathbf{H}} \in \mathbb{R}^{D \times K}$  is nonnegative and satisfies*

$$\bar{\mathbf{H}}^\top \bar{\mathbf{H}} = \sqrt{\frac{\lambda_W}{\lambda_H}} \frac{(1 - \sqrt{nK\lambda})_+}{\sqrt{n}} \mathbf{I}_K, \quad \mathbf{W}^\top \bar{\mathbf{H}} = (1 - \sqrt{nK\lambda})_+ \mathbf{I}_K, \quad \mathbf{W} = \sqrt{\frac{n\lambda_H}{\lambda_W}} \bar{\mathbf{H}} \quad (2.11)$$

where  $\lambda = \sqrt{\lambda_W \lambda_H}$  and  $\mathbf{Y}$  is defined in (2.2). Here (2.10) and (2.11) correspond to  $\mathcal{NC}_1$  and  $\mathcal{NC}_{2-3}$  of the UPFM under  $\ell_2$ -loss.

For part (a), the same result was obtained in Dang et al. (2024) for the unconstrained positive feature model under imbalanced datasets. The proof of Theorem 2(a) is provided in Section 4.1. Our approach relies on a convex relaxation of (2.6), where we derive the solution via the KKT conditions and show that the collapsed solution is in fact the global minimizer of the relaxation. Unlike the argument in Dang et al. (2024), our justification is conceptually different and can be easily extended to the imbalanced case; for this reason, we include a full proof here. For part (b), we impose the additional constraint  $\mathbf{H} \geq 0$  compared with Tirer and Bruna (2022). Using a relaxation argument, it suffices to identify a nonnegative  $\mathbf{H}$  within the original solution set. For completeness, we provide a short proof. From Theorem 2, it follows directly that the global minimizer exhibits both within-class variability collapse and the convergence of class means to an orthogonal frame under the unconstrained positive feature model.

Once we have a full characterization of the global minimizer to (2.6) and (2.7), we will study whether the global minimizer exhibits the  $\mathcal{NC}$  property for a given dataset  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$  and  $N = Kn$ . We begin with a two-layer neural network with ReLU activation, and the corresponding feature is exactly

$$\mathbf{H} = \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) \in \mathbb{R}^{D \times N}$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times D}$ , and  $d$  and  $D$  are the dimensions of the input data and output features respectively. Assume  $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_K] \in \mathbb{R}^{d \times Kn}$  with  $\mathbf{X}_k = [\mathbf{x}_{k1}, \dots, \mathbf{x}_{kn}] \in \mathbb{R}^{d \times n}$ , i.e.,  $\mathbf{X}_k$  is the  $k$ -th class consisting of  $n$  points in  $\mathbb{R}^d$ .

To see if the neural collapse occurs, Theorem 2 implies that we need to find out the existence of  $\mathbf{W}_1$  such that  $\mathbf{H} = \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) = \bar{\mathbf{H}}\mathbf{Y}$  with  $\bar{\mathbf{H}} \geq 0$  and  $\bar{\mathbf{H}}^\top \bar{\mathbf{H}} \propto \mathbf{I}_K$ . Note that for any nonnegative mean feature matrix  $\bar{\mathbf{H}}$  with orthogonal columns, it satisfies

$$\text{supp}(\bar{\mathbf{h}}_k) \cap \text{supp}(\bar{\mathbf{h}}_\ell) = \emptyset.$$

where  $\bar{\mathbf{h}}_k \geq 0$  and  $\|\bar{\mathbf{h}}_k\|$  is constant over  $1 \leq k \leq K$ . In other words, each row of  $\bar{\mathbf{H}} \in \mathbb{R}^{D \times K}$  is a one-hot vector in  $\mathbb{R}^K$  multiplied by a nonnegative scalar. This reduces to show the existence of a vector  $\beta_k \in \mathbb{R}^d$  such that

$$\mathbf{X}_k^\top \beta_k = \mathbf{1}_n, \quad \mathbf{X}_{k'}^\top \beta_k \leq 0, \quad \forall 1 \leq k' \neq k \leq K, \quad (2.12)$$

for each  $1 \leq k \leq K$ , which is a linear feasibility problem. The answer depends on  $K$ ,  $d$  and  $n$ .

Our first result is quite general and applies to any dataset, and it concerns whether the global minimizer to (2.4) contains the  $\mathcal{NC}$  configuration, i.e., whether there exists  $\mathbf{W}_1$  such that  $\mathbf{H} = \sigma(\mathbf{W}_1^\top \mathbf{X}) = \bar{\mathbf{H}} \otimes \mathbf{1}_n^\top$  for some  $\bar{\mathbf{H}} \in \mathbb{R}_+^{D \times K}$ .

**Theorem 3 (Neural collapse for general datasets).** *For any general dataset  $\mathbf{X}$ , we have the following results:*

- (a) *Suppose  $\mathbf{X}$  is not linearly separable, then the neural collapse does not occur.*
- (b) *Suppose  $d \geq Kn$  and moreover  $\mathbf{X}^\top \mathbf{X}$  is of rank  $Kn$ , and then (2.12) satisfies for any  $1 \leq k \leq K$ .*

(c) Suppose  $d < n$  and  $\mathbf{X}_k^\top$  does not contain  $\mathbf{1}_n$  in its range, then the neural collapse will not occur.

Theorem 3 specifies conditions on  $d$  and  $n$  under which the occurrence of  $\mathcal{NC}$  is deterministic. In particular, case (b) corresponds to the setting where the data are linearly separable—an assumption commonly adopted in the deep learning literature to simplify analysis (e.g., Wang et al. (2023)). In this case, even a linear classifier suffices to induce collapsed solutions. Another interesting implication of Theorem 3 is that even if the neural network is extremely wide,  $\mathcal{NC}$  may not happen if the input dimension  $d$  is too small, which answers Question 1. The proof is very simple and thus we present it here.

**Proof of Theorem 3.** For (a), we prove it by contradiction. Suppose the neural collapse occurs, then

$$\mathbf{X}_k^\top \boldsymbol{\beta}_k = \mathbf{1}_n, \quad \mathbf{X}_{k'}^\top \boldsymbol{\beta}_k \leq 0, \quad \forall k' \neq k.$$

It means the  $k$ -th cluster  $\mathbf{X}_k$  is separated from the rest by the hyperplane  $\{\mathbf{x} : \langle \mathbf{x}, \boldsymbol{\beta}_k \rangle = 1/2\}$  for all  $k$ . Therefore, if the linear separability fails to hold, then the neural collapse will not occur. For (b), suppose  $\mathbf{X}$  is of rank  $Kn$ , then the range of  $\mathbf{X}$  is of dimension  $Kn$  and definitely contains a vector such that (2.12) holds. For (c), suppose  $\mathbf{X}_k^\top$  does not  $\mathbf{1}_n$  in its column space, then the first equality in (2.12) cannot hold. ■

## 2.2 $\mathcal{NC}$ and SNR

Therefore, the more interesting scenario is  $n < d \leq Kn$  as when the input data dimension is greater than the data size, i.e.,  $d > Kn$ , it is very likely that the  $\mathcal{NC}$  will occur due to the linear feasibility in (2.12). To have a concrete discussion on the  $\mathcal{NC}$  with  $d < Kn$ , we consider the Gaussian mixture model with  $K$  classes:

$$\mathbf{X}_k = \boldsymbol{\mu}_k \mathbf{1}_n^\top + \sigma \mathbf{Z}_k \in \mathbb{R}^{d \times n}.$$

We start with a two-layer neural network with a Gaussian mixture model (GMM) of two clusters and then extend to the  $K$ -class scenario. The goal is to see how the signal-to-noise ratio (SNR) in the data affects the  $\mathcal{NC}$ . The data are assumed in the following form:

$$\mathbf{X}^\top := \begin{bmatrix} \mathbf{X}_1^\top \\ \mathbf{X}_2^\top \end{bmatrix} := \begin{bmatrix} \mathbf{1}_n \boldsymbol{\mu}_1^\top + \sigma \mathbf{Z}_1 \\ \mathbf{1}_n \boldsymbol{\mu}_2^\top + \sigma \mathbf{Z}_2 \end{bmatrix} \in \mathbb{R}^{2n \times d}$$

where  $\boldsymbol{\mu}_1$  and  $\boldsymbol{\mu}_2$  are the cluster centers, and each cluster consists of  $n$  samples in  $\mathbb{R}^d$ .

**Theorem 4 (Neural collapse for GMM with two clusters).** *Let  $0 < \varepsilon < 1$ ,  $2n \geq d > n + C\varepsilon^{-2} \log n$ , and  $\theta$  be the angle between  $\boldsymbol{\mu}_1$  and  $\boldsymbol{\mu}_2$ . With probability at least  $1 - O(n^{-1})$ , (2.12) is feasible, i.e., the global minimizer to the empirical risk is given by the  $\mathcal{NC}$  configuration under either scenario below:*

- For  $\cos \theta < -4\varepsilon/(1 - \varepsilon)^2$ , i.e.,  $\theta > \pi - \arccos(4\varepsilon/(1 - \varepsilon)^2)$ , then

$$\sigma \lesssim (1 - \varepsilon) \sqrt{\frac{d - n}{d \log n}} \min\{\|\boldsymbol{\mu}_1\|, \|\boldsymbol{\mu}_2\|\}. \quad (2.13)$$

- For  $\cos \theta > -4\varepsilon/(1 - \varepsilon)^2$ , i.e.,  $\theta < \pi - \arccos(4\varepsilon/(1 - \varepsilon)^2)$ , then

$$\sigma \lesssim (1 - \varepsilon) \sqrt{\frac{d - n}{d \log n} \left( 1 - \left( |\cos \theta| + \frac{4\varepsilon}{(1 - \varepsilon)^2} \right)^2 \right) \min\{\|\boldsymbol{\mu}_1\|, \|\boldsymbol{\mu}_2\|\}}. \quad (2.14)$$

In particular, if  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_2 = \boldsymbol{\mu}$ , i.e.,  $\theta = \pi$ , then

$$\frac{\sigma}{\|\boldsymbol{\mu}\|} \lesssim (1 - \varepsilon) \sqrt{\frac{d - n}{d \log n}},$$

then neural collapse occurs with probability at least  $1 - O(n^{-1})$ . Suppose  $d \geq 2n$ , (2.12) is feasible with probability 1, i.e., the neural collapse occurs.

In other words, if  $\sigma$  is sufficiently small, then even if  $d < 2n$ , the neural collapse still occurs. Now we proceed to consider data that satisfy the GMM with  $K$  classes, i.e.

$$\mathbf{X}^\top := \begin{bmatrix} \mathbf{X}_1^\top \\ \vdots \\ \mathbf{X}_K^\top \end{bmatrix} := \begin{bmatrix} \mathbf{1}_n \boldsymbol{\mu}_1^\top + \sigma \mathbf{Z}_1 \\ \vdots \\ \mathbf{1}_n \boldsymbol{\mu}_K^\top + \sigma \mathbf{Z}_K \end{bmatrix} \in \mathbb{R}^{Kn \times d} \quad (2.15)$$

where  $\{\mathbf{Z}_k\}_{k=1}^K$  are  $K$  i.i.d.  $n \times d$  Gaussian random matrices.

We define the mean vector matrix by

$$\boldsymbol{\Pi} = [\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots, \boldsymbol{\mu}_K]^\top \in \mathbb{R}^{K \times d}.$$

Also, we assume  $\boldsymbol{\Pi}$  is a matrix of full row rank. If it is not full rank, we simply pick the maximum independent set of  $\{\boldsymbol{\mu}\}_{k=1}^K$ , and then the theorem above with  $K = 2$  generalizes to the  $K$ -class case.

**Theorem 5 (Neural collapse for GMM with  $K$  clusters).** *For the GMM with  $K$  clusters, the following holds.*

- (a) Suppose  $\boldsymbol{\Pi}$  is full rank,  $d - n \geq CK^2 \log n$ , and

$$\sigma \lesssim \sqrt{\frac{d - n}{d \log(Kn)}} \cdot \frac{\sigma_{\min}(\boldsymbol{\Pi})}{\sqrt{K - 1}}, \quad (2.16)$$

then there exists a global minimizer to the empirical risk minimizer which is given by the  $\mathcal{NC}$  configuration with probability at least  $1 - O(n^{-1})$ .

- (b) Suppose

$$\frac{d}{n} \geq \frac{K + 1}{2} + 2\sqrt{\frac{(K - 1) \log n}{n}} + \frac{K + 2 \log n}{n}, \quad (2.17)$$

then the neural collapse occurs with a high probability of least  $1 - O(n^{-1})$ .

- (c) Suppose  $d \geq Kn$ , then the neural collapse occurs with probability one since  $\mathbf{X}$  is of rank  $Kn$ .

Theorem 4 and 5(a) provide an answer to Question 2. If the data has a cluster structure, i.e.,  $\sigma$  is small compared with  $\|\boldsymbol{\mu}\|$ , the  $\mathcal{NC}$  is more likely to occur. The proof follows from constructing a vector that satisfies (2.12) with high probability by using the union bound. Theorem 5(b) implies that for the GMM,  $d \geq Kn$  is not necessary to guarantee the emergence of the  $\mathcal{NC}$  and in fact  $d \geq (K + 1)n/2$  suffices. The main technique is to first reformulate the feasibility of (2.12) as a problem of finding the maximum of a Gaussian process. Then we apply Gordon’s bound to find a sharp upper bound of the Gaussian process and obtain Theorem 5(b). The proofs of Theorem 4 and 5 are provided in Section 4.2.

Building on GMM models, Theorems 4 and 5 establish precise statistical relationships among data SNR, sample size, and dimension, which cannot be derived under UFM settings. While the GMM model is admittedly specific, it illustrates the broader principle that  $\mathcal{NC}$  emerges more readily in data with higher SNR. Furthermore, the results of these theorems extend naturally to unbalanced datasets, since the proofs do not rely on class balance. The only change is that the global minimizer of the unconstrained positive feature model shifts slightly, depending on the relative sizes of the classes.

Note that Theorem 3 provides a negative answer to Question 1 for two-layer neural networks. However, things become interesting if we turn to three-layer neural networks with the feature map matrix equal to

$$\mathbf{H}(\mathbf{X}) = \sigma_{\text{ReLU}}\left(\mathbf{W}_2^\top \sigma_{\text{ReLU}}\left(\mathbf{W}_1^\top \mathbf{X}\right)\right) \quad (2.18)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times d_1}$  and  $\mathbf{W}_2 \in \mathbb{R}^{d_1 \times D}$ , and  $\mathbf{X} \in \mathbb{R}^{d \times N}$  represents a dataset of  $N$  points. The question of whether there exists  $\mathbf{W}_1$  and  $\mathbf{W}_2$  such that  $\mathbf{H}(\mathbf{X}) = \mathbf{H} \otimes \mathbf{1}_n^\top$  can be simplified to a two-layer scenario: it suffices to consider a random feature model: each  $\mathbf{W}_1 \in \mathbb{R}^{d \times d_1}$  is assumed to be i.i.d.  $\mathcal{N}(0, 1/d_1)$ . In other words, for a given input  $\mathbf{x} \in \mathbb{R}^d$ , the output through the first layer  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{x}) \in \mathbb{R}^{d_1}$  exactly satisfies truncated normal distribution.

Based on Theorem 3, to induce the  $\mathcal{NC}$ , it suffices to ensure that  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  is of rank  $N = Kn$  for a sufficiently large  $d_1$ , which is guaranteed by the following theorem.

**Theorem 6 (Neural collapse for three layer random feature network).** *Suppose  $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  is a given dataset with any pair of them non-parallel and  $\|\mathbf{x}_i\| = 1$ . Given a three-layer network with feature map (2.18) and  $\mathbf{W}_1$  is i.i.d.  $\mathcal{N}(0, 1/d_1)$ , if we have*

$$d_1 \gtrsim \frac{\|\mathbf{X}\|^4}{\lambda_{\min}^2(\mathbf{H}_\infty)} \cdot N \log N$$

where  $\mathbf{H}_\infty$  is the kernel matrix that equals:

$$[\mathbf{H}_\infty]_{ij} = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_d)} \left( \sigma_{\text{ReLU}}(\langle \mathbf{z}, \mathbf{x}_i \rangle) - \sqrt{\frac{2}{\pi}} \right) \left( \sigma_{\text{ReLU}}(\langle \mathbf{z}, \mathbf{x}_j \rangle) - \sqrt{\frac{2}{\pi}} \right).$$

Then, with probability at least  $1 - O(N^{-2})$ , the nonnegative matrix  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) \in \mathbb{R}^{d_1 \times N}$  is of full column rank. Hence, Theorem (3)(b) implies that neural collapse occurs.

In other words, Question 1 has a positive answer (the  $\mathcal{NC}$  will occur) as long as the neural network has a depth of more than 2 and the first layer is sufficiently wide with width  $d_1 \gtrsim N \log N$ . The proof of Theorem (6) consists of two main steps. First, we prove that

$\mathbf{H}_\infty$  is positive definite, i.e.,  $\lambda_{\min}(\mathbf{H}_\infty) > 0$ . Then, we apply concentration inequalities to control the spectral deviation between  $d_1^{-1}(\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}))^\top \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  and  $\mathbf{H}_\infty$  to prove  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  is full rank. The details are presented in Section 4.3.

Note that in Theorem 6, the random feature model serves to guarantee both a large ambient dimension and the presence of non-parallel pairs after the first layer. This condition should be naturally satisfied in a trained network with a comparable number of parameters, provided the dataset does not contain parallel training samples across different classes. Thus, the random feature model is used only as a constructive device for the proof; the results are expected to extend to general three-layer networks with a first layer of comparable width.

### 2.3 $\mathcal{NC}$ and generalization

This section is devoted to Question 3: if  $\mathcal{NC}$  occurs, does it necessarily imply good generalization? To understand this question, we consider the task of binary classification for a two-layer ReLU network. In particular, throughout the discussion in this section, we will focus on the best misclassification error a simplified two-neuron classifier can achieve in the presence of the  $\mathcal{NC}$ . We need to make some preparations before proceeding to our main results.

**Data model:** For the data generative model, we assume they are sampled from GMM with two clusters and the mean vectors are opposite:

$$\mathbf{x} = \xi \boldsymbol{\mu} + \sigma \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_d) \quad (2.19)$$

where  $\xi$  is a Rademacher random variable and the training data are

$$\mathbf{X}^\top = \begin{bmatrix} \mathbf{1}_n \boldsymbol{\mu}^\top \\ -\mathbf{1}_n \boldsymbol{\mu}^\top \end{bmatrix} + \sigma \mathbf{Z} \in \mathbb{R}^{2n \times d} \quad (2.20)$$

where  $\pm \boldsymbol{\mu} \in \mathbb{R}^d$  denotes the class mean and  $\mathbf{Z} \in \mathbb{R}^{2n \times d}$  is a Gaussian random matrix.

**Model under the  $\mathcal{NC}$ :** We now consider the generalization ability for a trained model that exhibits neural collapse:

$$f_\theta(\mathbf{x}) = \mathbf{W}^\top \sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{x}) \quad (2.21)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times D}$  and  $\mathbf{W} \in \mathbb{R}^{D \times 2}$ . We denote the  $i$ -th column of  $\mathbf{W}_1$  by  $\boldsymbol{\beta}_i$ ,  $1 \leq i \leq D$  and the  $k$ -th column of  $\mathbf{W}$  by  $\boldsymbol{\alpha}^k$ ,  $1 \leq k \leq 2$ . We first try to rewrite (2.21) by exploiting the information of  $\mathcal{NC}$ . Under the feature collapse  $\mathcal{NC}_1$ , we have

$$\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{x}_{ki}) = \bar{\mathbf{h}}_k, \quad \forall 1 \leq k \leq 2, 1 \leq i \leq n,$$

where  $\mathbf{x}_{ki}$  is the  $i$ -th sample in the  $k$ -th class and  $\bar{\mathbf{h}}_k \in \mathbb{R}_+^D$  is the mean feature for the class  $k$ . Based on Theorem 2(a), we have  $\langle \bar{\mathbf{h}}_1, \bar{\mathbf{h}}_2 \rangle = 0$  and  $\|\bar{\mathbf{h}}_1\| = \|\bar{\mathbf{h}}_2\|$ , and thus

$$\mathbf{W} = [\boldsymbol{\alpha}, -\boldsymbol{\alpha}], \quad \boldsymbol{\alpha} := \boldsymbol{\alpha}^1 = -\boldsymbol{\alpha}^2 \propto \bar{\mathbf{h}}_1 - \bar{\mathbf{h}}_2, \quad \text{supp}(\bar{\mathbf{h}}_1) \cap \text{supp}(\bar{\mathbf{h}}_2) = \emptyset \quad (2.22)$$

since  $\mathbf{W}^\top [\bar{\mathbf{h}}_1, \bar{\mathbf{h}}_2] \propto \mathbf{I}_2 - \mathbf{J}_2/2$  and  $\mathbf{W} \propto [\bar{\mathbf{h}}_1, \bar{\mathbf{h}}_2](\mathbf{I}_2 - \mathbf{J}_2/2) = [\bar{\mathbf{h}}_1 - \bar{\mathbf{h}}_2, \bar{\mathbf{h}}_2 - \bar{\mathbf{h}}_1] \in \mathbb{R}^{D \times 2}$ .

We denote  $S_1 = \{i : \mathbf{h}_{1i} \geq 0, \mathbf{h}_{2i} = 0, 1 \leq i \leq D\}$  and  $S_2 = \{1, \dots, D\} \setminus S_1$ , and then the classifier becomes

$$f_\theta(\mathbf{x}) = \sum_{i \in S_1} \alpha_i \sigma_{\text{ReLU}}(\langle \boldsymbol{\beta}_i, \mathbf{x} \rangle) - \sum_{i \in S_2} \alpha_i \sigma_{\text{ReLU}}(\langle \boldsymbol{\beta}_i, \mathbf{x} \rangle) \quad (2.23)$$

where  $\alpha_i > 0$  for  $i \in S_1$  and  $\alpha_i < 0$  otherwise. Suppose  $\mathbf{x}$  is sampled from the first class, i.e.,  $\mathbf{x} = \boldsymbol{\mu} + \sigma \mathbf{z}$ , then  $f_\theta(\mathbf{x})$  produces a correct classification if  $f_\theta(\boldsymbol{\mu} + \sigma \mathbf{z}) > 0$  and thus to upper bound the misclassification error, we need to control  $\Pr(f_\theta(\boldsymbol{\mu} + \sigma \mathbf{z}) < 0)$ . This quantity is too complicated to compute exactly because the actual  $\boldsymbol{\beta}_i$  is unknown and  $f_\theta(\cdot)$  involves the sum of  $D$  terms.

**Simplification to a two-neuron classifier:** We simplify it by considering the performance of a two-neuron classifier:

$$f(\mathbf{x}) = \sigma_{\text{ReLU}}(\boldsymbol{\beta}_1^\top \mathbf{x}) - \sigma_{\text{ReLU}}(\boldsymbol{\beta}_2^\top \mathbf{x}) \quad (2.24)$$

and study whether it is able to correctly classify a given data point. This simplification follows from two observations: (a) this two-neuron classifier is a reduced form of (2.23) by setting both  $|S_1| = |S_2| = 1$ ; (b) to achieve a good generalization performance on binary classification on the data sampled from GMM, it suffices to find one single hyperplane that is able to separate the two classes. Hence, understanding the generalization of this simple model sheds some light on the general  $D$ -neuron case.

Our analysis of the two-neuron classifier relies on our understanding on the output of a single neuron:  $g_\beta(\mathbf{x}) := \sigma_{\text{ReLU}}(\boldsymbol{\beta}^\top \mathbf{x})$  and consider

$$\mathcal{C}_1 := \{\boldsymbol{\beta} : \mathbf{X}_1^\top \boldsymbol{\beta} = \mathbf{1}_n, \mathbf{X}_2^\top \boldsymbol{\beta} \leq 0\}, \quad \mathcal{C}_2 := \{\boldsymbol{\beta} : \mathbf{X}_2^\top \boldsymbol{\beta} = \mathbf{1}_n, \mathbf{X}_1^\top \boldsymbol{\beta} \leq 0\}, \quad (2.25)$$

i.e., for any  $\boldsymbol{\beta} \in \mathcal{C}_1$ ,  $g_\beta(\mathbf{x})$  maps the first class of training data to 1 and the second class to 0 which means the feature is collapsed. A similar counterpart holds for  $\boldsymbol{\beta} \in \mathcal{C}_2$ . In other words, for  $\boldsymbol{\beta}_k \in \mathcal{C}_k$ , then  $f(\mathbf{x})$  in (2.24) achieves the feature variability collapse.

To estimate the misclassification of (2.24), we start with the misclassification of  $g_\beta(\mathbf{x})$  for  $\boldsymbol{\beta} \in \mathcal{C}_k$ . First note that for any  $\boldsymbol{\beta} \in \mathcal{C}_1$ , it holds that

$$\begin{aligned} & \frac{1}{2} \left[ \Pr(g_{\boldsymbol{\beta}_1}(\boldsymbol{\mu} + \sigma \mathbf{z}) \leq 0) + \Pr(g_{\boldsymbol{\beta}_1}(-\boldsymbol{\mu} + \sigma \mathbf{z}) > 0) \right] \\ &= \frac{1}{2} \left[ \Pr(\langle \boldsymbol{\mu} + \sigma \mathbf{z}, \boldsymbol{\beta}_1 \rangle \leq 0) + \Pr(\langle -\boldsymbol{\mu} + \sigma \mathbf{z}, \boldsymbol{\beta}_1 \rangle > 0) \right] \\ &= \frac{1}{2} \Pr_{z \sim \mathcal{N}(0,1)} \left( |z| \geq \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta}_1 \rangle}{\sigma \|\boldsymbol{\beta}_1\|} \right) = \Phi \left( -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta}_1 \rangle}{\sigma \|\boldsymbol{\beta}_1\|} \right) \end{aligned}$$

where  $\Phi(\cdot)$  is the c.d.f. of standard normal distribution. Similarly for  $\boldsymbol{\beta}_2 \in \mathcal{C}_2$ , then

$$\frac{1}{2} \left[ \Pr(g_{\boldsymbol{\beta}_2}(\boldsymbol{\mu} + \sigma \mathbf{z}) > 0) + \Pr(g_{\boldsymbol{\beta}_2}(-\boldsymbol{\mu} + \sigma \mathbf{z}) \leq 0) \right] = \frac{1}{2} \Pr_{z \sim \mathcal{N}(0,1)} \left( |z| \geq -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta}_2 \rangle}{\sigma \|\boldsymbol{\beta}_2\|} \right) = \Phi \left( \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta}_2 \rangle}{\sigma \|\boldsymbol{\beta}_2\|} \right).$$

Therefore, under the  $\mathcal{NC}$ , the best generalization and misclassification error of (2.24) is closely related to the global maximum of the following problems:

$$\max_{\boldsymbol{\beta} \in \mathcal{C}_1} \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\|\boldsymbol{\beta}\|} \quad \text{and} \quad \max_{\boldsymbol{\beta} \in \mathcal{C}_2} -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\|\boldsymbol{\beta}\|}. \quad (2.26)$$

The following theorems provide an estimation of (2.26) in two different regimes.

**Theorem 7.** *Consider the data sampled from (2.20) and the following statements hold true:*

(a) *Suppose  $d - n + 1 > C\varepsilon^{-2} \log n$  and*

$$\frac{\sigma}{\|\boldsymbol{\mu}\|} \lesssim (1 - \varepsilon) \sqrt{\frac{d - n + 1}{d \log n}}, \quad (2.27)$$

*then with probability at least  $1 - O(n^{-1})$  there exists  $\boldsymbol{\beta} \in \mathbb{R}^{d \times D}$  that can induce (2.25). Moreover, we have*

$$\min_{\boldsymbol{\beta} \in \mathcal{C}_1} \Phi \left( -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\sigma \|\boldsymbol{\beta}\|} \right) \lesssim n^{-2}, \quad \min_{\boldsymbol{\beta} \in \mathcal{C}_2} \Phi \left( \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\sigma \|\boldsymbol{\beta}\|} \right) \lesssim n^{-2}.$$

(b) *For  $d \geq 2n \log n$ , then the following holds*

$$\begin{aligned} & \min_{\boldsymbol{\beta} \in \mathcal{C}_1} \Phi \left( -\frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\sigma \|\boldsymbol{\beta}\|} \right), \quad \min_{\boldsymbol{\beta} \in \mathcal{C}_2} \Phi \left( \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\sigma \|\boldsymbol{\beta}\|} \right) \\ & \geq 1 - \Phi \left( \left( \frac{n}{2d} \left( \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}} + s^2 + 1 - (c_1 s^2 + c_2 s) \sqrt{\frac{\log n}{n}} \right) + s^2 \right)^{-1/2} \right), \end{aligned} \quad (2.28)$$

*where  $s = \sigma / \|\boldsymbol{\mu}\|$ , and  $c_1$  and  $c_2 > 0$ .*

Theorem 7(a) focuses on the estimation of (2.26) in the low-noise regime while Theorem 7(b) concerns the regime when  $d$  is sufficiently large and the  $\mathcal{NC}$  always occurs. Note that Theorem 7(a) does not rely on the simplified two-neuron model. In this case, we explicitly construct a parameter set that induces collapse while still achieving good generalization. To recover the original model, one can simply set the remaining parameters to zero, yielding an equivalent formulation. In contrast, part (b) establishes a lower bound on the misclassification rate. This requires searching over multiple values of  $\boldsymbol{\beta}$ , where the resulting objective function is nonlinear and nonsmooth. Hence, the simplification to the two-neuron setting is essential for this part.

Here we briefly describe the key steps of the proof. The key idea is to approximate the global maximum of (2.26). However, it is not straightforward to estimate it exactly. Therefore, we relax (2.26) by dropping the affine inequality constraints and solving it. Then we consider the sufficient conditions under which the global maximizer to the relaxed problem is also the global maximizer to (2.26). This leads to the proof and conclusion of Theorem 7. For Theorem 7(b) with  $d \geq 2n$ , it is different from Theorem 7(a) because Theorem 3 implies that the  $\mathcal{NC}_1$  occurs with probability 1. The feasibility of the  $\mathcal{NC}$  makes it possible to approximate the global maximizer to (2.26) directly by rewriting the affine constraints in  $\mathcal{C}_1$  and  $\mathcal{C}_2$ . The technical details are deferred to Section 4.4. With the theorem above, we can obtain a characterization of the generalization performance of a two-neuron classifier in presence of the  $\mathcal{NC}$ .

**Theorem 8 (Misclassification of a two-neuron classifier under the  $\mathcal{NC}$ ).** *Consider  $f(\mathbf{x}) = g_{\beta_1}(\mathbf{x}) - g_{\beta_2}(\mathbf{x})$  in (2.24) where  $\beta_k \in \mathcal{C}_k$  for  $k = 1, 2$ .*

- (a) Under the assumption of Theorem 7(a), there exists  $\beta_k \in \mathcal{C}_k$  such that the misclassification error of  $f(\mathbf{x})$  in (2.24) is bounded by  $O(n^{-2})$ ;
- (b) Under the assumption of Theorem 7(b), the misclassification error of  $f(\mathbf{x})$  in (2.24) for any  $\beta_k \in \mathcal{C}_k$  is at least

$$1 - \Phi \left( \left( \frac{n}{2d} \left( \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}} + s^2 + 1 - (c_1 s^2 + c_2 s) \sqrt{\frac{\log n}{n}} \right) + s^2 \right)^{-1/2} \right)$$

where  $s = \sigma/\|\boldsymbol{\mu}\|$  measures the noise level in (2.19).

The theorem above implies that in the low-noise regime and  $d - n + 1 \gtrsim \log n$ , there exists a two-neuron classifier that achieves the  $\mathcal{NC}$  and also enjoys an excellent generalization bound. For  $d \geq 2n \log n$  and also the  $\mathcal{NC}$  occurs, any two-neuron classifier that achieves the  $\mathcal{NC}$  does not have a small misclassification error if the noise level  $\sigma/\|\boldsymbol{\mu}\|$  is large. This provides a partial answer to Question 3.

**Proof of Theorem 8.** The misclassification error of a two-neuron classifier in (2.24) is given by

$$\frac{1}{2} \left[ \Pr(g_{\beta_1}(\boldsymbol{\mu} + \sigma \mathbf{z}) < g_{\beta_2}(\boldsymbol{\mu} + \sigma \mathbf{z})) + \Pr(g_{\beta_1}(-\boldsymbol{\mu} + \sigma \mathbf{z}) > g_{\beta_2}(-\boldsymbol{\mu} + \sigma \mathbf{z})) \right].$$

An upper bound of the error is given by

$$\begin{aligned} & \frac{1}{2} \left[ \Pr(g_{\beta_2}(\boldsymbol{\mu} + \sigma \mathbf{z}) > 0) + \Pr(g_{\beta_1}(-\boldsymbol{\mu} + \sigma \mathbf{z}) > 0) \right] \\ &= \frac{1}{2} \left[ \Pr(\sigma \langle \beta_1, \mathbf{z} \rangle > \langle \boldsymbol{\mu}, \beta_1 \rangle) + \Pr(\sigma \langle \beta_2, \mathbf{z} \rangle \geq -\langle \boldsymbol{\mu}, \beta_2 \rangle) \right] \\ &= \frac{1}{2} \left[ \Phi \left( -\frac{\langle \boldsymbol{\mu}, \beta_1 \rangle}{\|\beta_1\|} \right) + \Phi \left( \frac{\langle \boldsymbol{\mu}, \beta_2 \rangle}{\|\beta_2\|} \right) \right] \lesssim n^{-2}. \end{aligned}$$

Theorem 7(a) implies in the low noise regime, there exists  $\beta_k \in \mathcal{C}_k$  such that the corresponding two-neuron classifier enjoys a very small misclassification error.

On the other hand, a lower bound is given by

$$\begin{aligned} & \frac{1}{2} \left[ \Pr(g_{\beta_1}(\boldsymbol{\mu} + \sigma \mathbf{z}) < g_{\beta_2}(\boldsymbol{\mu} + \sigma \mathbf{z})) + \Pr(g_{\beta_1}(-\boldsymbol{\mu} + \sigma \mathbf{z}) > g_{\beta_2}(-\boldsymbol{\mu} + \sigma \mathbf{z})) \right] \\ & \geq \frac{1}{2} \left[ \Pr(g_{\beta_1}(\boldsymbol{\mu} + \sigma \mathbf{z}) \leq 0) + \Pr(g_{\beta_2}(-\boldsymbol{\mu} + \sigma \mathbf{z}) \leq 0) \right] \\ &= \frac{1}{2} \left[ \Pr(\langle \boldsymbol{\mu} + \sigma \mathbf{z}, \beta_1 \rangle \leq 0) + \Pr(\langle -\boldsymbol{\mu} + \sigma \mathbf{z}, \beta_2 \rangle \leq 0) \right] = \frac{1}{2} \left[ \Phi \left( -\frac{\langle \boldsymbol{\mu}, \beta_1 \rangle}{\sigma \|\beta_1\|} \right) + \Phi \left( \frac{\langle \boldsymbol{\mu}, \beta_2 \rangle}{\sigma \|\beta_2\|} \right) \right]. \end{aligned}$$

Theorem 7(b) implies that as noise level  $\sigma$  increases, even if  $d$  is large and  $\mathcal{NC}$  occurs (recall that if  $d > 2n$ , the  $\mathcal{NC}$  occurs), any  $\mathbf{W}_1$  that induces  $\mathcal{NC}$  will not achieve a misclassification error smaller than (2.28). For example, if  $s$  is large, i.e.,  $s \gtrsim \sqrt{n/d}$ , then (2.28) is roughly approximated by

$$1 - \Phi \left( \sqrt{\frac{d}{n}} \frac{1}{s^{3/2}} \right),$$

which is close to  $1/2$ . For any  $\beta_k \in \mathcal{C}_k$ , it holds that the misclassification error is at least  $1 - \Phi \left( \sqrt{\frac{d}{n}} \frac{1}{s^{3/2}} \right)$ . ■

## 2.4 Future Directions

We conclude this section by summarizing our findings and outlining directions for future work. In this paper, we have provided a rigorous study of the conditions under which two-layer and three-layer ReLU networks can induce neural collapse. We further presented a simplified generalization analysis for the collapsed model. Moving beyond UFM settings, our results connect  $\mathcal{NC}$  to fundamental data properties—signal-to-noise ratio (SNR), sample size, and feature dimension—which have not been systematically addressed in prior literature.

A key limitation of our work is the lack of algorithmic analysis. Because the loss landscape of nonlinear networks is inherently nonconvex, it is unclear whether standard optimization methods such as (stochastic) gradient descent will converge to an  $\mathcal{NC}$  state, even when such solutions exist. Current studies of  $\mathcal{NC}$  dynamics either impose strong assumptions (e.g., orthogonality of the dataset, reducing to linear models) Wang et al. (2023), or focus on regimes such as NTK and mean-field Jacot et al. (2024); Wu and Mondelli (2025). For two-layer nonlinear networks in particular, especially under mild over-parameterization or under-parameterization, convergence guarantees remain elusive. Understanding the algorithmic dynamics of  $\mathcal{NC}$  in shallow networks is therefore a natural and important direction for future work.

Another open question concerns the link between  $\mathcal{NC}$  and generalization. Our simplified analysis suggests that  $\mathcal{NC}$  may impair generalization in the presence of noisy data, but this remains far from fully understood. A promising direction is to design more refined models and to study how optimization dynamics can be guided toward solutions that combine the neural collapse with strong generalization guarantees.

## 3. Experiments

### 3.1 $\mathcal{NC}_1$ for two-layer neural networks

This subsection aims to verify the sufficient conditions for inducing the  $\mathcal{NC}$  derived in Theorem 4 and 5. For both theorems, we verify our bounds by sampling data from the GMM in (2.15). We solve the following linear programming to examine the feasibility of the  $\mathcal{NC}$ :

$$\text{Find } \beta \text{ such that } \left( \mathbf{1}_n \boldsymbol{\mu}_1^\top + \sigma \mathbf{Z}_1 \right) \beta = \mathbf{1}_n, \quad \left( \mathbf{1}_n \boldsymbol{\mu}_k^\top + \sigma \mathbf{Z}_k \right) \beta \leq 0, \quad (k \neq 1). \quad (3.1)$$

Note that here we only consider the linear feasibility problem regarding the first class because a similar conclusion can be drawn for other classes. For this test, the goal is to study the feasibility and its dependence on  $d, n, K$  and  $\sigma$ .

For  $K = 2$ , we let  $n = 300$  and  $d/n$  vary from 1 to 2. To simplify our discussion, we consider the mean vectors having the same norm  $\|\boldsymbol{\mu}_1\| = \|\boldsymbol{\mu}_2\| = \|\boldsymbol{\mu}\| = 1$  and denote their angle by  $\theta$ . For each pair of  $(d, \sigma)$ , we run 10 experiments and solve (3.1), and if the linear program is feasible, then we count it as a successful instance and we plot the successful rate of each set of parameters. From Figure 1, we can see that Theorem 4 and 5 does not exactly match the phase transition but provides a good approximation. For lower  $\sigma$  and larger  $\theta$ , then white regions are clearly larger, showing that well-clustered data are more likely to induce the  $\mathcal{NC}$ . For the case when the signal-to-noise ratio of the GMM model is

low, e.g., for  $\sigma$  is large or for  $\theta = 0$ , we can see  $d/n = 3/2 = (K + 1)/2$  provides an accurate description of the phase transition, which is guaranteed by the Gordon’s bound.

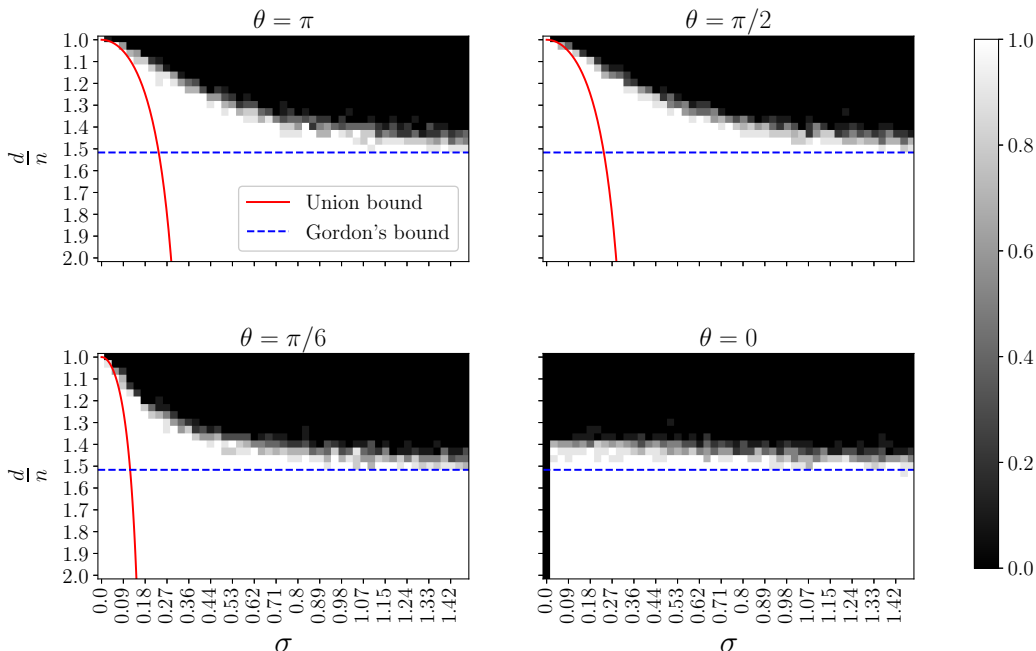


Figure 1: The  $\mathcal{NC}$  feasibility plot for a two-class GMM with mean vectors  $\mu_1$  and  $\mu_2$ , and  $\theta$  is their angle. The legends denote the sufficient conditions provided by Theorem 4 and 5 represented respectively by red solid line (union bound) and blue dashed line (Gordon’s bound).

For  $K \geq 2$ , the setup of experiments is similar to the  $K = 2$  case:  $K = 4, n = 250$ ,  $\mu_k = e_k$ , and  $d/n$  ranges from 1 to 4. For each  $(d, \sigma)$ , 10 experiments are carried out. Figure 2 shows a similar phase transition plot as in Figure 1. In the high and low  $\sigma$  regime, Gordon’s bound (blue dashed line) and union bound (red solid line) approximate the phase transition boundary respectively. Overall, our characterization given by Theorem 4 and 5 is not able to exactly capture the regime when  $d/n$  is larger than 1 and below  $(K + 1)/n$ . The further improvements of the bound will rely on a much more refined analysis to understand the feasibility of (3.1)

### 3.2 $\mathcal{NC}$ for two-layer neural network

Theorem 4 and 5 only provide an answer to whether a neural collapse configuration exists for a two-layer neural network. Note that the objective function (2.4) is non-convex for the training of two-layer neural networks with ReLU activation. This subsection is devoted to exploring whether such a network could converge to a neural collapse configuration by




Figure 2:  $\mathcal{NC}_1$  feasibility plot for GMM with four clusters and  $\boldsymbol{\mu}_k = \mathbf{e}_k$ ,  $1 \leq k \leq 4$ .

SGD. Similar to the previous subsection, we sample data with  $K = 2$  from GMM under different  $d$  and  $\sigma$  and train two-layer ReLU neural networks to do the classification task. We train each network for  $10^5$  epochs under the cross-entropy loss (2.4) with activation regularization: we use plain SGD starting with learning rate 0.1 and divide the rate by 10 at the 1/3 and 2/3 of the first epoch respectively and the regularization parameters are  $\lambda_W = 10^{-3}$  and  $\lambda_H = 10^{-6}$ . We fix  $\boldsymbol{\mu}_1 = -\boldsymbol{\mu}_2 = \mathbf{e}_1$  and let  $d/n \in [1.1, 1.3, 1.5, 2, 3, 4]$  and  $\sigma \in [0.18, 0.36, 0.53, 0.8, 1.07, 1.42]$  which correspond to the grid in Figure 1. We calculate the commonly adopted  $\mathcal{NC}_1$  metric to measure the degree of feature collapse for each network after training.

$$\mathcal{NC}_1 := \frac{1}{K} \text{Tr} \left( \boldsymbol{\Sigma}_W \boldsymbol{\Sigma}_B^\dagger \right), \quad (3.2)$$

where

$$\boldsymbol{\Sigma}_W := \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n (\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)(\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)^\top, \quad \boldsymbol{\Sigma}_B := \frac{1}{K} \sum_{k=1}^K (\bar{\mathbf{h}}_k - \mathbf{h}_G)(\bar{\mathbf{h}}_k - \mathbf{h}_G)^\top$$

and

$$\mathbf{h}_G := \frac{1}{K} \sum_{k=1}^K \bar{\mathbf{h}}_k, \quad \bar{\mathbf{h}}_k := \frac{1}{n} \sum_{i=1}^n \mathbf{h}_{ki}, \quad 1 \leq k \leq K.$$

We observe that when data is well clustered ( $\sigma = 0.18$ ), the  $\mathcal{NC}_1$  drops as the data dimension  $d$  increases. In particular, the  $\mathcal{NC}_1$ 's are all below  $10^{-4}$  except for the case when  $d/n = 1.1$ , which is consistent with our characterization in Figure 1 and under this noise level we find that large  $d$  facilitates the feature collapse. However, as  $\sigma$  increases, larger  $d$  does not yield to smaller  $\mathcal{NC}_1$  anymore. Especially, for cases where  $\sigma \geq 0.8$  (SNR is close or




Figure 3:  $\mathcal{NC}_1$  plot for two-layer ReLU networks after training for  $10^6$  epochs with varying  $d$  and variance  $\sigma$ . The legend denotes the variance  $\sigma$  in GMM models that generate data.

smaller than 1), we find the  $\mathcal{NC}_1$  almost all stay at the same level of magnitude. We think a higher  $\sigma$  makes the gradient less aligned among the SGD iterates. Also, increasing the dimension of the data means more neurons need to be aligned to induce feature collapse, which could result in slow convergence to neural collapse configurations or getting stuck at local minima. We verify this analysis by plotting the following quantities together with  $\mathcal{NC}_1$  along training epochs, which measure the convergence of mean and weight to neural collapse configurations (2.9) in Figures 4 and 5,

$$\begin{aligned}
 \mathcal{NC}_{2,\bar{\mathbf{H}}} &= \left\| \frac{\bar{\mathbf{H}}^\top \bar{\mathbf{H}}}{\|\bar{\mathbf{H}}^\top \bar{\mathbf{H}}\|_F} - \frac{1}{\sqrt{K}} \mathbf{I}_K \right\|_F, \\
 \mathcal{NC}_{2,\mathbf{W}} &= \left\| \frac{\mathbf{W}^\top \mathbf{W}}{\|\mathbf{W}^\top \mathbf{W}\|_F} - \frac{1}{\sqrt{K-1}} \mathbf{C}_K \right\|_F, \\
 \mathcal{NC}_3 &= \left\| \frac{\mathbf{W}^\top \bar{\mathbf{H}}}{\|\mathbf{W}^\top \bar{\mathbf{H}}\|_F} - \frac{1}{\sqrt{K-1}} \mathbf{C}_K \right\|_F,
 \end{aligned} \tag{3.3}$$

where  $\mathbf{C}_K = \mathbf{I}_K - \mathbf{J}_K/K$  is defined in (1.2)

Figures 4 and 5 respectively depict the change of  $\mathcal{NC}_{1-3}$  metrics of networks shown in Figure 3 when  $\sigma$  equals to 0.18 and 1.42. We observe that all the metrics converge for all the networks after  $10^5$  training epochs, which implies the stabilization of the training process. The major difference between the final trained parameters and the neural collapse configuration comes from  $\bar{\mathbf{H}}$  ( $\mathcal{NC}_{2,\bar{\mathbf{H}}}$ ) and the concentration of feature vectors around the

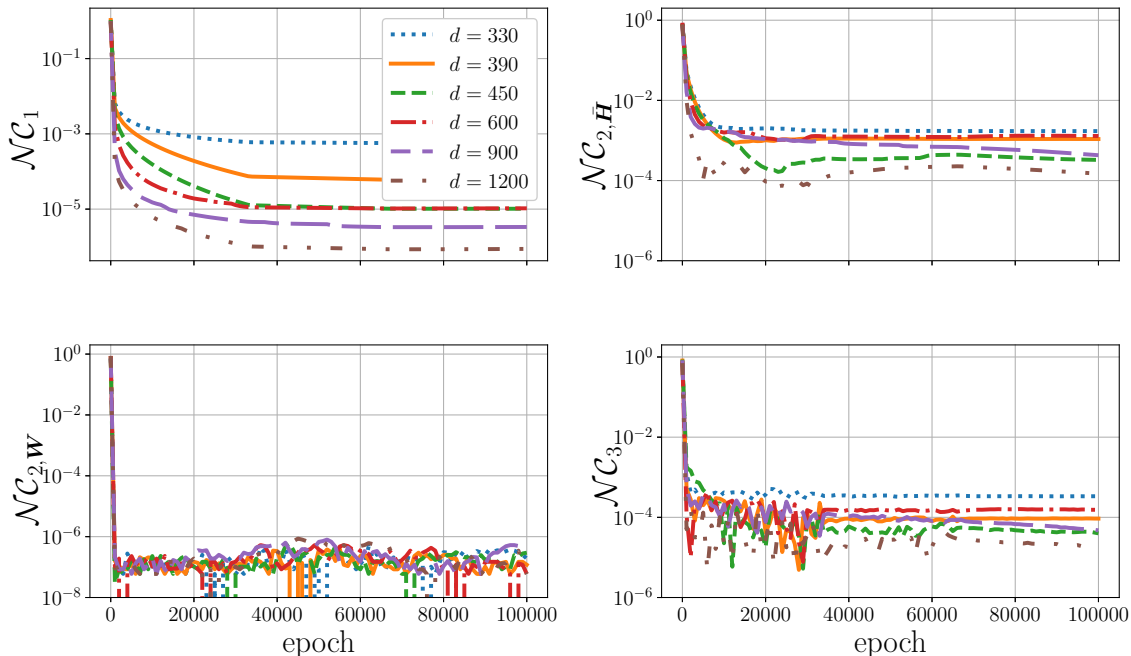


Figure 4:  $\mathcal{NC}_1$ - $\mathcal{NC}_3$  of two-layer networks along the training epochs in Figure 3 along the training epochs with  $\sigma = 0.18$ . The legend denotes the dimension  $d$  of the training data. Upper left plot:  $\mathcal{NC}_1$ . The left three plots record the alignment of mean feature matrix  $\bar{\mathbf{H}}$  and weight of classifier  $\mathbf{W}$  to orthogonal frame and ETF (described by (2.9)) respectively measured by the relative error (3.3) under Frobenius norm.

class mean feature vectors ( $\mathcal{NC}_1$ ) while  $\mathcal{NC}_{2,\mathbf{W}}$  drops to the level below  $10^{-6}$  for all the networks ( $\mathcal{NC}_{2,\mathbf{W}}$  actually becomes 0 under single-precision floating-point format at some points). Although adding dimension in general makes  $\mathcal{NC}_{2,\bar{\mathbf{H}}}$  smaller, it also makes the decay of  $\mathcal{NC}_1$  much harder in the case of  $\sigma = 1.42$ , as  $\mathcal{NC}_1$  quickly becomes stable for all networks in Figure 5. Additionally, by comparing Figures 4 and 5,  $\mathcal{NC}_1$  and  $\mathcal{NC}_{2,\bar{\mathbf{H}}}$  both become higher for  $\sigma = 1.42$ , which implies that the convergence becomes slower as the noise level rises. This observation leads us to conclude that when the GMM model has a low noise level, SGD could approach a neural collapse configuration as the global minimizer when it exists. However, when the noise level is high, it takes a very long time for the SGD to get close to a neural collapse configuration.

### 3.3 $\mathcal{NC}$ for three-layer neural network

As shown in Theorem 6, for a three-layer neural network with the first layer randomly initialized,  $\mathcal{NC}$  should emerge if the width of first layer  $d_1 \gtrsim N \log N$ :  $\mathbf{W}_1 \in \mathbb{R}^{d_1 \times d}$  to

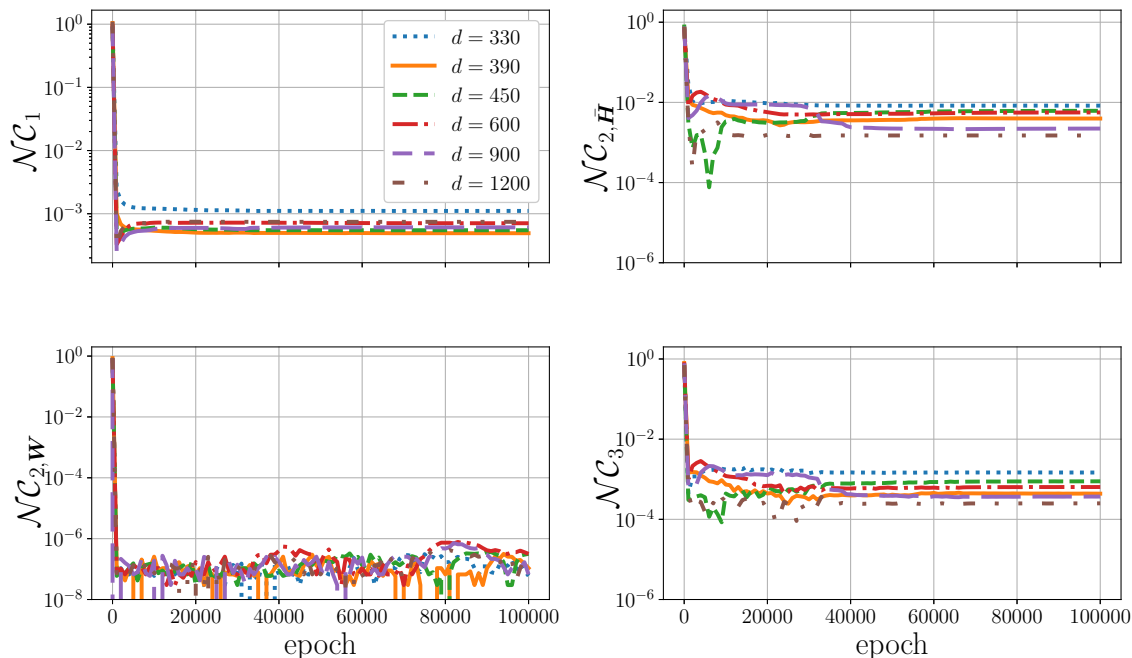


Figure 5:  $\mathcal{NC}_1$ - $\mathcal{NC}_3$  of two-layer networks along the training epochs in Figure 3 with  $\sigma = 1.42$ .

be Gaussian matrix containing i.i.d entries from  $\mathcal{N}(0, 1/\sqrt{d_1})$ . To better understand our theoretical result, we train three-layer neural networks with the weights on the first-layer ( $\mathbf{W}_1$ ) fixed. The dataset is FMnist containing  $n = 500$  samples from each class with  $K = 10$ .

We train each network under the regularized cross-entropy loss (2.5) with the same setting in terms of training epoch, regularization parameters, and stepsize schedule as those in the previous subsection. In Figure 6, we plot again the metrics  $\mathcal{NC}_{1-3}$  along the training epochs.

As  $d_1$  increases, the feature collapse and the convergence of  $\bar{\mathbf{H}}$  and  $\mathbf{W}$  to the  $\mathcal{NC}_2$  characterized in (2.9) are stronger in the terminal phase of training. The relative error proposed in (3.3) shrinks as the width  $d_1$  increases but it is still at the order of  $10^{-3}$  for  $d_1 = 2 \cdot 10^4$ . This implies despite the global optimality of  $\mathcal{NC}$  in the regularized ERM, it may still take much time for the actual training process to achieve  $\mathcal{NC}$ . Similar to cases presented in the previous subsection, the slow convergence is likely due to the nonconvex nature of the objective function and the random feature training. The gap between the experiments and theory exhibited in this section calls for further studies of the convergence of first-order iterative algorithms to the neural collapse configuration. Our code for all the experiments above is available on Github.




Figure 6:  $\mathcal{NC}_1$ - $\mathcal{NC}_3$  of three-layer networks along the training epoch. The legend denotes the width  $d_1$  of the first-layer weight.

## 4. Proofs

### 4.1 Unconstrained positive feature models

**Proof of Theorem 2(a).** Consider

$$R_N(\mathbf{W}, \mathbf{H}) = \frac{1}{N} \ell_{CE}(\mathbf{W}^\top \mathbf{H}, \mathbf{Y}) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2$$

subject to  $\mathbf{H} \geq 0$  and  $\mathbf{Y}$  is the label matrix defined in (2.2). This optimization problem is nonconvex. The proof idea is: we will find a convex relaxation of  $R_N(\mathbf{W}, \mathbf{H})$ , then find the corresponding global minimizer, and then prove that the solutions to the original nonconvex problem and convex relaxation are identical. Note that  $R_N(\mathbf{W}, \mathbf{H})$  only depends on  $\mathbf{Z} := \mathbf{W}^\top \mathbf{H}$ ,  $\mathbf{U} := \mathbf{W}^\top \mathbf{W}$ , and  $\mathbf{V} := \mathbf{H}^\top \mathbf{H}$ . Therefore, a convex relaxation is given by

$$\min \frac{1}{N} \ell_{CE}(\mathbf{Z}, \mathbf{Y}) + \frac{\lambda_W}{2} \text{Tr}(\mathbf{U}) + \frac{\lambda_H}{2} \text{Tr}(\mathbf{V})$$

subject to

$$\mathbf{Q} := \begin{bmatrix} \mathbf{U} & \mathbf{Z} \\ \mathbf{Z}^\top & \mathbf{V} \end{bmatrix} \succeq 0, \quad \mathbf{V} \geq 0$$

where  $\mathbf{U} \in \mathbb{R}^{K \times K}$  and  $\mathbf{V} \in \mathbb{R}^{N \times N}$ . We claim that

$$\mathbf{U} = \frac{a^2}{b} \mathbf{C}_K, \quad \mathbf{V} = b \mathbf{Y}^\top \mathbf{Y}, \quad \mathbf{Z} = a \mathbf{C}_K \mathbf{Y}$$

for some  $a > 0$  and  $b > 0$  where  $\mathbf{C}_K$  is the centering matrix in (1.2). Then it holds that

$$\mathbf{Q} = \begin{bmatrix} \mathbf{U} & \mathbf{Z} \\ \mathbf{Z}^\top & \mathbf{V} \end{bmatrix} = \begin{bmatrix} \frac{a^2}{b} \mathbf{C}_K & a \mathbf{C}_K \mathbf{Y} \\ a \mathbf{Y}^\top \mathbf{C}_K & b \mathbf{Y}^\top \mathbf{Y} \end{bmatrix} = \frac{1}{b} \begin{bmatrix} a \mathbf{C}_K \\ b \mathbf{Y}^\top \end{bmatrix} \begin{bmatrix} a \mathbf{C}_K \\ b \mathbf{Y}^\top \end{bmatrix}^\top \quad (4.1)$$

is exactly rank- $K$ .

The Lagrangian is

$$L(\mathbf{U}, \mathbf{V}, \mathbf{Z}, \mathbf{S}, \mathbf{B}) = \frac{1}{N} \ell_{CE}(\mathbf{Z}, \mathbf{Y}) + \frac{\lambda_W}{2} \text{Tr}(\mathbf{U}) + \frac{\lambda_H}{2} \text{Tr}(\mathbf{V}) - \left\langle \mathbf{S}, \begin{bmatrix} \mathbf{U} & \mathbf{Z} \\ \mathbf{Z}^\top & \mathbf{V} \end{bmatrix} \right\rangle - \frac{1}{2} \langle \mathbf{V}, \mathbf{B} \rangle$$

where  $\mathbf{B} \geq 0$  and  $\mathbf{S} \in \mathbb{R}^{(K+N) \times (K+N)}$  is positive semidefinite:

$$\mathbf{S} = \begin{bmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{bmatrix} \succeq 0.$$

The dimension of the blocks of  $\mathbf{S}$  matches that of  $\mathbf{U}$ ,  $\mathbf{V}$ , and  $\mathbf{Z}$ .

Setting  $\nabla_{\mathbf{U}, \mathbf{V}, \mathbf{Z}} L(\mathbf{U}, \mathbf{V}, \mathbf{Z}, \mathbf{S}, \mathbf{B}) = 0$  leads to

$$\lambda_W \mathbf{I}_K = 2\mathbf{S}_{11}, \quad \lambda_H \mathbf{I}_N = 2\mathbf{S}_{22} + \mathbf{B}, \quad 2\mathbf{S}_{12} = \frac{1}{N}(\mathbf{P} - \mathbf{Y})$$

where

$$\frac{\partial \ell_{CE}(\mathbf{Z}, \mathbf{Y})}{\partial z_{ki}} = \frac{1}{N}(\mathbf{p}_{ki} - \mathbf{e}_k), \quad \mathbf{p}_{ki} = \frac{\exp(z_{ki})}{\langle \exp(\mathbf{z}_{ki}), \mathbf{1}_K \rangle}.$$

We have the dual variable

$$\mathbf{S} = \frac{1}{2} \begin{bmatrix} \lambda_W \mathbf{I}_K & N^{-1}(\mathbf{P} - \mathbf{Y}) \\ N^{-1}(\mathbf{P} - \mathbf{Y})^\top & \lambda_H \mathbf{I}_N - \mathbf{B} \end{bmatrix}.$$

To ensure  $\mathbf{Q}$  is a global minimizer, then there exist  $\mathbf{S} \succeq 0$  and  $\mathbf{B} \geq 0$  such that the following KKT condition holds

$$\mathbf{Q} := \begin{bmatrix} \mathbf{U} & \mathbf{Z} \\ \mathbf{Z}^\top & \mathbf{V} \end{bmatrix} \succeq 0, \quad \mathbf{S}\mathbf{Q} = 0, \quad \langle \mathbf{B}, \mathbf{V} \rangle = 0. \quad (4.2)$$

Under  $\mathbf{Z} = a \mathbf{C}_K \mathbf{Y}$ , we have

$$\mathbf{P} = \bar{\mathbf{P}}\mathbf{Y} = \bar{\mathbf{P}} \otimes \mathbf{1}_n^\top, \quad \bar{\mathbf{P}} = \frac{\mathbf{J}_K + (e^a - 1)\mathbf{I}_K}{K - 1 + e^a}$$

and

$$\mathbf{I}_K - \bar{\mathbf{P}} = \frac{K\mathbf{I}_K - \mathbf{J}_K}{K - 1 + e^a} = \frac{K}{K - 1 + e^a}(\mathbf{I}_K - \mathbf{J}_K/K).$$

As a result, we have

$$\mathbf{S} = \frac{1}{2} \begin{bmatrix} \lambda_W \mathbf{I}_K & -\frac{K}{N(K-1+e^a)} \mathbf{C}_K \mathbf{Y} \\ -\frac{K}{N(K-1+e^a)} \mathbf{Y}^\top \mathbf{C}_K & \lambda_H \mathbf{I}_N - \mathbf{B} \end{bmatrix}.$$

Moreover, we choose

$$\mathbf{B} = t(\mathbf{J}_N - \mathbf{I}_K \otimes \mathbf{J}_n) = t(\mathbf{J}_K - \mathbf{I}_K) \otimes \mathbf{J}_n. \quad (4.3)$$

It remains to verify that (4.2) holds for some positive  $a$ ,  $b$ , and  $t$ . Note that  $\mathbf{Q} \succeq 0$  and  $\mathbf{B} \geq 0$  follow from its construction (4.1). For  $\langle \mathbf{B}, \mathbf{V} \rangle$ , we have

$$\langle \mathbf{B}, \mathbf{V} \rangle = bt \langle \mathbf{J}_N - \mathbf{I}_K \otimes \mathbf{J}_n, \mathbf{Y}^\top \mathbf{Y} \rangle = bt (\langle \mathbf{J}_N, \mathbf{I}_K \otimes \mathbf{J}_n \rangle - \|\mathbf{I}_K \otimes \mathbf{J}_n\|_F^2) = 0, \quad \forall b \geq 0, t \geq 0$$

where  $\mathbf{Y}^\top \mathbf{Y} = (\mathbf{I}_K \otimes \mathbf{1}_n)(\mathbf{I}_K \otimes \mathbf{1}_n^\top) = \mathbf{I}_K \otimes \mathbf{J}_n$ . It suffices to ensure  $\mathbf{S} \succeq 0$  and  $\mathbf{S}\mathbf{Q} = 0$ . Due to the factorization of  $\mathbf{Q}$  in (4.1), we have

$$\mathbf{S}\mathbf{Q} = 0 \iff \begin{bmatrix} \lambda_W \mathbf{I}_K & -\frac{K}{N(K-1+e^a)} \mathbf{C}_K \mathbf{Y} \\ -\frac{K}{N(K-1+e^a)} \mathbf{Y}^\top \mathbf{C}_K & \lambda_H \mathbf{I}_N - \mathbf{B} \end{bmatrix} \begin{bmatrix} a \mathbf{C}_K \\ b \mathbf{Y}^\top \end{bmatrix} = 0,$$

which is equivalent to

$$\begin{aligned} 0 &= a \lambda_W \mathbf{C}_K - \frac{b \mathbf{C}_K}{K-1+e^a}, \\ 0 &= -\frac{K a}{N(K-1+e^a)} \mathbf{Y}^\top \mathbf{C}_K + b \lambda_H \mathbf{Y}^\top - b \mathbf{B} \mathbf{Y}^\top \\ &= -\frac{K a}{N(K-1+e^a)} (\mathbf{Y}^\top - K^{-1} \mathbf{J}_{N \times K}) + b \lambda_H \mathbf{Y}^\top - b n t (\mathbf{J}_{N \times K} - \mathbf{Y}^\top) \end{aligned}$$

where  $\mathbf{Y} = \mathbf{I}_K \otimes \mathbf{1}_n^\top$ ,  $\mathbf{Y} \mathbf{Y}^\top = n \mathbf{I}_K = (N/K) \mathbf{I}_K$  and

$$t^{-1} \mathbf{B} \mathbf{Y}^\top = (\mathbf{J}_N - \mathbf{I}_K \otimes \mathbf{J}_n)(\mathbf{I}_K \otimes \mathbf{1}_n) = n \mathbf{J}_{N \times K} - n \mathbf{I}_K \otimes \mathbf{1}_n = n(\mathbf{J}_{N \times K} - \mathbf{Y}^\top).$$

To make  $\mathbf{S}\mathbf{Q} = 0$  hold, we need to have

$$\begin{aligned} \frac{b}{K-1+e^a} &= a \lambda_W, \\ \frac{K a}{N(K-1+e^a)} &= b \lambda_H + b n t, \\ \frac{a}{N(K-1+e^a)} &= b n t \end{aligned}$$

where  $n = N/K$ . Then the second and third equations determine  $t$ :

$$K = \frac{\lambda_H}{n t} + 1 \iff \lambda_H + n t - n t K = 0 \iff t = \frac{\lambda_H}{n(K-1)}. \quad (4.4)$$

The coefficients  $a$  and  $b$  are the solution to

$$\frac{b}{K-1+e^a} = a \lambda_W, \quad \frac{a}{N(K-1+e^a)} = \frac{b \lambda_H}{K-1} \implies \frac{a}{b} = \sqrt{\frac{nK}{K-1}} \cdot \frac{\lambda_H}{\lambda_W}. \quad (4.5)$$

Then

$$a = \log \left( (K-1) \left[ \sqrt{\frac{1}{nK(K-1)} \cdot \frac{1}{\lambda_H \lambda_W}} - 1 \right] \right), \quad b = \sqrt{\frac{K-1}{nK}} \cdot \frac{\lambda_W}{\lambda_H} a.$$

Finally, we verify  $\mathbf{S} \succeq 0$ :

$$\mathbf{S} = \frac{1}{2} \begin{bmatrix} \lambda_W \mathbf{I}_K & -\frac{K}{N(K-1+e^a)} \mathbf{C}_K \mathbf{Y} \\ -\frac{K}{N(K-1+e^a)} \mathbf{Y}^\top \mathbf{C}_K & \lambda_H \mathbf{I}_N - \frac{\lambda_H}{n(K-1)} (\mathbf{J}_K - \mathbf{I}_K) \otimes \mathbf{J}_n \end{bmatrix}$$

where

$$\frac{K}{N(K-1+e^a)} = \frac{bnKt}{a} = Kn \cdot \sqrt{\frac{K-1}{nK} \cdot \frac{\lambda_W}{\lambda_H}} \cdot \frac{\lambda_H}{n(K-1)} = \sqrt{\frac{K\lambda_W\lambda_H}{n(K-1)}}.$$

Note that  $\mathbf{S}_{11} = \frac{\lambda_W}{2} \mathbf{I}_K \succ 0$ , and thus  $\mathbf{S} \succeq 0$  equals  $\mathbf{S}_{22} - \mathbf{S}_{21} \mathbf{S}_{11}^{-1} \mathbf{S}_{12} \succeq 0$ , i.e.,

$$\begin{aligned} 2(\mathbf{S}_{22} - \mathbf{S}_{21} \mathbf{S}_{11}^{-1} \mathbf{S}_{12}) &= \lambda_H \mathbf{I}_N - \frac{\lambda_H}{n(K-1)} (\mathbf{J}_K - \mathbf{I}_K) \otimes \mathbf{J}_n - \frac{K^2}{N^2(K-1+e^a)^2 \lambda_W} \mathbf{Y}^\top \mathbf{C}_K \mathbf{Y} \\ &= \lambda_H \mathbf{I}_N - \frac{\lambda_H}{n(K-1)} (\mathbf{J}_K - \mathbf{I}_K) \otimes \mathbf{J}_n - \frac{K\lambda_H}{n(K-1)} \left( \mathbf{I}_K \otimes \mathbf{J}_n - \frac{\mathbf{J}_N}{K} \right) \\ &= \lambda_H \mathbf{I}_N - \frac{\lambda_H}{n} (\mathbf{I}_K \otimes \mathbf{J}_n) \succeq 0 \end{aligned}$$

where

$$\mathbf{Y}^\top \mathbf{C}_K \mathbf{Y} = \mathbf{Y}^\top \mathbf{Y} - \frac{1}{K} \mathbf{Y}^\top \mathbf{J}_K \mathbf{Y} = \mathbf{I}_K \otimes \mathbf{J}_n - \frac{\mathbf{J}_N}{K}.$$

To summarize, we have verified that (4.1) with coefficients (4.4) and (4.5) is the unique global minimizer where the uniqueness comes from the strict complementary slackness. Therefore, the global minimizer  $(\mathbf{W}, \mathbf{H})$  is given by

$$\mathbf{W}^\top \mathbf{W} = \frac{a^2}{b} \mathbf{C}_K, \mathbf{H}^\top \mathbf{H} = b \mathbf{Y}^\top \mathbf{Y} \implies \mathbf{W} = \frac{a}{\sqrt{b}} \mathbf{C}_K, \mathbf{H} = \sqrt{b} \mathbf{Y}$$

where  $a$  is strictly positive if

$$(K-1) \left[ \sqrt{\frac{1}{nK(K-1)} \cdot \frac{1}{\lambda_H \lambda_W}} - 1 \right] > 1 \iff \sqrt{\frac{K-1}{nK} \cdot \frac{1}{\lambda_H \lambda_W}} > K \iff \sqrt{\frac{n(K-1)}{K}} > N\lambda_Z. \quad \blacksquare$$

**Proof of Theorem 2(b).** Under  $\ell_2$ -loss function and unconstrained positive feature model, the empirical risk function is

$$R_N(\mathbf{H}_+, \mathbf{W}) = \frac{1}{2N} \|\mathbf{W}^\top \mathbf{H}_+ - \mathbf{Y}\|_F^2 + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}_+\|_F^2$$

where  $\mathbf{H}_+$  is a nonnegative matrix. We note that

$$\min_{\mathbf{W}, \mathbf{H}} R_N(\mathbf{W}, \mathbf{H}) \leq \min_{\mathbf{W}, \mathbf{H}_+} R_N(\mathbf{W}, \mathbf{H}_+).$$

Without any constraint on  $\mathbf{H}_+$ , the minimization of  $R_N(\mathbf{H}, \mathbf{W})$  is exactly singular value thresholding. Let  $\mathbf{Z} = \mathbf{W}^\top \mathbf{H}$ , and then

$$\min_{\mathbf{W}^\top \mathbf{H} = \mathbf{Z}} \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 = \sqrt{\lambda_W \lambda_H} \|\mathbf{Z}\|_*.$$

Thus  $R_N(\mathbf{W}, \mathbf{H})$  equals

$$\min_{\mathbf{Z}} \frac{1}{2N} \|\mathbf{Z} - \mathbf{Y}\|_F^2 + \lambda \|\mathbf{Z}\|_*$$

and the global minimum for  $R_N(\mathbf{W}, \mathbf{H})$  is attained at

$$\mathbf{Z} = \mathbf{W}^\top \mathbf{H} = \frac{(\sqrt{n} - N\lambda)_+}{\sqrt{n}} \mathbf{Y} = (1 - \sqrt{n}K\lambda)_+ \mathbf{Y}$$

where  $\mathbf{Y}$  is a binary matrix satisfying  $\mathbf{Y}\mathbf{Y}^\top = n\mathbf{I}_K$  and  $\lambda = \sqrt{\lambda_W \lambda_H}$ . Given  $\mathbf{Z}$ , (Zhu et al., 2021, Lemma A.3) implies that the global minimizer is given by  $(\mathbf{H}, \mathbf{W})$  that satisfies

$$\begin{aligned} \mathbf{H}^\top \mathbf{H} &= \sqrt{\frac{\lambda_W}{\lambda_H}} \frac{(1 - \sqrt{n}K\lambda)_+}{\sqrt{n}} \mathbf{Y}^\top \mathbf{Y} = \sqrt{\frac{\lambda_W}{\lambda_H}} \frac{(1 - \sqrt{n}K\lambda)_+}{\sqrt{n}} \mathbf{I}_K \otimes \mathbf{J}_n, \\ \mathbf{W}^\top \mathbf{W} &= \sqrt{\frac{\lambda_H}{\lambda_W}} n^{1/2} (1 - \sqrt{n}K\lambda)_+ \mathbf{I}_K, \end{aligned}$$

and we can select a nonnegative  $\mathbf{H}$  to be the global minimizer for (2.7), for example,

$$\mathbf{H} = \bar{\mathbf{H}} \otimes \mathbf{1}_n^\top, \quad \bar{\mathbf{H}} = \left( \frac{\lambda_W}{n\lambda_H} \right)^{\frac{1}{4}} \begin{bmatrix} (1 - \sqrt{n}K\lambda)_+^{\frac{1}{2}} \mathbf{I}_K \\ \mathbf{0}_{D-K, K} \end{bmatrix}.$$

■

## 4.2 Two-layer neural network

### 4.2.1 PROOF OF THEOREM 4: $\mathcal{NC}$ FOR GMM WITH TWO CLUSTERS FOR $2n \geq d > n$

**Proof of Theorem 4.** For now, we first try to establish a condition such that

$$\langle \boldsymbol{\mu}_1, \boldsymbol{\beta} \rangle + \sigma \mathbf{Z}_1 \boldsymbol{\beta} = 1, \quad \langle \boldsymbol{\mu}_1, \boldsymbol{\beta} \rangle = 1, \quad \langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle + \sigma \mathbf{Z}_2 \boldsymbol{\beta} \leq 0, \quad (4.6)$$

is feasible, i.e., the set is non-empty. Once it is done, the same result also applies to the second cluster. The second equality constraint makes the argument cleaner without compromising much. With (4.6), we have

$$\mathbf{Z}_1 \boldsymbol{\beta} = 0, \quad \langle \boldsymbol{\mu}_1, \boldsymbol{\beta} \rangle = 1, \quad \langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle + \sigma \mathbf{Z}_2 \boldsymbol{\beta} \leq 0,$$

which means  $\boldsymbol{\beta}$  is in the null space of  $\mathbf{Z}_1 \in \mathbb{R}^{n \times d}$ .

Let  $\boldsymbol{\Phi} \in \mathbb{R}^{d \times (d-n)}$  be a partial orthonormal matrix with columns being the basis of the null space of  $\mathbf{Z}_1$ , which belongs to the Grassmannian  $G_{d, d-n}$ . It is easy to see that  $\boldsymbol{\beta}$  has the following form:

$$\boldsymbol{\beta} = \boldsymbol{\Phi} \left( \frac{\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1}{\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1\|^2} + \mathbf{v} \right)$$

where  $\mathbf{v}$  is perpendicular to  $\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1$ , i.e.,  $\langle \mathbf{v}, \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1 \rangle = 0$ .

With the representation of  $\boldsymbol{\beta}$ , it remains to ensure the third constraint in (4.6) satisfies. Note that for any fixed  $\boldsymbol{\beta}$ ,

$$\langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_2 \boldsymbol{\beta} \sim \mathcal{N}(\langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle \mathbf{1}_n, \sigma^2 \|\boldsymbol{\beta}\|^2).$$

To guarantee a high probability of the event  $\langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_2 \boldsymbol{\beta} \leq 0$ , we need to maximize the ratio of  $\langle \boldsymbol{\mu}_2, \boldsymbol{\beta} \rangle$  over  $\|\boldsymbol{\beta}\|$  subject to the constraints on  $\boldsymbol{\beta}$ . Therefore, we define

$$f(\mathbf{v}) := \frac{\langle \boldsymbol{\mu}_2, \boldsymbol{\beta}(\mathbf{v}) \rangle}{\|\boldsymbol{\beta}(\mathbf{v})\|} = \frac{\langle \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1 / \|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1\|^2 + \mathbf{v}, \boldsymbol{\Phi}^\top \boldsymbol{\mu}_2 \rangle}{\sqrt{1/\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1\|^2 + \|\mathbf{v}\|^2}}$$

subject to  $\langle \mathbf{v}, \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1 \rangle = 0$ . By applying Lemma 10 with  $\mathbf{v}_1 := \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1 / \|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1\|^2$  and  $\mathbf{v}_2 := \boldsymbol{\Phi}^\top \boldsymbol{\mu}_2$ , we have

$$f^* := \min_{\langle \mathbf{v}, \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1 \rangle = 0} f(\mathbf{v}) = \begin{cases} -\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_2\| \sqrt{1 - \cos^2 \tilde{\theta}}, & \text{if } \cos \tilde{\theta} \geq 0, \\ -\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_2\|, & \text{if } \cos \tilde{\theta} < 0, \end{cases} \quad (4.7)$$

where  $\tilde{\theta}$  is the angle between  $\boldsymbol{\Phi}^\top \tilde{\boldsymbol{\mu}}_1$  and  $\boldsymbol{\Phi}^\top \tilde{\boldsymbol{\mu}}_2$ .

Therefore, it suffices to have i.i.d.  $n$  samples from  $\mathcal{N}(f^*, \sigma^2)$  such that all of them are nonpositive. By taking union bound over  $n$  samples, we have

$$2\sigma \sqrt{\log n} + f^* \leq 0 \quad (4.8)$$

with probability at least  $1 - O(n^{-1})$ . It remains to estimate  $f^*$ . From the lemma 11,

$$|\cos \tilde{\theta} - \cos \theta| \leq \frac{4\varepsilon}{(1 - \varepsilon)^2}.$$

For  $\cos \theta < -4\varepsilon/(1 - \varepsilon)^2$ , then  $\cos \tilde{\theta} < 0$ , and then

$$-f^* \geq (1 - \varepsilon) \sqrt{\frac{d - n}{d}} \|\boldsymbol{\mu}_2\|$$

where the bound on  $\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_2\| \geq (1 - \varepsilon) \sqrt{(d - n)/d} \|\boldsymbol{\mu}_2\|$  follows from the Johnson-Lindenstrauss lemma. For  $\cos \theta > -4\varepsilon/(1 - \varepsilon)^2$  and  $|\cos \tilde{\theta}| \leq |\cos \theta| + 4\varepsilon/(1 - \varepsilon)^2$ , then

$$\begin{aligned} -f^* &= \sqrt{\|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_2\|^2 - |\langle \boldsymbol{\Phi}^\top \boldsymbol{\mu}_1, \boldsymbol{\Phi}^\top \boldsymbol{\mu}_2 \rangle|^2 / \|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_1\|^2} \\ &= \|\boldsymbol{\Phi}^\top \boldsymbol{\mu}_2\| \sqrt{1 - \cos^2 \tilde{\theta}} \geq (1 - \varepsilon) \sqrt{\frac{d - n - 1}{d}} \|\boldsymbol{\mu}_2\| \sqrt{1 - \left( |\cos \theta| + \frac{4\varepsilon}{(1 - \varepsilon)^2} \right)^2}. \end{aligned}$$

All bounds above hold with probability at least  $1 - C \exp(-\varepsilon^2(d - n - 1))$ .

For  $d \geq 2n$ , we notice that  $\mathbf{X}^\top$  is full row rank for any  $\sigma > 0$ , and thus one can always find a  $\boldsymbol{\beta}$  such that  $\mathbf{X}_1^\top \boldsymbol{\beta} = \mathbf{1}$  and  $\mathbf{X}_2^\top \boldsymbol{\beta} \leq 0$ .  $\blacksquare$

#### 4.2.2 PROOF OF THEOREM 5(A): $\mathcal{NC}$ FOR GMM WITH MULTIPLE CLASSES FOR $Kn \geq d > n$

**Proof of Theorem 5(a).** To ensure the neural collapse occurs, we first show that for each fixed  $k$ , there exists  $\boldsymbol{\beta}$  such that

$$(\mathbf{1}_n \boldsymbol{\mu}_k^\top + \sigma \mathbf{Z}_k) \boldsymbol{\beta} = \mathbf{1}_n, \quad \langle \boldsymbol{\mu}_k, \boldsymbol{\beta} \rangle = 1, \quad \langle \boldsymbol{\mu}_{k'}, \boldsymbol{\beta} \rangle = -\gamma \quad (k' \neq k, \gamma > 0).$$

These equations reduce to the set such that

$$\mathbf{Z}_k \boldsymbol{\beta} = 0, \quad \mathbf{\Pi} \boldsymbol{\beta} = \mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k)$$

where  $\mathbf{\Pi} \in \mathbb{R}^{K \times d}$  and  $\mathbf{Z}_k \in \mathbb{R}^{n \times d}$ . Let  $\Phi$  be an  $d \times (d-n)$  orthogonal matrix whose columns span the null space of  $\mathbf{Z}_k$ . Let  $\boldsymbol{\beta} = \Phi \boldsymbol{\alpha}$  where  $\boldsymbol{\alpha} \in \mathbb{R}^{d-n}$  and  $\Phi \in \mathbb{R}^{d \times (d-n)}$  is a random projection matrix, and  $\mathbf{Z}_k \Phi = 0$ . Since  $\mathbf{\Pi}$  is assumed to be of rank- $K$ ,

$$\mathbf{\Pi} \Phi \boldsymbol{\alpha} = \mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k) \quad (4.9)$$

and  $\mathbf{\Pi} \Phi \in \mathbb{R}^{K \times (d-n)}$ . Under  $d-n > K$ , then  $\mathbf{\Pi} \Phi$  is of rank  $K$  with probability 1. Here we choose

$$\begin{aligned} \boldsymbol{\alpha} &= \Phi^\top \mathbf{\Pi}^\top (\mathbf{\Pi} \Phi \Phi^\top \mathbf{\Pi}^\top)^{-1} (\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k)), \\ \boldsymbol{\beta} &= \Phi \Phi^\top \mathbf{\Pi}^\top (\mathbf{\Pi} \Phi \Phi^\top \mathbf{\Pi}^\top)^{-1} (\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k)), \end{aligned}$$

where

$$\|\boldsymbol{\beta}\|^2 = (\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k))^\top (\mathbf{\Pi} \Phi \Phi^\top \mathbf{\Pi}^\top)^{-1} (\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k)).$$

Note that for each given  $\gamma$ ,

$$(\mathbf{1}_n \boldsymbol{\mu}_{k'}^\top + \sigma \mathbf{Z}_{k'}) \boldsymbol{\beta} \sim \mathcal{N}(-\gamma, \sigma^2 \|\boldsymbol{\beta}\|^2 \mathbf{I}_n)$$

for any  $k' \neq k$ , and  $\langle \boldsymbol{\mu}_{k'}, \boldsymbol{\beta} \rangle = -\gamma$ . We proceed to determine the choice of  $\gamma > 0$  by minimizing

$$\min_{\gamma > 0} -\frac{\gamma}{\|\boldsymbol{\beta}\|} \iff \max_{\gamma > 0} F(\gamma)$$

where

$$F(\gamma) := \frac{\gamma^2}{\|\boldsymbol{\beta}\|^2} = \frac{\gamma^2}{(\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k))^\top (\mathbf{\Pi} \Phi \Phi^\top \mathbf{\Pi}^\top)^{-1} (\mathbf{e}_k - \gamma(\mathbf{1}_K - \mathbf{e}_k))}. \quad (4.10)$$

Then we know that  $(\mathbf{1}_n \boldsymbol{\mu}_{k'}^\top + \sigma \mathbf{Z}_{k'}) \boldsymbol{\beta} \leq 0$  holds if

$$2\sigma \sqrt{\log(Kn)} \leq F(\gamma)$$

with probability at least  $1 - O((Kn)^{-1})$ .

By Lemma 12, it holds with probability at least  $1 - c \exp(-\varepsilon^2(d-n))$  that

$$\lambda_{\min}(\mathbf{\Pi} \Phi \Phi^\top \mathbf{\Pi}^\top) \geq ((1-\varepsilon)^2 - 4K\varepsilon) \frac{d-n}{d} \sigma_{\min}^2(\mathbf{\Pi})$$

where  $\sigma_{\min}(\mathbf{\Pi})$  is the smallest singular value of  $\mathbf{\Pi}$ . Therefore, we have

$$\|\boldsymbol{\beta}\|^2 \leq \frac{(1 + (K-1)\gamma^2)d}{((1-\varepsilon)^2 - 4K\varepsilon)(d-n)\sigma_{\min}^2(\mathbf{\Pi})},$$

which implies

$$\sup_{\gamma > 0} F(\gamma) \geq \sup_{\gamma > 0} \frac{\gamma^2}{\|\boldsymbol{\beta}\|^2} = \frac{((1-\varepsilon)^2 - 4K\varepsilon)(d-n)\sigma_{\min}^2(\mathbf{\Pi})}{(K-1)d}.$$

Then  $(\mathbf{1}_n \boldsymbol{\mu}_{k'}^\top + \sigma \mathbf{Z}_{k'}) \boldsymbol{\beta} \leq 0$  is implied by

$$\sigma < \frac{1}{2} \sqrt{\frac{((1-\varepsilon)^2 - 4K\varepsilon)(d-n)\sigma_{\min}^2(\mathbf{\Pi})}{(K-1)d \log(Kn)}}$$

with probability at least  $1 - c \exp(-\varepsilon^2(d-n)) - (Kn)^{-1}$ . Note that we need to choose  $\varepsilon < 1/(10K)$ , i.e.,  $d-n \geq CK^2 \log Kn$  such that the probability is at least  $1 - O(n^{-1})$ . ■

#### 4.2.3 PROOF OF THEOREM 5(B): $\mathcal{NC}$ FOR GMM WITH MULTIPLE CLASSES UNDER $d/n > (K+1)/2$

**Proof of Theorem 5(b).** In this subsection, we present a tighter bound for the linear feasibility when  $d/n > (K+1)/2$  holds. For class  $k$ , we consider the set of  $\boldsymbol{\beta}$ 's such that

$$(\mathbf{1}_n \boldsymbol{\mu}_k^\top + \sigma \mathbf{Z}_k) \boldsymbol{\beta} = \mathbf{1}_n, \quad \langle \boldsymbol{\mu}_k, \boldsymbol{\beta} \rangle = 1, \quad \langle \boldsymbol{\mu}_{k'}, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_{k'} \boldsymbol{\beta} \leq 0. \quad (4.11)$$

We consider  $\boldsymbol{\beta}$  in the following form:

$$\boldsymbol{\beta} = \Phi \left( \mathbf{v} + \frac{\Phi^\top \boldsymbol{\mu}_k}{\|\Phi^\top \boldsymbol{\mu}_k\|^2} \right) \quad (4.12)$$

where  $\Phi \in \mathbb{R}^{d \times (d-n)}$  represents the orthogonal basis of the null space of  $\mathbf{Z}_k$  and

$$\langle \mathbf{v}, \Phi^\top \boldsymbol{\mu}_k \rangle = 0, \quad 1 \leq k \leq K.$$

Such  $\mathbf{v} \in \mathbb{R}^{d-n}$  exists if  $d-n \geq K+1$ . Now for any  $k' \neq k$ , it holds that

$$\langle \boldsymbol{\mu}_{k'}, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_{k'} \boldsymbol{\beta} = \left( \langle \Phi^\top \boldsymbol{\mu}_{k'}, \Phi \boldsymbol{\mu}_k \rangle \mathbf{1}_n + \sigma \frac{\mathbf{Z}_{k'} \Phi \Phi^\top \boldsymbol{\mu}_k}{\|\Phi^\top \boldsymbol{\mu}_k\|^2} \right) + \sigma \mathbf{Z}_{k'} \Phi \mathbf{v}.$$

To have  $\langle \boldsymbol{\mu}_{k'}, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_{k'} \boldsymbol{\beta} < 0$  it suffices to find  $\mathbf{v}$  such that  $\mathbf{Z}_{k'} \Phi \mathbf{v} < 0$  for  $k'$ , since we can arbitrarily rescale the term  $\mathbf{Z}_{k'} \Phi \mathbf{v}$  by rescaling the norm of  $\mathbf{v}$ . The effective dimension of all  $\mathbf{v}$  is  $d-n-K$ . Therefore, this problem is equivalent to finding  $\bar{\mathbf{v}} \in \mathbb{R}^{d-n-K}$  such that

$$\mathbf{A} \bar{\mathbf{v}} < 0$$

where  $\mathbf{A} \in \mathbb{R}^{(K-1)n \times (d-n-K)}$  is a Gaussian random matrix. The next proposition provides sufficient conditions for constructing a high probability bound.

We aim to search for  $\bar{\mathbf{v}} \in \mathbb{R}^{d-n-K}$  such that  $\mathbf{A} \bar{\mathbf{v}} < 0$  holds entrywisely. By hyperplane separation theorem, if the convex set  $\{\mathbf{A}^\top \mathbf{s} : \|\mathbf{s}\|_1 = 1, \mathbf{s} \geq 0, \mathbf{s} \in \mathbb{R}^{(K-1)n}\}$  does not contain 0, then there exists  $\bar{\mathbf{v}}$  s.t.

$$\langle \bar{\mathbf{v}}, \mathbf{A}^\top \mathbf{s} \rangle < 0, \quad \forall \|\mathbf{s}\|_1 = 1, \mathbf{s} \geq 0 \iff \mathbf{A}^\top \bar{\mathbf{v}} < 0.$$

Hence it suffices to show that  $\min_{\mathbf{s} \geq 0, \|\mathbf{s}\|_1=1} \|\mathbf{A}^\top \mathbf{s}\| > 0$  holds with high probability. Note that

$$\begin{aligned} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|_1=1} \max_{\|\mathbf{u}\|=1} \langle \mathbf{u}, \mathbf{A}^\top \mathbf{s} \rangle &= \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|_1=1} \|\mathbf{A}^\top \mathbf{s}\| = \min_{\mathbf{s} \geq 0} \frac{\|\mathbf{A}^\top \mathbf{s}\|}{\|\mathbf{s}\|_1} = \min_{\mathbf{s} \geq 0} \frac{\|\mathbf{A}^\top \mathbf{s}\| \|\mathbf{s}\|_2}{\|\mathbf{s}\|_2 \|\mathbf{s}\|_1} \\ &\geq \frac{1}{\sqrt{(K-1)n}} \min_{\mathbf{s} \geq 0} \frac{\|\mathbf{A}^\top \mathbf{s}\|}{\|\mathbf{s}\|_2} = \frac{1}{\sqrt{(K-1)n}} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \|\mathbf{A}^\top \mathbf{s}\| \\ &= \frac{1}{\sqrt{(K-1)n}} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \max_{\|\mathbf{u}\|=1} \langle \mathbf{u}, \mathbf{A}^\top \mathbf{s} \rangle. \end{aligned}$$

Applying Proposition 13, we have,

$$\mathbb{E} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \max_{\|\mathbf{u}\|=1} \langle \mathbf{u}, \mathbf{A}^\top \mathbf{s} \rangle \geq \sqrt{d-n-K} - \sqrt{(K-1)n/2}. \quad (4.13)$$

Therefore, we have

$$\mathbb{E} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \|\mathbf{A}^\top \mathbf{s}\| \geq \frac{1}{\sqrt{(K-1)n}} \mathbb{E} \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \|\mathbf{A}^\top \mathbf{s}\| \geq \frac{\sqrt{d-n-K} - \sqrt{(K-1)n/2}}{\sqrt{(K-1)n}}.$$

By Lemma 15, we have

$$\Pr \left( \min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \|\mathbf{A}^\top \mathbf{s}\| \leq \sqrt{d-n-K} - \sqrt{(K-1)n/2} - t \right) \leq \exp(-t^2/2).$$

As a result, by taking  $t = \sqrt{2 \log n}$ , it holds with probability at least  $1 - 1/n$  that

$$\min_{\mathbf{s} \geq 0, \|\mathbf{s}\|=1} \|\mathbf{A}^\top \mathbf{s}\| \geq \frac{\sqrt{d-n-K} - \sqrt{(K-1)n/2} - \sqrt{2 \log n}}{\sqrt{(K-1)n}}.$$

Therefore, to ensure the existence of an  $\mathcal{NC}$  solution, it suffices to have the RHS above positive:

$$d - n - K \geq \frac{(K-1)n}{2} + 2 \log n + 2\sqrt{(K-1)n \log n}.$$

■

### 4.3 Proof of Theorem 6: neural collapse for three-layer neural network

Next, we consider the feasibility of inducing the neural collapse of a three-layer network with the first layer having random Gaussian weight, i.e. the format of the output is given by

$$f(\mathbf{W}_1, \mathbf{W}_2, \mathbf{W}_3) = \mathbf{W}_3^\top \sigma_{\text{ReLU}} \left( \mathbf{W}_2^\top \sigma_{\text{ReLU}} \left( \mathbf{W}_1^\top \mathbf{X} \right) \right) \quad (4.14)$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times d_1}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{d_1 \times d_2}$ ,  $\mathbf{W}_3 \in \mathbb{R}^{d_2 \times K}$  and  $\mathbf{X} \in \mathbb{R}^{d \times N}$  is some general data. One natural question is whether neural collapse occurs for a three-layer neural network. From Theorem 3(b), we know that if there exists  $\mathbf{W}_1$  such that  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) \in \mathbb{R}^{d_1 \times N}$  is rank  $N$  with  $d_1 \geq N$ , then there exists  $\mathbf{W}_2$  and  $\mathbf{W}_3$  such that the neural collapse occurs.

We will show that by setting  $\mathbf{W}_1$  to be a Gaussian random matrix with each entry being i.i.d Gaussian random variables, i.e.,  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  gives the ReLU random feature of  $N$  data points in  $\mathbb{R}^d$ , then the neural collapse occurs with high probability if  $d_1$  is sufficiently large.

Given  $N$  points  $\{\mathbf{x}_i\}_{i=1}^N$ , we define that

$$[\mathbf{H}_\infty]_{ij} := \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_d)} \left( \sigma_{\text{ReLU}}(\langle \mathbf{z}, \mathbf{x}_i \rangle) - \sqrt{\frac{2}{\pi}} \right) \left( \sigma_{\text{ReLU}}(\langle \mathbf{z}, \mathbf{x}_j \rangle) - \sqrt{\frac{2}{\pi}} \right)$$

as the kernel matrix. Provided that  $\mathbf{x}_i$  is not parallel to any  $\mathbf{x}_j$ , then it can be shown that  $\mathbf{H}_\infty(\mathbf{X})$  is full rank with  $\lambda_{\min}(\mathbf{H}_\infty) > 0$  as shown in the lemma below.

**Lemma 9.** Let  $\mathbf{X} \in \mathbb{R}^{d \times N}$  be a data matrix with columns  $\mathbf{x}_i \not\parallel \mathbf{x}_j$  for any  $(i, j)$  pair. The kernel matrix

$$\mathbf{H}_\infty := \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_d)} \sigma(\mathbf{X}^\top \mathbf{z}) \sigma(\mathbf{X}^\top \mathbf{z})^\top - \frac{2}{\pi} \mathbf{J}_N$$

has a strictly positive smallest singular value  $\lambda_{\min}(\mathbf{H}_\infty) > 0$ .

**Proof of Lemma 9** The main idea for this proof is adopted from (Du et al., 2018, Theorem 3.1). Define the feature map  $\phi_{\mathbf{x}}(\mathbf{w}) := \sigma(\mathbf{w}^\top \mathbf{x}) - \sqrt{2/\pi}$ ,  $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$ , which is a continuous function w.r.t.  $\mathbf{w} \in \mathbb{R}^d$  and the gradient  $\nabla_{\mathbf{w}} \sigma(\mathbf{w}^\top \mathbf{x}_i)$  is continuous everywhere except

$$D_i := \{\mathbf{w} : \langle \mathbf{x}_i, \mathbf{w} \rangle = 0\}.$$

Now to prove  $\mathbf{H}_\infty$  is strictly positive definite, it is equivalent to show  $\phi_{\mathbf{x}_i}(\mathbf{w})$  is linearly independent, i.e., for any  $\{a_i\}_{i=1}^N$  such that

$$\sum_{i=1}^N a_i \phi_{\mathbf{x}_i}(\mathbf{w}) = 0, \quad \forall \mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d),$$

then  $a_i = 0$  holds for  $1 \leq i \leq N$ . For every  $\{a_i\}_{i=1}^N$ , we define

$$f(\mathbf{w}) := \sum_{i=1}^N a_i \phi_{\mathbf{x}_i}(\mathbf{w}).$$

Assume  $f(\mathbf{w}) = 0$  holds for  $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$ . By definition,  $f(\mathbf{w})$  is a continuous function on  $\mathbb{R}^d$  and it means  $f(\mathbf{w}) \equiv 0$  and  $\nabla f(\mathbf{w}) \equiv \mathbf{0}_d$ .

Under  $\mathbf{x}_i \not\parallel \mathbf{x}_j$  for all  $(i, j)$  pair, it holds  $D_i \not\subseteq \cup_{j \neq i} D_j$ , i.e., there exists  $\mathbf{z} \in D_i \setminus \cup_{j \neq i} D_j$  such that  $\langle \mathbf{x}_i, \mathbf{z} \rangle = 0$  and  $\langle \mathbf{x}_j, \mathbf{z} \rangle \neq 0$  for any  $j \neq i$ . Since  $D_i$ 's are closed sets, there exists  $r_0 > 0$  such that for any  $r \leq r_0$ ,  $\mathcal{B}(\mathbf{z}, r) \cap (\cup_{j \neq i} D_j) = \emptyset$ . In other words,  $\phi_{\mathbf{x}_j}(\mathbf{w})$  is differentiable w.r.t.  $\mathbf{w}$  insides  $\mathcal{B}(\mathbf{z}, r_0)$  for all  $j \neq i$ . For  $i$ , the ball  $\mathcal{B}(\mathbf{z}, r_0)$  contains two disjoint parts:

$$\mathcal{B}^+(\mathbf{z}, r) := \{\mathbf{w} : \langle \mathbf{x}_i, \mathbf{w} \rangle > 0\} \cap \mathcal{B}(\mathbf{z}, r), \quad \mathcal{B}^-(\mathbf{z}, r) := \{\mathbf{w} : \langle \mathbf{x}_i, \mathbf{w} \rangle < 0\} \cap \mathcal{B}(\mathbf{z}, r)$$

and  $\mathbf{z}$  is on the boundary of both  $\mathcal{B}^+(\mathbf{z}, r)$  and  $\mathcal{B}^-(\mathbf{z}, r)$ .

Therefore, there exist two sequences  $\{\mathbf{w}_\ell^+\} \subseteq \mathcal{B}^+(\mathbf{z}, r)$  and  $\{\mathbf{w}_\ell^-\} \subseteq \mathcal{B}^-(\mathbf{z}, r)$  such that  $\lim_{\ell \rightarrow \infty} \mathbf{w}_\ell^+ = \lim_{\ell \rightarrow \infty} \mathbf{w}_\ell^- = \mathbf{z}$ . Due to the continuous differentiability of  $\phi_{\mathbf{x}_j}(\mathbf{w})$  in  $\mathcal{B}(\mathbf{z}, r)$  for  $j \neq i$ , we have

$$\lim_{\ell \rightarrow \infty} \nabla \phi_{\mathbf{x}_j}(\mathbf{w}_\ell^+) = \lim_{\ell \rightarrow \infty} \nabla \phi_{\mathbf{x}_j}(\mathbf{w}_\ell^-) = \nabla \phi_{\mathbf{x}_j}(\mathbf{z}). \quad (4.15)$$

For  $i$ , we note that  $\phi_{\mathbf{x}_i}(\mathbf{w})$  is not differentiable at  $\mathbf{z} \in D_i$  while the gradient exists on  $\mathcal{B}^+(\mathbf{z}, r)$  and  $\mathcal{B}^-(\mathbf{z}, r)$ . It holds

$$\lim_{\ell \rightarrow \infty} \nabla \phi_{\mathbf{x}_i}(\mathbf{w}_\ell^+) = \lim_{\ell \rightarrow \infty} \mathbf{w}_\ell = \mathbf{z}, \quad \lim_{\ell \rightarrow \infty} \nabla \phi_{\mathbf{x}_i}(\mathbf{w}_\ell^-) = 0 \quad (4.16)$$

where

$$\phi_{\mathbf{x}_i}(\mathbf{w}) = \begin{cases} \langle \mathbf{x}_i, \mathbf{w} \rangle - \sqrt{\frac{2}{\pi}}, & \mathbf{w} \in \mathcal{B}^+(\mathbf{z}, r), \\ -\sqrt{\frac{2}{\pi}}, & \mathbf{w} \in \mathcal{B}^-(\mathbf{z}, r). \end{cases}$$

Therefore,

$$\begin{aligned} 0 &= \lim_{\ell \rightarrow \infty} \nabla f(\mathbf{w}_\ell^+) - \lim_{\ell \rightarrow \infty} \nabla f(\mathbf{w}_\ell^-) \\ &= \lim_{\ell \rightarrow \infty} \sum_{\substack{j \\ j \neq i}}^N a_j (\nabla \phi_{\mathbf{x}_j}(\mathbf{w}_\ell^+) - \nabla \phi_{\mathbf{x}_j}(\mathbf{w}_\ell^-)) + a_i \lim_{\ell \rightarrow \infty} (\nabla \phi_{\mathbf{x}_i}(\mathbf{w}_\ell^+) - \nabla \phi_{\mathbf{x}_i}(\mathbf{w}_\ell^-)) = a_i \mathbf{z}. \end{aligned}$$

Note that  $\mathbf{z} \neq 0$  and it implies  $a_i = 0$ . ■

**Proof of Theorem 6.** By Theorem 3(b), when  $d_1 \geq N$ , it suffices to have  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) \in \mathbb{R}^{d_1 \times N}$  to be rank  $N$  to induce neural collapse. Without loss of generality, We assume each  $\mathbf{x}_i$  is a unit vector in  $\mathbb{R}^d$ . Here we define  $\boldsymbol{\varphi}_k$  as

$$\boldsymbol{\varphi}_k = \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) = \begin{bmatrix} \sigma_{\text{ReLU}}(\langle \mathbf{x}_1, \mathbf{z}_k \rangle) \\ \vdots \\ \sigma_{\text{ReLU}}(\langle \mathbf{x}_N, \mathbf{z}_k \rangle) \end{bmatrix}, \quad \mathbb{E} \boldsymbol{\varphi}_k = \sqrt{\frac{2}{\pi}} \mathbf{1}_d,$$

which is exactly the  $k$ -th row of  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  and it is a sub-gaussian random vector. To guarantee that  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X})$  is full rank, it suffices to have

$$\left\| \frac{1}{d_1} \sum_{k=1}^{d_1} (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)(\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)^\top - \mathbf{H}_\infty \right\| < \lambda_{\min}(\mathbf{H}_\infty).$$

The first step is directly implied by Lemma 9, i.e.  $\lambda_{\min}(\mathbf{H}_\infty) > 0$ . The estimation reduces to the covariance matrix estimation, which can be done by computing the sub-gaussian norm of the centered  $\boldsymbol{\varphi}_k$ . For any  $\mathbf{v} \in \mathbb{S}^{N-1}$  and define

$$f(\mathbf{z}) := \left\langle \mathbf{v}, \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}) - \sqrt{\frac{2}{\pi}} \mathbf{1}_N \right\rangle,$$

which is a random variable depending on  $\mathbf{z}$ . Then for any  $\mathbf{v}$ , it holds

$$|f(\mathbf{z}_1) - f(\mathbf{z}_2)| \leq \left\| \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_1) - \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_2) \right\| \leq \|\mathbf{X}\| \|\mathbf{z}_1 - \mathbf{z}_2\|. \quad (4.17)$$

Then by (Wainwright, 2019, Theorem 2.26), we know that  $f(\mathbf{z}) - \mathbb{E} f(\mathbf{z})$  is a subgaussian random variable with

$$\mathbb{E} e^{\lambda \langle \mathbf{v}, \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}) - \mathbb{E} \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}) \rangle} \leq e^{\lambda^2 \|\mathbf{X}\|^2 / 2}, \quad \forall \mathbf{v} \in \mathbb{S}^{N-1}, \lambda > 0.$$

Therefore, the random vector  $\sigma(\mathbf{X}^\top \mathbf{z}) - \mathbb{E} \sigma(\mathbf{X}^\top \mathbf{z})$  has a sub-gaussian norm bounded by  $\|\mathbf{X}\|$ .

By (Vershynin, 2018, Exercise 4.7.3) or (Wainwright, 2019, Theorem 6.5), it holds that

$$\left\| \frac{1}{d_1} \sum_k (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)(\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)^\top - \mathbf{H}_\infty \right\| \lesssim \|\mathbf{X}\|^2 \sqrt{\frac{N}{d_1}} < \lambda_{\min}(\mathbf{H}_\infty)$$

with probability at least  $1 - 2e^{-N}$  provided that  $d_1 \gtrsim \|\mathbf{X}\|^4 N \log N / \lambda_{\min}^2(\mathbf{H}_\infty)$ .

Note that

$$\begin{aligned}
 & \frac{1}{d_1} \sum_k \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top - \mathbb{E} \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top = \frac{1}{d_1} \sum_k \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top - (\mathbf{H}_\infty + \frac{2}{d_1 \pi} \mathbf{J}_N) \\
 & = \frac{1}{d_1} \sum_{k=1}^{d_1} (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k) (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)^\top - \mathbf{H}_\infty \\
 & \quad + \frac{1}{d_1} \sqrt{\frac{2}{\pi}} \left( \mathbf{1}_N \sum_{k=1}^{d_1} (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k)^\top + \sum_{k=1}^{d_1} (\boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k) \mathbf{1}_N^\top \right)
 \end{aligned}$$

where  $\mathbb{E} \boldsymbol{\varphi}_k = \sqrt{2/\pi} \mathbf{1}_{d_1}$ . It suffices to estimate  $\|d_1^{-1} \sum_k \boldsymbol{\varphi}_k - \mathbb{E} \boldsymbol{\varphi}_k\|$ . We know that

$$\mathbb{E} e^{\lambda \sum_k \langle \mathbf{v}, \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) - \mathbb{E} \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) \rangle} \leq e^{\lambda^2 d_1 \|\mathbf{X}\|^2 / 2}, \quad \forall \mathbf{v} \in \mathbb{S}^{N-1}, \lambda > 0$$

and thus  $d_1^{-1} \sum_k \langle \mathbf{v}, \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) - \mathbb{E} \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) \rangle$  is  $\|\mathbf{X}\|^2/d_1$ -subgaussian. Thus

$$\Pr \left( \left\| \frac{1}{d_1} \sum_k \sigma_{\text{ReLU}}(\mathbf{X}^\top \mathbf{z}_k) - \mathbb{E} \boldsymbol{\varphi}_k \right\| \geq t \right) \leq 2 \exp \left( -\frac{d_1 t^2}{2 \|\mathbf{X}\|^2} \right).$$

By picking  $t = 2 \|\mathbf{X}\| \sqrt{d_1^{-1} \log N}$ , we have

$$\left\| \frac{1}{d_1} \sum_k \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top - \mathbb{E} \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top \right\| \lesssim \|\mathbf{X}\|^2 \sqrt{\frac{N}{d_1}} + \|\mathbf{X}\| \sqrt{\frac{N \log N}{d_1}} \lesssim \lambda_{\min}(\mathbf{H}_\infty)$$

under  $d_1 \gtrsim \|\mathbf{X}\|^4 N \log N / \lambda_{\min}^2(\mathbf{H}_\infty)$ . Therefore, the smallest eigenvalue of  $d_1^{-1} \sum_k \boldsymbol{\varphi}_k \boldsymbol{\varphi}_k^\top$  is at least  $\lambda_{\min}(\mathbf{H}_\infty)/2$ , implying that  $\sigma_{\text{ReLU}}(\mathbf{W}_1^\top \mathbf{X}) \in \mathbb{R}^{d_1 \times N}$  is of full column rank.  $\blacksquare$

#### 4.4 Two-layer neural network: best generalization under $\mathcal{NC}$

##### 4.4.1 PROOF OF THEOREM 7(A)

**Proof of Theorem 7(a).** We consider the minimization for the term involving  $\boldsymbol{\beta}$ ,

$$\min_{\boldsymbol{\beta}} \Pr \left( \boldsymbol{\beta}^\top \boldsymbol{\mu} + \sigma \boldsymbol{\beta}^\top \mathbf{z} \leq 0 \right) \quad \text{s.t.} \quad \left( \mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_1 \right) \boldsymbol{\beta} = \mathbf{1}_n, \quad \left( -\mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_2 \right) \boldsymbol{\beta} \leq 0. \quad (4.18)$$

The difficulty of optimization mainly arises from incorporating the inequality constraints. For now, we drop the inequality constraints and consider the following simplified version that only involves the equality constraints, then the minimization is equivalent to solving the following maximization program.

$$\max_{\boldsymbol{\beta} \in \mathbb{R}^d} \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\|\boldsymbol{\beta}\|} \quad \text{s.t.} \quad \left( \mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_1 \right) \boldsymbol{\beta} = \mathbf{1}_n. \quad (4.19)$$

Note that for any  $\boldsymbol{\beta}$  satisfying  $(\mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_1) \boldsymbol{\beta} = \mathbf{1}_n$ , it holds that

$$(\mathbf{I}_n - \mathbf{J}_n/n) \mathbf{Z}_1 \boldsymbol{\beta} = 0, \quad \left\langle \boldsymbol{\mu} + \frac{\sigma \mathbf{Z}_1^\top \mathbf{1}_n}{n}, \boldsymbol{\beta} \right\rangle = 1.$$

Therefore, we let  $\Phi$  be the orthonormal basis of the null space of  $(\mathbf{I}_n - \mathbf{J}_n/n)\mathbf{Z}_1$ . Since  $\mathbf{Z}_1 \in \mathbb{R}^{n \times d}$  is a Gaussian random matrix,  $\Phi$  is a random matrix sampled from  $G_{d, d-n+1}$ .

Therefore, we have the following representation for  $\beta$ :

$$\beta(\mathbf{v}) = \Phi \left( \frac{\Phi^\top \hat{\boldsymbol{\mu}}}{\|\Phi^\top \hat{\boldsymbol{\mu}}\|^2} + \mathbf{v} \right)$$

where

$$\hat{\boldsymbol{\mu}} = \boldsymbol{\mu} + \frac{\sigma \mathbf{Z}_1^\top \mathbf{1}_n}{n}, \quad \langle \mathbf{v}, \Phi^\top \hat{\boldsymbol{\mu}} \rangle = 0.$$

Now we can see the maximization of (4.19) is equivalent to maximizing

$$f(\mathbf{v}) := \frac{\langle \boldsymbol{\mu}, \beta(\mathbf{v}) \rangle}{\|\beta(\mathbf{v})\|} = \frac{\langle \Phi^\top \boldsymbol{\mu}, \Phi^\top \hat{\boldsymbol{\mu}} / \|\Phi^\top \hat{\boldsymbol{\mu}}\|^2 + \mathbf{v} \rangle}{\sqrt{1 / \|\Phi^\top \hat{\boldsymbol{\mu}}\|^2 + \|\mathbf{v}\|^2}}$$

over  $\mathbf{v}$  such that  $\langle \mathbf{v}, \Phi^\top \hat{\boldsymbol{\mu}} \rangle = 0$ .

We apply Lemma 10 to  $-f(\mathbf{v})$  with  $\mathbf{v}_1 := \Phi^\top \hat{\boldsymbol{\mu}} / \|\Phi^\top \hat{\boldsymbol{\mu}}\|^2$  and  $\mathbf{v}_2 := -\Phi^\top \boldsymbol{\mu}$ . Then we obtain

$$f^* = \max_{\langle \mathbf{v}, \Phi^\top \hat{\boldsymbol{\mu}} \rangle = 0} f(\mathbf{v}) = \begin{cases} \|\Phi^\top \boldsymbol{\mu}\| \cdot |\sin \tilde{\theta}|, & \text{if } \cos \tilde{\theta} \geq 0, \\ \|\Phi^\top \boldsymbol{\mu}\|, & \text{if } \cos \tilde{\theta} < 0, \end{cases} \quad (4.20)$$

where  $\tilde{\theta}$  is the angle between  $\Phi^\top \boldsymbol{\mu}$  and  $-\Phi^\top \hat{\boldsymbol{\mu}}$ .

Now we aim to identify a regime for  $\sigma$  such that the global maximum of (4.20) matches the minimum of (4.18), i.e. the inequality constraints hold with high probability. It occurs if

$$\cos \tilde{\theta} < 0, \quad -\langle \boldsymbol{\mu}, \beta^* \rangle \mathbf{1}_n + \sigma \mathbf{Z}_2 \beta^* \leq 0$$

where  $\langle \boldsymbol{\mu}, \beta^* \rangle / \|\beta^*\| = \|\Phi^\top \boldsymbol{\mu}\|$ . For the first constraint, it holds

$$-\langle \Phi^\top \boldsymbol{\mu}, \Phi^\top \hat{\boldsymbol{\mu}} \rangle = -\|\Phi^\top \boldsymbol{\mu}\|^2 - \frac{\sigma}{n} \langle \mathbf{Z}_1^\top \mathbf{1}_n, \Phi^\top \boldsymbol{\mu} \rangle \sim \mathcal{N}(-\|\Phi^\top \boldsymbol{\mu}\|^2, \sigma^2 \|\Phi^\top \boldsymbol{\mu}\|^2 / n).$$

Therefore,

$$\Pr(\cos \tilde{\theta} < 0) = \Pr(Z < \sqrt{n} \sigma^{-1} \|\Phi^\top \boldsymbol{\mu}\|),$$

which means  $\cos \tilde{\theta} < 0$  holds with probability at least  $1 - O(n^{-1})$  if

$$\sigma \sqrt{\log n} < \sqrt{n} \|\Phi^\top \boldsymbol{\mu}\|. \quad (4.21)$$

For the second inequality to hold with probability at least  $1 - O(n^{-1})$ , we have

$$\Pr(-\langle \boldsymbol{\mu}, \beta^* \rangle + \sigma \langle \mathbf{z}, \beta^* \rangle < 0) = \Pr(Z < \sigma^{-1} \|\Phi^\top \boldsymbol{\mu}\|),$$

then we need

$$2\sigma \sqrt{\log n} < \|\Phi^\top \boldsymbol{\mu}\|, \quad (4.22)$$

by taking union bound over the  $n$  samples. As (4.22) is more strict than (4.21), we have the best possible generalization error bound is attained with probability at least  $1 - O(n^{-1})$  if

$$2\sigma \sqrt{\log n} < (1 - \varepsilon) \sqrt{\frac{d - n + 1}{d}} \|\boldsymbol{\mu}\|. \quad (4.23)$$

where we construct the lower bound (the RHS) for  $\|\Phi^\top \boldsymbol{\mu}\|$  by applying Johnson-Lindenstrauss lemma. Under condition (4.23), the best misclassification error is upper bounded by

$$\begin{aligned} \Pr\left(Z + \frac{f^*}{\sigma} < 0\right) &\leq \Pr\left(Z > \frac{1-\varepsilon}{\sigma} \sqrt{\frac{d-n+1}{d}} \|\boldsymbol{\mu}\|\right) \\ &\leq \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(1-\varepsilon)^2 \|\boldsymbol{\mu}\|^2}{2\sigma^2} \frac{d-n+1}{d}\right) \leq n^{-2} \end{aligned}$$

where  $f^* \geq (1-\varepsilon)\|\boldsymbol{\mu}\|\sqrt{(d-n+1)/d}$ . ■

#### 4.4.2 PROOF OF THEOREM 7(B): AN UPPER BOUND ON (4.18) WITH $d > 2n$

**Proof of Theorem 7(b).** Consider

$$\mathbf{X}_1^\top = \mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_1, \quad \mathbf{X}_2^\top = -\mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_2.$$

We assume  $\boldsymbol{\mu} = \|\boldsymbol{\mu}\| \mathbf{e}_1$  without loss of generality, and we will need to provide a lower bound of  $\min_{\boldsymbol{\beta}} \Pr_{z \sim \mathcal{N}(0,1)}(\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle + \sigma \|\boldsymbol{\beta}\| z \leq 0)$  subject to the constraints imposed by  $\mathcal{NC}$ . This lower bound leads us to consider program (4.18) again. The difference here is we assume  $d > 2n$ , i.e., the neural collapse occurs with probability 1. Our aim is to establish the best possible generalization bound v.s.  $\sigma/\|\boldsymbol{\mu}\|$ .

The idea to minimize (4.18) follows from two steps. We first consider

$$F(c, \gamma) := \max_{\boldsymbol{\beta} \in \mathbb{R}^d} \frac{\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle}{\|\boldsymbol{\beta}\|} \quad \text{s.t.} \quad (\mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_1) \boldsymbol{\beta} = \mathbf{1}_n, \quad (-\mathbf{1}_n \boldsymbol{\mu}^\top + \sigma \mathbf{Z}_2) \boldsymbol{\beta} = -\gamma, \quad \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle = c$$

for  $c > 0$  and  $\gamma \geq 0$ . Note that for any  $\gamma$  and  $c > 0$ , there exist feasible  $\boldsymbol{\beta}$ 's. To obtain the maximum of  $\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle / \|\boldsymbol{\beta}\|$  is equivalent to minimizing  $\|\boldsymbol{\beta}\|$ . Once we have that, the second step is to maximize over  $c$  and  $\gamma$ .

Now we compute  $F(c, \gamma)$  by minimizing  $\|\boldsymbol{\beta}\|$  subject to the constraints:

$$\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_1 \boldsymbol{\beta} = \mathbf{1}_n, \quad -\langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \mathbf{1}_n + \sigma \mathbf{Z}_2 \boldsymbol{\beta} = -\gamma, \quad \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle = c.$$

Note that the third constraint above implies  $\beta_1 = c \mathbf{e}_1 / \|\boldsymbol{\mu}\|$ . Substituting it into the first two equalities gives

$$\sigma \mathbf{Z}_1 \boldsymbol{\beta} = -c \mathbf{1}_n, \quad \sigma \mathbf{Z}_2 \boldsymbol{\beta} = -\gamma + c \mathbf{1}_n,$$

and then

$$\mathbf{Z}_{(-1)} \boldsymbol{\beta}_{(-1)} = \frac{1}{\sigma} \begin{bmatrix} (1-c) \mathbf{1}_n \\ c \mathbf{1}_n - \gamma \end{bmatrix} - \frac{c \mathbf{Z} \mathbf{e}_1}{\|\boldsymbol{\mu}\|}$$

where  $\mathbf{Z}_{(-1)}$  is a Gaussian matrix of size  $2n \times (d-1)$  that excludes the first column of  $\mathbf{Z}$ , and  $\boldsymbol{\beta}_{(-1)}$  is the same as  $\boldsymbol{\beta}$  after removing the first entry in  $\boldsymbol{\beta}$ . The minimum norm solution to  $\boldsymbol{\beta}_{(-1)}$  is given by

$$\boldsymbol{\beta}_{(-1)} = \mathbf{Z}_{(-1)}^\top (\mathbf{Z}_{(-1)} \mathbf{Z}_{(-1)}^\top)^{-1} \left( \frac{1}{\sigma} \begin{bmatrix} (1-c) \mathbf{1}_n \\ c \mathbf{1}_n - \gamma \end{bmatrix} - \frac{c \mathbf{Z} \mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right), \quad \beta_1 = \frac{c}{\|\boldsymbol{\mu}\|}$$

where  $d - 1 \geq 2n$  ensures the invertibility of  $\mathbf{Z}_{(-1)}$ . As a result, we have

$$F(c, \gamma) = \frac{c}{\sqrt{\left\| \mathbf{Z}_{(-1)}^\top (\mathbf{Z}_{(-1)} \mathbf{Z}_{(-1)}^\top)^{-1} \left( \frac{1}{\sigma} \begin{bmatrix} (1-c)\mathbf{1}_n \\ c\mathbf{1}_n - \gamma \end{bmatrix} - \frac{c\mathbf{Z}\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right) \right\|^2 + \frac{c^2}{\|\boldsymbol{\mu}\|^2}}}.$$

Maximizing  $F(c, \gamma)$  is equivalent to minimizing  $1/F^2(c, \gamma)$ :

$$\begin{aligned} \frac{1}{F^2(c, \gamma)} &= \frac{1}{c^2} \left\| \mathbf{Z}_{(-1)}^\top (\mathbf{Z}_{(-1)} \mathbf{Z}_{(-1)}^\top)^{-1} \left( \frac{1}{\sigma} \begin{bmatrix} (1-c)\mathbf{1}_n \\ c\mathbf{1}_n - \gamma \end{bmatrix} - \frac{c\mathbf{Z}\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right) \right\|^2 + \frac{1}{\|\boldsymbol{\mu}\|^2} \\ &= \left\| \mathbf{Z}_{(-1)}^\top (\mathbf{Z}_{(-1)} \mathbf{Z}_{(-1)}^\top)^{-1} \left( \frac{1}{\sigma} \begin{bmatrix} (1/c-1)\mathbf{1}_n \\ \mathbf{1}_n - \gamma/c \end{bmatrix} - \frac{\mathbf{Z}\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right) \right\|^2 + \frac{1}{\|\boldsymbol{\mu}\|^2}. \end{aligned}$$

For  $\mathbf{Z}_{(-1)} \in \mathbb{R}^{2n \times (d-1)}$  Gaussian random matrix, we have

$$\frac{d}{2} \mathbf{I} \preceq \mathbf{Z}_{(-1)} \mathbf{Z}_{(-1)}^\top$$

holds with high probability provided that  $d \geq 2n \log n$ . Let

$$G(c, \gamma) = \left\| \frac{1}{\sigma} \begin{bmatrix} (1/c-1)\mathbf{1}_n \\ \mathbf{1}_n - \gamma/c \end{bmatrix} - \frac{\mathbf{Z}\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right\|^2 = \left\| \frac{1-c}{c\sigma} \mathbf{1}_n + \frac{\mathbf{Z}_1\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right\|^2 + \left\| \frac{1}{\sigma} \left( \mathbf{1}_n - \frac{\gamma}{c} \right) - \frac{\mathbf{Z}_2\mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right\|^2$$

and then

$$\frac{G(c, \gamma)}{2d} + \frac{1}{\|\boldsymbol{\mu}\|^2} \leq \frac{1}{F^2(c, \gamma)}.$$

We proceed to compute the minimum of  $G(c, \gamma)$ . For any positive  $c > 0$ , the minimizer of the second term is attained at

$$\gamma = c \left[ \mathbf{1}_n - \frac{\sigma \mathbf{Z}_2 \mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right]_+$$

with minimum value

$$\min_{\gamma \geq 0} \left\| \frac{1}{\sigma} \left( \mathbf{1}_n - \frac{\gamma}{c} \right) - \frac{\mathbf{Z}_2 \mathbf{e}_1}{\|\boldsymbol{\mu}\|} \right\|^2 = \left\| \left[ \frac{\mathbf{Z}_2 \mathbf{e}_1}{\|\boldsymbol{\mu}\|} - \frac{1}{\sigma} \mathbf{1}_n \right]_+ \right\|^2,$$

which is independent of  $c$ . So we are left to minimize the first term, which is a quadratic function in  $1/c$ ,

$$\begin{aligned} \operatorname{argmin}_{c > 0} \frac{1}{\sigma^2} \left\| \frac{1}{c} \mathbf{1}_n + \left[ \frac{\sigma \mathbf{Z}_1 \mathbf{e}_1}{\|\boldsymbol{\mu}\|} - \mathbf{1}_n \right] \right\|^2 &= \operatorname{argmin}_{c > 0} \frac{n}{c^2} + \frac{2}{c} \left\langle \mathbf{1}_n, \frac{\sigma \mathbf{Z}_1 \mathbf{e}_1}{\|\boldsymbol{\mu}\|} - \mathbf{1}_n \right\rangle \\ &= \begin{cases} \frac{n\|\boldsymbol{\mu}\|}{n\|\boldsymbol{\mu}\| - \sigma \mathbf{1}_n^\top \mathbf{Z}_1 \mathbf{e}_1}, & \text{if } n\|\boldsymbol{\mu}\| - \sigma \mathbf{1}_n^\top \mathbf{Z}_1 \mathbf{e}_1 > 0, \\ +\infty, & \text{else,} \end{cases} \end{aligned} \tag{4.24}$$

and thus

$$\min_{c, \gamma} G(c, \gamma) = \frac{1}{\sigma^2} \|[s\mathbf{Z}_2\mathbf{e}_1 - \mathbf{1}_n]_+\|^2 + \begin{cases} \frac{1}{\|\boldsymbol{\mu}\|^2} \left\| \left( \mathbf{I}_n - \frac{\mathbf{1}_n\mathbf{1}_n^\top}{n} \right) \mathbf{Z}_1\mathbf{e}_1 \right\|^2, & \text{if } n\|\boldsymbol{\mu}\| - \sigma\mathbf{1}_n^\top\mathbf{Z}_1\mathbf{e}_1 > 0, \\ \frac{1}{\sigma^2} \|\mathbf{1}_n + s\mathbf{Z}_1\mathbf{e}_1\|^2, & \text{else} \end{cases}$$

where  $s = \sigma/\|\boldsymbol{\mu}\|$ . Now we provide a high probability value bound for the second term in two cases respectively.

**Case 1:**  $n\|\boldsymbol{\mu}\| - \sigma\mathbf{1}_n^\top\mathbf{Z}_1\mathbf{e}_1 \leq 0$ : The term is a sum of exponential random variable:  $\mathbb{E}|1 + sZ|^2 = 1 + s^2$  for  $Z \sim \mathcal{N}(0, 1)$  and

$$\mathbb{E}(|1 + sZ|^2 - 1 - s^2)^2 = s^2 \mathbb{E}(2Z + sZ^2 - s)^2 = s^2(4 + 2s^2).$$

Then we have

$$\left| \|s\mathbf{Z}_1\mathbf{e}_1 + \mathbf{1}_n\|^2 - n(1 + s^2) \right| \leq Cs\sqrt{(4 + 2s^2)n \log n} \quad (4.25)$$

holds with probability at least  $1 - O(n^{-2})$ .

**Case 2:**  $n\|\boldsymbol{\mu}\| - \sigma\mathbf{1}_n^\top\mathbf{Z}_1\mathbf{e}_1 > 0$ : In this case, we want to obtain concentration bound for quadratic form  $\mathbf{z}^\top \mathbf{C}_n \mathbf{z}$  where  $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_n)$ . Note that  $\mathbf{z}^\top \mathbf{C}_n \mathbf{z} \sim \chi_{n-1}^2$ , and then we have

$$\left| \mathbf{z}^\top \mathbf{C}_n \mathbf{z} - \mathbb{E} \mathbf{z}^\top \mathbf{C}_n \mathbf{z} \right| > c\sqrt{n \log n} \quad (4.26)$$

with probability as least  $1 - O(n^{-1})$  where we use the fact that  $\|\mathbf{C}_n\|_{op} = 1$  and we have  $\mathbb{E} \mathbf{z}^\top \mathbf{C}_n \mathbf{z} = \text{Tr}(\mathbf{C}_n) = n - 1$ .

For the first term, we notice it is also a sum of exponential random variables, so we have

$$\left| \|[-\mathbf{1}_n + s\mathbf{Z}_2\mathbf{e}_1]_+\|^2 - n \mathbb{E}[-1 + sZ]_+^2 \right| \leq Cs\sqrt{(4 + 2s^2)n \log n} \quad (4.27)$$

holds with probability at least  $1 - O(n^{-2})$ .

In expectation, it holds that

$$\begin{aligned} \mathbb{E}[-1 + sZ]_+^2 &= \frac{1}{\sqrt{2\pi s^2}} \int_0^\infty x^2 \exp\left(-\frac{(x+1)^2}{2s^2}\right) dx \\ &= \frac{1}{\sqrt{2\pi s^2}} \int_1^\infty (x-1)^2 \exp\left(-\frac{x^2}{2s^2}\right) dx \\ &= \frac{1}{\sqrt{2\pi s^2}} \left( (s^2 + 1) \int_1^\infty e^{-\frac{x^2}{2s^2}} dx - s^2 e^{-\frac{1}{2s^2}} \right). \end{aligned}$$

Note that

$$\begin{aligned} \int_1^\infty x \exp\left(-\frac{x^2}{2s^2}\right) dx &= -s^2 \int_1^\infty de^{-\frac{x^2}{2s^2}} = s^2 \exp\left(-\frac{1}{2s^2}\right), \\ \int_1^\infty x^2 \exp\left(-\frac{x^2}{2s^2}\right) dx &= -s^2 \int_1^\infty x de^{-\frac{x^2}{2s^2}} = s^2 e^{-\frac{1}{2s^2}} + s^2 \int_1^\infty e^{-\frac{x^2}{2s^2}} dx, \\ \int_1^\infty \exp\left(-\frac{x^2}{2s^2}\right) dx &\leq \int_1^\infty x \exp\left(-\frac{x^2}{2s^2}\right) dx = s^2 e^{-\frac{1}{2s^2}}. \end{aligned}$$

As a result, we have

$$\mathbb{E}[-1 + sZ]_+^2 \leq \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}}.$$

Hence, collecting all the concentration bounds (4.25),(4.26),(4.27) we have derived above, it holds with at least probability  $1 - O(n^{-1})$ ,

$$\min G(c, \gamma) \geq \begin{cases} \frac{n}{\sigma^2} \left( \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}} + s^2 - (c'_1 s^2 + c'_2 s) \sqrt{\frac{\log n}{n}} \right), & \text{if } s \mathbf{1}_n^\top \mathbf{Z}_1 \mathbf{e}_1 < n, \\ \frac{n}{\sigma^2} \left( \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}} + s^2 + 1 - (c'_3 s^2 + c'_4 s) \sqrt{\frac{\log n}{n}} \right), & \text{else,} \end{cases}$$

for some positive constants  $c'_1, c'_2, c'_3$ , and  $c'_4$ . Recall that we have

$$F^* = \max F(c, \gamma) \leq \left( \frac{\min G(c, \gamma)}{2d} + \frac{1}{\|\boldsymbol{\mu}\|^2} \right)^{-1/2}.$$

As a result, the best possible misclassification error is lower bounded by:

$$\Pr \left( \frac{F^*}{\sigma} + Z < 0 \right) \geq \Pr \left( Z > \left( \frac{n}{2d} \left( \frac{s^3 e^{-\frac{1}{2s^2}}}{\sqrt{2\pi}} + s^2 + 1 - (c_1 s^2 + c_2 s) \sqrt{\frac{\log n}{n}} \right) + s^2 \right)^{-1/2} \right) \quad (4.28)$$

where  $c_1 > 0$  and  $c_2 > 0$ . ■

## 5. Appendix

**Lemma 10.** *Let  $\mathbf{v}_1$  and  $\mathbf{v}_2 \in \mathbb{R}^d$  be two arbitrary vectors. Then it holds*

$$\min_{\langle \mathbf{v}, \mathbf{v}_1 \rangle = 0} h(\mathbf{v}) := \frac{\langle \mathbf{v}_1 + \mathbf{v}, \mathbf{v}_2 \rangle}{\sqrt{\|\mathbf{v}_1\|^2 + \|\mathbf{v}\|^2}} = \begin{cases} -\sqrt{\|\mathbf{v}_2\|^2 - |\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|^2 / \|\mathbf{v}_1\|^2}, & \text{if } \langle \mathbf{v}_1, \mathbf{v}_2 \rangle \geq 0, \\ -\|\mathbf{v}_2\|, & \text{if } \langle \mathbf{v}_1, \mathbf{v}_2 \rangle < 0. \end{cases}$$

**Proof.** Note that  $\mathbf{v}$  is in a linear subspace, and thus the global minimum is certainly negative. Let  $c = \langle \mathbf{v}_2, \mathbf{v}_1 + \mathbf{v} \rangle < 0$  be fixed and then  $h(\mathbf{v})$  is increasing w.r.t.  $\|\mathbf{v}\|$ . As a result, we need to search for a  $\mathbf{v}$  with the minimum norm subject to

$$\langle \mathbf{v}, \mathbf{v}_1 \rangle = 0, \quad \langle \mathbf{v}, \mathbf{v}_2 \rangle = c - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle.$$

Let  $\mathbf{P} = \mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^\top / \|\mathbf{v}_1\|^2$  be the projection matrix onto the complement of  $\text{span}(\mathbf{v}_1)$ . The minimum is given by

$$\mathbf{v} = (c - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle) \frac{\mathbf{P} \mathbf{v}_2}{\|\mathbf{P} \mathbf{v}_2\|^2}.$$

Therefore, the minimization of  $h(\mathbf{v})$  is reduced to minimizing  $h(c)$  for  $c < 0$ :

$$\begin{aligned} h(c) &:= \frac{c}{\sqrt{\|\mathbf{v}_1\|_2^2 + (c - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle)^2 / \|\mathbf{P} \mathbf{v}_2\|^2}} \\ &= \frac{c \|\mathbf{P} \mathbf{v}_2\|}{\sqrt{\|\mathbf{v}_1\|^2 \|\mathbf{P} \mathbf{v}_2\|^2 + (c - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle)^2}} \\ &= \frac{-\|\mathbf{P} \mathbf{v}_2\|}{\sqrt{(\|\mathbf{v}_1\|^2 \|\mathbf{P} \mathbf{v}_2\|^2 + \langle \mathbf{v}_1, \mathbf{v}_2 \rangle^2) / c^2 + 2 \langle \mathbf{v}_1, \mathbf{v}_2 \rangle / c + 1}}. \end{aligned}$$

If  $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle \geq 0$ , then function  $h(c)$  satisfies

$$\inf_{c < 0} h(c) = -\|\mathbf{P}\mathbf{v}_2\| = -\sqrt{\|\mathbf{v}_2\|^2 - \frac{|\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|^2}{\|\mathbf{v}_1\|^2}}.$$

If  $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle < 0$ , then the denominator is a quadratic function of  $1/|c|$ . The minimum is achieved at

$$\frac{1}{c} = \frac{\langle \mathbf{v}_1, \mathbf{v}_2 \rangle}{\|\mathbf{v}_1\|^2 \|\mathbf{P}\mathbf{v}_2\|^2 + \langle \mathbf{v}_1, \mathbf{v}_2 \rangle^2} = \frac{\langle \mathbf{v}_1, \mathbf{v}_2 \rangle}{\|\mathbf{v}_1\|^2 \|\mathbf{v}_2\|^2}$$

and then  $\min_{c < 0} h(c) = -\|\mathbf{v}_2\|$ . ■

**Lemma 11 (Johnson-Lindenstrauss lemma for angles).** *Let  $\Phi \in \mathbb{R}^{n \times m}$  be a random orthogonal matrix with  $m > n$ . For any two unit vectors  $\mathbf{v}_1$  and  $\mathbf{v}_2 \in \mathbb{R}^n$ , we denote the angle between  $\mathbf{v}_1$  and  $\mathbf{v}_2$  and  $\Phi^\top \mathbf{v}_1$  and  $\Phi^\top \mathbf{v}_2$  by  $\theta$  and  $\tilde{\theta}$  respectively. Then with probability at least  $1 - 4 \exp(-c\varepsilon^2 m)$ ,*

$$|\cos \tilde{\theta} - \cos \theta| \leq \frac{4\varepsilon}{(1 - \varepsilon)^2}.$$

**Proof.** Without loss of generality, we assume  $\|\mathbf{v}_1\|_2 = \|\mathbf{v}_2\|_2 = 1$ . By JL lemma, we have with probability at least  $1 - 2 \exp(-c\varepsilon^2 m)$ , for  $k = 1, 2$ ,

$$(1 - \varepsilon) \sqrt{\frac{m}{n}} \|\mathbf{v}_k\|_2 \leq \|\Phi^\top \mathbf{v}_k\|_2 \leq (1 + \varepsilon) \sqrt{\frac{m}{n}} \|\mathbf{v}_k\|_2. \quad (5.1)$$

By applying union bound over  $\|\Phi^\top \mathbf{v}_k\|_2$ , ( $k = 1, 2$ ), we have with probability at least  $1 - 4 \exp(-c\varepsilon^2 m)$ , the following bounds hold. We proceed to estimate  $\cos \tilde{\theta}$ :

$$\begin{aligned} \cos \tilde{\theta} &= \frac{\|\Phi^\top(\mathbf{v}_1 + \mathbf{v}_2)\|^2 - \|\Phi^\top(\mathbf{v}_1 - \mathbf{v}_2)\|^2}{4\|\Phi^\top \mathbf{v}_1\| \|\Phi^\top \mathbf{v}_2\|} \\ &\leq \frac{(1 + \varepsilon)^2(1 + \cos \theta) - (1 - \varepsilon)^2(1 - \cos \theta)}{2(1 - \varepsilon)^2} = \frac{2\varepsilon + (1 + \varepsilon^2) \cos \theta}{(1 - \varepsilon)^2}. \end{aligned}$$

As a result, it holds

$$\cos \tilde{\theta} - \cos \theta \leq \frac{2\varepsilon + (1 + \varepsilon^2) \cos \theta - (1 - \varepsilon)^2 \cos \theta}{(1 - \varepsilon)^2} = \frac{2\varepsilon + 2\varepsilon \cos \theta}{(1 - \varepsilon)^2}.$$

Similarly, for the lower bound, we have

$$\begin{aligned} \cos \tilde{\theta} &\geq \frac{\|\Phi^\top(\mathbf{v}_1 + \mathbf{v}_2)\|^2 - \|\Phi^\top(\mathbf{v}_1 - \mathbf{v}_2)\|^2}{4\|\Phi^\top \mathbf{v}_1\| \|\Phi^\top \mathbf{v}_2\|} \\ &\geq \frac{(1 - \varepsilon)^2(1 + \cos \theta) - (1 + \varepsilon)^2(1 - \cos \theta)}{2(1 + \varepsilon)^2} = -\frac{2\varepsilon}{(1 + \varepsilon)^2} + \cos \theta. \end{aligned}$$

Therefore, for  $\cos \theta > 0$ , we have

$$-\frac{2\varepsilon}{(1 + \varepsilon)^2} \leq \cos \tilde{\theta} - \cos \theta \leq \frac{4\varepsilon}{(1 - \varepsilon)^2} \iff |\cos \tilde{\theta} - \cos \theta| \leq \frac{4\varepsilon}{(1 - \varepsilon)^2}.$$

For  $\cos \theta < 0$ , we consider  $-\mathbf{v}_1$  and  $\mathbf{v}_2$  instead, and the same bound holds. ■

**Lemma 12 (Johnson-Lindenstrauss lemma for singular values).** *Let  $\mathbf{V} \in \mathbb{R}^{K \times d}$  and  $\Phi \in \mathbb{R}^{d \times m}$  is a random subspace sampled uniformly from  $G_{d,m} \subset \mathbb{R}^{d \times m}$  with  $m < d$  where  $G_{d,m}$  stands for the Grassmannian consisting of all  $m$ -dimensional subspaces in  $\mathbb{R}^d$ , then with probability at least  $1 - K^2 \exp(-c\varepsilon^2 m)$  for some constant  $c > 0$ , the following inequalities hold,*

$$\begin{aligned}\sigma_{\max}(\mathbf{V}\Phi) &\leq \sqrt{\frac{m}{d}(1 + \varepsilon^2 + 2K\varepsilon)}\sigma_{\max}(\mathbf{V}), \\ \sigma_{\min}(\mathbf{V}\Phi) &\geq \sqrt{\frac{m}{d}(1 - \varepsilon^2 - 2K\varepsilon)}\sigma_{\min}(\mathbf{V}),\end{aligned}\tag{5.2}$$

where  $\sigma_{\max}(\mathbf{V})$  and  $\sigma_{\min}(\mathbf{V})$  denote the largest and smallest singular value of matrix  $\mathbf{V}$ . In other words, under  $m > c^{-1}\varepsilon^{-2} \log(K^2 d)$ , the probability of success is at least  $1 - O(d^{-1})$ .

**Proof.** The idea is to apply the Johnson-Lindenstrauss lemma to bound the quadratic form  $\mathbf{a}^\top (\mathbf{V}\Phi\Phi^\top \mathbf{V}^\top) \mathbf{a}$  for  $\mathbf{a} \in \mathbb{S}^{K-1}$  where  $\Phi \in \mathbb{R}^{d \times m}$  is a random orthogonal matrix with  $\Phi^\top \Phi = \mathbf{I}_m$ . Without loss of generality, we can assume  $\mathbf{V}\mathbf{V}^\top = \mathbf{I}_K$ , as we can perform the SVD on  $\mathbf{V}$  and only take care of the orthogonal parts.

Let  $\mathbf{V}^\top = [\mathbf{v}_1, \dots, \mathbf{v}_K] \in \mathbb{R}^{d \times K}$  with orthogonal columns. Then by applying Johnson-Lindenstrauss lemma on  $\mathbf{v}_i$ , see (Vershynin, 2018, Lemma 5.3.2), we have with probability at least  $1 - K^2 \exp(-c\varepsilon^2 m)$ ,

$$\begin{aligned}(1 - \varepsilon)\sqrt{\frac{m}{d}} &\leq \|\Phi^\top \mathbf{v}_i\| \leq (1 + \varepsilon)\sqrt{\frac{m}{d}}, \\ \sqrt{2}(1 - \varepsilon)\sqrt{\frac{m}{d}} &\leq \|\Phi^\top (\mathbf{v}_i + \mathbf{v}_j)\| \leq \sqrt{2}(1 + \varepsilon)\sqrt{\frac{m}{d}},\end{aligned}\tag{5.3}$$

where we use the fact that  $\|\mathbf{v}_i\| = 1$  and  $\|\mathbf{v}_i \pm \mathbf{v}_j\| = \sqrt{2}$  ( $i \neq j$ ). We have

$$\begin{aligned}\mathbf{a}^\top (\mathbf{V}\Phi\Phi^\top \mathbf{V}^\top) \mathbf{a} &= \sum_{i=1}^K a_i^2 \|\Phi^\top \mathbf{v}_i\|^2 + \sum_{i=1}^K \sum_{j \neq i}^K a_i a_j \langle \Phi^\top \mathbf{v}_i, \Phi^\top \mathbf{v}_j \rangle \\ &= \sum_{i=1}^K a_i^2 \|\Phi^\top \mathbf{v}_i\|^2 + \sum_{i=1}^K \sum_{j \neq i}^K \frac{a_i a_j}{2} \left( \|\Phi^\top (\mathbf{v}_i + \mathbf{v}_j)\|^2 - \|\Phi^\top \mathbf{v}_i\|^2 - \|\Phi^\top \mathbf{v}_j\|^2 \right).\end{aligned}\tag{5.4}$$

Applying (5.3), we obtain

$$\|\Phi^\top (\mathbf{v}_i + \mathbf{v}_j)\|^2 - \|\Phi^\top \mathbf{v}_i\|^2 - \|\Phi^\top \mathbf{v}_j\|^2 \leq \frac{2m}{d} (1 + \varepsilon^2 - (1 - \varepsilon)^2) \leq \frac{4m\varepsilon}{d}.\tag{5.5}$$

Therefore by (5.3), (5.4) and (5.5), it holds

$$\begin{aligned}\mathbf{a}^\top (\mathbf{V}\Phi\Phi^\top \mathbf{V}^\top) \mathbf{a} &\leq \frac{m}{d} \sum_{i=1}^K a_i^2 (1 + \varepsilon)^2 + \frac{2\varepsilon m}{d} \sum_{i=1}^K \sum_{j \neq i}^K |a_i a_j| \\ &\leq \frac{m}{d} ((1 + \varepsilon)^2 + 2(K - 1)\varepsilon) = \frac{m}{d} (1 + \varepsilon^2 + 2K\varepsilon)\end{aligned}\tag{5.6}$$

where  $\sum_{i \neq j} |a_i a_j| \leq K - 1$ . Similarly,

$$\begin{aligned} \mathbf{a}^\top (\mathbf{V} \Phi \Phi^\top \mathbf{V}^\top) \mathbf{a} &\geq \frac{m}{d} \sum_{i=1}^K a_i^2 (1 - \varepsilon)^2 - \frac{2\varepsilon m}{d} \sum_{i=1}^K \sum_{j \neq i}^K |a_i a_j| \\ &\geq \frac{m}{d} ((1 - \varepsilon)^2 - 2(K - 1)\varepsilon) = \frac{m}{d} (1 + \varepsilon^2 - 2K\varepsilon). \end{aligned} \quad (5.7)$$

■

**Proposition 13.** *Let  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  be a standard Gaussian random matrix with i.i.d. entries. Then it holds*

$$\mathbb{E} \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} \langle \mathbf{s}, \mathbf{Z} \mathbf{u} \rangle \geq \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n), \mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_d)} [\|\mathbf{h}\| - \|\mathbf{g}_+\|] \sim \sqrt{d} - \sqrt{n/2}$$

where  $\mathbf{s} \in \mathbb{R}^n$  and  $\mathbf{u} \in \mathbb{R}^d$ , and  $\mathbf{g}_+$  is the positive part of  $\mathbf{g}$ . Here

$$\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n)} \|\mathbf{g}_+\| \leq \sqrt{\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n)} \|\mathbf{g}_+\|^2}$$

and

$$\mathbb{E}_{\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_d)} \|\mathbf{h}\| = \frac{\sqrt{2}\Gamma(\frac{d+1}{2})}{\Gamma(\frac{d}{2})} \sim \sqrt{\frac{d-1}{d-2}} \left(1 + \frac{1}{d-2}\right)^{\frac{d-2}{2}} \frac{\sqrt{d-1}}{\sqrt{e}} \sim \sqrt{d}.$$

The core idea of the proof relies on Gordon's inequality.

**Theorem 14 (Gordon's inequality).** *Suppose  $X_{u,t}$  and  $Y_{u,t}$  are two Gaussian processes indexed by  $(u, t)$ . Assume that*

$$\begin{aligned} \mathbb{E}(X_{u,t} - X_{u,s})^2 &= \mathbb{E}(Y_{u,t} - Y_{u,s})^2, \quad \forall u, t, s, \\ \mathbb{E}(X_{u,t} - X_{v,s})^2 &\geq \mathbb{E}(Y_{u,t} - Y_{v,s})^2, \quad \forall u \neq v, \text{ and } t, s, \end{aligned}$$

then

$$\mathbb{E} \inf_u \sup_t X_{u,t} \leq \mathbb{E} \inf_u \sup_t Y_{u,t}.$$

**Proof of Proposition 13.** Consider two Gaussian processes:

$$X_{s,u} = \langle \mathbf{s}, \mathbf{Z} \mathbf{u} \rangle, \quad Y_{s,u} = \langle \mathbf{s}, \mathbf{g} \rangle + \langle \mathbf{h}, \mathbf{u} \rangle$$

where  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n)$  and  $\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_d)$ , and all Then

$$\begin{aligned} \mathbb{E} |X_{s,u} - X_{t,v}|^2 &= \|\mathbf{u} \mathbf{s}^\top - \mathbf{v} \mathbf{t}^\top\|_F^2 = 2 - 2\langle \mathbf{u}, \mathbf{v} \rangle \langle \mathbf{s}, \mathbf{t} \rangle \\ &\leq \|\mathbf{s} - \mathbf{t}\|^2 + \|\mathbf{u} - \mathbf{v}\|^2 = \mathbb{E} \|Y_{s,u} - Y_{t,v}\|^2. \end{aligned}$$

For  $\mathbf{s} = \mathbf{t}$ , it holds

$$\mathbb{E} \|X_{s,u} - X_{s,v}\|^2 = \|\mathbf{u} - \mathbf{v}\|^2 = \mathbb{E} \|Y_{s,u} - Y_{s,v}\|^2.$$

Note that

$$\mathbb{E} \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} \{\langle \mathbf{s}, \mathbf{g} \rangle + \langle \mathbf{h}, \mathbf{u} \rangle\} = \mathbb{E}[-\|\mathbf{g}_+\| + \|\mathbf{h}\|].$$

By Gordon's inequality, it holds that

$$\begin{aligned} \mathbb{E} \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} \langle \mathbf{s}, \mathbf{Z}\mathbf{u} \rangle &= \mathbb{E} \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} X_{\mathbf{s},\mathbf{u}} \\ &\geq \mathbb{E} \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} Y_{\mathbf{s},\mathbf{u}} = \mathbb{E} [\|\mathbf{h}\| - \|\mathbf{g}_+\|] \end{aligned}$$

where  $\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_d)$  and  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n)$ . ■

**Lemma 15.** *Let  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  be a standard Gaussian random matrix. Define*

$$f(\mathbf{Z}) := \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \max_{\mathbf{u} \in \mathbb{S}^{d-1}} \langle \mathbf{s}, \mathbf{Z}\mathbf{u} \rangle,$$

we have

$$\Pr(|f(\mathbf{Z}) - \mathbb{E} f(\mathbf{Z})| \geq t) \leq 2e^{-\frac{t^2}{2}}. \quad (5.8)$$

**Proof.** First note that

$$f(\mathbf{Z}) = \min_{\mathbf{s} \in \mathbb{S}_+^{n-1}} \|\mathbf{Z}^\top \mathbf{s}\|.$$

Given two Gaussian random matrices  $\mathbf{Z}_1$  and  $\mathbf{Z}_2$ , and we let  $\mathbf{s}_1$  and  $\mathbf{s}_2$  be two positive unit vectors such that

$$f(\mathbf{Z}_\ell) = \|\mathbf{Z}_\ell^\top \mathbf{s}_\ell\|, \quad \ell = 1, 2.$$

Then

$$\|\mathbf{Z}_2^\top \mathbf{s}_2\| - \|\mathbf{Z}_1^\top \mathbf{s}_2\| \leq f(\mathbf{Z}_2) - f(\mathbf{Z}_1) = \|\mathbf{Z}_2^\top \mathbf{s}_2\| - \|\mathbf{Z}_1^\top \mathbf{s}_1\| \leq \|\mathbf{Z}_2^\top \mathbf{s}_1\| - \|\mathbf{Z}_1^\top \mathbf{s}_1\|$$

and it implies

$$|f(\mathbf{Z}_2) - f(\mathbf{Z}_1)| \leq \|\mathbf{Z}_2 - \mathbf{Z}_1\|_F,$$

i.e.,  $f(\mathbf{Z})$  is a Lipschitz-1 continuous function. Then following from (Wainwright, 2019, Theorem 2.26), we have  $\Pr(|f(\mathbf{Z}) - \mathbb{E} f(\mathbf{Z})| \geq t) \leq 2e^{-\frac{t^2}{2}}$ . ■

## References

- Tina Behnia, Ganesh Ramachandra Kini, Vala Vakilian, and Christos Thrampoulidis. On the implicit geometry of cross-entropy parameterizations for label-imbalanced data. In *International Conference on Artificial Intelligence and Statistics*, pages 10815–10838. PMLR, 2023.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.
- Hien Dang, Tan Nguyen, Tho Tran, Hung Tran, and Nhat Ho. Neural collapse in deep linear network: From balanced to imbalanced data. *arXiv preprint arXiv:2301.00437*, 2023a.

- Hien Dang, Tho Tran, Stanley Osher, Hung Tran-The, Nhat Ho, and Tan Nguyen. Neural collapse in deep linear networks: from balanced to imbalanced data. In *Proceedings of the 40th International Conference on Machine Learning*, pages 6873–6947, 2023b.
- Hien Dang, Tho Tran, Tan Nguyen, and Nhat Ho. Neural collapse for cross-entropy class-imbalanced learning with unconstrained ReLU feature model. *arXiv preprint arXiv:2401.02058*, 2024.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2018.
- Weinan E and Stephan Wojtowytsch. On the emergence of simplex symmetry in the final and penultimate layers of neural network classifiers. In *Mathematical and Scientific Machine Learning*, pages 270–290. PMLR, 2022.
- Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.
- Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning. In *International Conference on Learning Representations*, 2021.
- Connall Garrod and Jonathan P Keating. Unifying low dimensional observations in deep learning through the deep linear unconstrained feature model. *arXiv preprint arXiv:2404.06106*, 2024.
- XY Han, Vardan Papyan, and David L Donoho. Neural collapse under MSE loss: Proximity to and dynamics on the central path. In *International Conference on Learning Representations*, 2021.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- Wanli Hong and Shuyang Ling. Neural collapse for unconstrained feature model under cross-entropy loss with imbalanced data. *Journal of Machine Learning Research*, 25(192):1–48, 2024.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- Like Hui, Mikhail Belkin, and Preetum Nakkiran. Limitations of neural collapse for understanding generalization in deep learning. *arXiv preprint arXiv:2202.08384*, 2022.
- Arthur Jacot, Peter Sůkeník, Zihan Wang, and Marco Mondelli. Wide neural networks trained with weight decay provably exhibit neural collapse. *arXiv preprint arXiv:2410.04887*, 2024.

- Wenlong Ji, Yiping Lu, Yiliang Zhang, Zhun Deng, and Weijie J Su. An unconstrained layer-peeled perspective on neural collapse. In *International Conference on Learning Representations*, 2022.
- Vignesh Kothapalli, Ebrahim Rasromani, and Vasudev Awatramani. Neural collapse: A review on modelling principles and generalization. *arXiv preprint arXiv:2206.04041*, 2022.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015.
- Xiao Li, Sheng Liu, Jinxin Zhou, Xinyu Lu, Carlos Fernandez-Granda, Zhihui Zhu, and Qing Qu. Principled and efficient transfer learning of deep models via neural collapse. *arXiv preprint arXiv:2212.12206*, 2022.
- Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with ReLU activation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Jianfeng Lu and Stefan Steinerberger. Neural collapse under cross-entropy loss. *Applied and Computational Harmonic Analysis*, 59:224–241, 2022.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33): E7665–E7671, 2018.
- Dustin G Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *Sampling Theory, Signal Processing, and Data Analysis*, 20(2):11, 2022.
- Duc Anh Nguyen, Ron Levie, Julian Lienen, Eyke Hüllermeier, and Gitta Kutyniok. Memorization-dilation: Modeling neural collapse under noise. In *The Eleventh International Conference on Learning Representations*, 2023.
- Vardan Pappayan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- Tomaso Poggio and Qianli Liao. Explicit regularization and implicit bias in deep network classifiers trained with the square loss. *arXiv preprint arXiv:2101.00072*, 2021.
- Akshay Rangamani and Andrzej Banburski-Fahey. Neural collapse in deep homogeneous classifiers and the role of weight decay. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4243–4247. IEEE, 2022.
- Grant Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of artificial neural networks: An interacting particle system approach. *Communications on Pure and Applied Mathematics*, 75(9):1889–1935, 2022.
- Mariia Seleznova, Dana Weitzner, Raja Giryes, Gitta Kutyniok, and Hung-Hsu Chou. Neural (tangent kernel) collapse. volume 36, 2024.

- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Peter SÚkeník, Marco Mondelli, and Christoph H Lampert. Deep neural collapse is provably optimal for the deep unconstrained features model. *Advances in Neural Information Processing Systems*, 36, 2024.
- Christos Thrampoulidis, Ganesh Ramachandra Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry. *Advances in Neural Information Processing Systems*, 35:27225–27238, 2022.
- Tom Tirer and Joan Bruna. Extended unconstrained features model for exploring deep neural collapse. In *International Conference on Machine Learning*, pages 21478–21505. PMLR, 2022.
- Tom Tirer, Haoxiang Huang, and Jonathan Niles-Weed. Perturbation analysis of neural collapse. In *International Conference on Machine Learning*, pages 34301–34329. PMLR, 2023.
- Roman Vershynin. *High-dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge University Press, 2018.
- Martin J Wainwright. *High-dimensional Statistics: A Non-asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Peng Wang, Xiao Li, Can Yaras, Zhihui Zhu, Laura Balzano, Wei Hu, and Qing Qu. Understanding deep representation learning via layerwise feature compression and discrimination. *arXiv preprint arXiv:2311.02960*, 2023.
- Diyuan Wu and Marco Mondelli. Neural collapse beyond the unconstrained features model: Landscape, dynamics, and generalization in the mean-field regime. *arXiv preprint arXiv:2501.19104*, 2025.
- Yongyi Yang, Jacob Steinhardt, and Wei Hu. Are neurons actually collapsed? on the fine-grained structure in neural representations. In *International Conference on Machine Learning*, pages 39453–39487. PMLR, 2023.
- Can Yaras, Peng Wang, Zhihui Zhu, Laura Balzano, and Qing Qu. Neural collapse with normalized features: A geometric analysis over the riemannian manifold. *Advances in Neural Information Processing Systems*, 35:11547–11560, 2022.
- Jinxin Zhou, Xiao Li, Tianyu Ding, Chong You, Qing Qu, and Zhihui Zhu. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *International Conference on Machine Learning*, pages 27179–27202. PMLR, 2022.
- Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.