

Cross Entropy versus Label Smoothing: A Neural Collapse Perspective

Anonymous authors

Paper under double-blind review

Abstract

Label smoothing loss is a widely adopted technique to mitigate overfitting in deep neural networks. This paper studies label smoothing from the perspective of Neural Collapse (NC), a powerful empirical and theoretical framework which characterizes model behavior during the terminal phase of training. We first show empirically that models trained with label smoothing converge faster to neural collapse solutions and attain a stronger level of neural collapse compared to those trained with cross-entropy loss. Furthermore, we show that at the same level of NC1, models under label smoothing loss exhibit intensified NC2. These findings provide valuable insights into the impact of label smoothing on model performance and calibration. Then, leveraging the unconstrained feature model, we derive closed-form solutions for the global minimizers under both label smoothing and cross-entropy losses. We show that models trained with label smoothing have a lower conditioning number and, therefore, theoretically converge faster. Our study, combining empirical evidence and theoretical results, not only provides nuanced insights into the differences between label smoothing and cross-entropy losses, but also serves as an example of how the powerful neural collapse framework can be used to improve our understanding of DNNs.

1 Introduction

The effectiveness of a deep neural network (DNN) hinges significantly on the choice of the loss function during training. While cross-entropy loss is one of the most popular choices for classification tasks, many alternatives with improved empirical performance have been proposed. Among these, label smoothing (Szegedy et al., 2016; Müller et al., 2019; Lukasik et al., 2020) has emerged as a common technique to enhance the performance of DNNs. Instead of supervising the model training with one-hot key labels, label smoothing introduces a soft target label by blending the hard target label with a uniform distribution over the labels. This procedure is generally understood as a means of regularisation for improving the model’s generalizability.

In this paper, we examine the benefits of label smoothing from the perspective of neural collapse (Papayan et al., 2020), a new and powerful framework for obtaining an improved understanding of DNNs.

Background and Related Work

Deep neural networks (DNNs) typically consist of a non-linear feature extractor and a linear classification layer. Neural collapse, first observed in (Papayan et al., 2020), occurs in DNNs for classification tasks where the data is balanced and cross-entropy loss is employed. It characterizes the geometric properties of the features produced by the feature extractor and the weight vectors of the classifier during the terminal phase of training (TPT):

- NC1: The learned features of samples from the same class approach their respective class means.
- NC2: The collapsed features from different classes and the classification weight vectors approach the vertices of a simplex equiangular tight frame (ETF).
- NC3: Up to rescaling, the linear classifier weights approach the corresponding class means, giving a self-dual configuration.

Other than cross-entropy loss, neural collapse phenomena have been observed and studied under the mean-squared loss as well (Han et al., 2021; Poggio & Liao, 2020; Zhou et al., 2022a). Furthermore, research has extended neural collapse investigations to scenarios involving imbalanced data (Fang et al., 2021; Hong & Ling, 2023; Thrampoulidis et al., 2022; Dang et al., 2023; Hong & Ling, 2023) and cases with a large number of classes (Jiang et al., 2023).

In addition to being observed empirically, neural collapse has also been proven to arise mathematically. Beginning with [Mixon et al. \(2020\)](#), a series of studies have employed approximation models such as the Unconstrained Feature Models (UFMs) [Mixon et al. \(2020\)](#) or layer-peeled models ([Fang et al., 2021](#)) to provide theoretical evidence for the emergence of neural collapse. These optimization models simplify a DNN by treating the last-layer features as free variables to optimize over, which is justified due to the expressiveness of DNNs ([Hornik, 1991](#); [Cybenko, 1989](#); [Lu et al., 2017](#); [Shaham et al., 2018](#)). Theoretical advancements in neural collapse not only enhance our understanding of DNNs but also inspire new techniques to improve their performance in diverse applications, such as imbalanced learning ([Xie et al., 2023](#); [Liu et al., 2023](#)), transfer learning ([Galanti et al., 2021](#); [Li et al., 2022](#)), and continual learning ([Yu et al., 2022](#); [Yang et al., 2023](#)), etc.

Several studies have explored UFMs under various loss functions and regularization strategies ([Wojtowysch et al., 2020](#); [Zhu et al., 2021](#); [Dang et al., 2023](#); [Lu & Steinerberger, 2022](#); [Tirer & Bruna, 2022](#); [Tirer et al., 2023](#); [Yaras et al., 2022](#); [Zhou et al., 2022b](#)). Particularly, [Zhou et al. \(2022b\)](#) established that under the UFM model, the global minimizers for a wide range of loss functions, including cross-entropy loss and label smoothing loss, have the idealized neural collapse properties and that the UFM model has a benevolent landscape allowing the global minimizer to be effectively attained using iterative algorithms. However, their investigation does not provide insights into why label smoothing consistently outperforms cross-entropy loss, or why label smoothing converges faster during training. In this paper, leveraging the neural collapse framework, we conduct an in-depth investigation of training under these loss functions, aiming to explain the reasons behind the observed superiority of label smoothing loss over cross-entropy loss.

Our Contributions

Our study begins with a thorough empirical comparison between cross-entropy (CE) loss and label smoothing (LS) loss during the training process. Specifically, we carefully study how the last layer features and linear classifiers evolve during training. Our findings are as follows:

1. Compared with CE loss, models trained under LS loss exhibit accelerated convergence in terms of training error, testing error, and neural collapse metrics. Furthermore, they converge to a more pronounced level of NC1 and NC2.
2. Along with accelerated convergence, label smoothing loss maintains a distinct balance between NC1 and NC2. Notably, as compared with CE loss, LS loss results in a more pronounced level of NC2 when reaching a comparable level of NC1. We argue that this phenomenon originates from the implicit inductive bias introduced by LS loss, which equalizes the logits of all non-target classes and thus promotes the emergence of a simplex ETF structure in both the learned features and classification weight. We posit that the emphasis on NC2 in label smoothing (LS), which promotes maximally separable features between classes, enhances the model’s generalization performance. Conversely, an excessive level of NC1 may cause the features to overly specialize in the training data, hindering the model’s ability to generalize effectively.
3. Models trained with LS loss exhibit improved calibration [Guo et al. \(2017\)](#) by implicitly regularizing classification weights and features during training. However, if temperature scaling [Guo et al. \(2017\)](#) is applied as post-processing to counteract the regularization effect, label smoothing can lead to deteriorated model calibration. This is because an excessive level of NC1 under label smoothing can cause the model to be overconfident in its predictions, even when they are incorrect, thereby negatively impacting model calibration.

Of equal importance to the empirical results, we perform a mathematical analysis of the convergence properties of the UFM models under CE and LS losses. While [Zhou et al. \(2022a\)](#) demonstrate that, for a broad class of loss functions, including CE and LS losses, the global minimizers exhibit neural collapse

properties, the authors neither derive the exact form of these global minimizers nor thoroughly examine the landscape around them.

1. We first derive closed-form solutions for the global minimizers under both CE and LS loss functions, which explicitly depend on the smoothing hyperparameter δ .
2. Utilizing these closed-form solutions, we conduct a second-order theoretical analysis of the optimization landscape around their respective global optimizers. Within the UFM context, our mathematical analysis reveals that LS exhibits a more well-conditioned landscape around the global minimum, which facilitates the faster convergence observed in our empirical study.

This paper provides a significantly deeper understanding of why LS provides better convergence and performance than CE loss. Additionally, the paper illustrates how the powerful framework of neural collapse and its associated mathematical models can be employed to gain a more nuanced understanding of the “why” of DNNs. However, we acknowledge that theoretical evidence about how neural collapse impacts model generalization is lacking in this study. We hope that our work will inspire future research into the intricate interplay between neural collapse, convergence speed, and model generalizability.

2 Preliminaries

2.1 The Problem Setup

A deep neural network is comprised of two key components: a feature extractor and a linear classifier. The feature extractor $\phi_{\theta}(\cdot)$ is a nonlinear mapping that maps the input \mathbf{x} to the corresponding feature embedding $\mathbf{h} := \phi_{\theta}(\mathbf{x}) \in \mathbb{R}^d$. Meanwhile, the linear classifier involves a weight matrix $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_K] \in \mathbb{R}^{d \times K}$ and a bias vector $\mathbf{b} \in \mathbb{R}^K$. Consequently, the architecture of a deep neural network is captured by the following equation:

$$f_{\Theta}(\mathbf{x}) := \mathbf{W}^{\top} \phi_{\theta}(\mathbf{x}) + \mathbf{b}, \quad (1)$$

where $\Theta := \{\theta, \mathbf{W}, \mathbf{b}\}$ represents the set of all model parameters. In this work, we consider training a deep neural network using a balanced dataset denoted as $\{(\mathbf{x}_{ki}, \mathbf{y}_{ki})\}_{1 \leq k \leq K, 1 \leq i \leq n}$. This dataset consists of samples distributed across K distinct classes, with n samples allocated per class. Here, \mathbf{x}_{ki} represents the i -th sample from the k -th class, and \mathbf{y}_{ki} is a one-hot vector with unity solely in the k -th entry. Our objective is to learn the parameters Θ by minimizing the empirical risk over the total $N = nK$ training samples:

$$\min_{\Theta} \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n l(f_{\Theta}(\mathbf{x}_{ki}), \mathbf{y}_{ki}) + \frac{\lambda}{2} \|\Theta\|_F^2, \quad (2)$$

where $l(\cdot, \cdot)$ denotes the chosen loss function, and $\lambda > 0$ is the regularization parameter (i.e., the weight decay parameter).

2.2 Training Losses

To simplify the notation, we use $\mathbf{z} = \mathbf{W}^{\top} \phi_{\theta}(\mathbf{x}) + \mathbf{b}$ to represent the network’s output logit vector for a given input \mathbf{x} , and $\mathbf{p} = \text{Softmax}(\mathbf{z})$ to denote the predicted distribution from the model. The cross-entropy between the target distribution \mathbf{y} and the predicted distribution \mathbf{p} is defined as $l_{CE}(\mathbf{y}, \mathbf{p}) = -\sum_k y_k \log(p_k)$, where \mathbf{y} is a one-hot vector with a value of 1 in the dimension corresponding to the target class. In contrast, LS loss minimizes the cross-entropy between the smoothed soft label \mathbf{y}^{δ} and the predicted distribution \mathbf{p} , denoted as $l_{CE}(\mathbf{y}^{\delta}, \mathbf{p})$, where the soft label $\mathbf{y}^{\delta} = (1 - \delta)\mathbf{y} + (\delta/K)\mathbf{1}_K$ combines the hard ground truth label \mathbf{y} with a uniform distribution over the labels. Here $\mathbf{1}_K$ denotes the K -dimensional vector of all ones. The hyper-parameter δ determines the degree of smoothing.

For simplicity, we use the following formulation to represent both CE loss and LS loss:

$$l_{CE}(\mathbf{y}^{\delta}, \mathbf{p}) = -\sum_k y_k^{\delta} \log(p_k), \quad (3)$$

where $y_k^\delta = (1 - \delta)y_k + \delta/K$. The provided loss corresponds to CE loss when $\delta = 0$, and for any other value of $\delta \in (0, 1)$, it represents LS loss.

3 Empirical Analysis of Cross-Entropy and Label Smoothing Losses

This section conducts a comprehensive empirical comparison between cross-entropy (CE) loss and label smoothing (LS) loss from the perspective of neural collapse. In Section 3.1, we examine the convergence of the model to neural collapse solutions under both CE and LS losses. Our results demonstrate that while models under both loss functions converge to neural collapse, LS loss induces faster convergence and reaches a more pronounced level of neural collapse. Beyond convergence rates, Section 3.2 delves into the dynamics of convergence, revealing that LS loss introduces a bias toward solutions with a symmetric simplex ETF structure, thereby enforcing NC2. In Section 3.3, we attempt to understand the impact of label smoothing on model calibration from the perspective of neural collapse. In addition, we also investigate how the smoothing hyperparameter δ influences neural collapses during the model training in the appendix B.4.

Experiment Setup. We conducted experiments on CIFAR-10 (Krizhevsky et al., 2009), CIFAR-100, and STL-10 (Coates et al., 2011). For consistency with prior studies (Papayan et al., 2020; Zhu et al., 2021), we adopt ResNet-18 (He et al., 2016) as the backbone network for CIFAR-10 and ResNet-50 (He et al., 2016) as the backbone for CIFAR-100 and STL-10.

To focus on the behaviors associated with neural collapse and minimize the influence of other factors, we adopt standard preprocessing techniques without introducing any data augmentation for the training data. To comprehensively analyze the models' behavior during TPT, we extend the training period to 800 epochs. For all datasets, we employ a default batch size of 128, stochastic gradient descent with a momentum of 0.9, and a learning rate that is initialized at 0.05 and undergoes multi-step decay, decreasing by a factor of 0.1 at epochs 150 and 350. We use a default weight decay of 5×10^{-4} , except for the experiments in Section 3.3, where a weight decay value of 1×10^{-4} is used.

Metrics for Measuring NC. We assess neural collapse in the last-layer features and the classifiers using metrics based on the properties introduced in Section 1, with metrics similar to those presented in (Papayan et al., 2020). For convenience, we denote the global mean and class mean of the last-layer features as:

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

Within class variability (NC1) measures the relative magnitude of the within-class covariance matrix $\Sigma_W := \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n (\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)(\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)^\top$ compared to the between-class covariance matrix $\Sigma_B := \frac{1}{K} \sum_{k=1}^K (\bar{\mathbf{h}}_k - \mathbf{h}_G)(\bar{\mathbf{h}}_k - \mathbf{h}_G)^\top$ of the last-layer features. It is formulated as:

$$NC_1 = \frac{1}{K} \text{trace} \left(\Sigma_W \Sigma_B^\dagger \right), \quad (4)$$

where Σ_B^\dagger denotes the pseudo inverse of Σ_B .

Distance to simplex ETF (NC2) quantifies the difference between the product of the classifier weight matrix and the centered class mean feature matrix, and a simplex ETF, defined as follows¹:

$$NC_2 := \left\| \frac{\mathbf{W}^\top \bar{\mathbf{H}}}{\|\mathbf{W}^\top \bar{\mathbf{H}}\|_F} - \frac{1}{\sqrt{K-1}} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right) \right\|_F, \quad (5)$$

where $\bar{\mathbf{H}} = [\bar{\mathbf{h}}_1 - \mathbf{h}_G, \dots, \bar{\mathbf{h}}_K - \mathbf{h}_G] \in \mathbb{R}^{d \times K}$ represents centered class mean matrix.

Self-duality (NC3) measures the distance between the classifier weight matrix \mathbf{W} and the centered class-means $\bar{\mathbf{H}}$:

$$NC_3 := \left\| \frac{\mathbf{W}}{\|\mathbf{W}\|_F} - \frac{\bar{\mathbf{H}}}{\|\bar{\mathbf{H}}\|_F} \right\|_F. \quad (6)$$

¹When the bias \mathbf{b} is an all-zero vector or a constant vector, $NC_2 = 0$ implies that the average logit matrix, defined as $\bar{\mathbf{Z}} = \mathbf{W}^\top \bar{\mathbf{H}} + \mathbf{b} \mathbf{1}_K^\top$, satisfies $\bar{\mathbf{Z}} = a \left(\mathbf{I} - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right)$ for some constant a , i.e., $\bar{\mathbf{Z}}$ is a simplex ETF.

It is evident that when NC_2 in (5) and NC_3 in (6) both reach zero, the matrices \mathbf{W} and \mathbf{H} form the same simplex ETF up to some scaling factor. Thus, our definitions of NC1-NC3 capture the same concepts as the definitions in (Papayan et al., 2020) and (Zhu et al., 2021). We say that neural collapse occurs if NC1, NC2, and NC3 collectively approach zero during the Terminal Phase of Training (TPT).

3.1 Terminal Phase Training under Label Smoothing and Cross-Entropy Loss

In this section, we investigate the distinct behaviors exhibited by models trained under cross-entropy (CE) and label smoothing (LS) losses during TPT. We conduct experiments on CIFAR10, CIFAR100, and STL10 datasets, with LS loss employing a default smoothing hyperparameter of $\delta = 0.05$. We visualize the training process in Figure 1.

The leftmost column in Figure 1 (with a detailed view provided in Figure 5 in the Appendix) illustrates the progression of training and testing errors throughout the training process. As expected, models trained under LS loss exhibit lower testing errors compared to those under CE loss for all datasets, showcasing an improvement in model generalizability with LS loss. Furthermore, models trained with LS loss exhibit faster convergence for both training and testing errors. Figure 1 additionally presents the three neural collapse metrics for models trained under both cross-entropy and label smoothing losses. While the values of NC3, representing the alignment of \mathbf{W} and \mathbf{H} , remain consistently low and comparable under both loss functions, models with LS loss exhibit faster convergence in both NC1 and NC2 and eventually reach lower levels for both NC1 and NC2. We provide a theoretical explanation for these empirical findings in Section 4. Additionally, in Section 3.3, we establish a connection between NC1 and model calibration under both loss functions.

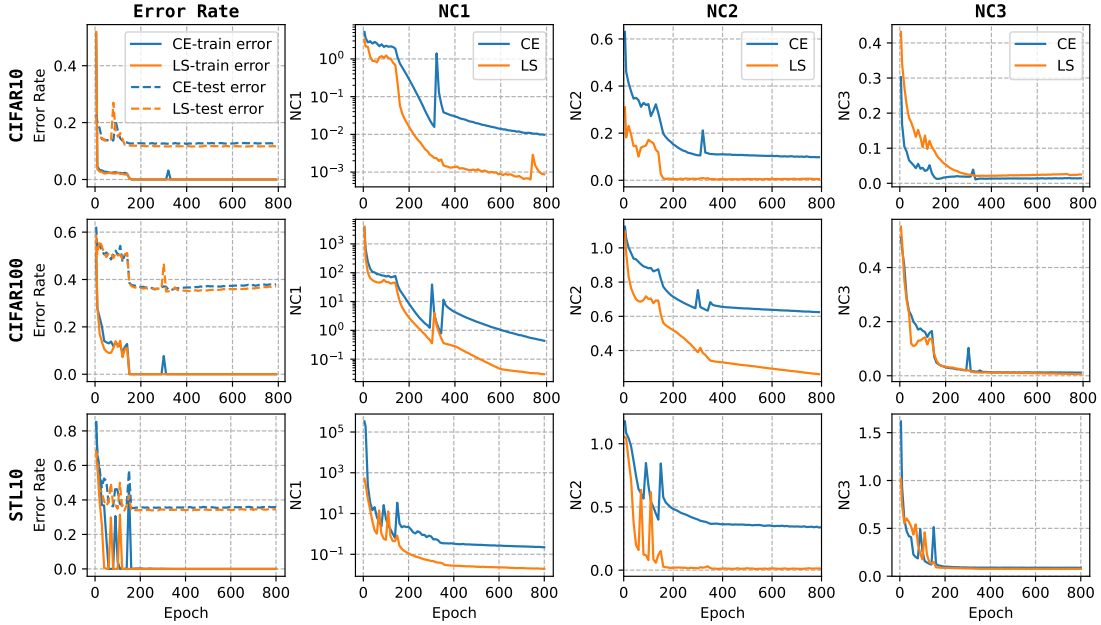


Figure 1: **Neural collapse with cross-entropy and label smoothing losses.** Columns from left to right represent the model’s error rate (Train/Test), NC1, NC2, and NC3 under CE and LS losses.

3.2 Label Smoothing Induces Enhanced NC2

Along with the faster convergence to neural collapse under LS loss, we further observe that LS loss maintains a distinct balance between NC1 and NC2 throughout training. Specifically, we find that for the same level of NC1, LS consistently provides a more pronounced manifestation of NC2 compared to CE Loss. Figure 2 provides scatter plots of NC1 and NC2 under both loss functions, with the NC metrics recorded every 10 training




Figure 2: Scatter plots for **NC1 vs. NC2 under CE/LS losses** with colors indicating the testing error rate, for CIFAR10 (left), CIFAR100 (middle), and STL10 (right) datasets.

epochs. Across all three datasets, the data points under LS loss consistently position to the left of those for CE loss. This observation suggests that at equivalent levels of NC1, LS loss induces an intensified level of NC2.

To gain insight into this phenomenon, we now closely examine the formulations of cross-entropy and label smoothing loss. Given an input \mathbf{x} , the output logit and the predicted distribution are equal to $\mathbf{z} = f_{\Theta}(\mathbf{x})$ and $\mathbf{p} = \text{Softmax}(\mathbf{z})$, respectively. As per Equation 3, assuming the observation \mathbf{x} is from the k -th class, CE loss is formulated as $l_{CE} = -\log(p_k)$. To minimize cross-entropy loss, the emphasis is solely on making the logit of the target class larger than the logits of non-target classes *without* constraining the logit variation among the non-target classes.

On the other hand, LS with smoothing hyperparameter δ is given by

$$l_{LS} = - \left(\left(1 - \frac{K-1}{K} \delta \right) \log(p_k) + \frac{\delta}{K} \sum_{l \neq k} \log(p_l) \right).$$

According to Jensen’s inequality, we have

$$\frac{1}{K-1} \sum_{l \neq k} \log(p_l) \leq \log \left(\frac{1}{K-1} \sum_{l \neq k} p_l \right) = \log(1 - p_k),$$

with equality achieved only if $p_l = p_{l'}$ (for $l, l' \neq k$). This implies that LS loss reaches its minimum only if the predicted probabilities for non-target classes are all equal. *Thus, LS loss strengthens the inductive bias towards equalizing the logits of the non-target classes.* This property aligns with the definition of NC_2 in (5), explaining why label smoothing loss reinforces NC2. As demonstrated in Section B.1 in the appendix, for deep neural networks with L2 regularization, the convergence of NC_2 as defined in (5) to zero indicates that both the classification weight \mathbf{W} and class mean features $\bar{\mathbf{H}}$ converge to a simplex ETF structure. The convergence of classifier weights and class mean features towards a simplex ETF promotes the development of maximally separable features and classifiers, a critical factor for enhancing overall model performance. We posit that the improvement in model generalizability observed with LS loss can be attributed to the prioritization of enhancing NC2 over NC1 during the training process. Nonetheless, this observation primarily serves as supporting evidence. Theoretical evidence and more rigorous validation of the benefits of NC2 remain topics for future research.

3.3 Label Smoothing and Model Calibration

DNN models often suffer from poor calibration, where the assigned probability values to class labels tend to overestimate the actual likelihood of correctness. This issue arises from the high capacity of DNN models, making them prone to overfitting the training data. Guo et al. (Guo et al., 2017) introduced the expected calibration error (ECE) as a metric for assessing model calibration. Additionally, they proposed temperature scaling as an effective post-processing technique for improving calibration, which involves dividing a network’s

logits by a scalar $T > 0$ before applying softmax, thereby softening (for $T > 1$) or sharpening (for $T < 1$) the predicted probability to make it more aligned with the true confidence levels.

Previous studies (Müller et al., 2019) have highlighted the effectiveness of label smoothing in improving model calibration. Conversely, our empirical study reveals intriguing findings: without temperature scaling, label smoothing (with properly tuned smoothing parameter δ) enhances model calibration compared to cross-entropy loss, *but when temperature scaling is applied as post-processing, label smoothing leads to inferior model calibration*. This inconsistency motivates us to investigate the underlying reason. In this section, we provide an explanation from the perspective of neural collapse.

Regularization effect of label smoothing on weight and feature norms. We trained ResNet18 on CIFAR10 and ResNet50 on CIFAR100 using label smoothing loss with various smoothing hyperparameters ranging from 0 to 0.3. Then, we evaluated models’ testing ECE both before and after temperature scaling, as depicted in Figure 3. Our observations reveal that compared to cross-entropy loss ($\delta = 0$), LS

loss with properly chosen δ lead to lower model calibration error (ECE), while larger values of δ may degrade model calibration. This is because tuning the smoothing hyperparameter δ has a similar effect to tuning temperature T in temperature scaling. Particularly, as shown empirically in the right column Figure 3 and theoretically in Theorem 4.1, increasing δ decreases both the feature and classification vector norms², thereby reducing the confidence level of model predictions. Consequently, the corresponding optimal temperature T decreases alongside the increase in δ . In fact, there exists a smoothing hyperparameter δ^* for which the optimal T equals 1, making temperature scaling post-processing unnecessary. Selecting $\delta > \delta^*$ results in an excessively low confidence level in the model’s predictions, leading to a notable increase in test ECE.

However, if we apply temperature scaling to counteract the regularization effect of LS loss on feature and classification weight, models trained with LS loss demonstrate higher test ECE (Post ECE as shown in the left column of Figure 3) compared to models trained with CE loss. We further investigate how this phenomenon is attributed to the more pronounced level of NC1 under LS loss.

Excessive NC1 adversely impacts model calibration. To investigate the impact of NC1 on model calibration, we analyze how miscalibration occurs during training, using CIFAR10 data with CE and LS loss ($\delta = 0.05$) as examples. We specifically partition the test set into two subgroups: those correctly classified and those incorrectly classified by the model. In Figure 4, the left column illustrates the average CE loss and the average entropy of the model predictions for both subgroups of test samples. The middle section of the figure showcases the NC1 metric for the entire test dataset, as well as for each of the subgroups individually. Meanwhile, the right column of Figure 4 presents the training and testing misclassification errors, accompanied by the test set ECE.

In Figure 4, models trained with CE and LS losses show similar trends. For correctly classified test samples, both test loss and average entropy exhibit a consistent decreasing trend, indicating the model’s increasing confidence in its correct predictions. Conversely, for incorrectly classified test samples, although the average loss starts to increase after 350 epochs, the overall downward trend of entropy suggests the model’s growing confidence in its incorrect predictions. Moreover, the middle plot in Figure 4 indicates that NC1 for both




Figure 3: Plots of ECE Pre/Post temperature scaling (left column), alongside feature and classification vector norms and optimal T (right column) for CIFAR10 and CIFAR100 datasets with varying values of δ .

²The feature norm is computed as the average norm of each class mean feature, represented by $\sum_{k=1}^K \|\bar{\mathbf{h}}_k - \mathbf{h}_G\|/K$. The classification vector norm is determined as the mean norm of the weight vector for each class given by $\sum_{k=1}^K \|\mathbf{w}_k\|/K$, where $\mathbf{w}_k \in \mathbb{R}^d$ represents the k -th column of the classifier weight \mathbf{W} .

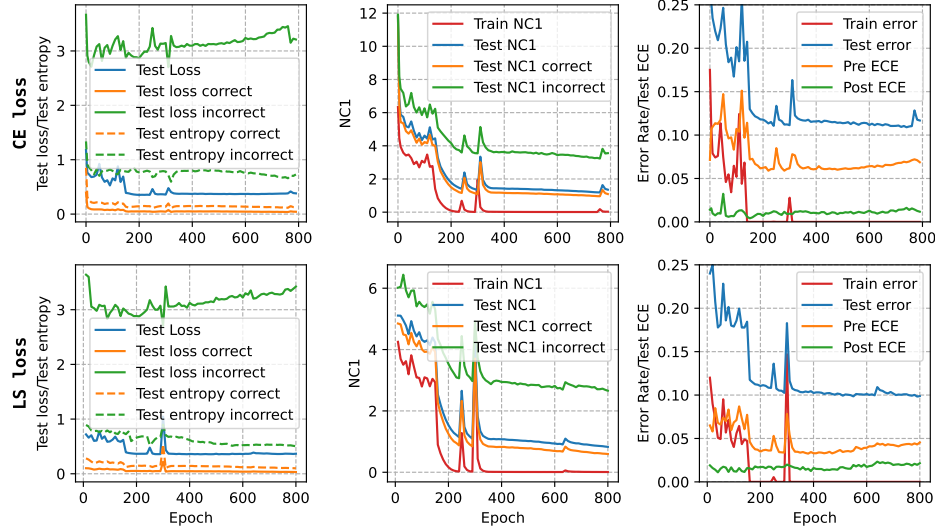


Figure 4: Leftmost plot: average error and entropy for correctly and incorrectly classified test samples. Middle plot: NC1 for the training set, testing set, and correctly/incorrectly classified testing samples. Right plot: model classification error rate and the test ECE before and after temperature scaling.

correctly and incorrectly classified test samples consistently decreases during training. For misclassified samples, a smaller NC1 indicates that their feature vectors align more closely with the mean of the incorrectly predicted class, making the model more confident about its incorrect predictions.³ As a result, the entropy for misclassified test samples continues to decrease, and the test ECE starts to rise during TPT. And if we apply temperature scaling to counteract the regularization effect of LS loss on feature and classification vector norms, LS loss results in higher ECE due to its stronger NC1.

4 Theoretical Analysis

4.1 Unconstrained Feature Model

Analyzing deep neural networks poses significant challenges because of the non-linearities and complex interactions between layers in the feature mapping $\mathbf{h} := \phi_{\theta}(\mathbf{x})$. The Unconstrained Feature Model (UFM) (Fang et al., 2021) simplifies DNN models by treating the features of the last layer as free optimization variables. This choice is motivated by the idea that over-parameterized DNN models are able to approximate any continuous functions (Hornik, 1991; Cybenko, 1989; Lu et al., 2017; Shaham et al., 2018).

Recall that the dimension of a feature vector \mathbf{h} is denoted by d . Let $\mathbf{H} = \{\mathbf{h}_{ki}\}_{1 \leq k \leq K, 1 \leq i \leq n} \in \mathbb{R}^{d \times N}$ be the feature matrix of all training samples with \mathbf{h}_{ki} denoting the $((k-1)n + i)$ -th column of \mathbf{H} . Under the UFM, we investigate regularized empirical risk minimization, a variant of the formulation in Equation 2:

$$\min_{\mathbf{W}, \mathbf{H}, \mathbf{b}} \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n l_{CE}(\mathbf{W}^{\top} \mathbf{h}_{ki} + \mathbf{b}, \mathbf{y}_k^{\delta}) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 + \frac{\lambda_b}{2} \|\mathbf{b}\|^2, \quad (7)$$

where $\mathbf{y}_k^{\delta} = (1 - \delta)\mathbf{e}_k + (\delta/K)\mathbf{1}_K$ is the soft label for class k with a smoothing parameter δ . Here \mathbf{e}_k a one-hot vector with the k -th element equal to 1, and $\lambda_W, \lambda_H, \lambda_b > 0$ are the regularization parameters.

Let $\mathbf{Y} = [\mathbf{e}_1 \mathbf{1}_n^{\top}, \dots, \mathbf{e}_K \mathbf{1}_n^{\top}] \in \mathbb{R}^{K \times N}$ represent the matrix form of the (hard) ground truth labels. Consequently, the matrix form of the soft labels can be represented as $\mathbf{Y}^{\delta} = (1 - \delta)\mathbf{Y} + \frac{\delta}{K}\mathbf{1}_K \mathbf{1}_N^{\top}$. The empirical

³For misclassified test samples, we assess the within-class variance relative to the incorrectly predicted class.

risk minimization in (7) has an equivalent matrix form:

$$\min_{\mathbf{W}, \mathbf{H}, \mathbf{b}} \mathcal{L}(\mathbf{W}, \mathbf{H}, \mathbf{b}) := \frac{1}{N} l_{CE}(\mathbf{W}^\top \mathbf{H} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 + \frac{\lambda_b}{2} \|\mathbf{b}\|^2, \quad (8)$$

where $l_{CE}(\mathbf{W}^\top \mathbf{H} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta)$ computes the cross-entropy column-wise and takes the sum.

4.2 Theoretical Results

Within the UFM framework, Zhou et al. (Zhou et al., 2022b) demonstrate that, under a wide range of loss functions, including CE loss and LS loss, all the global minimizers satisfy neural collapses properties, and in particular \mathbf{H} and \mathbf{W} form aligned simplex ETFs. Moreover, they show that every critical point is either a global minimizer or a strict saddle point, which implies that these global minimizers can be effectively attained through iterative algorithms. However, their work does not provide the exact expression of the global minimizers, nor does it provide a landscape analysis (condition number) across different loss functions, which is critical for understanding the model's convergence behavior. In contrast, our work derives the exact solutions of both \mathbf{W} and \mathbf{H} for UFM models with CE and LS losses. Furthermore, based on these solutions, we employ conditioning number analysis to closely compare the optimization landscape of the models surrounding their respective global minimizers, and thereby provide an explanation for the accelerated convergence observed under LS loss. The proofs of our theorems can be found in the Appendix.

Let $\lambda_Z := \sqrt{\lambda_W \lambda_H}$ and let a^δ be defined by

(i) if $\sqrt{KN}\lambda_Z + \delta \geq 1$, then $a^\delta = 0$;

(ii) if $\sqrt{KN}\lambda_Z + \delta < 1$, then

$$a^\delta = \log \left(\frac{K}{\sqrt{KN}\lambda_Z + \delta} - K + 1 \right).$$

Theorem 4.1. (*Global Optimizer*). Assume that the feature dimension d is greater than or equal to the number of classes K , i.e., $d \geq K$, and the dataset is balanced. Then any global optimizer of $(\mathbf{W}, \mathbf{H}, \mathbf{b})$ of (8) satisfies the following condition:

There is a semi-orthogonal matrix $\mathbf{P} \in \mathbb{R}^{d \times K}$ (i.e., $\mathbf{P}^\top \mathbf{P} = \mathbf{I}_K$) such that:

(i) The classification weight matrix \mathbf{W} is given by

$$\mathbf{W} = \left(\frac{n\lambda_H}{\lambda_W} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K). \quad (9)$$

(ii) The matrix of last-layer features \mathbf{H} is given by

$$\mathbf{H} = \left(\frac{\lambda_W}{n\lambda_H} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K) \mathbf{Y}. \quad (10)$$

(iii) The bias \mathbf{b} is a zero vector, i.e. $\mathbf{b} = \mathbf{0}$.

The above theorem provides an explicit closed-form solution for the global minimizer $(\mathbf{W}, \mathbf{H}, \mathbf{b})$. Notably, our findings indicate that an increase in δ is associated with a decrease in the norm of \mathbf{W} and \mathbf{H} . This observation aligns closely with the empirical results detailed in Section 3.3.

Our observations in Section 3 reveal that LS loss demonstrates accelerated convergence in terms of the NC metrics and training error. To better understand this phenomenon, we analyze the optimization landscape under the UFM. Note that condition number of the Hessian matrix, representing the ratio of the largest to the smallest eigenvalue, plays a crucial role in convergence rate analysis (Nocedal & Wright, 1999; Trefethen & Bau, 2022). In the vicinity of the local minimizer, a smaller condition number typically signifies a faster convergence rate. We now present our main theoretical result, which, to the best of our knowledge, is the first result providing insight into the convergence rate to the optimal solution under UFM models.

Theorem 4.2. (*Optimization Landscape*). Assume the feature dimension d is greater than or equal to the number of classes K ($d \geq K$) and the dataset is balanced. In addition, assume the regularization parameters and the smoothing parameter satisfy $\sqrt{KN\lambda_W\lambda_H} + \delta < 1$. Then, at the global minimizer (as provided in Equation 9-10), the Hessian of the empirical loss with respect to \mathbf{W} when \mathbf{H} is fixed (also the Hessian w.r.t. \mathbf{H} when \mathbf{W} is fixed) has smaller condition number under label smoothing loss ($0 < \delta < 1 - \sqrt{KN\lambda_W\lambda_H}$) than under cross-entropy loss ($\delta = 0$).

This theorem suggests that the LS loss results in a better-conditioned optimization landscape in the vicinity of the global minimizers, facilitating faster convergence behavior, as observed in Section 3.1. It is noteworthy that while the theorem primarily analyzes the local landscape around the global minimizer, experimental results demonstrate that gradient descent with random initialization indeed converges faster to the global solution of neural collapse under label smoothing loss.

5 Discussion

We conducted a comprehensive empirical comparison of CE loss and LS loss during the training process. We found that models trained under LS loss exhibit accelerated convergence in terms of training error and neural collapse metrics. Furthermore, they converge to a more pronounced level of NC1 and NC2. Along with the accelerated convergence, we found that label smoothing maintains a distinct balance between NC1 and NC2. We posit that the emphasis on NC2 in LS enhances the model’s generalization performance. Additionally, we investigated the impact of LS on model calibration from the perspective of neural collapse, revealing that label smoothing has a regularization effect on the classifier weight and feature norms, and excessively small NC1 values may adversely affect model calibration.

We performed a mathematical analysis of the convergence properties of the UFM models under CE and LS losses. We first derived closed-form solutions for the global minimizers under both loss functions. Then we conducted a second-order theoretical analysis of the optimization landscape around their respective global optimizers, which reveals that LS exhibits a better-conditioned optimization landscape around the global minimum, which facilitates the faster convergence observed in our empirical study.

This paper provides a significantly deeper understanding of why LS excels in terms of convergence, performance, and model calibration compared to CE loss. Additionally, it illustrates how the powerful framework of neural collapse and its associated mathematical models can be employed to gain a more nuanced understanding of the "why" of DNNs. We expect that these results will inspire future research into the interplay between neural collapse, convergence speed, and model generalizability.

References

- Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, 2011.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- 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. *arXiv preprint arXiv:2301.00437*, 2023.
- 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. *arXiv preprint arXiv:2112.15121*, 2021.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.

- XY Han, Vardan Papyan, and David L Donoho. Neural collapse under mse loss: Proximity to and dynamics on the central path. *arXiv preprint arXiv:2106.02073*, 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*, pp. 770–778, 2016.
- Wanli Hong and Shuyang Ling. Neural collapse for unconstrained feature model under cross-entropy loss with imbalanced data. *arXiv preprint arXiv:2309.09725*, 2023.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- Jiachen Jiang, Jinxin Zhou, Peng Wang, Qing Qu, Dustin Mixon, Chong You, and Zhihui Zhu. Generalized neural collapse for a large number of classes. *arXiv preprint arXiv:2310.05351*, 2023.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- 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.
- Xuantong Liu, Jianfeng Zhang, Tianyang Hu, He Cao, Yuan Yao, and Lujia Pan. Inducing neural collapse in deep long-tailed learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 11534–11544. PMLR, 2023.
- Jianfeng Lu and Stefan Steinerberger. Neural collapse under cross-entropy loss. *Applied and Computational Harmonic Analysis*, 59:224–241, 2022.
- Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The expressive power of neural networks: A view from the width. *Advances in neural information processing systems*, 30, 2017.
- Michal Lukasik, Srinadh Bhojanapalli, Aditya Menon, and Sanjiv Kumar. Does label smoothing mitigate label noise? In *International Conference on Machine Learning*, pp. 6448–6458. PMLR, 2020.
- Dustin G Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *arXiv preprint arXiv:2011.11619*, 2020.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019.
- Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 625–632, 2005.
- Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999.
- Vardan Papyan, 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*, 2020.
- Uri Shoham, Alexander Cloninger, and Ronald R Coifman. Provable approximation properties for deep neural networks. *Applied and Computational Harmonic Analysis*, 44(3):537–557, 2018.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- 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*, pp. 21478–21505. PMLR, 2022.
- Tom Tirer, Haoxiang Huang, and Jonathan Niles-Weed. Perturbation analysis of neural collapse. In *International Conference on Machine Learning*, pp. 34301–34329. PMLR, 2023.
- Lloyd N Trefethen and David Bau. *Numerical linear algebra*, volume 181. Siam, 2022.
- Stephan Wojtowytsch et al. On the emergence of simplex symmetry in the final and penultimate layers of neural network classifiers. *arXiv preprint arXiv:2012.05420*, 2020.
- Liang Xie, Yibo Yang, Deng Cai, and Xiaofei He. Neural collapse inspired attraction–repulsion-balanced loss for imbalanced learning. *Neurocomputing*, 527:60–70, 2023.
- Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning. *arXiv preprint arXiv:2302.03004*, 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.
- Longhui Yu, Tianyang Hu, Lanqing Hong, Zhen Liu, Adrian Weller, and Weiyang Liu. Continual learning by modeling intra-class variation. *arXiv preprint arXiv:2210.05398*, 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*, pp. 27179–27202. PMLR, 2022a.
- Jinxin Zhou, Chong You, Xiao Li, Kangning Liu, Sheng Liu, Qing Qu, and Zhihui Zhu. Are all losses created equal: A neural collapse perspective. *Advances in Neural Information Processing Systems*, 35:31697–31710, 2022b.
- 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.

A Appendix

B Additional Experiments

B.1 Metrics for Measuring NC

We introduced the metrics for measuring neural collapse in Section 3. For convenience, we restate them here, providing additional details about the subtle differences between our definitions and those used in prior works (Zhu et al., 2021; Zhou et al., 2022b). We denote the global mean and classwise means of the last-layer features as:

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

where \mathbf{h}_G is the global mean and $\bar{\mathbf{h}}_k$ represents the mean of class k . The metrics for neural collapse are then defined as follows:

(NC1) Within class variability collapse measures the relative magnitude of the within-class covariance $\Sigma_W := \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n (\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)(\mathbf{h}_{ki} - \bar{\mathbf{h}}_k)^\top$ compared to the between-class covariance matrix $\Sigma_B := \frac{1}{K} \sum_{k=1}^K (\bar{\mathbf{h}}_k - \mathbf{h}_G)(\bar{\mathbf{h}}_k - \mathbf{h}_G)^\top$ of the last-layer features. It is formulated as:

$$NC_1 = \frac{1}{K} \text{trace} \left(\Sigma_W \Sigma_B^\dagger \right),$$

where Σ_B^\dagger denotes the pseudo inverse of Σ_B .

(NC2) Convergence to simplex ETF quantifies the difference between the normalized classifier weight matrix and the centered class mean features in comparison to a normalized simplex ETF, defined as follows:

$$NC_2 := \left\| \frac{\mathbf{W}^\top \bar{\mathbf{H}}}{\|\mathbf{W}^\top \bar{\mathbf{H}}\|_F} - \frac{1}{\sqrt{K-1}} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right) \right\|_F, \quad (11)$$

where $\bar{\mathbf{H}} = [\bar{\mathbf{h}}_1 - \mathbf{h}_G, \dots, \bar{\mathbf{h}}_K - \mathbf{h}_G] \in \mathbb{R}^{d \times K}$ represents centered class mean matrix.

This definition differs slightly from those in (Zhu et al., 2021; Zhou et al., 2022b). However, it can be shown that the convergence of NC_2 in Equation (11) to zero indicates the simultaneous convergence of both \mathbf{W} and \mathbf{H} towards the simplex ETF structure. Specifically, when the bias \mathbf{b} is an all-zero vector or a constant vector, NC_2 as defined in (11) approaching zero indicates that the average logit matrix, formulated as $\bar{\mathbf{Z}} = \mathbf{W}^\top \bar{\mathbf{H}} + \mathbf{b} \mathbf{1}_K^\top$, satisfies the condition $\bar{\mathbf{Z}} = a \left(\mathbf{I} - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right)$ for some constant a . Remarkably, this matrix aligns with the simplex-encoding label (SEL) matrix introduced in (Thrampoulidis et al., 2022), up to a scaling factor a .

From Proposition C.4, we have

$$\min_{\mathbf{W}^\top \mathbf{H} = \bar{\mathbf{Z}}} \frac{1}{2} (\lambda_W \|\mathbf{W}\|_F^2 + \lambda_H \|\mathbf{H}\|_F^2) \geq \|\bar{\mathbf{Z}}\|_*$$

where $\|\bar{\mathbf{Z}}\|_*$ represents the nuclear norm of $\bar{\mathbf{Z}}$ and the equality holds only when $\mathbf{H} = \bar{\mathbf{H}} \mathbf{Y}$ and $\bar{\mathbf{H}} = \sqrt{\lambda_W / n \lambda_H} \mathbf{W}$, indicating the self-duality of the class mean feature and the classification vector. Consequently, during the model training with an L2 penalty on both \mathbf{W} and \mathbf{H} (with the norm of \mathbf{H} implicitly penalized by penalizing the model parameters Θ), the convergence of NC_2 as defined in (11) to zero indicates the simultaneous convergence of both \mathbf{W} and \mathbf{H} towards the simplex ETF structure.

(NC3) Convergence to self-duality measures the distance between the classifier weight matrix \mathbf{W} and the centered class-means $\bar{\mathbf{H}}$:

$$NC_3 := \left\| \frac{\mathbf{W}}{\|\mathbf{W}\|_F} - \frac{\bar{\mathbf{H}}}{\|\bar{\mathbf{H}}\|_F} \right\|_F. \quad (12)$$

B.2 Experiment Setup

In this section, we provide with more details for reproducing the experiments presented in the paper. In particular, we emphasize that the datasets involved in the paper, namely CIFAR10, CIFAR100, and STL10, are all publicly available for academic usage: CIFAR10 and CIFAR100 datasets are made available under the MIT license. All experiments are conducted on a single RTX3090 GPU with 24GB memory. For consistency with prior studies (Papayan et al., 2020; Zhu et al., 2021), we adopt ResNet-18 (He et al., 2016) as the backbone network for CIFAR-10 and ResNet-50 (He et al., 2016) as the backbone for CIFAR-100 and STL-10.

To focus on behaviors associated with neural collapse and minimize the influence of other factors, we adopted standard preprocessing techniques without introducing any data augmentation for the training data. We extended the training period to 800 epochs to analyze the model’s behavior during the terminal phase of training. For all datasets, we used a default batch size of 128, stochastic gradient descent with a momentum of 0.9, and a learning rate initialized at 0.05, which undergoes multi-step decay, decreasing by a factor of 0.1 at epochs 150 and 350. We considered two different scenarios for model regularization:

- **Model with Weight Decay.** For the default setting, we trained the models with weight decay regularization set at 5×10^{-4} . It’s worth noting that the norm of \mathbf{H} is implicitly penalized by the weight decay on the model parameters Θ .
- **Simulated UFM Model.** To simulate the unconstrained feature model, we turn off the weight decay and add an L2 penalty on the classifier parameters \mathbf{W} and \mathbf{b} and the last-layer feature \mathbf{H} to the cross-entropy (CE) or label smoothing (LS) loss. This setting was only used in Section B.6 to validate the theoretical results.

B.3 Neural Collapses under Cross-Entropy Loss and Label Smoothing Loss

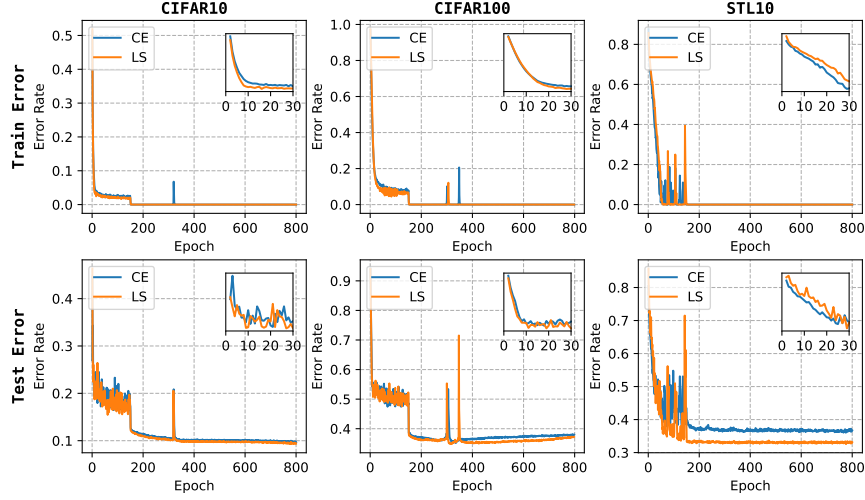


Figure 5: Comparison of training (top row) and testing (bottom row) error rates using cross-entropy (CE) and label smoothing (LS) losses.

This section presents additional visualizations supporting the conclusions drawn in the main paper. Figure 5 provides a detailed zoom-in plot illustrating the progression of training and testing errors throughout the model training process for the CIFAR-10, CIFAR-100, and STL-10 datasets, under both cross-entropy and label smoothing losses. Notably, the visualization suggests that models trained with label smoothing exhibit a more rapid convergence in both training and testing errors compared to their cross-entropy counterparts.

B.4 Impact of the Smoothing Hyperparameter

In Section 3.1 and 3.2, a default smoothing hyperparameter of $\delta = 0.05$ is utilized. This section explores how the choice of δ impacts the model’s convergence to neural collapse and its generalizability. Specifically, we consider various values for δ within the interval $[0, 1)$.

Figure 6 presents the results of the experiments. The experiments yield several important insights. Firstly, the smoothing hyperparameter significantly influences the model’s convergence to neural collapse. The curves of NC1 and NC2 as a function of δ (2nd column in Figure 6) reveal a U-shaped trend, with the levels of NC3 remaining relatively consistent across different values of δ . Exceptionally small and large values of δ result in reduced collapse for both NC1 and NC2.

Secondly, the average norms of the classifier vectors and the class mean features decrease as δ increases. This observation aligns with intuition: given features and classifier vectors that form a simplex ETF structure, decreasing their norms softens the output probabilities. For LS loss, a higher smoothing parameter δ corresponds to smoother target labels, and consequently, the features and classification vectors with lower norms can achieve close-to-zero label smoothing loss. This observation was also supported by the closed-form expressions for \mathbf{W} and \mathbf{H} provided in Theorem 4.1. Thirdly, the trend of the test error also exhibits a U-shape, underscoring the necessity of selecting an appropriate δ . When δ approaches 1, the nearly uniform smoothed labels do not provide effective training signals to update the parameters. Additionally, the norm of both classification vectors and the last-layer features decrease towards zero, causing the classifier to be dominated by noise, leading to rapid deterioration in model performance. Finally, we observe a surprising robustness in both neural collapse metrics and model performance for the choice of $\delta \in [0.02, 0.8]$. For classification tasks with a larger number of classes, such as CIFAR-100, opting for a relatively large δ is viable.

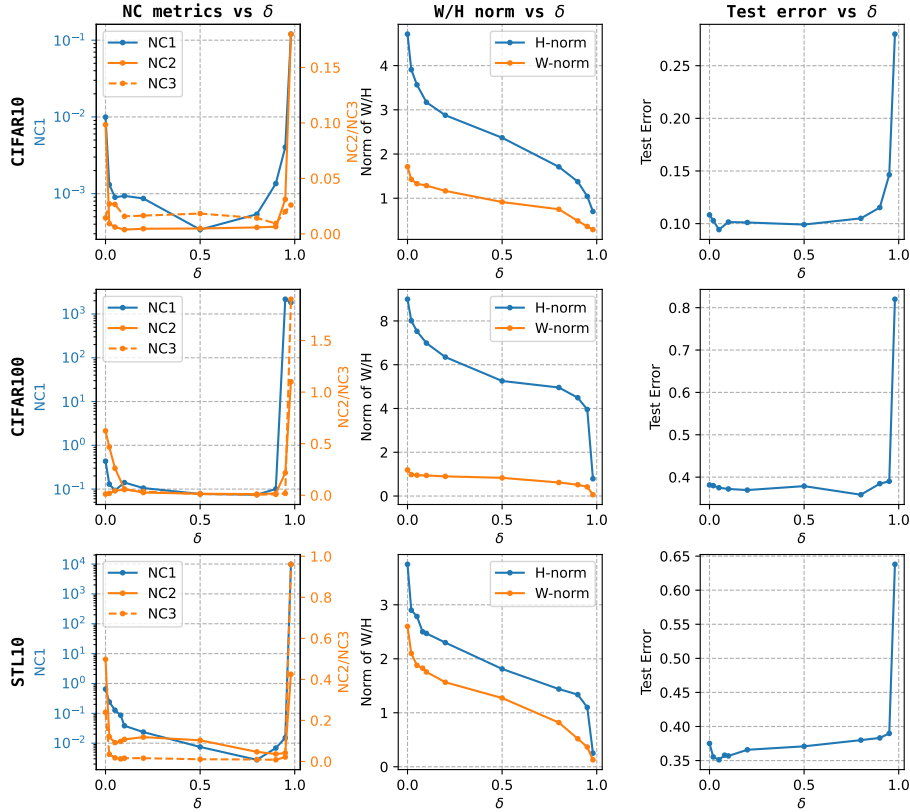


Figure 6: **The effect of the smoothing hyperparameter δ .** The columns from left to right visualize (a) NC metrics vs. δ , (b) Average norm of classification weight vectors and mean features vs. δ , and (c) Testing performance (test error rate) vs. δ .

B.5 Weight and Feature Regularization Improve Model Calibration

As mentioned in Section 3.3, label smoothing loss (with properly chosen δ) can improve model calibration. This improvement is due to the regularization effect of label smoothing on the classification vector norm and last-layer feature norm, which leads to predictions with lower confidence. To support this, Figure 7 presents a 20-bin reliability plot (Niculescu-Mizil & Caruana, 2005) for models trained under CE and LS losses for CIFAR-10, with a default smoothing hyperparameter of 0.05. Additionally, we include a plot showing the percentage of samples in each confidence bin. Clearly, LS loss exhibits less confidence in its predictions and better model calibration compared to CE loss without temperature scaling. In the right column of Figure 7, we employ temperature scaling to calibrate the model, with the hyperparameter T selected through cross-validation. Notably, after temperature scaling, models trained with CE loss achieve even better calibration than those trained with LS loss. This is because temperature scaling counteracts the regularization effect of label smoothing on the classification vector norm and feature norm. Consequently, LS loss can result in worse test expected calibration error (ECE) due to an excessive level of NC1.

B.6 Empirical Validation of Theoretical Results

In Theorem 4.1, we derive closed-form solutions for both the mean feature matrix $\bar{\mathbf{H}}$ and the classifier weight matrix \mathbf{W} , both exhibiting a simplex ETF structure. To validate these theoretical findings, we train ResNet18 models on CIFAR10 dataset with varying smoothing parameters δ . To ensure consistency with the UFM model, we adopted the simulated UFM model as introduced in Section B.2, where weight decay is disabled, and L2 regularization is applied to the last-layer features and classification parameters, ensuring




Figure 7: Reliability plots with 20 bins (top row) and the percentage of samples in each bin (bottom row) for models trained with cross-entropy (CE) and label smoothing (LS) losses, both pre- (left column) and post- (right column) temperature scaling. The results are based on the CIFAR-10 dataset.

exact regularization effects. We set the regularization parameters as follows: $\lambda_W = 10^{-3}$, $\lambda_H = 10^{-6}$, and $\lambda_b = 10^{-2}$. Under these settings, we compare the trained average feature norm $\|\bar{\mathbf{h}}\| = \sum_{k=1}^K \|\bar{\mathbf{h}}_k - \mathbf{h}_G\|/K$ and average classification vector norm $\|\bar{\mathbf{w}}\| = \sum_{k=1}^K \|\mathbf{w}_k\|/K$ with their theoretical counterparts derived from Theorem 4.1 which can be represented as

$$\|\bar{\mathbf{h}}\|^* = \sqrt{(K-1)a^\delta/K} \left(\frac{n\lambda_H}{\lambda_W} \right)^{-1/4}, \quad \|\bar{\mathbf{w}}\|^* = \sqrt{(K-1)a^\delta/K} \left(\frac{n\lambda_H}{\lambda_W} \right)^{1/4}.$$

Figure 8 illustrates the differences between the trained predictions $\|\bar{\mathbf{h}}\|$ and $\|\bar{\mathbf{w}}\|$ and their theoretical solutions for models trained under LS loss with different δ . The plot shows a clear decreasing trend in the relative error between the empirical and theoretical values for both the feature and classification vector norms, with differences reducing to below 0.01 across all cases. These results demonstrate that the empirical findings align with the theoretical predictions, thereby verifying Theorem 4.1.




Figure 8: Relative error of **classification vector norm** $\|\bar{\mathbf{w}}\|$ and **feature norm** $\|\bar{\mathbf{h}}\|$ compared to their theoretical results over epochs for ResNet-18 on CIFAR10.

According to Theorem 4.2, models trained with label smoothing loss (satisfying $\sqrt{KN}\lambda_Z + \delta < 1$) converges faster than models trained with cross-entropy loss. In Figure 9, we train ResNet18 and ResNet50 on CIFAR10 and CIFAR100, respectively, with a default weight decay value of 5×10^{-4} . Comparing models trained with

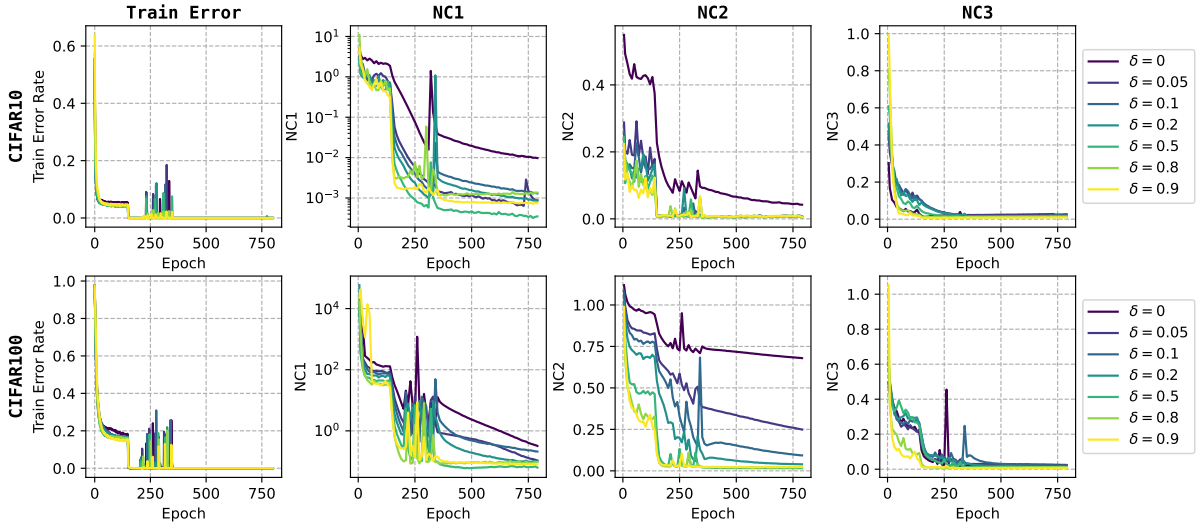


Figure 9: Comparison of model convergence under weight decay regularization. From left to right, the figures display the models’ training error rate, NC1, NC2, and NC3 with varying smoothing hyperparameters δ .

different smoothing hyperparameters, we observe that for a wide range of δ values, models with larger δ converge faster in terms of both training error and NC metrics.

Furthermore, we investigate model convergence under the simulated UFM model, where weight decay is disabled, and L2 regularization is added. As illustrated in Figure 10, across a wide range of δ such that $\sqrt{KN}\lambda_Z + \delta < 1$, we observe faster convergence for larger δ values.




Figure 10: Comparison of model convergence under simulated UFM model. From left to right, the figures display the models’ training error rate, NC1, NC2, and NC3 with varying smoothing hyperparameters δ .

C Proofs

The first key observation of the phenomenon of neural collapse, that is (NC1), refers to a type of collapse that involves the convergence of the feature of samples from the same class to a unique mean feature vector.

The second key observation of neural collapse, namely (NC2), involves these unique mean feature vectors (after recentering by their global mean) as they form an equiangular tight frame (ETF), i.e., they share the same pairwise angles and length. Before providing the theoretical poof, we first formally define the rank K canonical simplex ETF in the definition below.

C.1 Basics

Definition C.1. (K-simplex ETF) A K-simplex ETF is a collection of points in \mathbb{R}^d specified by the columns of the matrix

$$\mathbf{M} = \sqrt{\frac{K}{K-1}} \mathbf{P} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right),$$

where $\mathbf{I}_K \in \mathbb{R}^{K \times K}$ is the identity matrix and $\mathbf{1}_K \in \mathbb{R}^K$ is the ones vector, and $\mathbf{P} \in \mathbb{R}^{d \times K}$ ($d \geq K$) is a partial-orthogonal matrix such that $\mathbf{P}^\top \mathbf{P} = \mathbf{I}_K$.

Note that the matrix \mathbf{M} satisfies:

$$\mathbf{M}^\top \mathbf{M} = \frac{K}{K-1} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^\top \right).$$

Next, we prove a series of lemmas that will prove crucial upon establishing our main theorems.

Lemma C.2. (Young's Inequality) Let p, q be positive numbers satisfying $\frac{1}{p} + \frac{1}{q} = 1$. Then for any $a, b \in \mathbb{R}$, we have

$$|ab| \leq \frac{|a|^p}{p} + \frac{|b|^q}{q},$$

where the equality holds if and only if $|a|^p = |b|^q$. The case for $p = q = 2$ is just the AM-GM inequality which is $|ab| \leq \frac{1}{2} (a^2 + b^2)$, where the equality holds if and only if $|a| = |b|$

Lemma C.3. For any fixed $\mathbf{Z} \in \mathbb{R}^{K \times N}$ and $\alpha > 0$, we have

$$\min_{\mathbf{Z}=\mathbf{W}^\top \mathbf{H}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2) = \|\mathbf{Z}\|_*. \quad (13)$$

Here $\|\mathbf{Z}\|_*$ denotes the nuclear norm of \mathbf{Z} :

$$\|\mathbf{Z}\|_* := \sum_k \sigma_k(\mathbf{Z}) = \text{trace}(\Sigma), \text{ with } \mathbf{Z} = \mathbf{U} \Sigma \mathbf{V}^\top,$$

where $\{\sigma_k\}_{k=1}^{\min(K,N)}$ denote the singular values of \mathbf{Z} , and $\mathbf{Z} = \mathbf{U} \Sigma \mathbf{V}^\top$ is the singular value decomposition (SVD) of \mathbf{Z} .

Proof. Let $\mathbf{Z} = \mathbf{U} \Sigma \mathbf{V}^\top$ be the SVD of \mathbf{Z} . From the fact that $\mathbf{U} \mathbf{U}^\top = \mathbf{I}$, $\mathbf{V} \mathbf{V}^\top = \mathbf{I}$, and $\text{trace}(\mathbf{A}^\top \mathbf{A}) = \|\mathbf{A}\|_F^2$, we have

$$\begin{aligned} \|\mathbf{Z}\|_* &= \text{trace}(\Sigma) = \frac{1}{2\sqrt{\alpha}} \text{trace}(\sqrt{\alpha} \mathbf{U}^\top \mathbf{U} \Sigma) + \frac{\sqrt{\alpha}}{2} \text{trace}\left(\frac{1}{\sqrt{\alpha}} \Sigma \mathbf{V}^\top \mathbf{V}\right) \\ &= \frac{1}{2\sqrt{\alpha}} \left(\left\| \alpha^{1/4} \mathbf{U} \Sigma^{1/2} \right\|_F^2 + \alpha \left\| \alpha^{-1/4} \Sigma^{1/2} \mathbf{V}^\top \right\|_F^2 \right). \end{aligned}$$

This implies that there exists some $\mathbf{W} = \alpha^{1/4} \Sigma^{1/2} \mathbf{U}^\top$ and $\mathbf{H} = \alpha^{-1/4} \Sigma^{1/2} \mathbf{V}^\top$, such that $\|\mathbf{Z}\|_* = \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2)$, which further indicates that

$$\|\mathbf{Z}\|_* \geq \min_{\mathbf{Z}=\mathbf{W}^\top \mathbf{H}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2). \quad (14)$$

On the other hand, for any $\mathbf{W}^\top \mathbf{H} = \mathbf{Z}$, we have

$$\begin{aligned} \|\mathbf{Z}\|_* &= \text{trace}(\Sigma) = \text{trace}(\mathbf{U}^\top \mathbf{Z} \mathbf{V}) = \text{trace}(\mathbf{U}^\top \mathbf{W}^\top \mathbf{H} \mathbf{V}) \\ &\leq \frac{1}{2\sqrt{\alpha}} \|\mathbf{U}^\top \mathbf{W}^\top\|_F^2 + \frac{\sqrt{\alpha}}{2} \|\mathbf{H} \mathbf{V}\|_F^2 = \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2), \end{aligned}$$

where the first inequality is guaranteed by Young's inequality in Lemma C.2, and equality only holds when $\mathbf{W} \mathbf{U} = \sqrt{\alpha} \mathbf{H} \mathbf{V}$. The last equality follows because $\mathbf{U} \mathbf{U}^\top = \mathbf{I}$ and $\mathbf{V} \mathbf{V}^\top = \mathbf{I}$. Therefore, we have

$$\|\mathbf{Z}\|_* \leq \min_{\mathbf{Z}=\mathbf{W}^\top \mathbf{H}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2). \quad (15)$$

Combining the results in (14) and (15), we complete the proof. \square

Proposition C.4. Consider matrices $\mathbf{H} = [\mathbf{H}_1, \dots, \mathbf{H}_n] \in \mathbb{R}^{d \times N}$ and $\mathbf{Z} = [\mathbf{Z}_1, \dots, \mathbf{Z}_n] \in \mathbb{R}^{K \times N}$, where $\mathbf{H}_i \in \mathbb{R}^{d \times K}$ and $\mathbf{Z}_i \in \mathbb{R}^{K \times K}$ with $N = nK$. Let $\bar{\mathbf{H}} = \frac{1}{n} \sum_i \mathbf{H}_i$ and $\bar{\mathbf{Z}} = \frac{1}{n} \sum_i \mathbf{Z}_i$. Then $\mathbf{Z} = \mathbf{W}^\top \mathbf{H}$ indicates $\bar{\mathbf{Z}} = \mathbf{W}^\top \bar{\mathbf{H}}$. If $\bar{\mathbf{Z}}$ is a symmetric matrix, then we have

$$\min_{\mathbf{Z}=\mathbf{W}^\top \mathbf{H}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2) = \min_{\bar{\mathbf{Z}}=\mathbf{W}^\top \bar{\mathbf{H}}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha n \|\bar{\mathbf{H}}\|_F^2) = \|\bar{\mathbf{Z}}\|_*, \quad (16)$$

with the minimum is reached only if $\mathbf{H}_i = \bar{\mathbf{H}}$ (for $\forall i = 1, \dots, n$) and $\mathbf{W} = \sqrt{\alpha n} \bar{\mathbf{H}}$.

Proof. From Lemma C.2, we have $\|\mathbf{H}\|_F^2 = \sum_i \|\mathbf{H}_i\|_F^2 \geq n \|\bar{\mathbf{H}}\|_F^2$, with the equality hold only when $\mathbf{H}_i = \bar{\mathbf{H}}$ for any $i \neq j$. Consequently, this yields the following result:

$$\min_{\mathbf{Z}=\mathbf{W}^\top \mathbf{H}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2) = \min_{\bar{\mathbf{Z}}=\mathbf{W}^\top \bar{\mathbf{H}}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha n \|\bar{\mathbf{H}}\|_F^2).$$

Utilizing Lemma C.3, we further deduce:

$$\min_{\bar{\mathbf{Z}}=\mathbf{W}^\top \bar{\mathbf{H}}} \frac{1}{2\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha n \|\bar{\mathbf{H}}\|_F^2) = \|\bar{\mathbf{Z}}\|_*.$$

This minimum is achieved only when $\mathbf{W} = \sqrt{\alpha n} \bar{\mathbf{H}}$, thus completing the proof. \square

D Proof of Theorem 4.1

D.1 Proof of Theorem 4.1

In this section, we present the proof of Theorem 4.1 in Section 4, which we restate as follows.

Theorem D.1. (Global Optimizer). Assume that the feature dimension d is no less than the number of classes K , i.e., $d \geq K$, and the dataset is balanced. Then any global optimizer of $(\mathbf{W}, \mathbf{H}, \mathbf{b})$ of

$$\min_{\mathbf{W}, \mathbf{H}, \mathbf{b}} \mathcal{L}(\mathbf{W}, \mathbf{H}, \mathbf{b}) := \frac{1}{N} l_{CE}(\mathbf{W}^\top \mathbf{H} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) + \frac{\lambda_W}{2} \|\mathbf{W}\|_F^2 + \frac{\lambda_H}{2} \|\mathbf{H}\|_F^2 + \frac{\lambda_b}{2} \|\mathbf{b}\|^2 \quad (17)$$

satisfies the following properties:

The classification weight matrix \mathbf{W} is given by

$$\mathbf{W} = \left(\frac{\lambda_H n}{\lambda_W} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K). \quad (18)$$

The matrix of last-layer feature \mathbf{H} can be represented as

$$\mathbf{H} = \bar{\mathbf{H}} \mathbf{Y}, \quad \bar{\mathbf{H}} = \left(\frac{\lambda_W}{\lambda_H n} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K). \quad (19)$$

The bias \mathbf{b} is a zero vector, i.e. $\mathbf{b} = \mathbf{0}$.

Here, $\mathbf{P} \in \mathbb{R}^{d \times K}$ ($d \geq K$) is a partial orthogonal matrix such that $\mathbf{P}^\top \mathbf{P} = \mathbf{I}_K$ and a^δ satisfies:

(i) if $\sqrt{KN}\lambda_Z + \delta \geq 1$, then $a^\delta = 0$;

(ii) if $\sqrt{KN}\lambda_Z + \delta < 1$, then

$$a^\delta = \log \left(\frac{K}{\sqrt{KN}\lambda_Z + \delta} - K + 1 \right).$$

with $\lambda_Z = \sqrt{\lambda_W \lambda_H}$.

Proof. The main idea of proving Theorem 4.1 is first to connect the problem (17) to its corresponding convex counterpart. This allows us to derive the precise form of the global minimizer for the convex optimization problem. Subsequently, we can further characterize the specific structures of \mathbf{W} and \mathbf{H} based on the acquired global minimizer.

Connection of (17) to a Convex Problem. Let $\mathbf{Z} = \mathbf{W}^\top \mathbf{H} \in \mathbb{R}^{K \times N}$ represent the output logit matrix with $N = nK$ and $\alpha = \frac{\lambda_H}{\lambda_W}$. Utilizing Lemma C.3, we get:

$$\begin{aligned} \min_{\mathbf{W}^\top \mathbf{H} = \mathbf{Z}} \lambda_W \|\mathbf{W}\|_F^2 + \lambda_H \|\mathbf{H}\|_F^2 &= \sqrt{\lambda_W \lambda_H} \min_{\mathbf{W}^\top \mathbf{H} = \mathbf{Z}} \frac{1}{\sqrt{\alpha}} (\|\mathbf{W}\|_F^2 + \alpha \|\mathbf{H}\|_F^2) \\ &= 2\sqrt{\lambda_W \lambda_H} \|\mathbf{Z}\|_*, \end{aligned} \quad (20)$$

where $\|\mathbf{Z}\|_*$ represent nuclear norm of \mathbf{Z} . Additionally from Lemma C.3, the minimum is attained only when $\mathbf{H} = \bar{\mathbf{H}}\mathbf{Y}$ and $\mathbf{W} = \sqrt{\alpha n} \bar{\mathbf{H}} = \sqrt{\frac{\lambda_H}{\lambda_W n}} \bar{\mathbf{H}}$.

Let $\lambda_Z := \sqrt{\lambda_W \lambda_H}$, then Equation (17) becomes:

$$\min_{\mathbf{Z}, \mathbf{b}} \mathcal{L}(\mathbf{Z}, \mathbf{b}) := \frac{1}{N} l_{CE}(\mathbf{Z} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) + \lambda_Z \|\mathbf{Z}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2, \quad (21)$$

which is a convex optimization problem.

Characterizing the Optimal Solution of (17) based on the Convex Program (21): We first derive the exact form of the global minimizer for the convex optimization problem (21). Particularly, we first establish that the predicted logit vectors within each class collapse to their sample means, i.e., $\mathbf{Z} = \bar{\mathbf{Z}}\mathbf{Y}$ (Lemma D.3). Subsequently, we derive the closed-form solution of $\bar{\mathbf{Z}}$ in Lemma D.5.

Furthermore, from Lemma C.3, we establish that the minimum in (20) is attained only if $\mathbf{H} = \bar{\mathbf{H}}\mathbf{Y}$ and $\mathbf{W} = \sqrt{\frac{\lambda_H}{\lambda_W n}} \bar{\mathbf{H}}$. Since $\mathbf{Z} = \mathbf{W}^\top \mathbf{H}$, combining the above result with Lemma D.5 yields the global minimizer $(\mathbf{W}, \mathbf{H}, \mathbf{b})$ of (17), satisfying:

$$\mathbf{W} = \left(\frac{\lambda_H n}{\lambda_W} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K), \quad (22)$$

$$\mathbf{H} = \bar{\mathbf{H}}\mathbf{Y}, \quad \bar{\mathbf{H}} = \left(\frac{\lambda_W}{\lambda_H n} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K), \quad (23)$$

and

$$\mathbf{b} = \mathbf{0},$$

where $\mathbf{P} \in \mathbb{R}^{d \times K}$ ($d \geq K$) is a partial orthogonal matrix such that $\mathbf{P}^\top \mathbf{P} = \mathbf{I}_K$ and a^δ is defined per the specifications in Lemma D.5.

Hence, we conclude the proof. \square

D.2 Supporting Lemma

Lemma D.2. (*Optimality Condition*) The first-order optimality condition of $\mathcal{L}(\mathbf{Z}, \mathbf{b})$ in Equation (21) is

$$N^{-1}(\mathbf{Y}^\delta - \mathbf{P}) \in \lambda_Z \partial \|\mathbf{Z}\|_*, \quad N^{-1}(\mathbf{Y}^\delta - \mathbf{P}) \mathbf{1}_N = \lambda_b \mathbf{b}, \quad (24)$$

where \mathbf{P} is the prediction matrix defined as

$$\mathbf{P} = \{\mathbf{p}_{ki}\}_{1 \leq k \leq K, 1 \leq i \leq n} \in \mathbb{R}^{K \times N}, \quad \mathbf{p}_{ki} := \frac{\exp(\mathbf{z}_{ki} + \mathbf{b})}{\langle \exp(\mathbf{z}_{ki} + \mathbf{b}), \mathbf{1}_K \rangle}, \quad (25)$$

\mathbf{Y}^δ is the matrix of the smoothed soft target defined as

$$\mathbf{Y}^\delta = (1 - \delta)\mathbf{Y} + \frac{\delta}{K} \mathbf{1}_K \mathbf{1}_N^\top, \quad \mathbf{Y} = [\mathbf{e}_1 \mathbf{1}_n^\top, \dots, \mathbf{e}_K \mathbf{1}_n^\top] \in \mathbb{R}^{K \times N}, \quad (26)$$

and $\partial \|\mathbf{Z}\|_*$ represents the subdifferential of the nuclear norm of \mathbf{Z} .

Proof. Consider

$$\mathcal{L}(\mathbf{Z}, \mathbf{b}) = \frac{1}{N} l_{CE}(\mathbf{Z} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) + \lambda_Z \|\mathbf{Z}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2$$

in (21). Define

$$\phi(\mathbf{Z}, \mathbf{b}) = \frac{1}{N} l_{CE}(\mathbf{Z} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) = \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n l_{CE}(\mathbf{z}_{ki} + \mathbf{b}, \mathbf{y}_k^\delta), \quad (27)$$

where $\mathbf{y}_k^\delta = (1 - \delta)\mathbf{e}_k + \frac{\delta}{K} \mathbf{1}_K$ is the smoothed target. The gradient of ϕ is

$$\frac{\partial \phi}{\partial \mathbf{z}_{ki}} = \frac{1}{N} (\mathbf{p}_{ki} - \mathbf{y}_k^\delta), \quad \frac{\partial \phi}{\partial \mathbf{b}} = \frac{1}{N} \sum_{k,i} (\mathbf{p}_{ki} - \mathbf{y}_k^\delta),$$

whose matrix form is:

$$\frac{\partial \phi}{\partial \mathbf{Z}} = \frac{1}{N} (\mathbf{P} - \mathbf{Y}^\delta), \quad \frac{\partial \phi}{\partial \mathbf{b}} = \frac{1}{N} (\mathbf{P} - \mathbf{Y}^\delta) \mathbf{1}_N.$$

Hence the gradient (subgradient) of \mathcal{L} is

$$\frac{\partial \mathcal{L}}{\partial \mathbf{Z}} = N^{-1}(\mathbf{P} - \mathbf{Y}^\delta) + \lambda_Z \partial \|\mathbf{Z}\|_*, \quad \frac{\partial \mathcal{L}}{\partial \mathbf{b}} = N^{-1}(\mathbf{P} - \mathbf{Y}^\delta) \mathbf{1}_N + \lambda_b \mathbf{b},$$

where $\partial \|\mathbf{Z}\|_*$ is the subdifferential of the nuclear norm at \mathbf{Z} . Thus, (\mathbf{Z}, \mathbf{b}) is a global minimizer of \mathcal{L} if its gradient (subgradient) is equal to zero, i.e.

$$N^{-1}(\mathbf{Y}^\delta - \mathbf{P}) \in \lambda_Z \partial \|\mathbf{Z}\|_*, \quad N^{-1}(\mathbf{Y}^\delta - \mathbf{P}) \mathbf{1}_N = \lambda_b \mathbf{b}.$$

□

Lemma D.3. Assume that the number of classes K is less than the feature dimension d , i.e., $K \leq d$, and the dataset is balanced. Then the prediction vectors, formulated as $\mathbf{z}_{ki} = f_\Theta(\mathbf{x}_{ki})$, $(1 \leq k \leq K, 1 \leq i \leq n)$ within each class collapse to their sample means $\bar{\mathbf{z}}_k$:

$$\mathbf{z}_{ki} = \bar{\mathbf{z}}_k, \quad 1 \leq i \leq n, \quad (28)$$

In other words, the prediction matrix \mathbf{Z} can be written as the following factorized form:

$$\mathbf{Z} = \bar{\mathbf{Z}} \mathbf{Y} \in \mathbb{R}^{K \times N} \quad (29)$$

where

$$\bar{\mathbf{Z}} = [\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_K], \quad \mathbf{Y} = [\mathbf{e}_1 \mathbf{1}_n^\top, \dots, \mathbf{e}_K \mathbf{1}_n^\top].$$

Proof. The proof follows from the convexity of the loss function. Recall that the loss function in (21) was given by

$$\mathcal{L}(\mathbf{Z}, \mathbf{b}) := \phi(\mathbf{Z}, \mathbf{b}) + \lambda_Z \|\mathbf{Z}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2,$$

where $\phi(\mathbf{Z}, \mathbf{b})$ is the Cross-Entropy (CE) or Label Smoothing (LS) loss (depending on the value of the smoothing parameter δ) as defined in (27). Recalling that $\{\mathbf{z}_{ik}\}_{i=1}^n$ belong to the same class and $\bar{\mathbf{z}}_k = n^{-1} \sum_{i=1}^n \mathbf{z}_{ki}$ is their respective sample mean. From Jensen's inequality, we have

$$\begin{aligned} \phi(\mathbf{Z}, \mathbf{b}) &= \frac{1}{Kn} \sum_{k=1}^K \sum_{i=1}^n l_{CE}(\mathbf{z}_{ki} + \mathbf{b}, \mathbf{y}_k^\delta) \\ &\geq \frac{1}{K} \sum_{k=1}^K l_{CE} \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{z}_{ki} + \mathbf{b}), \mathbf{y}_k^\delta \right) \\ &= \frac{1}{K} \sum_{k=1}^K l_{CE}(\bar{\mathbf{z}}_k + \mathbf{b}, \mathbf{y}_k^\delta) \end{aligned}$$

where the inequality becomes equality only when $\mathbf{z}_{ki} = \bar{\mathbf{z}}_k$.

In the rest of the proof, we employ a permutation argument. Let $\mathbf{Z}_l = [\mathbf{z}_{1l}, \dots, \mathbf{z}_{Kl}]$ for $1 \leq l \leq n$ and $\tilde{\mathbf{Z}} = [\mathbf{Z}_1, \dots, \mathbf{Z}_n]$. Let Γ_i represent a distinct permutation of n . Consider $\tilde{\mathbf{Z}}_{\Gamma_i} = \tilde{\mathbf{Z}} \Pi_{\Gamma_i}$, where Π_{Γ_i} is a permutation matrix rearranging $\tilde{\mathbf{Z}}$ so that the elements \mathbf{Z}_l are ordered according to Γ_i . Then

$$\|\tilde{\mathbf{Z}}_{\Gamma_i}\|_* = \|\tilde{\mathbf{Z}}\|_* = \|\mathbf{Z}\|_*.$$

Since $\|\cdot\|_*$ is a convex function, we deduce

$$\|\mathbf{Z}\|_* = \frac{1}{n!} \left(\sum_i \|\tilde{\mathbf{Z}}_{\Gamma_i}\|_* \right) \geq \left\| \frac{1}{n!} \sum_i \tilde{\mathbf{Z}}_{\Gamma_i} \right\|_*,$$

where the inequality becomes equality only when $\mathbf{Z}_l = \bar{\mathbf{Z}} (1 \leq l \leq n)$ or equivalently $\mathbf{z}_{ki} = \bar{\mathbf{z}}_k$. As a result, it holds that

$$\begin{aligned} \mathcal{L}(\mathbf{Z}, \mathbf{b}) &= \phi(\mathbf{Z}, \mathbf{b}) + \lambda_Z \|\mathbf{Z}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2 \\ &\geq \frac{1}{K} \sum_{k=1}^K l_{CE}(\bar{\mathbf{z}}_k + \mathbf{b}, \mathbf{y}_k^\delta) + \lambda_Z \|\bar{\mathbf{Z}}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2 \\ &= \frac{1}{N} l_{CE}(\bar{\mathbf{Z}} \mathbf{Y} + \mathbf{b} \mathbf{1}_N^\top, \mathbf{Y}^\delta) + \lambda_Z \|\bar{\mathbf{Z}}\|_* + \frac{\lambda_b}{2} \|\mathbf{b}\|^2. \end{aligned}$$

The global optimality of $\mathcal{L}(\mathbf{Z}, \mathbf{b})$ implies that $\mathbf{Z} = \bar{\mathbf{Z}} \mathbf{Y}$, i.e., $\mathbf{z}_{ki} = \bar{\mathbf{z}}_k$ for $1 \leq k \leq n$. □

Lemma D.4. Assume $\mathbf{z}_{ki} = \bar{\mathbf{z}}_k$, for $1 \leq k \leq K, 1 \leq i \leq n$ and $\langle \bar{\mathbf{z}}_k, \mathbf{1}_K \rangle = 0$, then it holds that

$$\begin{aligned} N^{-1} \left((1 - \delta) \mathbf{I}_K + \frac{\delta}{K} \mathbf{J}_K - \bar{\mathbf{P}} \right) &= \lambda_Z \left(n^{-1/2} \left[(\bar{\mathbf{Z}} \bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}} + \bar{\mathbf{R}} \right), \\ \frac{n}{N} \left((1 - \delta) \mathbf{I}_K + \frac{\delta}{K} \mathbf{J}_K - \bar{\mathbf{P}} \right) \mathbf{1}_K &= \lambda_b \mathbf{b}, \end{aligned}$$

where $\mathbf{J}_K \in \mathbb{R}^{K \times K}$ is the ones matrix, $\bar{\mathbf{Z}} = [\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_K] \in \mathbb{R}^{K \times K}$, and $\bar{\mathbf{R}}$ satisfies $\bar{\mathbf{R}} \bar{\mathbf{Z}}^\top = 0$, $\bar{\mathbf{Z}}^\top \bar{\mathbf{R}} = 0$, and $\|\bar{\mathbf{R}}\| \leq n^{-1/2}$.

In particular, if $\bar{\mathbf{Z}}$ is of rank $K - 1$, then

$$N^{-1} \left((1 - \delta) \mathbf{I}_K + \frac{\delta}{K} \mathbf{J}_K - \bar{\mathbf{P}} \right) = \frac{\lambda_Z}{\sqrt{n}} \left(\left[(\bar{\mathbf{Z}} \bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}} \right).$$

Proof. Under the assumption, $\mathbf{z}_{ki} = \bar{\mathbf{z}}_k$, $1 \leq i \leq n$, and $\langle \bar{\mathbf{z}}_k, \mathbf{1}_K \rangle = 0$, we have $\mathbf{Z} = \bar{\mathbf{Z}}\mathbf{Y}$ and $\mathbf{P} = \bar{\mathbf{P}}\mathbf{Y}$, where $\bar{\mathbf{P}}$ is defined as

$$\bar{\mathbf{P}} = [\bar{\mathbf{p}}_1, \dots, \bar{\mathbf{p}}_K]$$

with $\bar{\mathbf{p}}_k$ as the probability vector w.r.t. the $\bar{\mathbf{z}}_k + \mathbf{b}$. Then the optimality condition in (24) reduces to

$$N^{-1} \left((1-\delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} \right) \mathbf{Y} \in \lambda_Z \partial \|\mathbf{Z}\|_*, \quad (30)$$

Let $\mathbf{Z} = \mathbf{U}\Sigma\mathbf{V}^\top$ be the SVD of \mathbf{Z} , then we have:

$$\partial \|\mathbf{Z}\|_* = \{\mathbf{U}\mathbf{V}^\top + \mathbf{R} : \|\mathbf{R}\| \leq 1, \mathbf{U}^\top \mathbf{R} = 0, \mathbf{R}\mathbf{V} = 0\},$$

where

$$\mathbf{U}\mathbf{V}^\top = \left[(\mathbf{Z}\mathbf{Z}^\top)^\dagger \right]^{1/2} \mathbf{Z}.$$

Since $\mathbf{Z} = \bar{\mathbf{Z}}\mathbf{Y}$ and $\mathbf{Y}\mathbf{Y}^\top = n\mathbf{I}_K$, we further get:

$$\mathbf{U}\mathbf{V}^\top = n^{-1/2} \left[(\bar{\mathbf{Z}}\bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}}\mathbf{Y}.$$

Then (30) is equivalent to:

$$N^{-1} \left((1-\delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} \right) \mathbf{Y} = \lambda_Z \left(n^{-1/2} \left[(\bar{\mathbf{Z}}\bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}}\mathbf{Y} + \mathbf{R} \right),$$

where \mathbf{R} is in the form of

$$\mathbf{R} = \bar{\mathbf{R}}\mathbf{Y}, \quad \bar{\mathbf{R}} = [\bar{\mathbf{r}}_1, \dots, \bar{\mathbf{r}}_K]$$

such that

$$\bar{\mathbf{R}}\mathbf{Y} (\bar{\mathbf{Z}}\mathbf{Y})^\top = n\bar{\mathbf{R}}\bar{\mathbf{Z}}^\top = 0, \quad \bar{\mathbf{Z}}^\top \bar{\mathbf{R}} = 0, \quad \|n^{1/2}\bar{\mathbf{R}}\| \leq 1.$$

This further leads to

$$N^{-1} \left((1-\delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} \right) = \lambda_Z \left(n^{-1/2} \left[(\bar{\mathbf{Z}}\bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}} + \bar{\mathbf{R}} \right)$$

where $\bar{\mathbf{R}}\bar{\mathbf{Z}}^\top = 0$, $\bar{\mathbf{Z}}^\top \bar{\mathbf{R}} = 0$ and $\|\bar{\mathbf{R}}\| \leq n^{-1/2}$.

For \mathbf{b} , it is easy to see that the optimality condition in equation 24 reduces to

$$\lambda_b \mathbf{b} = N^{-1} (\mathbf{Y}^\delta - \mathbf{P}) \mathbf{1}_N = \frac{n}{N} \left((1-\delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} \right) \mathbf{1}_K.$$

Now, if $\bar{\mathbf{Z}}$ is of rank $K-1$, since $\mathbf{1}_K^\top \bar{\mathbf{Z}} = 0$, the columns of $\bar{\mathbf{R}}$ are parallel to $\mathbf{1}_K$. Moreover, $\bar{\mathbf{P}}$ is a positive left stochastic matrix with $\bar{\mathbf{P}}^\top \mathbf{1}_K = \mathbf{1}_K$, and therefore $\mathbf{I}_K - \bar{\mathbf{P}}$ is also of rank $K-1$. The left stochasticity of $\bar{\mathbf{P}}$ further deduces that $(\mathbf{I}_K - \bar{\mathbf{P}})^\top \mathbf{1}_K = 0$. In other words, $\mathbf{1}_K$ is in the left null space of $\mathbf{I}_K - \bar{\mathbf{P}}$, which in turn implies $\bar{\mathbf{R}} = 0$. This leads to

$$N^{-1} \left((1-\delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} \right) = \frac{\lambda_Z}{\sqrt{n}} \left(\left[(\bar{\mathbf{Z}}\bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}} \right).$$

□

Lemma D.5. Assume that the number of classes K is less than or equal to the feature dimension d , i.e., $K \leq d$, and the dataset is balanced. Then the global minimizer (\mathbf{Z}, \mathbf{b}) of (21) satisfies the following properties:

$$\mathbf{Z} = \bar{\mathbf{Z}}\mathbf{Y}, \quad \bar{\mathbf{Z}} = a^\delta (\mathbf{I}_K - \mathbf{J}_K/K), \quad \mathbf{b} = 0. \quad (31)$$

In Particular, we have

(i) if $\sqrt{KN}\lambda_Z + \delta \geq 1$, then $a^\delta = 0$;

(ii) if $\sqrt{KN}\lambda_Z + \delta < 1$, then

$$a^\delta = \log \left(\frac{K}{\sqrt{KN}\lambda_Z + \delta} - K + 1 \right)$$

Proof. From Lemma D.3, we have $\mathbf{Z} = \bar{\mathbf{Z}}\mathbf{Y}$. Moreover, according to Lemma 8 in (Zhou et al., 2022b), there exists a constant a such that the global optimum of (21) satisfies:

$$\mathbf{Z} = a (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K) \mathbf{Y}, \quad \mathbf{b} = 0.$$

Equivalently, we have $\bar{\mathbf{Z}} = a (\mathbf{I}_K - \mathbf{1}_K \mathbf{1}_K^\top / K)$. Further, we define $\bar{\mathbf{P}} = [\bar{\mathbf{p}}_1, \dots, \bar{\mathbf{p}}_K]$ with $\bar{\mathbf{p}}_k$ as the probability vector w.r.t. the logit $\mathbf{z}_k + \mathbf{b}$. Then,

$$\bar{\mathbf{P}} = \frac{\mathbf{J}_K + (e^a - 1)\mathbf{I}_K}{K - 1 + e^a}, \quad (32)$$

and

$$\begin{aligned} (1 - \delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} &= (1 - \delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \frac{\mathbf{J}_K + (e^a - 1)\mathbf{I}_K}{K - 1 + e^a} \\ &= \left(\frac{K}{K - 1 + e^a} - \delta \right) \left(\mathbf{I}_K - \frac{1}{K}\mathbf{J}_K \right) \end{aligned}$$

Recall that the first-order optimality condition as expressed in Lemma D.2 indicates that:

$$(1 - \delta)\mathbf{I}_K + \frac{\delta}{K}\mathbf{J}_K - \bar{\mathbf{P}} = N\lambda_Z \left(n^{-1/2} \left[(\bar{\mathbf{Z}}\bar{\mathbf{Z}}^\top)^\dagger \right]^{1/2} \bar{\mathbf{Z}} + \bar{\mathbf{R}} \right) \quad (33)$$

$$= N\lambda_Z \left(\text{sign}(a) \sqrt{\frac{1}{n}} (\mathbf{I}_K - \mathbf{J}_K/K) + \bar{\mathbf{R}} \right). \quad (34)$$

Suppose $a \neq 0$, then $\bar{\mathbf{Z}}$ has rank $K - 1$ and we have $\bar{\mathbf{R}} = 0$. Thus the above optimality condition implies:

$$\left(\frac{K}{K - 1 + e^a} - \delta \right) (\mathbf{I}_K - \mathbf{J}_K/K) = \text{sign}(a) \sqrt{KN}\lambda_Z (\mathbf{I}_K - \mathbf{J}_K/K),$$

which is equivalent to

$$\frac{K}{K - 1 + e^a} - \delta = \text{sign}(a) \sqrt{KN}\lambda_Z.$$

For $a > 0$, the above equation has the following solution:

$$a = \log \left(\frac{K}{\sqrt{KN}\lambda_Z + \delta} - K + 1 \right) \quad \text{if } \sqrt{KN}\lambda_Z + \delta < 1.$$

If $\sqrt{KN}\lambda_Z + \delta \geq 1$, then we select $a = 0$, and noting that in that case $\bar{\mathbf{P}} = K^{-1}\mathbf{J}_K$, we have that the optimality condition in (33) implies that $\bar{\mathbf{R}}$ satisfies

$$(1 - \delta)\mathbf{I}_K + \frac{\delta - 1}{K}\mathbf{J}_K = N\lambda_Z \bar{\mathbf{R}}.$$

We get

$$\bar{\mathbf{R}} = \frac{1}{N\lambda_Z} \left[(1 - \delta)\mathbf{I}_K + \frac{\delta - 1}{K}\mathbf{J}_K \right],$$

where $\|\bar{\mathbf{R}}\| \leq n^{-1/2}$ meets the requirement of the optimality condition.

□

E Proof of Theorem 4.2

In this section of the appendices, we prove Theorem 4.2.

Proof. To provide insight into the accelerated convergence of the model under label smoothing loss, we examine the Hessian matrix of the empirical loss function, defined as follows:

$$\min_{\mathbf{W}, \mathbf{H}} \phi(\mathbf{W}, \mathbf{H}) := \frac{1}{N} l_{CE}(\mathbf{W}^\top \mathbf{H}, \mathbf{Y}^\delta). \quad (35)$$

Under commonly used deep learning frameworks, the model parameters are updated iteratively, and it is common and often practical to analyze the Hessian matrix with respect to \mathbf{H} , \mathbf{W} individually rather than considering the full Hessian. Particularly, we demonstrate that the Hessian of the empirical loss concerning \mathbf{W} when \mathbf{H} is fixed and the Hessian w.r.t. \mathbf{H} when \mathbf{W} is fixed are positive semi-definite at the global optimizer under both cross-entropy and label smoothing losses. Furthermore, we establish that the condition numbers of these Hessian matrices are notably lower under the LS loss in comparison to the CE loss.

Hessian matrix with respect to \mathbf{Z} . Let $\mathbf{Z} = \mathbf{W}^\top \mathbf{H} \in \mathbb{R}^{K \times N}$ represent the prediction logit matrix. In the proof of Lemma D.2, we obtained the first-order partial derivatives of the loss ϕ :

$$\frac{\partial \phi}{\partial \mathbf{z}_{ki}} = \frac{1}{N} (\mathbf{p}_{ki} - \mathbf{y}_k^\delta), \quad \frac{\partial \phi}{\partial \mathbf{b}} = \frac{1}{N} \sum_{k,i} (\mathbf{p}_{ki} - \mathbf{y}_k^\delta),$$

where \mathbf{p}_{ki} is the prediction for the i -th sample that belongs to the k -th class and \mathbf{y}_k^δ is the soft label for class k with a smoothing parameter δ . These partial derivatives lead to the corresponding second-order partial derivatives of the loss ϕ :

$$\begin{aligned} \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} &= \frac{1}{N} (\text{diag}(\mathbf{p}_{ki}) - \mathbf{p}_{ki} \mathbf{p}_{ki}^\top), \quad \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki} \partial \mathbf{z}_{k'i'}} = 0, \forall (k, i) \neq (k', i'), \\ \frac{\partial^2 \phi}{\partial \mathbf{b}^2} &= \frac{1}{N} \sum_{k,i} (\text{diag}(\mathbf{p}_{ki}) - \mathbf{p}_{ki} \mathbf{p}_{ki}^\top), \quad \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki} \partial \mathbf{b}} = \frac{1}{N} (\text{diag}(\mathbf{p}_{ki}) - \mathbf{p}_{ki} \mathbf{p}_{ki}^\top). \end{aligned}$$

From Theorem 4.1, we have the global optimizer satisfies $\mathbf{z}_{ki}^2 = \bar{\mathbf{z}}_k$ and $\mathbf{p}_{ki}^2 = \bar{\mathbf{p}}_k$, $1 \leq k \leq K, 1 \leq i \leq n$. Consequently, it follows that $\frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} = \frac{\partial^2 \phi}{\partial \mathbf{z}_{k'i'}^2}$ ($\forall 1 \leq i, i' \leq n$). To simplify the notation, we denote

$$\mathbf{D}_k = \text{diag}(\bar{\mathbf{p}}_k) - \bar{\mathbf{p}}_k \bar{\mathbf{p}}_k^\top. \quad (36)$$

Then we have $\frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} = \mathbf{D}_k / N$.

The closed-form solution for $\bar{\mathbf{P}}$ in (32) yields the expression:

$$\bar{\mathbf{p}}_k = \frac{1}{K - 1 + e^{a^\delta}} \left((e^{a^\delta} - 1) \mathbf{e}_k + \mathbf{1}_K \right). \quad (37)$$

To simplify the notation, we denote

$$p_t = e^{a^\delta} / (K - 1 + e^{a^\delta}), \quad p_n = 1 / (K - 1 + e^{a^\delta}), \quad (38)$$

where p_t (p_n) is the predicted probability for the target (non-target) class at the global optimal solution. Under this notation, we have $p_t + (K - 1)p_n = 1$ and $\bar{\mathbf{p}}_k = (p_t - p_n) \mathbf{e}_k + p_n \mathbf{1}_K$.

For any $1 \leq k \leq K$, $\mathbf{p}_k \mathbf{p}_k^\top$ is a positive matrix and the associated Laplacian \mathbf{D}_k is positive-semidefinite with all eigenvalues non negative. The smallest eigenvalue of the Laplacian matrix \mathbf{D}_k is $\sigma_1 = 0$ with the corresponding eigenvector $\mathbf{v}_1 = \mathbf{1}_K / \|\mathbf{1}_K\|$.

Define

$$\mathbf{v}_2 = (K\mathbf{e}_1 - \mathbf{1}_K) / \|K\mathbf{e}_1 - \mathbf{1}_K\|,$$

then we have

$$\begin{aligned} \mathbf{D}_k \mathbf{v}_2 &= \text{diag}(\bar{\mathbf{p}}_k) \mathbf{v}_2 - \bar{\mathbf{p}}_k \bar{\mathbf{p}}_k^\top \mathbf{v}_2 \\ &= K p_t p_n \mathbf{v}_2, \end{aligned}$$

from which we get $K p_t p_n$ is an eigenvalue of \mathbf{D}_k with the corresponding eigenvector \mathbf{v}_2 .

Consider the outer-product matrix $\mathbf{p}_k \mathbf{p}_k^\top$, its largest eigenvalue is $\|\mathbf{p}_k\|^2$ with the eigenvector $\mathbf{p}_k / \|\mathbf{p}_k\|$. Its null space is of dimension $K - 1$. Particularly, we can find a set of the basis vectors for $\text{Null}(\mathbf{p}_k \mathbf{p}_k^\top)$ as follows:

$$\{\mathbf{o}_1, \dots, \mathbf{o}_{K-1} : \text{span}(\mathbf{o}_1, \mathbf{p}_k) = \text{span}(\mathbf{e}_k, \mathbf{p}_k)\}, \quad (39)$$

which means the vectors \mathbf{o}_1 and \mathbf{p}_k span the same 2D space as \mathbf{e}_k and \mathbf{p}_k .

Note that

$$\text{span}(\mathbf{e}_k, \mathbf{p}_k) = \text{span}(\mathbf{v}_1, \mathbf{v}_2),$$

where $\mathbf{v}_1, \mathbf{v}_2$ are the eigenvector for \mathbf{D}_k .

Then we can show that $\mathbf{o}_2 \dots, \mathbf{o}_{K-1}$ are eigenvectors of \mathbf{D}_k with the corresponding eigenvalue as $\sigma_3 = p_n$. Particularly, with \mathbf{o}_l orthonormal to both $\bar{\mathbf{p}}_k$ and \mathbf{e}_k , we have

$$\begin{aligned} \mathbf{D}_k \mathbf{o}_l &= \text{diag}(\bar{\mathbf{p}}_k) \mathbf{o}_l - \bar{\mathbf{p}}_k \bar{\mathbf{p}}_k^\top \mathbf{o}_l \\ &= \text{diag}(\bar{\mathbf{p}}_k) \mathbf{o}_l \\ &= p_n \mathbf{o}_l, \end{aligned}$$

for $\mathbf{o}_l (2 \leq l \leq K - 1)$ defined in (39).

To summarize the unique eigenvalues of the matrix \mathbf{D} are

$$\sigma_1 = 0, \quad \sigma_2 = p_n, \quad \sigma_3 = K p_t p_n \quad (40)$$

with p_t and p_n defined in (38).

Hessian matrix with respect to \mathbf{H} . The gradient of ϕ with respect to \mathbf{h}_{ki} is

$$\frac{\partial \phi}{\partial \mathbf{h}_{ki}} = \frac{\partial \mathbf{z}_{ki}}{\partial \mathbf{h}_{ki}} \frac{\partial \phi}{\partial \mathbf{z}_{ki}} = \mathbf{W} \frac{\partial \phi}{\partial \mathbf{z}_{ki}}.$$

Further, we can easily get the corresponding Hessian:

$$\frac{\partial^2 \phi}{\partial \mathbf{h}_{ki}^2} = \mathbf{W} \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} \mathbf{W}^\top = \frac{1}{N} \mathbf{W} \mathbf{D}_k \mathbf{W}^\top, \quad \frac{\partial^2 \phi}{\partial \mathbf{h}_{ki} \partial \mathbf{h}_{k'i'}} = 0, \quad (41)$$

where \mathbf{D}_k is the Laplacian matrix as defined in (36).

From Theorem 4.1, we have

$$\mathbf{W} = \left(\frac{\lambda_H n}{\lambda_W} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{J}_K / K).$$

For simplicity, we denote $a_W^\delta = (\lambda_H n / \lambda_W)^{1/4} \sqrt{a^\delta}$, then $\mathbf{W} = a_W^\delta \mathbf{P} (\mathbf{I}_K - \mathbf{J}_K / K)$. Under this notation, we have

$$\frac{1}{N} \mathbf{W} \mathbf{D}_k \mathbf{W}^\top = \frac{(a_W^\delta)^2}{N} \mathbf{P} (\mathbf{I}_K - \mathbf{J}_K / K) \mathbf{D}_k (\mathbf{I}_K - \mathbf{J}_K / K) \mathbf{P}^\top \quad (42)$$

$$= \frac{(a_W^\delta)^2}{N} \mathbf{P} \mathbf{D}_k \mathbf{P}^\top. \quad (43)$$

Note that $\mathbf{P} \in \mathbb{R}^{d \times K}$, ($d > K$) is a partial orthogonal matrix. Given the eigenvalues of \mathbf{D}_k in (40), the eigenvalues for (42) are:

$$\lambda_1 = 0, \quad \lambda_2 = \frac{(a_W^\delta)^2}{N} p_n, \quad \lambda_3 = \frac{(a_W^\delta)^2}{N} K p_t p_n \quad (44)$$

with the corresponding multiplicities $m(\lambda_1) = 1 + d - K$, $m(\lambda_2) = K - 2$, $m(\lambda_3) = 1$.

From (47), the hessian of ϕ with respect to $\mathbf{h} = \text{vec}(\mathbf{H})$ is a block diagonal matrix which can be expressed as

$$\frac{\partial^2 \phi}{\partial \mathbf{h}^2} = \frac{1}{N} \text{blkdiag}(\mathbf{W} \mathbf{D}_1 \mathbf{W}^\top, \dots, \mathbf{W} \mathbf{D}_1 \mathbf{W}^\top, \dots, \mathbf{W} \mathbf{D}_K \mathbf{W}^\top, \dots, \mathbf{W} \mathbf{D}_K \mathbf{W}^\top). \quad (45)$$

The unique eigenvalues of the Hessian matrix $\frac{\partial^2 \phi}{\partial \mathbf{h}^2}$ are the same as provided in (44). Given that the Hessian matrix contains a zero eigenvalue, our analysis centers on its condition number within the non-zero eigenvalue space. This is calculated as follows:

$$\kappa(\nabla_{\mathbf{h}}^2 \phi) = \lambda_3 / \lambda_2 = K p_t \quad (46)$$

Considering the formula for p_t given in (38), an increase in δ leads to a decrease in p_t , subsequently resulting in a smaller condition number.

Hessian matrix with respect to \mathbf{W} . The gradient of ϕ with respect to \mathbf{w}_l ($l = 1, \dots, K$) is

$$\frac{\partial \phi}{\partial \mathbf{w}_l} = \frac{\partial \mathbf{z}_{ki}}{\partial \mathbf{w}_l} \frac{\partial \phi}{\partial \mathbf{z}_{ki}}.$$

Further, we get the corresponding Hessian:

$$\begin{aligned} \frac{\partial^2 \phi}{\partial \mathbf{w}_l \partial \mathbf{w}_{l'}} &= \sum_{k,i} \frac{\partial \mathbf{z}_{ki}}{\partial \mathbf{w}_{l'}} \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} \left(\frac{\partial \mathbf{z}_{ki}}{\partial \mathbf{w}_l} \right)^\top \\ &= \sum_{k,i} \mathbf{h}_{ki} \mathbf{e}_{l'}^\top \frac{\partial^2 \phi}{\partial \mathbf{z}_{ki}^2} \mathbf{e}_l \mathbf{h}_{ki}^\top \\ &= \sum_{k,i} \mathbf{D}_k(l', l) \mathbf{h}_{ki} \mathbf{h}_{ki}^\top \\ &= n \sum_k \mathbf{D}_k(l', l) \bar{\mathbf{h}}_k \bar{\mathbf{h}}_k^\top \end{aligned}$$

where $\mathbf{D}_k(l', l)$ is (l', l) -th element in \mathbf{D}_k and $\bar{\mathbf{h}}_k$ is the k -th column vector in $\bar{\mathbf{H}}$ as defined in 19. From the definition of \mathbf{D}_k in (36), we have

$$\frac{\partial^2 \phi}{\partial \mathbf{w}_l \partial \mathbf{w}_{l'}} = \begin{cases} n \bar{\mathbf{H}} [\text{diag}(\bar{\mathbf{p}}_l) - \text{diag}(\bar{\mathbf{p}}_l)^2] \bar{\mathbf{H}}^\top & \text{if } l = l' \\ n \bar{\mathbf{H}} [-\text{diag}(\bar{\mathbf{p}}_l) \text{diag}(\bar{\mathbf{p}}_{l'})] \bar{\mathbf{H}}^\top & \text{if } l \neq l' \end{cases}$$

where $\bar{\mathbf{p}}_l$ is the average prediction vector as defined in (37). Particularly, the l -th element of $\bar{\mathbf{p}}_l$ equals p_t and the others equal p_n .

From (19), we have

$$\bar{\mathbf{H}} = \left(\frac{\lambda_W}{\lambda_H n} \right)^{1/4} \sqrt{a^\delta} \mathbf{P} (\mathbf{I}_K - \mathbf{J}_K / K).$$

Hence, the Hessian of ϕ w.r.t. the $\mathbf{w} = \text{vec}(\mathbf{W})$ can be written as:

$$\frac{\partial^2 \phi}{\partial \mathbf{w}^2} = n \left(\frac{\lambda_W}{\lambda_H n} \right)^{1/2} a^\delta \mathbf{D}_P \mathbf{D}_\Pi \mathbf{S} \mathbf{D}_\Pi \mathbf{D}_P^\top \quad (47)$$

where

$$\mathbf{D}_P = \text{blkdiag}(\mathbf{P}, \dots, \mathbf{P}) \in \mathbb{R}^{Kd \times K^2}, \quad (48)$$

$$\mathbf{D}_\Pi = \text{blkdiag}(\mathbf{\Pi}, \dots, \mathbf{\Pi}) \in \mathbb{R}^{K^2 \times K^2}, \quad \mathbf{\Pi} = \mathbf{I}_K - \frac{\mathbf{1}_K \mathbf{1}_K^\top}{K} \quad (49)$$

and

$$\mathbf{S} = \begin{bmatrix} \Lambda_1 & & \\ & \ddots & \\ & & \Lambda_K \end{bmatrix} - \begin{bmatrix} \Lambda_1 \\ \vdots \\ \Lambda_K \end{bmatrix} [\Lambda_1 \quad \dots \quad \Lambda_K] \quad (50)$$

with $\Lambda_l = \text{diag}(\bar{p}_l)$.

Since \mathbf{P} is a partial orthogonal matrix, the condition number of (47) is the same as the condition number of the following matrix:

$$\mathbf{B} = \mathbf{D}_\Pi \mathbf{S} \mathbf{D}_\Pi. \quad (51)$$

Subsequently, we proceed to derive the eigenvalues for the above matrix \mathbf{B} .

First, considering an eigenvector \mathbf{v} of \mathbf{B} , let Δ be a $K \times K$ matrix such that $\mathbf{v} = \text{vec}(\Delta)$. Considering any Δ satisfying $\Delta \mathbf{\Pi} = 0$ or $\Delta^\top \mathbf{\Pi} = 0$, for the corresponding $\mathbf{v} = \text{vec}(\Delta)$ we get $\mathbf{D}_\Pi \mathbf{v} = 0$ and $\mathbf{v}^\top \mathbf{D}_\Pi = 0$, which further yields $\mathbf{B} \mathbf{v} = 0$. Since $\text{rank}(\mathbf{\Pi}) = K - 1$, it follows that $\lambda_1 = 0$ is an eigenvalue of \mathbf{B} with multiplicity $2K - 1$.

On the other hand, consider any Δ that satisfies the conditions:

$$\text{diag}(\Delta) = 0, \quad \Delta \mathbf{1}_K = \Delta^\top \mathbf{1}_K = 0. \quad (52)$$

It is noteworthy that from $\Delta \mathbf{1}_K = \Delta^\top \mathbf{1}_K = 0$, the corresponding vector \mathbf{v} satisfies $\mathbf{D}_\Pi \mathbf{v} = \mathbf{v}$. If Δ further satisfies $\text{diag}(\Delta) = 0$ then $\mathbf{S} \mathbf{v} = p_n \mathbf{v}$. Consequently, we deduce that

$$\mathbf{B} \mathbf{v} = \mathbf{D}_\Pi \mathbf{S} \mathbf{D}_\Pi \mathbf{v} = \mathbf{D}_\Pi \mathbf{S} \mathbf{v} = \mathbf{D}_\Pi p_n \mathbf{v} = p_n \mathbf{v}$$

We can conclude that p_n is an eigenvalue of \mathbf{B} with multiplicity $K^2 - K - (2K - 1) = K^2 - 3K + 1$.

Next, let us consider the matrix $\Delta = \mathbf{\Pi} \text{diag}(\mathbf{u}) \mathbf{\Pi}$, where \mathbf{u} is a vector satisfying $\mathbf{u}^\top \mathbf{1}_K = 0$. The vectorized form of Δ can be denoted as

$$\mathbf{v} = \text{vec}(\Delta) = \begin{bmatrix} \mathbf{\Pi}^\top \text{diag}(\mathbf{u}) \mathbf{\Pi}_1 \\ \mathbf{\Pi}^\top \text{diag}(\mathbf{u}) \mathbf{\Pi}_2 \\ \vdots \\ \mathbf{\Pi}^\top \text{diag}(\mathbf{u}) \mathbf{\Pi}_K \end{bmatrix},$$

where $\mathbf{\Pi}_k$ represents the k -th column of the matrix $\mathbf{\Pi}$. Due to the properties $\mathbf{\Pi} = \mathbf{\Pi}^\top$ and $\mathbf{\Pi}^2 = \mathbf{\Pi}$, it follows that $\mathbf{D}_\Pi \mathbf{v} = \mathbf{v}$. Additionally, given $\mathbf{u}^\top \mathbf{1}_K = 0$, we obtain $\mathbf{S} \mathbf{v} = \lambda_2 \mathbf{v}$ with $\lambda_2 = (1 - p_t + p_n) \cdot \frac{p_n + (K-1)p_t}{K}$. Consequently, we conclude that $\lambda_2 = (1 - p_t + p_n) \cdot \frac{p_n + (K-1)p_t}{K}$ is an eigenvalue of \mathbf{B} with a multiplicity equal to the degrees of freedom of \mathbf{u} , which is $K - 1$.

Lastly, considering $\mathbf{v} = \text{vec}(\mathbf{\Pi})$, we observe that $\mathbf{D}_\Pi \mathbf{v} = \mathbf{v}$. Further, we have $\mathbf{S} \mathbf{v} = K p_t p_n \mathbf{v}$. Hence, we obtain

$$\mathbf{B} \mathbf{v} = \mathbf{D}_\Pi \mathbf{S} \mathbf{D}_\Pi \mathbf{v} = \mathbf{D}_\Pi \mathbf{S} \mathbf{v} = K p_t p_n \mathbf{v}.$$

Therefore, $\mathbf{v} = \text{vec}(\mathbf{\Pi})$ is an eigenvector of \mathbf{B} with eigenvalue $\lambda_3 = K p_t p_n$.

In summary, we identify the distinct eigenvalues of the matrix \mathbf{B} as defined in (51) as follows:

$$\lambda_0 = 0, \quad \lambda_1 = p_n, \quad \lambda_2 = (1 - p_t + p_n) \cdot \frac{p_n + (K-1)p_t}{K}, \quad \lambda_3 = K p_t p_n \quad (53)$$

with the corresponding multiplicities $2K - 1$, $K^2 - 3K + 1$, $K - 1$, and 1, respectively.

Since the matrix \mathbf{B} has zero eigenvalues, we consider its condition number within the non-zero eigenvalue subspace, which is given by $\lambda_3/\lambda_1 = K p_t$. Consequently, we straightforwardly determine the condition number of the Hessian $\frac{\partial^2 \phi}{\partial \mathbf{w}^2}$ as defined in (47) is

$$\kappa(\nabla_{\mathbf{w}}^2 \phi) = K p_t. \quad (54)$$

From the formula of p_t given in (38), it is evident that increasing δ leads to a decrease of p_t , and consequently a reduction in the condition number.

By combining (46) and (54), we demonstrate that the Hessian matrices, $\nabla_{\mathbf{W}}^2 \phi(\mathbf{W}, \mathbf{H})$ and $\nabla_{\mathbf{H}}^2 \phi(\mathbf{W}, \mathbf{H})$, are positive semi-definite at the global optimizer. Moreover, the condition numbers of these Hessian matrices are notably lower under the label smoothing loss (with $0 < \delta < 1 - \sqrt{KN\lambda_W\lambda_H}$) compared to the cross-entropy loss (with $\delta = 0$). This observation suggests that the optimization landscape of the empirical loss function is better conditioned around its global minimizer under label smoothing loss, thus completing our proof.

□