
Investigating Why Contrastive Learning Benefits Robustness against Label Noise

Yihao Xue¹ Kyle Whitecross¹ Baharan Mirzasoleiman¹

Abstract

Self-supervised Contrastive Learning (CL) has been recently shown to be very effective in preventing deep networks from overfitting noisy labels. Despite its empirical success, the theoretical understanding of the effect of contrastive learning on boosting robustness is very limited. In this work, we rigorously prove that the representation matrix learned by contrastive learning boosts robustness, by having: (i) one prominent singular value corresponding to each sub-class in the data, and significantly smaller remaining singular values; and (ii) a large alignment between the prominent singular vectors and the clean labels of each sub-class. The above properties enable a linear layer trained on such representations to effectively learn the clean labels without overfitting the noise. We further show that the low-rank structure of the Jacobian of deep networks pre-trained with contrastive learning allows them to achieve a superior performance initially, when fine-tuned on noisy labels. Finally, we demonstrate that the initial robustness provided by contrastive learning enables robust training methods to achieve state-of-the-art performance under extreme noise levels, e.g., an average of 27.18% and 15.58% increase in accuracy on CIFAR-10 and CIFAR-100 with 80% symmetric noisy labels, and 4.11% increase in accuracy on WebVision.

1. Introduction

Large datasets have enabled deep neural networks to achieve a remarkable success in various domains, such as vision and natural language processing (Deng et al., 2009; Floridi & Chiriatti, 2020). However, this success is highly dependent

¹Department of Computer Science, University of California, Los Angeles, CA 90095, USA. Correspondence to: Yihao Xue <yihaoxue@g.ucla.edu>.

on the quality of the training labels. As datasets grow, manual labeling of data becomes prohibitive and the commonly used web-crawling, crowd-sourcing, and automated data labeling techniques result in noisy labels being ubiquitous in large real-world datasets (Krishna et al., 2016). Over-parameterized networks trained with first-order gradient methods can fit any (even random) labeling of the training data (Zhang et al., 2016). Hence, noisy labels drastically degrade the generalization performance of deep models. To address this, techniques that can robustly learn from noisy labeled data has attracted a lot of attention in recent years (Zheltonozhskii et al., 2022; Li et al., 2020; Zhang et al., 2020; Cao et al., 2020; Mirzasoleiman et al., 2020).

Classical work on robust learning from noisy labels is mainly focused on estimating the noise transition matrix (Goldberger & Ben-Reuven, 2016; Patrini et al., 2017), designing robust loss functions (Ghosh et al., 2017; Van Rooyen et al., 2015; Wang et al., 2019; Zhang & Sabuncu, 2018), correcting noisy labels (Ma et al., 2018; Reed et al., 2014; Tanaka et al., 2018; Li et al., 2020), using explicit regularization techniques (Cao et al., 2020; Zhang et al., 2020; 2017; Liu et al., 2020), and selecting or reweighting training examples (Chen et al., 2019; Han et al., 2018; Jiang et al., 2018; Malach & Shalev-Shwartz, 2017; Ren et al., 2018; Wang et al., 2019; Mirzasoleiman et al., 2020). However as the level of noise increases, these techniques become highly ineffective.

Very recently, self-supervised contrastive learning has shown a lot of promise in boosting robustness of deep networks against noisy labels. Contrastive learning discards all the labels, and learns representations by maximizing agreement between differently augmented views of the same data point via a contrastive loss in the latent space (Chen et al., 2020). Then a linear layer is trained on the representations with the (potentially noisy) labels in a supervised manner. Empirically, networks trained in this way enjoy a superior degree of robustness against noisy labels (Zheltonozhskii et al., 2022; Hendrycks et al., 2019; Ghosh & Lan, 2021).

Despite its empirical success, the theoretical understanding of the effect of contrastive learning on improving robustness of deep networks against noisy labels is very limited. To the best of our knowledge, the only existing theoretic

cal result is on training a binary classifier on pre-trained embeddings obeying a Gaussian distribution (Cheng et al., 2021). The corresponding theory is, however, derived under very limited assumptions, and does not use any properties of self-supervised or contrastive learning.

In this work, we address the above limitations by theoretically characterizing the beneficial properties of representations obtained by contrastive learning for enhancing robustness against noisy labels. We prove that contrastive learning produces a representation matrix that has: (i) a prominent singular value corresponding to each sub-class in the data, and significantly smaller remaining singular values; and (ii) a large alignment between the prominent singular vectors and the ground-truth labels. Then we analyze the case where a linear model is trained on the obtained representations with labels that are either perturbed with Gaussian noise, or flipped at random to other classes. We show that noise has minimal effect on learning the clean labels and the model can hardly memorize the wrong labels.

We further show that deep networks pre-trained with contrastive learning and fine-tuned on noisy labels can achieve a superior performance initially, before overfitting the noise. This is attributed to the initial low-rank structure of the Jacobian. Contrastive pre-training produces a Jacobian matrix with a larger gap between the prominent singular values and the remaining smaller ones, compared to a randomly initialized network. This gap effectively slows down overfitting at the early phase of training. Finally, we demonstrate that the initial robustness provided by contrastive learning can be further leveraged by robust methods to achieve state-of-the-art performance under extreme levels of noise. Such methods do not let the low-rank Jacobian matrix to overfit the noise, even after a long number of training iterations.

We conduct extensive experiments on noisy CIFAR-10 and CIFAR-100 (Krizhevsky & Hinton, 2009), where noisy labels are generated by random flipping the original ones, and the mini Webvision datasets (Li et al., 2017) which is a benchmark consisting of images crawled from websites, containing real-world noisy labels. We show that contrastive learning enables robust training methods to achieve state-of-the-art performance, e.g., an average of 27.18% and 15.58% increase in accuracy on CIFAR-10 and CIFAR-100 with 80% symmetric noisy labels, and 4.11% increase in accuracy on WebVision.

2. Additional Related Work

Contrastive learning and robustness against noise. Recent empirical results demonstrated the effectiveness of self-supervised learning in improving robustness of deep models against adversarial examples, label corruption and input corruption (Hendrycks et al., 2019). Contrastive learning has

been also shown to boost robustness of existing supervised methods (Ghosh & Lan, 2021; Zheltonozhskii et al., 2022) to learn with noisy labels. Notably, Zheltonozhskii et al. (2022) found a large improvement by combining contrastive learning with two state-of-the-art methods, namely ELR (Liu et al., 2020) and DivideMix (Li et al., 2020).

Despite the recent success of contrastive learning in improving robustness of deep networks, a theoretical explanation is yet to be found. Very recently, Cheng et al. (2021) analyzed the performance of a linear binary classifier trained on the embeddings obtained by self-supervised learning. However, their results are based on the assumption that the embeddings follow a Gaussian distribution. Nevertheless, the validity of such assumption and its relation to self-supervised learning is not justified. In contrast, we rigorously prove that contrastive learning extracts the underlying sub-class structure from the augmented data distribution and encodes it into the embeddings. This guarantees the robustness of the downstream supervised learning task.

Theoretical works on self-supervised learning. A recent line of theoretical works have studied self-supervised learning (Arora et al., 2019; Tosh et al., 2021; HaoChen et al., 2021). In particular, it is shown that under conditional independence between positive pairs given the label and/or additional latent variables, representations learned by reconstruction-based self-supervised learning algorithms can achieve small errors in the downstream linear classification task (Arora et al., 2019; Tosh et al., 2021). More closely related to our work is the recent result of HaoChen et al. (2021) that analyzed contrastive learning without assuming conditional independence of positive pairs. Based on the concept of augmentation graph, they showed that spectral decomposition on the augmented distribution leads to embeddings with provable accuracy guarantees under linear probe evaluation. Here, we further leverage the properties of the augmentation graph and provide rigorous robustness guarantees for the performance of linear models trained with representations learned by self-supervised contrastive learning on noisy labels.

3. Problem Formulation and Background

Suppose we have a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, where $(\mathbf{x}_i, \mathbf{y}_i)$ denotes the i -th sample with input $\mathbf{x}_i \in \mathbb{R}^d$ and its clean one-hot encoded label $\mathbf{y}_i \in \mathbb{R}^K$ corresponding to one of the K classes. For example, for a data point \mathbf{x}_i from class $j \in [K]$, we have $\mathbf{y}_i = \mathbf{e}_j$ where \mathbf{e}_j denotes the vector with a 1 in the j th coordinate and 0's elsewhere. We further assume that there are $\bar{K} \geq K$ sub-classes in the data. Sub-classes of a class share the same label, but are distinguishable from each other. For example, apple and orange could be two sub-classes of the class fruit.

We assume that for every data point \mathbf{x}_i , we only observe a noisy version of its label $\hat{\mathbf{y}}_i$. The noise $\Delta\mathbf{y}_i$ can be either generated from a Gaussian distribution $\Delta\mathbf{y}_i = \mathcal{N}(0, \sigma^2 \mathbf{I}_n/K)$, or by randomly flipping the label to one of the other classes. For example for a data point \mathbf{x}_i whose label is flipped from class j to k , we have $\Delta\mathbf{y}_i = \mathbf{e}_k - \mathbf{e}_j$. We denote by $\mathbf{Y}, \hat{\mathbf{Y}} \in \mathbb{R}^{n \times K}$ the matrices of all the one-hot encoded clean and noisy labels of the training data points.

We consider the case where the representations are learned with self-supervised contrastive learning, and then a linear layer is trained with the representations on the noisy labels.

3.1. Self-supervised Contrastive Learning

Self-supervised contrastive learning learns representations of different data points by maximizing agreement between differently augmented views of the same example and minimizing agreement between differently augmented views of different examples. This is achieved via a contrastive loss in the latent space, as we discuss below.

Augmentation graph. The augmentations of different data points can be used to construct the *population augmentation graph* (HaoChen et al., 2021), whose vertices are all the augmented data points in the population distribution, and two vertices are connected with an edge if they are augmentations of the same natural (original) example. Hence, ground-truth classes naturally form connected sub-graphs. Formally, let P be the distribution of all natural data points (raw inputs without augmentation). For a natural data point $\mathbf{x}^* \sim P$, let $\mathcal{A}(\cdot|\mathbf{x}^*)$ be the distribution of \mathbf{x}^* 's augmentations. For instance, when \mathbf{x}^* represents an image, $\mathcal{A}(\cdot|\mathbf{x}^*)$ can be the distribution of common augmentations (Chen et al., 2020) including Gaussian blur, color distortion and random cropping. Then, for an augmented data point \mathbf{x} , $\mathcal{A}(\mathbf{x}|\mathbf{x}^*)$ is the probability of generating \mathbf{x} from \mathbf{x}^* . The edge weights $w_{\mathbf{x}_i \mathbf{x}_j} = \mathbb{E}_{\mathbf{x}^* \sim P}[\mathcal{A}(\mathbf{x}_i|\mathbf{x}^*)\mathcal{A}(\mathbf{x}_j|\mathbf{x}^*)]$ can be interpreted as the marginal probability of generating \mathbf{x}_i and \mathbf{x}_j from a random natural data point.

Contrastive loss. The embeddings produced by contrastive learning can be viewed as a low-rank approximation of the normalized augmentation graph. Effectively, minimizing a loss that performs spectral decomposition on the population augmentation graph can be succinctly written as a contrastive learning objective $\mathcal{C}(f)$ on neural network representations (HaoChen et al., 2021):

$$\mathcal{C}(f) = -2\mathbb{E}_{\mathbf{x}, \mathbf{x}^+}[f(\mathbf{x})^\top; f(\mathbf{x}^+)] + \mathbb{E}_{\mathbf{x}, \mathbf{x}^-}[((f(\mathbf{x})^\top; f(\mathbf{x}^-))^2)], \quad (1)$$

where $f(\mathbf{x}) \in \mathbb{R}^p$ is the neural network representation for an input \mathbf{x} , and \mathbf{x}, \mathbf{x}^+ are drawn from the augmentations of the same natural data point, and \mathbf{x}, \mathbf{x}^- are two augmentations generated independently either from the same data point or two different data points. The above loss function is similar

to many standard contrastive loss functions (Oord et al., 2018; Sohn, 2016; Wu et al., 2018), including SimCLR (Chen et al., 2020) that we will use in our experiments. Minimizing this objective leads to representations with provable accuracy guarantees under linear probe evaluation. We use f_{\min} to denote the minimizer, i.e., $f_{\min} = \arg \min_f \mathcal{C}(f)$.

3.2. Training the Linear Head with Label Noise

Here, we introduce the notations for training a linear classifier on the representations learned by contrastive learning, based on which we perform theoretical analysis. In Section 4.2, we discuss how our idea can be extended to understand the performance of fine-tuning all the layers of the neural network.

We assume the representations are given by f_{\min} , the global minimizer of the contrastive loss. In practice, this is easier to be achieved by larger networks trained for longer (Chen et al., 2020). In Section 5.1, we confirm superior robustness of representations learned by larger networks against noisy labels by our experiments. Given a matrix $\mathbf{F} \in \mathbb{R}^{n \times p}$ where each row $\mathbf{F}_i = f_{\min}(\mathbf{x}_i)^\top$ is the learned representation of a data point \mathbf{x}_i , we consider the downstream task of training a linear model, parameterized by $\mathbf{W} \in \mathbb{R}^{p \times K}$, to minimize the MSE loss with l_2 regularization with parameter β

$$\min_{\mathbf{W} \in \mathbb{R}^{p \times K}} \frac{1}{2} \|\hat{\mathbf{Y}} - \mathbf{F}\mathbf{W}\|_{\mathbb{F}}^2 + \frac{\beta}{2} \|\mathbf{W}\|. \quad (2)$$

Let $\hat{\mathbf{W}}^*$ denote the solution that has the following closed-form expression

$$\hat{\mathbf{W}}^* = (\mathbf{F}^\top \mathbf{F} + \beta \mathbf{I})^{-1} \mathbf{F}^\top \hat{\mathbf{Y}}. \quad (3)$$

While we use MSE in our analysis, we empirically show that our results hold for other losses, such as cross-entropy.

4. Contrastive learning Boosts Robustness

In this section we first show that training a linear head on representations learned by contrastive learning is provably robust to label noise. Then we look into the phenomenon that fine-tuning the deep network pre-trained by contrastive learning achieves a superior performance at early phase of training. Finally, we discuss how the initial robustness provided by contrastive learning boosts robust training methods, and corroborate this with extensive experiments in Section 5.

4.1. Provable Robustness of the Linear Head

To understand the robustness provided by contrastive learning, we assume certain properties of the augmentation graph and analyze the low-rank structure of the the resulting representation matrix. In particular, we utilize the following natural assumptions that formalize the following two properties on the data augmentation: (1) the augmented examples




Figure 1. Training accuracy w.r.t. ground-truth labels and test accuracy of a linear classifier trained on representations learned by contrastive learning (SimCLR). Experiments are conducted on CIFAR-10 (C10) and CIFAR-100 (C100) under different noise levels. Dashed lines show loss and accuracy on training set w.r.t. ground-truth labels and solid lines show test loss and accuracy.

of one sub-class are similar to each other; and (2) the augmented examples of one sub-class are different from the augmented examples of other sub-classes.

Assumption 4.1 (Compact sub-class structure). For a triple of augmented examples \mathbf{x}_j , \mathbf{x}_s and \mathbf{x}_t from the same sub-class, the marginal probability of \mathbf{x}_s , \mathbf{x}_j being generated from a natural data point is close to that of \mathbf{x}_t , \mathbf{x}_j . Formally, we have $w_{\mathbf{x}_s \mathbf{x}_j} / w_{\mathbf{x}_i \mathbf{x}_j} \in [\frac{1}{1+\delta}, 1 + \delta]$, for small $\delta \in [0, 1)$.

Assumption 4.2 (Distinguishable sub-class structure). For two pairs of augmented examples $(\mathbf{x}_i, \mathbf{x}_j)$ and $(\mathbf{x}_s, \mathbf{x}_t)$ where \mathbf{x}_i , \mathbf{x}_j are from different sub-classes and \mathbf{x}_s , \mathbf{x}_t are from the same sub-class, the marginal probability of \mathbf{x}_i , \mathbf{x}_j being generated from a natural data is much smaller than that of \mathbf{x}_s , \mathbf{x}_t . Formally, we have $w_{\mathbf{x}_i \mathbf{x}_j} / w_{\mathbf{x}_s \mathbf{x}_t} \leq \xi$, for small $\xi \in [0, 1)$.

The above assumptions result in an augmentation graph where augmented data points from different subclasses form nearly disconnected subgraphs with similar edge weights. In particular for $\xi = 0$, we get disconnected subgraph structure.

4.1.1. DESIRABLE PROPERTIES OF CL REPRESENTATIONS

The key to our analysis is that, based on compact and distinguishable sub-class structure assumptions 4.1, 4.2, contrastive learning produces a low-rank representation matrix \mathbf{F} that captures the sub-class structure. More formally, the representation matrix has \bar{K} singular vectors that align well with the ground-truth labels, and the corresponding \bar{K} singular values are significantly larger than the other singular values. The following theorem is a summary of Lemmas A.2 A.3 B.2 B.3 and Corollary A.6 in the Appendix which details the desirable properties of the representation matrix.

Theorem 4.3. *Having \bar{K} compact and distinguishable sub-classes in the data, the representation matrix \mathbf{F} learned by contrastive learning has \bar{K} prominent singular values of magnitude $\mathcal{O}(1)$. At the same time, the sum of the remaining singular values is significantly smaller, i.e., $\mathcal{O}(\sqrt{\delta} + \xi)$. Furthermore, the most prominent \bar{K} singular vectors and the ground-truth labels has a $\mathcal{O}(1)$ alignment, measured by*

the normalized projection of the clean labels \mathbf{Y} onto the span of the singular vectors.

Intuitively, the above three properties of the representation matrix affect the downstream training in the following sense: (1) the magnitude of largest singular values determines the speed at which the model evolves as well as the extent to which the model can fit the training data; (2) the alignment between prominent singular vectors and clean labels indicates whether the model evolves in the right direction; and (3) the magnitude of smaller singular values dictates the amount of overfitting. As a result, Theorem 4.3 implies that the model trained on such representation learns mainly the correct information from the training data, which we formally show in Theorems 4.4 and 4.5.

4.1.2. TRAINING PERFORMANCE W.R.T. GROUND-TRUTH LABELS REFLECTS ROBUSTNESS

To simplify the theoretical analysis, instead of studying the generalization performance (usually measured by the expected loss over the data distribution), we will examine the loss and accuracy on the training data w.r.t. ground-truth labels. This strongly correlates with the test accuracy, especially under large noise. We empirically confirm this correlation in Figure 1, where the dashed lines show training loss and training accuracy w.r.t. ground-truth labels, and solid lines show test loss and test accuracy. We clearly see the high correlation between training and performance, in particular under significant levels of label noise.

4.1.3. GAUSSIAN LABEL NOISE

We first consider the case where label noise is generated from a Gaussian distribution. Formally, $\hat{\mathbf{Y}} = \mathbf{Y} + \Delta\mathbf{Y}$, where \mathbf{Y} is the clean label matrix containing all the one-hot encoded labels, and $\Delta\mathbf{Y}$ is the label noise matrix, where each column drawn independently from $\mathcal{N}(0, \sigma^2 \mathbf{I}_n / K)$. We consider this setting first, as it provides the most convenient way to analyze robustness. Here, our analysis mainly aims at breaking down the effect of label perturbations on

training dynamics, in terms of bias and variance. This could provide theoretical insights into the benefits of contrastive learning for boosting robustness.

The following theorem bounds the expected error on training data w.r.t. ground-truth labels, and shows how contrastive learning exploits the augmented sub-class structure to improve robustness.

Theorem 4.4. *For a dataset of size n with K classes, \bar{K} balanced compact and distinguishable sub-classes (c.f. assumptions 4.2, 4.1) and labels corrupted with Gaussian noise $\mathcal{N}(0, \sigma^2 \mathbf{I}_n / K)$, a linear model trained by minimizing the objective in Eq. (2) with the representations obtained by minimizing contrastive loss in Eq. (1) has the following expected error on the training set w.r.t. the ground-truth labels \mathbf{Y} :*

$$\begin{aligned} & \mathbb{E}_{\Delta \mathbf{Y}} \frac{1}{n} \|\mathbf{Y} - \mathbf{F}\hat{\mathbf{W}}^*\|_F^2 \quad (4) \\ & \leq \underbrace{\left(\frac{\beta}{\beta+1}\right)^2 + \mathcal{O}(\delta + \xi)}_{\text{bias}^2} + \underbrace{\sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1}\right)^2 + \sigma^2 \mathcal{O}\left(\frac{\sqrt{\delta} + \xi}{\beta}\right)}_{\text{variance}}. \end{aligned}$$

We note that the above results can be easily extended to imbalanced sub-class structure.

The proof can be found in the Appendix. The proof follows the intuition discussed in Section 4.1.1 that the desirable properties of the learned representation benefit the downstream training. In a nutshell, we derive the bound by writing the error in terms of singular values and vectors of \mathbf{F} and then applying Theorem 4.3.

In Eq. (4), the error is decomposed into bias and variance. The bias captures the mismatch between the average prediction of the model and the ground-truth labels. It depends on the magnitude of the prominent singular values as well as the alignment of the corresponding singular vectors with the ground-truth labels. Contrastive learning reduces the bias by aligning the first \bar{K} singular vectors with ground-truth labels (Theorem 4.3), thus producing a small second term in the bias. The variance quantifies the sensitivity to label noise, and is controlled by the magnitude of the non-prominent singular values, which is guaranteed to be small by Theorem 4.3. The regularization parameter β appears in both terms and can be tuned as a trade-off between underfitting and overfitting.

With small enough δ and ξ , one can select a small β to not explicitly penalize the variance much. This results in a small bias, and subsequently a small total error. For example, when there exists a β such that $\sqrt{\delta} + \xi \ll \beta \ll 1$, the error $\approx \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1}\right)^2$, which is the inevitable cost of achieving a small bias, when the representation matrix \mathbf{F} has \bar{K} prominent singular values.

4.1.4. RANDOM LABEL FLIPPING

Next, we study the case where the label noise $\Delta \mathbf{Y} = \hat{\mathbf{Y}} - \mathbf{Y}$ is generated by flipping a fraction of the clean labels at random. Formally, for an example \mathbf{x}_i belongs to class j with $\mathbf{y}_i = \mathbf{e}_j$, if its label is flipped to class k , we have $\Delta \mathbf{y}_i = \mathbf{e}_k - \mathbf{e}_j$. We introduce the following notations to analyze the case of asymmetric label noise which mimics the real-world noise, where wrong labels are generated in a (sub)class-dependent way. Let $m_{\bar{k}}$ be the number of mislabeled examples in subclass \bar{k} , $m_{\bar{k},k}$ be the number of mislabeled examples in subclass \bar{k} that are labeled as class k , and Z_k be the set of sub-classes in class k . We define $c_{k|\bar{k}} := \frac{m_{\bar{k},k}}{m_{\bar{k}}}$ for all \bar{k}, k such that $\bar{k} \notin Z_k$, to be the fraction of mislabeled examples in sub-class \bar{k} that are mislabeled as k . We have that $\sum_{k: \bar{k} \notin Z_k} c_{k|\bar{k}} = 1, \forall \bar{k} \in [\bar{K}]$. When $c_{k|\bar{k}} = \frac{1}{K-1} \forall k \in [K], \bar{k} \in [\bar{K}]$, the noise is symmetric.

For simplicity we assume $\xi = 0$. The general case of $\xi \geq 0$ requires more involved analysis which we discuss in the Appendix. The following theorem shows that for a dataset with compact and distinguishable sub-class structure the linear classifier trained on the representations obtained by contrastive learning can recover the clean label for all training data.

Theorem 4.5 (Asymmetric Noise). *For a dataset with K classes and \bar{K} compact and distinguishable sub-class structure (c.f. assumptions 4.2 4.1) with $\xi = 0$, let n_{\min}, n_{\max} be the size of the smallest and largest sub-class, and α be the fraction of mislabeled examples in the training set. Let $c_{\max} := \max_{k \in [K], \bar{k} \in [\bar{K}]} c_{k|\bar{k}} \in [\frac{1}{K-1}, 1]$ be the maximum fraction of wrong labels in a subclass that are flipped to another class. Then as long as*

$$\alpha < \frac{1}{1 + \frac{n_{\max}}{n_{\min}} c_{\max}} - \mathcal{O}(\sqrt{\delta}), \quad (5)$$

a linear model trained by minimizing the objective in Eq. (2) with the representations obtained by minimizing contrastive loss in Eq. (1) can predict the ground-truth labels for all training examples, i.e.,

$$\forall i, \arg \max_{j \in [K]} (\mathbf{F}\hat{\mathbf{W}})_{i,j} = \arg \mathfrak{S}_{j \in [K]} (\mathbf{Y}_{i,j} = 1)$$

In other words contrastive learning can prevent the linear model from memorizing any wrong label even under large noise. Theorem 4.5 also shows that the model can tolerate more noise when the sub-class structure is more compact, i.e., δ is smaller, or the noise is more symmetric, or the sub-classes are more balanced. The following corollary for symmetric noise is simply obtained by setting $c_{\max} = \frac{1}{K-1}$ in Theorem 4.5.

Corollary 4.6 (Symmetric Noise). *For symmetric noise,*

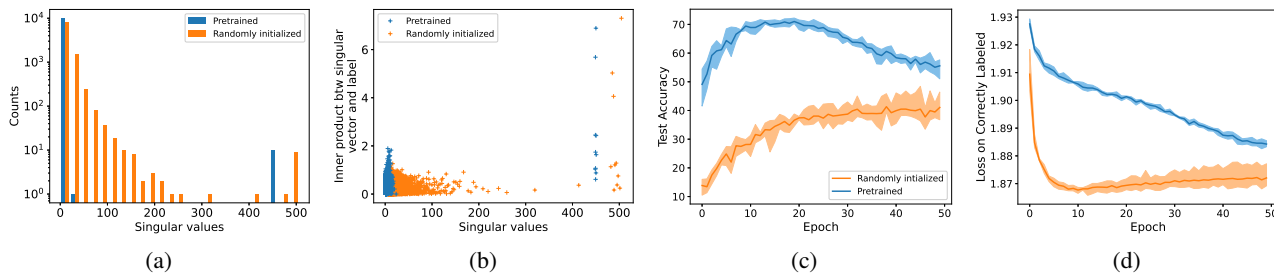


Figure 2. Jacobian spectrum and dynamics of training a randomly initialized vs. fine-tuning a pre-trained ResNet32 on CIFAR-10 with 80% randomly flipped labels. (a) distribution of singular values of the initial Jacobian, (b) alignment of the clean labels with the initial Jacobian, (c) test accuracy and (d) loss on noisy labeled data points within the first 50 epochs. While pre-training does not improve the alignment of the Jacobian with the labels, it significantly slows down overfitting at the beginning by shrinking the smaller singular values of the Jacobian matrix.

under the same assumption as in Theorem 4.5, as long as

$$\alpha < \frac{K-1}{K + \frac{n_{\max}}{n_{\min}} - 1} - \mathcal{O}(\sqrt{\delta}), \quad (6)$$

a linear model trained by minimizing the objective in Eq. (2) with the representations obtained by minimizing contrastive loss in Eq. (1) can predict the ground-truth labels for all training examples.

If we further let $\delta = 0$ and $n_{\max}/n_{\min} = 1$, we get $(K-1)/K$ noise tolerance. We note that this, however, does not imply that a dataset with more classes necessarily has a higher noise tolerance. In Appendix B.1, we show that less distinguishable sub-class structure, i.e. $\xi > 0$, introduces a $\mathcal{O}(\bar{K}^{5/2}\xi)$ perturbation to the singular values and a $\mathcal{O}(\bar{K}^{5/2}\xi)$ rotation in the direction of singular vectors of the representation matrix. Datasets with more classes usually contains more sub-classes, which greatly reduces the noise tolerance. This is also reflected by our empirical results (Figure 1) where the performance of the linear model is worse on CIFAR-100 compared to CIFAR-10 under the same noise level.

4.2. Fine-tuning the Network to Boost Performance

As discussed, representations obtained by contrastive learning provably boosts robustness of a linear classifier trained on noisy labels. Here, we first discuss the benefits of fine-tuning deep networks trained by contrastive learning over that of training randomly initialized networks, on boosting robustness against noisy labels. Then, we show that the initial robustness provided by contrastive learning enables robust training methods that are effective under small to moderate amount of noisy labels (Liu et al., 2020; Zhang et al., 2017; Mirzasoleiman et al., 2020) to achieve the state-of-the-art performance under extreme noise levels.

Contrastive learning slows down overfitting. In the previous section, we showed that training a linear model on

representations learned by contrastive learning is provably robust. Here, we study fine-tuning all layers of the deep network. Interestingly, we show that fine-tuning all layers achieves a superior performance initially, but overfits the noise later.

Recall that the theoretical guarantee for linear model (theorems 4.4 and 4.5) is obtained by examining singular values and singular vectors of \mathbf{F} . Here, we use a similar idea to understand benefits of contrastive learning on robustness when all the layers are trained. Intuitively, during the early stage of training, it is natural to assume that the gradient does not considerably change, and therefore the model is nearly linear. In this case, the initial Jacobian matrix plays the same role as the representation matrix, \mathbf{F} , to the linear model. This is supported by the recent studies suggesting the following two properties: the early learning dynamics can mimicked by training a linear model (Hu et al., 2020), and a network that provides a better alignment between prominent directions of the Jacobian matrix and the label vector is more likely to generalize well (Oymak et al., 2019).

We examine the spectrum of the Jacobian of a network pre-trained with contrastive learning and compare it to that of a randomly initialized network. Fig. 2a, 2b compares the Jacobian spectrum and alignment of clean label vector with Jacobian of ResNet32 on a random sample of 1000 data points from CIFAR10. Interestingly, Fig. 2b shows that while pre-training does not considerably improve (in Appendix C we show a slight improvement) the alignment between singular vectors of the Jacobian and the clean label vector, it greatly shrinks the smaller singular values of the Jacobian, as is illustrated by Fig. 2a. As a result, it takes substantially longer for the pre-trained network to overfit the noisy labels. As Fig. 2d shows, while a randomly initialized network experience a sharp drop in loss of noisy labeled data points during the first few epochs of training, it takes much longer for a pre-trained network to overfit the noise. In particular, Fig. 2c demonstrates the superior performance

of a pre-trained network over that of a randomly initialized before it overfits the noise. Details of the experiment can be found in Appendix D.

4.3. Contrastive learning boosts robust methods

As discussed, pre-training the network with contrastive learning effectively shrinks the smaller singular values of the Jacobian and slows down overfitting the noisy labels. The initial level of robustness provided by contrastive learning can be leveraged by existing robust training methods to achieve a superior performance under extreme noise levels. Next, we briefly discuss three methods that prevent the pre-trained network from overfitting the noisy labels, through regularization (Liu et al., 2020; Zhang et al., 2017), or identifying clean examples (Mirzasoleiman et al., 2020).

ELR (Liu et al., 2020) regularizes the loss by $\frac{1}{n} \sum_{i=1}^n \log(1 - \mathbf{p}(x_i)^\top \mathbf{t}(x_i))$ to encourage the alignment between the model prediction $\mathbf{p}(x)$ and the running average of the model outputs in previous rounds $\mathbf{t}(x)$. The effectiveness of ELR is attributed to the early-learning phenomenon where the model first fits the correct labels and then memorizes the noisy ones (Oymak et al., 2019). Effectively, the regularization term stretches the prediction toward the clean labels predicted early by the model. However, under extreme label noise, the memorization phase starts very early, and does not let the model to learn clean labels and high-quality targets. As discussed, contrastive learning makes a large separation between learning and memorization and gives the network enough time to learn high-quality targets. As we show in our experiments, applying ELR to fine-tune the network learned by contrastive learning significantly boosts the generalization performance.

Mixup (Zhang et al., 2017) extends the training distribution by linear interpolations of feature vectors and their associated labels: $\hat{\mathbf{x}} = \lambda \mathbf{x}_i + (1 - \lambda) \mathbf{x}_j$, $\hat{\mathbf{y}} = \lambda \mathbf{y}_i + (1 - \lambda) \mathbf{y}_j$, where $\lambda \sim \text{Beta}(\alpha, \alpha)$. In doing so, mixup makes linear transition in the decision boundary between classes and provide a smoother estimate of uncertainty. Larger α prevents overfitting by generating examples that are less similar to the training examples and are more difficult for the network to memorize. In our experiments, we show that the network learned by contrastive learning can be robustly fine-tuned by mixup to achieve a superior generalization performance.

CRUST (Mirzasoleiman et al., 2020) provides provable robustness guarantees by extracting clean examples that cluster closely in the gradient space based on the following observation: as the nuisance space is very high dimensional, data points with noisy labels spread out in the gradient space. In contrast, the information space is low-dimensional and data points with clean labels that have similar gradients cluster closely together. Central clean examples in the gradient

space can be efficiently extracted by maximizing a submodular function. To enable Crust to find the clean examples under extreme noise, we first fine-tune the entire network on noisy labels for around 20 epochs and then randomly label half of the examples with the prediction of the model. As discussed, pre-training the network with contrastive learning shrinks the smaller singular values of the Jacobian. This allows the clean examples to make clear clusters around the large singular directions and be easily extracted. In our experiments, we show that the pre-trained network can significantly boost Crust’s performance under extreme noise.

5. Experiments

We evaluate the effectiveness of contrastive learning in boosting the robustness of deep networks under various levels of label noise. We first consider fine-tuning all layers of a network pre-trained with contrastive learning on noisy labels, and show that it can achieve a comparable generalization performance to the state-of-the-art robust methods. Then, we show that the structure of the representation matrix obtained by contrastive learning can be leveraged by robust methods to achieve a superior generalization performance under extreme noise levels.

For our evaluation, we use artificially corrupted versions of CIFAR-10 and CIFAR-100 (Krizhevsky & Hinton, 2009), as well as a subset of the real-world dataset Webvision (Li et al., 2017), which naturally contains noisy labels. Our method was developed using PyTorch (Paszke et al., 2017). We use 1 Nvidia A40 for all experiments.

Baselines. We compare our results with many commonly used baselines for robust training against label noise: (1) F-correction (Patrini et al., 2017) is a two step process, where a neural network is first trained on noisily-labelled data, then retrained using a corrected loss function based on an estimation of the noise transition matrix. (2) Decoupling (Malach & Shalev-Shwartz, 2017) is a meta-algorithm that trains two networks concurrently, only training on examples where the two networks disagree. (3) Co-teaching (Han et al., 2018) also trains two networks simultaneously. Each network selects subsets of clean data with high probability for the other network to train on. (4) MentorNet (Jiang et al., 2018) uses two neural networks, a student and a mentor. The mentor dynamically creates a curriculum based on the student, while the student trains on the curriculum provided by the mentor. (5) D2L (Ma et al., 2018) learns the training data distribution, then dynamically adapts the loss function based on the changes in dimensionality of subspaces during training. (6) INCV (Chen et al., 2019) identifies random subsets of the training data with fewer noisy labels, then applies Co-teaching to iteratively train on subsets found with the most clean labels. (7) T-Revision (Xia et al., 2019) learns the transition matrix efficiently using an algorithm that does

Table 1. Average test accuracy (3 runs) on CIFAR-10 and CIFAR-100. The best test accuracy is marked in bold. We note the higher performance of methods that use SimCLR (CL) pretraining, especially under higher noise levels. In particular, under 80% noise, methods see an average of 27.18%, and 15.58% increase in test accuracy for CIFAR-10 and CIFAR-100 respectively. Results marked with (*) are reproduced from publicly available code. E2E refers to end to end fine-tuning the pre-trained network.

Dataset	CIFAR-10				CIFAR-100			
	Noise Type	Sym		Asym	Sym		Asym	
Noise Ratio	20	50	80	40	20	50	80	40
F-correction	85.1 ± 0.4	76.0 ± 0.2	34.8 ± 4.5	83.6 ± 2.2	55.8 ± 0.5	43.3 ± 0.7	–	42.3 ± 0.7
Decoupling	86.7 ± 0.3	79.3 ± 0.6	36.9 ± 4.6	75.3 ± 0.8	57.6 ± 0.5	45.7 ± 0.4	–	43.1 ± 0.4
Co-teaching	89.1 ± 0.3	82.1 ± 0.6	16.2 ± 3.2	84.6 ± 2.8	64.0 ± 0.3	52.3 ± 0.4	–	47.7 ± 1.2
MentorNet	88.4 ± 0.5	77.1 ± 0.4	28.9 ± 2.3	77.3 ± 0.8	63.0 ± 0.4	46.4 ± 0.4	–	42.4 ± 0.5
D2L	86.1 ± 0.4	67.4 ± 3.6	10.0 ± 0.1	85.6 ± 1.2	12.5 ± 4.2	5.6 ± 5.4	–	14.1 ± 5.8
INCV	89.7 ± 0.2	84.8 ± 0.3	52.3 ± 3.5	86.0 ± 0.5	60.2 ± 0.2	53.1 ± 0.4	–	50.7 ± 0.2
T-Revision	79.3 ± 0.5	78.5 ± 0.6	36.2 ± 1.6	76.3 ± 0.8	52.4 ± 0.3	37.6 ± 0.3	–	32.3 ± 0.4
L_DMI	84.3 ± 0.4	78.8 ± 0.5	20.9 ± 2.2	84.8 ± 0.7	56.8 ± 0.4	42.2 ± 0.5	–	39.5 ± 0.4
Crust*	85.3 ± 0.5	86.8 ± 0.3	33.8 ± 1.3	76.7 ± 3.4	62.9 ± 0.3	55.5 ± 1.1	18.5 ± 0.8	52.5 ± 0.4
Mixup	89.7 ± 0.7	84.5 ± 0.3	40.7 ± 1.1	86.3 ± 0.1	64.0 ± 0.4	53.4 ± 0.5	15.1 ± 0.1	54.4 ± 2.0
ELR*	90.6 ± 0.6	87.7 ± 1.0	69.5 ± 5.0	86.6 ± 2.9	63.6 ± 1.7	52.5 ± 4.2	23.4 ± 1.9	59.7 ± 0.1
CL+E2E*	88.8 ± 0.5	82.8 ± 0.2	72.0 ± 0.3	83.5 ± 0.5	63.5 ± 0.2	56.1 ± 0.3	36.7 ± 0.3	52.4 ± 0.2
CL+Crust*	86.5 ± 0.7	87.6 ± 0.3	77.9 ± 0.3	85.9 ± 0.4	63.0 ± 0.8	58.3 ± 0.1	34.8 ± 1.5	53.3 ± 0.7
CL+Mixup*	90.8 ± 0.2	84.6 ± 0.4	74.8 ± 0.3	87.5 ± 1.3	64.4 ± 0.4	55.5 ± 0.1	30.3 ± 0.4	55.5 ± 0.8
CL+ELR*	91.3 ± 0.0	89.1 ± 0.1	77.7 ± 0.2	89.7 ± 0.3	64.7 ± 0.2	55.6 ± 0.2	35.9 ± 0.3	63.6 ± 0.1

not rely on known points with clean labels. (8) L_DMI (Xu et al., 2019) uses a novel information-theoretic loss function based on determinant based mutual information. (9) ELR (Liu et al., 2020) uses semi-supervised learning techniques to regularize based on the early-learning phase of training, to ensure the noisy labels are not overfit. (10) CRUST (Mirzasoileiman et al., 2020) dynamically selects subsets of clean data points by clustering in the gradient space. (11) Mixup (Zhang et al., 2017) smooths the decision boundary by adding linear interpolations of feature vectors and their labels to the dataset.

5.1. Empirical Results on Artificially Corrupted CIFAR

We first evaluate our method on CIFAR-10 and CIFAR-100, which each contain 50,000 training images, and 10,000 test images of size $32 \times 32 \times 3$. CIFAR-10 and CIFAR-100 have 10 and 100 classes respectively. We use the same testing protocol as (Xu et al., 2019; Liu et al., 2020; Xia et al., 2019), by evaluating our method on symmetric and asymmetric label noise. For both CIFAR-10 and CIFAR-100, we use symmetric noise ratios of 0.2, 0.5, 0.8, and an asymmetric noise ratio of 0.4.

In our experiments, we first pre-train ResNet-32 (He et al., 2016) using SimCLR (Chen et al., 2020; SimCLR) for 1000 epochs using the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 3×10^{-4} , a weight decay of 1×10^{-6} and a batch size of 128. When pre-training, the last linear layer of ResNet-32 is replaced with a 2-layer projection head with an output dimensionality of 64. When pretraining is finished, we replace the projection head with a new,

randomly initialized classification layer, and begin training normally. We also report the results when ELR, Mixup, and Crust are applied to fine-tune the pre-trained network. For ELR, we use $\beta = 0.7$ for the temporal ensembling parameter, and $\lambda = 3$ for the regularization strength. For mixup, we use a mixup strength of $\alpha = 1$. For Crust, we choose a coreset ratio of 0.5.

The results are shown in Table 1. We note that SimCLR pretraining leads to an across the board improvement for Crust, ELR, and Mixup. For lower noise ratios, the improvement is marginal. However, for extreme noise ratios, the improvement is more dramatic. In particular, pre-training boosts the performance of Crust by up to 44.1%, ELR by up to 8.2%, and Mixup by up to 34.1% under 80% noise. We also note that under 80% noise, SimCLR pretraining alone outperforms all methods without SimCLR pretraining on CIFAR-10 and CIFAR-100.

Effects of Network Size Empirically, larger networks trained for longer can achieve smaller contrastive loss (Chen et al., 2020), thus providing a representation closer to the optimal. In this regard, increasing network size should further improve the robustness to label noise. We confirm this by comparing pretrained ResNet34 (with 46x params) and ResNet32 in Table 2. Both models are finetuned with Mixup. We see that increasing the network size can greatly improve the performance.

Table 2. Comparison between ResNet34 and ResNet32. We fine-tune the contrastively pretrained network with Mixup on CIFAR10/100 with different fractions of noise.

Dataset	CIFAR-10		CIFAR-100	
	Sym 80	Asym 40	Sym 80	Asym 40
ResNet-32	74.8 ± 3	87.5 ± 1.3	30.3 ± 4	55.5 ± 8
ResNet-34	90.8 ± 6	90.4 ± 4	69.4 ± 3	65.2 ± 2

Table 3. Test accuracy on mini WebVision. The best test accuracy is marked in bold. SimCLR (CL) pre-training leads to average improvements of 4.11% and 3.20% for mini Webvision and ImageNet respectively.

Method	WebVision		ImageNet	
	Top-1	Top-5	Top-1	Top-5
F-correction	61.12	82.68	57.36	82.36
Decoupling	62.54	84.74	58.26	82.26
Co-teaching	63.58	85.20	61.48	84.70
MentorNet	63.00	81.40	57.80	79.92
D2L	62.68	84.00	57.80	81.36
INCV	65.24	85.34	61.60	84.98
Crust	72.40	89.56	67.36	87.84
Mixup	71.38	87.36	68.34	88.44
ELR	76.26	91.26	68.71	87.84
CL + E2E	71.84	88.84	68.48	89.32
CL + Mixup	76.34	90.52	72.25	89.72
CL + ELR	79.52	93.80	71.20	90.80

5.2. Empirical Results on WebVision

WebVision is large scale image dataset with noisy labels (Li et al., 2017). It contains 2.4 million images crawled from Google Images search and Flickr that share the same 1000 classes as the ImageNet dataset. The noise ratio in classes varies from 0.5% to 88%, and the number of images per class varies from 300 to more than 10,000 (Fig. 4 in (Li et al., 2017) shows the noise distribution). We follow the setting in (Jiang et al., 2018) and create a mini WebVision dataset that consists of the top 50 classes in the Google subset with 66,000 images. We use both WebVision and ImageNet test sets for testing the performance of the model. We train InceptionResNet-v2 (Szegedy et al., 2017) for 120 epochs with a starting learning rate of 0.02, which we anneal by a factor of 0.01 at epochs 40 and 80. We use the SGD optimizer with a weight decay of 1×10^{-3} , and a minibatch size of 32. For Mixup and ELR, we use the same hyperparameters as CIFAR.

Table 3 shows the Top-1 and Top-5 accuracy of different methods evaluated on WebVision and ImageNet. We see that for both ELR and Mixup, SimCLR pretraining leads to average improvements of 4.11% and 3.20% for mini Webvision and ImageNet respectively. Furthermore, we note that SimCLR pretraining on its own outperforms every method without SimCLR pretraining, except ELR and Crust.

6. Conclusion

We showed that representations learned by contrastive learning provably boosts robustness against noisy labels. In particular, contrastive learning provides a representation matrix that has: (i) a significant gap between the prominent singular values and the remaining ones; (ii) a large alignment between the prominent singular vectors and the clean labels. The above properties allow a linear layer trained on the representations to effectively learn the clean labels well while barely overfitting the noise. Then we showed that fine-tuning all layers of a network pre-trained with contrastive learning achieves a superior performance initially, before overfitting the noise. Crucially, contrastive learning only slightly improves the alignment between the Jacobian matrix and the true labels, but greatly reduces the smaller singular values, which slows down the overfitting. Finally, we demonstrated that the initial robustness provided by contrastive learning can greatly boost robust methods and achieve state-of-the-art performance under extreme noise levels. Our results confirm benefits of contrastive pretraining for robust machine learning.

Acknowledgements

This research was supported in part by Cisco Systems and UCLA-Amazon Science Hub for Humanity and Artificial Intelligence.

References

- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Cao, K., Chen, Y., Lu, J., Arechiga, N., Gaidon, A., and Ma, T. Heteroskedastic and imbalanced deep learning with adaptive regularization. *arXiv preprint arXiv:2006.15766*, 2020.
- Chen, P., Liao, B. B., Chen, G., and Zhang, S. Understanding and utilizing deep neural networks trained with noisy labels. In *International Conference on Machine Learning*, pp. 1062–1070, 2019.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Cheng, H., Zhu, Z., Sun, X., and Liu, Y. Demystifying how self-supervised features improve training from noisy labels. *arXiv preprint arXiv:2110.09022*, 2021.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-

- Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- Eckart, C. and Young, G. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- Floridi, L. and Chiriatti, M. Gpt-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4):681–694, 2020.
- Ghosh, A. and Lan, A. Contrastive learning improves model robustness under label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2703–2708, 2021.
- Ghosh, A., Kumar, H., and Sastry, P. Robust loss functions under label noise for deep neural networks. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Goldberger, J. and Ben-Reuven, E. Training deep neural networks using a noise adaptation layer. 2016.
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I., and Sugiyama, M. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in neural information processing systems*, pp. 8527–8537, 2018.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *arXiv preprint arXiv:2106.04156*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. *arXiv preprint arXiv:1906.12340*, 2019.
- Hu, W., Xiao, L., Adlam, B., and Pennington, J. The surprising simplicity of the early-time learning dynamics of neural networks. *Advances in Neural Information Processing Systems*, 33:17116–17128, 2020.
- Jiang, L., Zhou, Z., Leung, T., Li, L.-J., and Fei-Fei, L. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pp. 2309–2318, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Krishna, R. A., Hata, K., Chen, S., Kravitz, J., Shamma, D. A., Fei-Fei, L., and Bernstein, M. S. Embracing error to enable rapid crowdsourcing. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pp. 3167–3179, 2016.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Li, J., Socher, R., and Hoi, S. C. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*, 2020.
- Li, W., Wang, L., Li, W., Agustsson, E., and Van Gool, L. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017.
- Liu, S., Niles-Weed, J., Razavian, N., and Fernandez-Granda, C. Early-learning regularization prevents memorization of noisy labels. *arXiv preprint arXiv:2007.00151*, 2020.
- Ma, X., Wang, Y., Houle, M. E., Zhou, S., Erfani, S., Xia, S., Wijewickrema, S., and Bailey, J. Dimensionality-driven learning with noisy labels. In *International Conference on Machine Learning*, pp. 3355–3364, 2018.
- Malach, E. and Shalev-Shwartz, S. Decoupling” when to update” from” how to update”. In *Advances in Neural Information Processing Systems*, pp. 960–970, 2017.
- Minc, H. On the maximal eigenvector of a positive matrix. *SIAM Journal on Numerical Analysis*, 7(3):424–427, 1970.
- Mirzasoleiman, B., Cao, K., and Leskovec, J. Coresets for robust training of deep neural networks against noisy labels. *Advances in Neural Information Processing Systems*, 33, 2020.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Oymak, S., Fabian, Z., Li, M., and Soltanolkotabi, M. Generalization guarantees for neural networks via harnessing the low-rank structure of the jacobian. *arXiv preprint arXiv:1906.05392*, 2019.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. Automatic differentiation in pytorch. 2017.
- Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1944–1952, 2017.

- Reed, S., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., and Rabinovich, A. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*, 2014.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pp. 4334–4343, 2018.
- SimCLR. <https://github.com/spijkervet/simclr>.
- Sohn, K. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in neural information processing systems*, pp. 1857–1865, 2016.
- Stewart, G. W. Matrix perturbation theory. 1990.
- Szegedy, C., Ioffe, S., Vanhoucke, V., and Alemi, A. A. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- Tanaka, D., Ikami, D., Yamasaki, T., and Aizawa, K. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5552–5560, 2018.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 22(281): 1–31, 2021.
- Van Rooyen, B., Menon, A., and Williamson, R. C. Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems*, pp. 10–18, 2015.
- Wang, X., Hua, Y., Kodirov, E., and Robertson, N. M. Imae for noise-robust learning: Mean absolute error does not treat examples equally and gradient magnitude’s variance matters. *arXiv preprint arXiv:1903.12141*, 2019.
- Wedin, P.-Å. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3733–3742, 2018.
- Xia, X., Liu, T., Wang, N., Han, B., Gong, C., Niu, G., and Sugiyama, M. Are anchor points really indispensable in label-noise learning? In *Advances in Neural Information Processing Systems*, pp. 6838–6849, 2019.
- Xu, Y., Cao, P., Kong, Y., and Wang, Y. L_{dmi}: A novel information-theoretic loss function for training deep nets robust to label noise. In *Advances in Neural Information Processing Systems*, pp. 6225–6236, 2019.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Zhang, H., Lee, H., Arik, S., Pfister, T., and Zhang, Z. Distilling effective supervision from severe label noise. 2020.
- Zhang, Z. and Sabuncu, M. Generalized cross entropy loss for training deep neural networks with noisy labels. In *Advances in neural information processing systems*, pp. 8778–8788, 2018.
- Zheltonozhskii, E., Baskin, C., Mendelson, A., Bronstein, A. M., and Litany, O. Contrast to divide: Self-supervised pre-training for learning with noisy labels. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1657–1667, 2022.

A. Analysis for Disconnected Subclasses

In this section we consider the case where $\xi = 0$ in assumption 4.2, which implies that the probability of two augmentation data from different subclasses being generated from the same random natural datum is exactly zero. And in section B we extend the results to any $\xi \in [0, 1)$ via eigenvalue and eigenvector perturbation. We use $\|\cdot\|_1$, $\|\cdot\|_2$ and $\|\cdot\|_F$ to denote the 1-norm, operator norm and Frobenius norm, respectively.

A.1. Spectral Decomposition of Augmentation Graph

An important technical idea we use to formalize the representations obtained by contrastive learning is augmentation graph (HaoChen et al., 2021), which is an undirected graph with all augmentation data $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ as its vertices and $w_{\mathbf{x}_i \mathbf{x}_j}$ as the weight for edge $(\mathbf{x}_i, \mathbf{x}_j)$. Let \mathbf{A} denote the adjacency matrix of the augmentation graph, that is, each entry $a_{ij} = w_{\mathbf{x}_i \mathbf{x}_j}$. And the normalized adjacency matrix $\bar{\mathbf{A}}$ is defined as

$$\bar{\mathbf{A}} := \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2},$$

where $\mathbf{D} = \text{diag}(w_{\mathbf{x}_1}, w_{\mathbf{x}_2}, \dots, w_{\mathbf{x}_n})$ with $w_{\mathbf{x}_i} = \sum_{j=1}^n w_{\mathbf{x}_i \mathbf{x}_j}$. For simplicity we index the augmentation data in the following way: the first n_1 data are from subclass 1, the next n_2 data are from subclass 2, \dots , the last $n_{\bar{K}}$ data are from subclass \bar{K} . Lemma A.1 states an important property of $\bar{\mathbf{A}}$.

Lemma A.1. *Assumption 4.2 with $\xi = 0$ implies that $\bar{\mathbf{A}}$ is a block diagonal matrix $\text{diag}(\bar{\mathbf{A}}_1, \bar{\mathbf{A}}_2, \dots, \bar{\mathbf{A}}_{\bar{K}})$ where $\bar{\mathbf{A}}_{\bar{k}} \in \mathbb{R}^{n_{\bar{k}} \times n_{\bar{k}}}$. This combined with 4.1 gives us: for each block $\bar{\mathbf{A}}_{\bar{k}}$, the ratio between two entries in the same column is bounded as follows*

$$\forall \bar{k} \in [\bar{K}], \max_{j,s,t} \frac{\bar{a}_{\bar{k},s,j}}{\bar{a}_{\bar{k},t,j}} \leq 1 + \delta',$$

where $\bar{a}_{\bar{k},i,j}$ are the entries of $\bar{\mathbf{A}}_{\bar{k}}$ and $\delta' = (1 + \delta)^{3/2} - 1 = \frac{3}{2}\delta + \mathcal{O}(\delta^2)$.

Let $f_{\min} = \arg \min_f \mathfrak{C}(f)$ and $\mathbf{F}_{\min} = [f_{\min}(\mathbf{x}_1) f_{\min}(\mathbf{x}_2) \dots f_{\min}(\mathbf{x}_n)]^\top$, according the theorem in (HaoChen et al., 2021), \mathbf{F}_{\min} is also the minimizer of the following matrix factorization problem

$$\min_{\mathbf{F} \in \mathbb{R}^{n \times p}} \|\bar{\mathbf{A}} - \mathbf{F} \mathbf{F}^\top\|_F^2, \quad (7)$$

and therefore can be further decomposed as

$$\mathbf{F}_{\min} \bar{\mathbf{A}} = \mathbf{F}^* \Sigma \mathbf{R}, \quad (8)$$

by Eckart–Young–Mirsky theorem (Eckart & Young, 1936), where $\mathbf{F}^* \in \mathbb{R}^{n \times p} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p] \in \mathbb{R}^{n \times p}$, $\Sigma = \text{diag}(\sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_p})$, $\mathbf{R} \in \mathbb{R}^{p \times p}$ is some orthogonal matrix, $\lambda_1, \lambda_2, \dots, \lambda_p$ are the p largest eigenvalues of $\bar{\mathbf{A}}$ and $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p$ are the corresponding unit-norm eigenvectors. W.l.o.g., we assume $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$. Our following proofs are based on this decomposition. To avoid cluttered notation we drop the subscript of \mathbf{F}_{\min} , i.e., we use \mathbf{F} for the learned representation.

A.2. Properties of Singular Values/Vectors of the Representation Matrix

From the above section we know that the singular values/vectors of \mathbf{F} are the first p eigenvalues/vectors of $\bar{\mathbf{A}}$. For each block $\bar{\mathbf{A}}_{\bar{k}}$, let $\lambda_{\bar{k},1}, \lambda_{\bar{k},2}, \dots, \lambda_{\bar{k},n_{\bar{k}}}$ denote the eigenvalues (in descending order) and $\mathbf{v}_{\bar{k},1}, \mathbf{v}_{\bar{k},2}, \dots, \mathbf{v}_{\bar{k},n_{\bar{k}}}$ denote the corresponding eigenvectors. The eigenvalues of $\bar{\mathbf{A}}$ are the list of the eigenvalues of all blocks. The corresponding eigenvectors are the block vectors $(\vec{0}_1, \vec{0}_2, \dots, \vec{0}_{\bar{k}-1}, \mathbf{v}_{\bar{k},i}, \vec{0}_{\bar{k}+1}, \dots, \vec{0}_{\bar{K}}) := \hat{\mathbf{v}}_{\bar{k},i}$ with each $\vec{0}_j$ being a zero vector of length n_j . Since $\bar{\mathbf{A}}$ is a normalized adjacency matrix, each block $\bar{\mathbf{A}}_{\bar{k}}$ is also normalized. Then the largest eigenvalue for each block is exactly 1, i.e., $\lambda_{\bar{k},1} = 1$. It immediately follow Lemma A.2.

Lemma A.2. *The \bar{K} largest eigenvalues of $\bar{\mathbf{A}}$ are all 1, i.e., $\lambda_1 = \lambda_2 = \dots = \lambda_{\bar{K}} = 1$.*

As long as $p \geq \bar{K}$, all $\lambda_{\bar{k},1}$ and $\mathbf{v}_{\bar{k},1}$ appear in the decomposition of \mathbf{F} . Let $p_{\bar{k}} \geq 1$ be the number of eigenvalues/eigenvectors of block $\bar{\mathbf{A}}_{\bar{k}}$ that also appear in the decomposition of \mathbf{F} . The following Lemmas and Corollaries states other important properties of eigenvalues/eigenvectors of $\bar{\mathbf{A}}$.

Lemma A.3. *By assumption 4.1, the 1-norm of $\mathbf{v}_{\bar{k},1}$ has the following lower bound.*

$$\|\mathbf{v}_{\bar{k},1}\|_1^2 \geq \frac{n_{\bar{k}}^2}{1 + (n_{\bar{k}} - 1)(1 + \delta')^2} = n_{\bar{k}} - 2(n_{\bar{k}} - 1)\delta' - \mathcal{O}(\delta'^2).$$

Proof. Write $\mathbf{v}_{\bar{k},1}$ as $[e_1, e_2, \dots, e_{n_{\bar{k}}}]^T$. By Perron-Frobenius theorem, all the elements here are positive since $\bar{\mathbf{A}}_{\bar{k}}$ is a positive matrix. Then the quotient of any two elements in $\mathbf{v}_{\bar{k},1}$ can be bounded in terms of the entries of $\bar{\mathbf{A}}_{\bar{k}}$ (Minc, 1970) and therefore $1 + \delta'$ by lemma A.1:

$$\max_{i,j} \frac{e_i}{e_j} \leq \max_{j,s,t} \frac{\bar{a}_{\bar{k},s,j}}{\bar{a}_{\bar{k},t,j}} \leq 1 + \delta'.$$

Let $e_{min} = \min_i e_i$, we have

$$1 = \|\mathbf{v}_{\bar{k},1}\|_2^2 = \sum_{i=1}^{n_{\bar{k}}} e_i^2 \leq e_{min}^2 + (n_{\bar{k}} - 1)(1 + \delta')^2 e_{min}^2,$$

and

$$\|\mathbf{v}_{\bar{k},1}\|_1^2 = \left(\sum_{i=1}^{n_{\bar{k}}} e_i\right)^2 \geq n_{\bar{k}}^2 e_{min}^2.$$

Combining the preceding two equations yields

$$\|\mathbf{v}_{\bar{k},1}\|_1^2 \geq \frac{n_{\bar{k}}^2}{1 + (n_{\bar{k}} - 1)(1 + \delta')^2}.$$

□

Lemma A.4. *The sum of squared eigenvalues of each block $\bar{\mathbf{A}}_{\bar{k}}$ can be bounded.*

$$\sum_{i=1}^{n_{\bar{k}}} \lambda_{\bar{k},i}^2 \leq \frac{(1 + (n_{\bar{k}} - 1)(1 + \delta')^2)(1 + \delta')^2}{n_{\bar{k}}}.$$

Proof. First we have

$$\sum_{i=1}^{n_{\bar{k}}} \lambda_{\bar{k},i}^2 = \|\bar{\mathbf{A}}_{\bar{k}}\|_F^2 = \sum_{j=1}^{n_{\bar{k}}} \|\mathbf{c}_{\bar{k},j}\|_2^2, \quad (9)$$

where $\mathbf{c}_{\bar{k},j} = [\bar{a}_{\bar{k},1,j}, \bar{a}_{\bar{k},2,j}, \dots, \bar{a}_{\bar{k},n_{\bar{k}},j}]^T$ denotes the j -th column in $\bar{\mathbf{A}}_{\bar{k}}$. Analogous to lemma A.3, here we can bound $\|\mathbf{c}_{\bar{k},j}\|_2^2$ in terms of $\|\mathbf{c}_{\bar{k},j}\|_1^2$ by lemma A.1

$$\|\mathbf{c}_{\bar{k},j}\|_2^2 \leq \frac{(1 + (n_{\bar{k}} - 1)(1 + \delta')^2)\|\mathbf{c}_{\bar{k},j}\|_1^2}{n_{\bar{k}}}. \quad (10)$$

We also have

$$\|\mathbf{c}_{\bar{k},j}\|_1^2 \leq \max_j \|\mathbf{c}_{\bar{k},j}\|_1^2 \leq (1 + \delta')^2 \min_j \|\mathbf{c}_{\bar{k},j}\|_1^2 \leq (1 + \delta')^2 \lambda_{\bar{k},1}^2 = (1 + \delta')^2, \quad (11)$$

where the second inequality holds because of assumption 4.1 and the third inequality holds because of Perron-Frobenius theorem. Combining equations 9, 10 and 11 completes the proof. □

Corollary A.5. *The eigenvalues except the \bar{K} largest ones are each upper bounded by*

$$\lambda_{\bar{k},i} \leq \sqrt{\frac{(1 + (n_{\bar{k}} - 1)(1 + \delta')^2)(1 + \delta')^2}{n_{\bar{k}}} - 1} = \mathcal{O}(\sqrt{\delta'}), \quad \forall i = 2, 3, \dots, n_{\bar{k}}, \quad \forall \bar{k} \in [\bar{K}].$$

Proof. By Lemmas A.2 and A.4

$$\begin{aligned}
 \sum_{i=2}^{n_{\bar{k}}} \lambda_{\bar{k},i}^2 &= \sum_{i=1}^{n_{\bar{k}}} \lambda_{\bar{k},i}^2 - \lambda_{\bar{k},1}^2 \\
 &\leq \frac{(1 + (n_{\bar{k}} - 1)(1 + \delta')^2)(1 + \delta')^2}{n_{\bar{k}}} - 1 \\
 &= 4\delta' + \mathcal{O}(\delta'^2).
 \end{aligned} \tag{12}$$

Then

$$\begin{aligned}
 \lambda_{\bar{k},i} &\leq \sqrt{\sum_{i=2}^{n_{\bar{k}}} \lambda_{\bar{k},i}^2} \\
 &\leq \sqrt{\frac{(1 + (n_{\bar{k}} - 1)(1 + \delta')^2)(1 + \delta')^2}{n_{\bar{k}}} - 1}
 \end{aligned}$$

□

Corollary A.6. For each block $\bar{\mathbf{A}}_{\bar{k}}$, the sum of the eigenvalues from the second to the $p_{\bar{k}}$ -th is bounded by

$$\sum_{i=2}^{p_{\bar{k}}} \lambda_{\bar{k},i} \leq 2\sqrt{(p_{\bar{k}} - 1)\delta'} + \mathcal{O}(\delta').$$

And the sum of eigenvalues of $\bar{\mathbf{A}}$ from the $\bar{K} + 1$ -th to the p -th is bounded by

$$\sum_{i=\bar{K}+1}^p \lambda_i \leq 2\sqrt{(p - \bar{K})\bar{K}\delta'} + \mathcal{O}(\delta').$$

Proof. Applying Cauchy–Schwarz inequality to equation 12 yields the bound for $\sum_{i=2}^{p_{\bar{k}}} \lambda_{\bar{k},i}$. Summing both sides of equation 12 over $\bar{k} \in [\bar{K}]$ and then applying Cauchy–Schwarz inequality give the bound for $\sum_{i=\bar{K}+1}^p \lambda_i$. □

A.3. Error under Gaussian Noise when $\xi = 0$

With the decomposition in Equation 8, the learned parameter of the linear model in Equation 3 can be rewritten as

$$\hat{\mathbf{W}}^* = \mathbf{R}^\top \mathbf{diag}\left(\frac{\sqrt{\lambda_1}}{\lambda_1 + \beta}, \frac{\sqrt{\lambda_2}}{\lambda_2 + \beta}, \dots, \frac{\sqrt{\lambda_p}}{\lambda_p + \beta}\right) \mathbf{F}^{*\top} \hat{\mathbf{Y}}.$$

The output on the training set $\mathbf{F}\hat{\mathbf{W}}^*$ is

$$\mathbf{F}\hat{\mathbf{W}}^* = \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} \hat{\mathbf{Y}},$$

where $\mathbf{B} = \mathbf{diag}(b_1, b_2, \dots, b_p)$ with $b_i = \frac{\lambda_i}{\lambda_i + \beta}$. And the error on training set w.r.t. ground-truth labels can be therefore written in terms of the eigenvalues and eigenvectors of $\bar{\mathbf{A}}$

$$\begin{aligned}
 \mathbb{E}_{\Delta \mathbf{Y}} \frac{1}{n} \|\mathbf{Y} - \mathbf{F}\hat{\mathbf{W}}^*\|_F^2 &= \mathbb{E}_{\Delta \mathbf{Y}} \left[\frac{1}{n} \|\mathbf{Y} - \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} (\mathbf{Y} + \Delta \mathbf{Y})\|_F^2 \right] \\
 &= \frac{1}{n} \|\mathbf{Y} - \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} \mathbf{Y}\|_F^2 + \mathbb{E}_{\Delta \mathbf{Y}} [\|\mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} \Delta \mathbf{Y}\|_F^2] \\
 &= \frac{1}{n} \|\mathbf{Y}\|_F^2 + \frac{1}{n} \sum_{i=1}^p \sum_{j=1}^K (b_i^2 - 2b_i) (\mathbf{v}_i^\top \mathbf{y}_j)^2 + \frac{\sigma^2}{n} \sum_{i=1}^p b_i^2 \\
 &= 1 + \underbrace{\frac{1}{n} \sum_{i=1}^p \sum_{j=1}^K (b_i^2 - 2b_i) (\mathbf{v}_i^\top \mathbf{y}_j)^2}_{\text{bias}^2} + \underbrace{\frac{\sigma^2}{n} \sum_{i=1}^p b_i^2}_{\text{variance}},
 \end{aligned} \tag{13}$$

where $\mathbf{y}_j \in \mathbb{R}^n$ is the j -th column of \mathbf{Y} .

We first calculate the \mathbf{bias}^2 term. We have $b_i^2 - 2b_i \leq 0$ since $b_i \in (0, 1]$. Then we drop items with $i \geq \bar{K} + 1$ in the summation and apply Lemma A.2 to get an upper bound

$$\begin{aligned} \mathbf{bias}^2 &\leq 1 - \frac{1}{n} \sum_{i=1}^{\bar{K}} \sum_{j=1}^K (2b_i - b_i^2) (\mathbf{v}_i^\top \mathbf{y}_j)^2 \\ &= 1 - \frac{1}{n} \left(\frac{2}{1+\beta} - \left(\frac{1}{1+\beta} \right)^2 \right) \sum_{i=1}^{\bar{K}} \sum_{j=1}^K (\mathbf{v}_i^\top \mathbf{y}_j)^2 \\ &= 1 - \frac{1}{n} \left(\frac{2}{1+\beta} - \left(\frac{1}{1+\beta} \right)^2 \right) \sum_{\bar{k}=1}^{\bar{K}} \sum_{j=1}^K (\mathbf{v}_{\bar{k},1}^\top \mathbf{y}_j)^2 \end{aligned}$$

By Perron-Frobenius theorem all elements in $\mathbf{v}_{\bar{k},1}$ are positive, thus the sum of elements of $\mathbf{v}_{\bar{k},1}$ can be written as $\|\mathbf{v}_{\bar{k},1}\|_1$. With the observation that $\mathbf{v}_{\bar{k},1}^\top \mathbf{y}_j = \|\mathbf{v}_{\bar{k},1}\|_1$ when \bar{k} is a subclass within class j and otherwise $\mathbf{v}_{\bar{k},1}^\top \mathbf{y}_j = 0$, the above equation can be rewritten as

$$\mathbf{bias}^2 \leq 1 - \frac{1}{n} \left(\frac{2}{1+\beta} - \left(\frac{1}{1+\beta} \right)^2 \right) \sum_{\bar{k}=1}^{\bar{K}} \|\mathbf{v}_{\bar{k},1}\|_1^2. \quad (14)$$

Then by Lemma A.3,

$$\begin{aligned} \mathbf{bias}^2 &\leq 1 - \frac{1}{n} \left(\frac{2}{1+\beta} - \left(\frac{1}{1+\beta} \right)^2 \right) \sum_{\bar{k}=1}^{\bar{K}} (n_{\bar{k}} - 2(n_{\bar{k}} - 1)\delta' - \mathcal{O}(\delta'^2)) \\ &= \left(\frac{\beta}{1+\beta} \right)^2 + 2 \left(1 - \frac{\bar{K}}{n} \right) \frac{(2\beta+1)}{(\beta+1)^2} \delta' + \mathcal{O}(\delta'^2) \\ &= \left(\frac{\beta}{1+\beta} \right)^2 + 3 \left(1 - \frac{\bar{K}}{n} \right) \frac{(2\beta+1)}{(\beta+1)^2} \delta + \mathcal{O}(\delta^2) \end{aligned} \quad (15)$$

Now we bound the **variance** term. By Lemma A.2

$$\begin{aligned} \mathbf{variance} &= \frac{\sigma^2}{n} \sum_{i=1}^p b_i^2 \\ &= \frac{\sigma^2}{n} \sum_{i=1}^{\bar{K}} b_i^2 + \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p b_i^2 \\ &= \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1} \right)^2 + \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p \left(1 - \frac{\beta}{\lambda_i + \beta} \right) \\ &= \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1} \right)^2 + \frac{\sigma^2}{n} (p - \bar{K}) - \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p \frac{\beta}{\lambda_i + \beta}. \end{aligned} \quad (16)$$

Apply Cauchy–Schwarz inequality and Corollary A.6 to bound the summation in the last term

$$\begin{aligned} \sum_{i=\bar{K}+1}^p \frac{\beta}{\lambda_i + \beta} &\geq \frac{\beta(p - \bar{K})^2}{\sum_{i=\bar{K}+1}^p (\lambda_i + \beta)} \\ &\geq \frac{\beta(p - \bar{K})^2}{2\sqrt{(p - \bar{K})} \sqrt{\delta'} + \mathcal{O}(\delta') + (p - \bar{K})\beta} \\ &= p - \bar{K} + \frac{2\sqrt{p - \bar{K}}}{\beta} \sqrt{\bar{K}\delta'} + \mathcal{O}(\delta'). \end{aligned}$$

Plugging the preceding into Equation 16 yields

$$\begin{aligned} \text{variance} &\leq \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1} \right)^2 + \sigma^2 \frac{2\sqrt{\bar{K}(p-\bar{K})}}{n} \frac{\sqrt{\delta'}}{\beta} + \mathcal{O}(\delta') \\ &= \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta+1} \right)^2 + \sigma^2 \frac{\sqrt{6\bar{K}(p-\bar{K})}}{n} \frac{\sqrt{\delta}}{\beta} + \mathcal{O}(\delta) \end{aligned}$$

A.4. Accuracy under Label Flipping (Proof for Theorem 4.5)

We study the accuracy by looking at the entries of the output $\mathbf{F}\hat{\mathbf{W}}^*$.

$$\begin{aligned} \mathbf{F}\hat{\mathbf{W}}^* &= \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} (\mathbf{Y} + \Delta \mathbf{Y}) \\ &= [\mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} (\mathbf{y}_1 + \Delta \mathbf{y}_1), \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} (\mathbf{y}_2 + \Delta \mathbf{y}_2), \dots, \mathbf{F}^* \mathbf{B} \mathbf{F}^{*\top} (\mathbf{y}_K + \Delta \mathbf{y}_K)] \\ &= \left[\sum_{i=1}^p b_i \mathbf{v}_i \mathbf{v}_i^\top (\mathbf{y}_1 + \Delta \mathbf{y}_1), \sum_{i=1}^p b_i \mathbf{v}_i \mathbf{v}_i^\top (\mathbf{y}_2 + \Delta \mathbf{y}_2), \dots, \sum_{i=1}^p b_i \mathbf{v}_i \mathbf{v}_i^\top (\mathbf{y}_K + \Delta \mathbf{y}_K) \right] \\ &:= [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_K] \end{aligned}$$

For convenience we define the notations C_k and $S_{\bar{k}}$ as the sets of indices of examples from class k and subclass \bar{k} , respectively

$$\begin{aligned} C_k &:= \{i : \mathbf{x}_i \text{ belongs to class } k\} \\ S_{\bar{k}} &:= \{i : \mathbf{x}_i \text{ belongs to subclass } \bar{k}\} = \left\{ i : \sum_{j=1}^{\bar{k}-1} n_j < i \leq \sum_{j=1}^{\bar{k}} n_j \right\}. \end{aligned}$$

Let the notation $\boldsymbol{\mu}^{(j)}$ denote the j -th element of vector $\boldsymbol{\mu}$. Then $\mathbf{z}_k^{(j)}$ can be written as $\sum_{i=1}^p b_i \mathbf{v}_i^{(j)} \mathbf{v}_i^\top (\mathbf{y}_k + \Delta \mathbf{y}_k)$. Let \bar{k}_j denote the subclass that \mathbf{x}_j belongs to, i.e., $j \in S_{\bar{k}}$ and define $e_{\min, s} := \min_j \mathbf{v}_{\bar{k}_j, 1}^{(s)}$ and $e_{\max, s} := \max_s \mathbf{v}_{\bar{k}_j, 1}^{(s)}$. We have the following two lemmas.

Lemma A.7. $\mathbf{z}_k^{(j)}$ can be bounded

$$\begin{cases} \mathbf{z}_k^{(j)} \geq \frac{1}{\beta+1} e_{\min, j} \dot{\mathbf{v}}_{\bar{k}_j, 1}^{(j)} n_{\min} (1 - \alpha) - \sqrt{n_{\max}} \frac{2\sqrt{p-1}}{\beta} \sqrt{\delta'} - \mathcal{O}(\delta'), & j \in C_k \\ \mathbf{z}_k^{(j)} \leq \frac{1}{\beta+1} e_{\max, j} \dot{\mathbf{v}}_{\bar{k}_j, 1}^{(j)} n_{\max} \alpha c_{\max} + \sqrt{n_{\max}} c_{\max} \frac{2\sqrt{p-1}}{\beta} \sqrt{\delta'} + \mathcal{O}(\delta'), & j \notin C_k. \end{cases}$$

Proof. Let $b_{\bar{k}, i}$ denote $\frac{\lambda_{\bar{k}, i}}{\beta + \lambda_{\bar{k}, i}}$. Recalling that one property of the block vector $\dot{\mathbf{v}}_{\bar{k}, i}^{(j)}$ is that $\dot{\mathbf{v}}_{\bar{k}, i}^{(j)} = 0$ when $j \notin S_{\bar{k}}$, we have

$$\begin{aligned} \mathbf{z}_k^{(j)} &= \sum_{i=1}^p b_i \mathbf{v}_i^{(j)} \mathbf{v}_i^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) \\ &= \sum_{\bar{k}=1}^{\bar{K}} \sum_{i=1}^{p_{\bar{k}}} b_{\bar{k}, i} \dot{\mathbf{v}}_{\bar{k}, i}^{(j)} \dot{\mathbf{v}}_{\bar{k}, i}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) \\ &= \sum_i^{p_{\bar{k}_j}} b_{\bar{k}_j, i} \dot{\mathbf{v}}_{\bar{k}_j, i}^{(j)} \dot{\mathbf{v}}_{\bar{k}_j, i}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) \\ &= b_{\bar{k}_j, 1} \dot{\mathbf{v}}_{\bar{k}_j, 1}^{(j)} \dot{\mathbf{v}}_{\bar{k}_j, 1}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) + \sum_{i=2}^{p_{\bar{k}_j}} b_{\bar{k}_j, i} \dot{\mathbf{v}}_{\bar{k}_j, i}^{(j)} \dot{\mathbf{v}}_{\bar{k}_j, i}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) \\ &= \frac{1}{\beta+1} \dot{\mathbf{v}}_{\bar{k}_j, 1}^{(j)} \dot{\mathbf{v}}_{\bar{k}_j, 1}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) + \sum_{i=2}^{p_{\bar{k}_j}} b_{\bar{k}_j, i} \dot{\mathbf{v}}_{\bar{k}_j, i}^{(j)} \dot{\mathbf{v}}_{\bar{k}_j, i}^\top (\mathbf{y}_k + \Delta \mathbf{y}_k) \end{aligned} \tag{17}$$

For the nonzero elements in $\dot{\mathbf{v}}_{\bar{k}_j,1}$, if $j \in C_k$, there are at least $n_{\min}(1 - \alpha)$ elements being 1 at corresponding positions in $\mathbf{y}_k + \Delta\mathbf{y}_k$; if $j \notin C_k$, there are at most $n_{\max}\alpha c_{\max}$ elements being 1 at corresponding positions in $\mathbf{y}_k + \Delta\mathbf{y}_k$. Then the inner product in the first term in equation 17 can be bounded by

$$\dot{\mathbf{v}}_{\bar{k}_j,1}^\top (\mathbf{y}_k + \Delta\mathbf{y}_k) = \sum_{s=1}^{n_{\bar{k}_j}} \dot{\mathbf{v}}_{\bar{k}_j,1}^{(s)} (\mathbf{y}_k + \Delta\mathbf{y}_k)^{(s)} \begin{cases} \geq e_{\min,j} n_{\min}(1 - \alpha), & j \in C_k \\ \leq e_{\max,j} \dot{\mathbf{v}}_{\bar{k}_j,1}^{(j)} n_{\max}\alpha c_{\max}, & j \notin C_k \end{cases}$$

For the inner product in the second term in equation 17, if $j \in C_k$, then $\dot{\mathbf{v}}_{\bar{k}_j,i}^\top (\mathbf{y}_k + \Delta\mathbf{y}_k)$ is the sum of at most n_{\max} elements in $\dot{\mathbf{v}}_{\bar{k}_j,i}$. Since $\|\dot{\mathbf{v}}_{\bar{k}_j,i}\|_2^2 = 1$, the sum is bounded by $[-\sqrt{n_{\max}}, \sqrt{n_{\max}}]$. Similarly we can get the bound $[-\sqrt{n_{\max}c_{\max}}, \sqrt{n_{\max}c_{\max}}]$ for $j \notin C_k$. We also know that $|\dot{\mathbf{v}}_{\bar{k}_j,i}^{(j)}| \leq 1$. Now it remains to bound $\sum_{i=2}^{p_{\bar{k}_j}} b_{\bar{k}_j,i}$ by applying Cauchy–Schwarz inequality and Corollary A.6

$$\begin{aligned} \sum_{i=2}^{p_{\bar{k}_j}} b_{\bar{k}_j,i} &= \sum_{i=2}^{p_{\bar{k}_j}} \frac{\lambda_{\bar{k},i}}{\beta + \lambda_{\bar{k},i}} \\ &= p_{\bar{k}_j} - 1 - \sum_{i=2}^{p_{\bar{k}_j}} \frac{\beta}{\beta + \lambda_{\bar{k},i}} \\ &\leq p_{\bar{k}_j} - 1 - \frac{\beta(p_{\bar{k}_j} - 1)^2}{\sum_{i=2}^{p_{\bar{k}_j}} (\lambda_{\bar{k},i} + \beta)} \\ &\leq p_{\bar{k}_j} - 1 - \frac{\beta(p_{\bar{k}_j} - 1)^2}{2\sqrt{(p_{\bar{k}} - 1)\sqrt{\delta'}} + \mathcal{O}(\delta') + \beta(p_{\bar{k}_j} - 1)} \\ &\leq 2\sqrt{p-1} \frac{\sqrt{\delta'}}{\beta} + \mathcal{O}(\delta') \end{aligned}$$

□

Lemma A.8. $e_{\min,j}$, $e_{\max,j}$ and $e_{\max,j}/e_{\min,j}$ can be bounded.

$$\begin{aligned} e_{\min,j} &\geq \sqrt{\frac{1}{1 + (n_{\max} - 1)(1 + \delta')^2}} \\ e_{\max,j} &\leq \sqrt{\frac{1}{1 + (n_{\min} - 1)/(1 + \delta')^2}} \\ \frac{e_{\max,j}}{e_{\min,j}} &\leq (1 + \delta') \sqrt{\frac{(1 + (1 + \delta')^2(n_{\max} - 1))}{n_{\max} - 1 + (1 + \delta')^2}} \end{aligned}$$

Proof. The proof is analogous to that for lemma A.3. □

Now we are ready to calculate the maximum noise level that allows correct prediction on all training examples, i.e., $\mathbf{z}_{k:j \in C_k}^j > \mathbf{z}_{k':j \notin C_{k'}}^j$. Let the preceding hold, then by lemma A.7 we get

$$\alpha < \frac{1 - \frac{2\sqrt{n_{\max}}(1 + \sqrt{c_{\max}})\sqrt{p-1}}{\beta + 1} \sqrt{\delta'}}{1 + \frac{n_{\max}e_{\max,j}}{n_{\min}e_{\min,j}} c_{\max}} - \mathcal{O}(\delta') \quad (18)$$

Plugging lemma A.8 into equation 18 with some algebraic manipulation yields

$$\alpha < \frac{1}{1 + \frac{n_{\max}}{n_{\min}} c_{\max}} - \mathcal{O}(\sqrt{\delta'}) = \alpha < \frac{1}{1 + \frac{n_{\max}}{n_{\min}} c_{\max}} - \mathcal{O}(\sqrt{\delta}). \quad (19)$$

B. Considering Off-Diagonal Entries in the Adjacency Matrix (Connected Subclasses)

For here on we assume $n_{\bar{k}} = \frac{n}{\bar{K}}, \forall \bar{k} \in \bar{K}$ for simplicity, despite that our results can easily extend to unbalanced dataset.

Lemma B.1. *Under assumption 4.2, the off-diagonal entries in A is no longer zero. Let $\tilde{\mathbf{A}}$ denote the new normalized matrix, which also contains non-zero off-diagonal entries. With a bit abuse of notation, in the following we use $\bar{\mathbf{A}}$ to denote the matrix obtained by normalizing A with off-diagonal elements ignored. Then all the properties of eigenvectors and eigenvalues of $\bar{\mathbf{A}}$ stated before including those lemma A.1, A.3, A.4 still hold. And $\tilde{\mathbf{A}}$ can be written as a perturbation of $\bar{\mathbf{A}}$*

$$\tilde{\mathbf{A}} = \bar{\mathbf{A}} + \mathbf{E},$$

with $\|\mathbf{E}\|_F = \mathcal{O}(\bar{K}^{5/2}\xi)$.

Proof. Let $\mathring{\mathbf{A}}$ be a matrix in the same shape of $\bar{\mathbf{A}}$ containing all elements of $\tilde{\mathbf{A}}$ in the diagonal blocks. Let \mathbf{H} be a matrix that collects the remaining off-diagonal elements. Therefore $\tilde{\mathbf{A}} = \mathring{\mathbf{A}} + \mathbf{H}$, which can be rewritten as

$$\tilde{\mathbf{A}} = \bar{\mathbf{A}} + \mathring{\mathbf{A}} - \bar{\mathbf{A}} + \mathbf{H}. \quad (20)$$

For all off-diagonal elements $h_{i,j}$ in \mathbf{H}

$$h_{i,j} \leq \frac{\bar{K}\xi}{n}.$$

Since there are $n^2 - \sum_{\bar{k}=1}^{\bar{K}} n_{\bar{k}}^2$ entries outside of the diagonal blocks, the norm of \mathbf{H} can be bounded by

$$\|\mathbf{H}\|_F \leq \sqrt{1 - \frac{1}{\bar{K}}} \bar{K}\xi \quad (21)$$

Each element in the diagonal blocks of $\bar{\mathbf{A}} - \mathring{\mathbf{A}}$ is non-negative. Also, supposing \mathbf{x}_i and \mathbf{x}_j are from subclass \bar{k} , we have

$$\begin{aligned} (\bar{\mathbf{A}} - \mathring{\mathbf{A}})_{i,j} &= \frac{w_{\mathbf{x}_i \mathbf{x}_j}}{\sqrt{\sum_{s: \mathbf{x}_s \in C_{\bar{k}}} w_{\mathbf{x}_i \mathbf{x}_s}} \sqrt{\sum_{t: \mathbf{x}_t \in C_{\bar{k}}} w_{\mathbf{x}_t \mathbf{x}_j}}} - \frac{w_{\mathbf{x}_i \mathbf{x}_j}}{\sqrt{\sum_{s=1}^n w_{\mathbf{x}_i \mathbf{x}_s}} \sqrt{\sum_{t=1}^n w_{\mathbf{x}_t \mathbf{x}_j}}} \\ &\leq \frac{n(\bar{K}-1)\xi}{\frac{n}{\bar{K}(1+\delta)} \left(\frac{n}{\bar{K}(1+\delta)} + n(\bar{K}-1)\xi \right)} \\ &= \mathcal{O}\left(\frac{\bar{K}^3(1+\delta)^2\xi}{n}\right), \end{aligned}$$

by which the norm of $\bar{\mathbf{A}} - \mathring{\mathbf{A}}$ is bounded

$$\|\bar{\mathbf{A}} - \mathring{\mathbf{A}}\|_F = \mathcal{O}\left(\bar{K}^{5/2}(1+\delta)^2\xi\right) = \mathcal{O}\left(\bar{K}^{5/2}\xi\right). \quad (22)$$

Combining equations 20, 21 and 22 completes the proof. \square

$\bar{\mathbf{A}}$ has the following eigendecomposition

$$\bar{\mathbf{A}} = [\mathbf{V}_I \mathbf{V}_N] \begin{bmatrix} \mathbf{\Sigma}_I & 0 \\ 0 & \mathbf{\Sigma}_N \end{bmatrix} \begin{bmatrix} \mathbf{V}_I^\top \\ \mathbf{V}_N^\top \end{bmatrix},$$

where $\mathbf{\Sigma}_I$ collects the \bar{K} largest eigenvalues $\lambda_1, \dots, \lambda_{\bar{K}}$ on the diagonal and $\mathbf{\Sigma}_N$ collects the remaining $\lambda_{\bar{K}+1}, \dots, \lambda_{\bar{n}}$. \mathbf{V}_I and \mathbf{V}_N collects the corresponding \bar{K} and $n - \bar{K}$ eigenvectors, respectively. Let $\tilde{\mathbf{A}}$ has analogous decomposition

$$\tilde{\mathbf{A}} = [\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_N] \begin{bmatrix} \tilde{\mathbf{\Sigma}}_I & 0 \\ 0 & \tilde{\mathbf{\Sigma}}_N \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{V}}_I^\top \\ \tilde{\mathbf{V}}_N^\top \end{bmatrix},$$

with eigenvalues $\tilde{\lambda}_1, \dots, \tilde{\lambda}_n$ and eigenvectors $\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_n$. Eigenvalues of both matrices are indexed in descending order.

B.1. Perturbation in Eigenvalues and Eigenvectors

The following two lemmas bound the changes in eigenvalues, eigenvectors and the alignment between labels and eigenvectors caused by ξ .

Lemma B.2. *We have the following bound for eigenvalues of $\tilde{\mathbf{A}}$:*

$$\begin{cases} 1 - \mathcal{O}(K^{5/2}\xi) \leq \tilde{\lambda}_i \leq 1, & i = 1, 2, \dots, \bar{K} \\ \sum_{i=\bar{K}+1}^p \lambda_i \leq \mathcal{O}(\sqrt{\delta} + \bar{K}^{5/2}\xi) \end{cases}$$

Proof. From Lemma A.2 and Corollary A.6 we know that

$$\begin{cases} \lambda_i = 1, & i = 1, 2, \dots, \bar{K}, \\ \sum_{i=\bar{K}+1}^p \lambda_i \leq \mathcal{O}(\sqrt{\delta}). \end{cases} \quad (23)$$

By Weyl's inequality on perturbation, we have

$$|\tilde{\lambda}_i - \lambda_i| \leq \|\mathbf{E}\|_2.$$

The right-hand-side is $\leq \|\mathbf{E}\|_F$ and therefore $\mathcal{O}(\bar{K}^{5/2}\xi)$ by lemma B.1. Combining the preceding with equation 23 completes the proof. \square

Lemma B.3. *The norm of the projection of \mathbf{Y} onto the range of $\tilde{\mathbf{V}}_I$ is bounded from below, i.e.,*

$$\|\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T \mathbf{Y}\|_F^2 \geq \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 - \mathcal{O}(n\bar{K}^2\xi).$$

Proof. By Lemma A.2 and Corollary A.5 we have

$$\lambda_{\bar{K}} - \lambda_{\bar{K}+1} \geq 1 - \mathcal{O}(\sqrt{\delta}).$$

By Wedin's Theorem (Wedin, 1972; Stewart, 1990), we have the following bound on the principle angle between the range of $\tilde{\mathbf{V}}_I$ and the range of \mathbf{V}_I as long as $1 \geq \mathcal{O}(\sqrt{\delta} + \bar{K}^{5/2}\xi)$

$$\|\mathbf{V}_I \mathbf{V}_I^T (\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T - \mathbf{I})\|_F \leq \mathcal{O}\left(\frac{\|\mathbf{E}\|_F}{\lambda_{\bar{K}} - \lambda_{\bar{K}+1}}\right) \leq \mathcal{O}\left(\frac{\mathcal{O}(\bar{K}^{5/2}\xi)}{1 - \mathcal{O}(\sqrt{\delta})}\right) = \mathcal{O}(\bar{K}^{5/2}\xi).$$

Thus

$$\begin{aligned} \|\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T \mathbf{Y}\|_F^2 &= \|\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T \mathbf{Y} - \mathbf{V}_I \mathbf{V}_I^T \mathbf{Y} + \mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 \\ &\geq \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 + 2\langle \tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T \mathbf{Y} - \mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}, \mathbf{V}_I \mathbf{V}_I^T \mathbf{Y} \rangle_F \\ &\geq \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 + 2\text{Tr}((\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T - \mathbf{V}_I \mathbf{V}_I^T) \mathbf{Y} \mathbf{Y}^T \mathbf{V}_I \mathbf{V}_I^T) \\ &= \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 + 2\text{Tr}(\mathbf{V}_I \mathbf{V}_I^T (\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T - \mathbf{I}) \mathbf{Y} \mathbf{Y}^T) \\ &= \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 + 2\langle \mathbf{V}_I \mathbf{V}_I^T (\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T - \mathbf{I}), \mathbf{Y} \mathbf{Y}^T \rangle_F \\ &\geq \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 - 2\|\mathbf{V}_I \mathbf{V}_I^T (\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^T - \mathbf{I})\|_F \|\mathbf{Y} \mathbf{Y}^T\|_F \\ &\geq \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 - \mathcal{O}(\bar{K}^{5/2}\xi) \frac{n}{\sqrt{K}} \\ &= \|\mathbf{V}_I \mathbf{V}_I^T \mathbf{Y}\|_F^2 - \mathcal{O}(n\bar{K}^2\xi). \end{aligned}$$

\square

B.2. Error under Gaussian Noise (Proof for Theorem 4.4)

Considering ξ , rewrite \mathbf{bias}^2 as

$$\begin{aligned}
 \mathbf{bias}^2 &= 1 + \frac{1}{n} \sum_{i=1}^p \sum_{j=1}^K (\tilde{b}_i^2 - 2\tilde{b}_i) (\tilde{\mathbf{v}}_i^\top \mathbf{y}_j)^2 \\
 &\leq 1 - \frac{1}{n} \sum_{i=1}^{\bar{K}} \sum_{j=1}^K (2\tilde{b}_i - \tilde{b}_i^2) (\tilde{\mathbf{v}}_i^\top \mathbf{y}_j)^2 \\
 &\leq 1 - \frac{1}{n} (2\tilde{b}_{\bar{K}} - \tilde{b}_{\bar{K}}^2) \sum_{i=1}^{\bar{K}} \sum_{j=1}^K (\tilde{\mathbf{v}}_i^\top \mathbf{y}_j)^2 \\
 &\leq 1 - \frac{1}{n} (2\tilde{b}_{\bar{K}} - \tilde{b}_{\bar{K}}^2) \|\tilde{\mathbf{V}}_I \tilde{\mathbf{V}}_I^\top \mathbf{Y}\|_F^2,
 \end{aligned} \tag{24}$$

where $\tilde{b}_i = \frac{\tilde{\lambda}_i}{\tilde{\lambda}_i + \beta}$. Also, Lemma A.3 gives us the lower bound for $\|\mathbf{V}_I \mathbf{V}_I^\top \mathbf{Y}\|_F^2$

$$\begin{aligned}
 \|\mathbf{V}_I \mathbf{V}_I^\top \mathbf{Y}\|_F^2 &= \sum_{\bar{k}}^{\bar{K}} \|\mathbf{v}_{\bar{k},1}\|_1^2 \\
 &\geq n - \mathcal{O}(\delta).
 \end{aligned} \tag{25}$$

Combining lemma B.2, lemma B.3, equation 24 and equation 25 yields the bound for the bias

$$\mathbf{bias}^2 = \left(\frac{\beta}{1 + \beta}\right)^2 + \mathcal{O}(\delta + \xi).$$

We bound the variance in the same manner as in Section A.3 by applying Cauchy–Schwarz inequality, Corollary A.6 and Lemma B.2

$$\begin{aligned}
 \mathbf{variance} &= \frac{\sigma^2}{n} \sum_{i=1}^p \tilde{b}_i^2 \\
 &= \frac{\sigma^2}{n} \sum_{i=1}^{\bar{K}} \tilde{b}_i^2 + \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p \tilde{b}_i^2 \\
 &\leq \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta + 1}\right)^2 + \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p \left(1 - \frac{\beta}{\tilde{\lambda}_i + \beta}\right) \\
 &= \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta + 1}\right)^2 + \frac{\sigma^2}{n} (p - \bar{K}) - \frac{\sigma^2}{n} \sum_{i=\bar{K}+1}^p \frac{\beta}{\tilde{\lambda}_i + \beta} \\
 &\leq \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta + 1}\right)^2 + \frac{\sigma^2}{n} (p - \bar{K}) - \frac{\sigma^2}{n} \frac{\beta(p - \bar{K})^2}{\sum_{i=\bar{K}+1}^p (\tilde{\lambda}_i + \beta)} \\
 &\leq \sigma^2 \frac{\bar{K}}{n} \left(\frac{1}{\beta + 1}\right)^2 + \sigma^2 \mathcal{O}\left(\frac{\sqrt{\delta} + \xi}{\beta}\right).
 \end{aligned}$$

C. Contrastive Learning Slightly Improves the Alignment Between Jacobian Matrix and Ground-truth Labels

We compare the alignments between the clean label vector and the initial Jacobian matrix of (1) network pretrained using SimCLR for 1000 epochs, (2) network pretrained using SimCLR for 100 epochs and (3) randomly initialized network. $\mathbf{y} \in \mathbb{R}^{nK}$ is the vector obtained by flattening the label matrix \mathbf{Y} , i.e., concatenating the n rows of \mathbf{Y} . Let $\mathbf{z}(\mathbf{x}_i, \mathbf{W}) \in \mathbb{R}^K$




Figure 3. Test accuracy of fine-tuning a pre-trained network with frozen encoder v.s. unfrozen encoder on CIFAR-10 (top) and CIFAR-100 (bottom) under 20%, 50%, 80% symmetric noise and 40% asymmetric noise (left to right).

	$\ \Pi_{\mathcal{I}}(\mathbf{y})\ _F$	$\ \Pi_{\mathcal{N}}(\mathbf{y})\ _F$	$\ \mathbf{J}\mathbf{J}^T\mathbf{y}\ _F/\ \mathbf{J}\mathbf{J}^T\ _F$
Pretrained for 1000 epochs	10.063	29.979	3.184
Pretrained for 100 epochs	10.036	29.988	3.175
Randomly initialized	10.014	29.995	3.055

Table 4. Alignment between the Jacobian matrix and the clean labels.

be the output of the network given example \mathbf{x}_i and parameters $\mathbf{W} \in \mathbb{R}^d$ (we see the parameters of the network as a vector). Then the Jacobian \mathbf{J} is defined as

$$\mathbf{J}(\mathbf{W}) = \left[\frac{\partial \mathbf{z}(\mathbf{x}_1, \mathbf{W})}{\mathbf{W}} \dots \frac{\partial \mathbf{z}(\mathbf{x}_n, \mathbf{W})}{\mathbf{W}} \right]^\top.$$

Note that $\frac{\partial \mathbf{z}(\mathbf{x}_i, \mathbf{W})}{\mathbf{W}} \in \mathbb{R}^{d \times K}$, therefore $\mathbf{J}(\mathbf{W}) \in \mathbb{R}^{nK \times d}$. In table 4 $\Pi_{\mathcal{I}}(\mathbf{y})$ is the projection of \mathbf{y} onto the span of the 10 singular vectors of $\mathbf{J}(\mathbf{W}_0)$ with largest singular values and $\Pi_{\mathcal{N}}(\mathbf{y})$ is the projection of \mathbf{y} onto the span of the remaining singular vectors. Interestingly, pretraining for more epochs leads to larger $\Pi_{\mathcal{I}}(\mathbf{y})$ and smaller $\Pi_{\mathcal{N}}(\mathbf{y})$ and therefore larger $\|\mathbf{J}\mathbf{J}^T\mathbf{y}\|_F/\|\mathbf{J}\mathbf{J}^T\|_F$. How much this slight improvement in the alignment contributes to the robustness deserves further investigation.

D. Training Only the Linear Layer v.s. Training All Layers

Figure 3 compares the performance of training only the linear layer (i.e., with the encoder frozen) and fine-tuning all layers (i.e., with the encoder unfrozen). For both CIFAR-10 and CIFAR-100 we first pretrain a Res-Net 32 using SimCLR for 1000 epochs and using the Adam optimizer with a learning rate of 3×10^{-4} , a weight decay of 1×10^{-6} and a batch size of 128. For downstream tasks, we use the SGD optimizer with a learning rate of 5×10^{-3} , a weight decay of 1×10^{-3} , a batch size of 64. We see that in most cases fine-tuning achieves a higher test accuracy. However, finetuning will eventually overfit if trained for longer. Also, we note that training all layers is more likely to overfit, especially under large noise level (column 3 in figure 3).