

# Unraveling Projection Heads in Contrastive Learning: Insights from Expansion and Shrinkage

Yu Gui<sup>1</sup>, Cong Ma<sup>1</sup>, and Yiqiao Zhong<sup>2</sup>

<sup>1</sup>Department of Statistics, University of Chicago

<sup>2</sup>Department of Statistics, University of Wisconsin-Madison

June 5, 2023

## Abstract

We investigate the role of projection heads, also known as projectors, within the encoder-projector framework (e.g., SimCLR) used in contrastive learning. We aim to demystify the observed phenomenon where representations learned before projectors outperform those learned after—measured using the downstream linear classification accuracy, even when the projectors themselves are linear.

In this paper, we make two significant contributions towards this aim. Firstly, through empirical and theoretical analysis, we identify two crucial effects—expansion and shrinkage—induced by the contrastive loss on the projectors. In essence, contrastive loss either expands or shrinks the signal direction in the representations learned by an encoder, depending on factors such as the augmentation strength, the temperature used in contrastive loss, etc. Secondly, drawing inspiration from the expansion and shrinkage phenomenon, we propose a family of linear transformations to accurately model the projector’s behavior. This enables us to precisely characterize the downstream linear classification accuracy in the high-dimensional asymptotic limit. Our findings reveal that linear projectors operating in the shrinkage (or expansion) regime hinder (or improve) the downstream classification accuracy. This provides the first theoretical explanation as to why (linear) projectors impact the downstream performance of learned representations. Our theoretical findings are further corroborated by extensive experiments on both synthetic data and real image data.

## 1 Introduction

Representation learning (Bengio et al., 2013) is a fundamental task in machine learning and statistics with the aim of extracting representations from the data that are useful for building future classifiers or predictors. While supervised learning is effective for this purpose (for instance, deep neural networks such as ResNet (He et al., 2016) have achieved remarkable performance in image classification), it is limited by the availability of massive labeled data.

Self-supervised learning (SSL) (Balestriero et al., 2023) has recently emerged as a novel paradigm to learn meaningful representations from huge *unlabeled* datasets (Misra and van der Maaten, 2019; Chen et al., 2020a; He et al., 2020; Dwibedi et al., 2021; HaoChen et al., 2021; Jing et al., 2021; Wang and Isola, 2020; Ji et al., 2021). Among SSL methods (Chen et al., 2020a; Zbontar et al., 2021; Bardes et al., 2021), contrastive learning (Chen et al., 2020a) is arguably the most popular one, which is also the focus of this paper. In essence, contrastive learning learns representations by

encouraging proximity between the representations of similar inputs (also known as positive pairs), while forcing the representations of dissimilar inputs (i.e., negative pairs) to be far from each other.

Below, we compare contrastive learning with the classical representation learning methods to help readers better understand the former one. Readers familiar with contrastive learning can jump directly to Section 1.3.

## 1.1 Classical unsupervised learning based on encoders and decoders

Principal component analysis (PCA), dating back to Pearson (1901); Hotelling (1933), is perhaps the oldest unsupervised representation learning method. In a nutshell, PCA aims to find a linear function  $\mathbf{f}_{\mathbf{W}}(\mathbf{x}) = \mathbf{W}\mathbf{x}$  of the input  $\mathbf{x}$  that preserves as much information about the original data as possible. Mathematically, PCA can be formulated as minimizing the *reconstruction loss*: denoting  $\mathbf{h} = \mathbf{f}_{\mathbf{W}}(\mathbf{x})$ , we search for another linear function  $\mathbf{g}_{\mathbf{T}}(\mathbf{h}) = \mathbf{T}\mathbf{h}$  such that the empirical reconstruction loss

$$\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{g}_{\mathbf{T}}(\mathbf{f}_{\mathbf{W}}(\mathbf{x}_i))\|^2$$

is minimized over all possible linear maps  $\mathbf{f}_{\mathbf{W}}$  and  $\mathbf{g}_{\mathbf{T}}$  of fixed dimensions. Here,  $\{\mathbf{x}_i\}_{i \leq n}$  denotes the input data.

Using the machine learning terminology, the linear function  $\mathbf{f}_{\mathbf{W}}(\cdot)$  is called an *encoder* that maps input data to latent representations while the linear function  $\mathbf{g}_{\mathbf{T}}(\cdot)$  is called a *decoder* that reproduces the original input as accurately as possible. More generally, this *encoder-decoder* approach forms the core principle of many other representation learning methods (Ghojogh et al., 2023) including autoencoders (Bouillard and Kamp, 1988) and variational autoencoders (Kingma and Welling, 2013).

This encoder-decoder approach serves as the precursor of modern deep learning. Before 2010, learning features from large unlabeled data (also called pretraining), often followed by fine-tuning on a smaller labeled dataset, is known to be beneficial for downstream tasks (Hinton and Salakhutdinov, 2006). However, it is realized that the decoder component is not essential if we are not asking for a generative model, which can be time-consuming to train (Chen et al., 2020a). This gives rise to new attempts and ideas of learning with no or limited labeled data. Contrastive learning is one emerging approach that achieves outstanding empirical performance, producing state-of-the-art models such as CLIP (Radford et al., 2021).

## 1.2 An encoder-projector framework for contrastive learning

Contrastive learning marks its departure from the classical representation learning methods (e.g., autoencoders (Bouillard and Kamp, 1988; Kingma and Welling, 2013) in two aspects: (1) first, instead of the encoder-decoder framework, contrastive learning generally follows the *encoder-projector* framework (cf. Figure 1); (2) second, instead of minimizing the reconstruction loss, contrastive learning minimizes a *contrastive loss* with the aim to pull positive pairs closer and push negative ones farther. Below we detail these two modifications and other essential components of contrastive learning using images as a running example for the input data.

**Encoder.** Similar to autoencoders, contrastive learning starts with an encoder  $\mathbf{h} = \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x})$  that maps an input image  $\mathbf{x}$  to its feature representation  $\mathbf{h}$ . Oftentimes, the encoder  $\mathbf{f}_{\boldsymbol{\theta}}$  is a deep neural network, e.g., ResNet-50 (He et al., 2016).

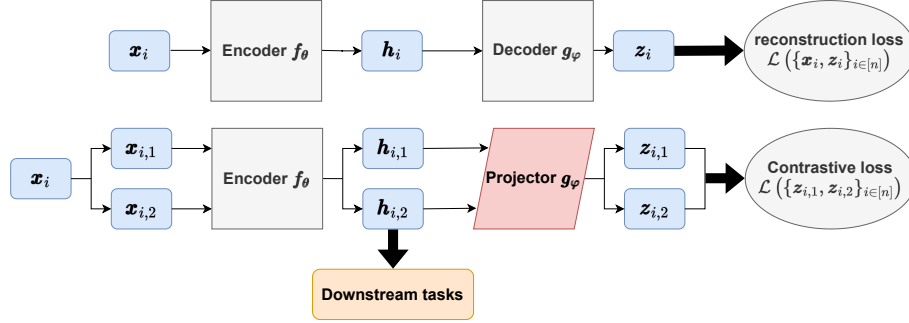


Figure 1: Encoder-decoder framework vs. encoder-projector framework.

**Projector.** Contrastive learning has a unique component  $\mathbf{z} = \mathbf{g}_\varphi(\mathbf{h})$  that projects the representation  $\mathbf{h}$  to its embedding  $\mathbf{z}$ , used later for calculating loss functions. This map  $\mathbf{g}_\varphi$  is usually called the *projection head* or *projector*. One often uses a simple network, e.g., a multilayer perceptron (MLP) with one or two hidden layers for the projector.

**Augmentation and contrastive loss.** As we have mentioned, contrastive learning aims to learn representations that are close for positive pairs and far for negative pairs. Here we provide one example for building positive and negative pairs from unlabeled data. Each sample  $\mathbf{x}_i$  is augmented by transformations (e.g., color distortion) to produce semantically similar positive samples  $\{\mathbf{x}_{i,k}\}_{1 \leq k \leq K}$ , where  $K$  denotes the number of augmentations (a.k.a. views). Two augmentations  $(\mathbf{x}_{i,k}, \mathbf{x}_{j,k'})$  are called a positive pair if  $i = j$ , and a negative pair if  $i \neq j$ . Figure 2 gives an example of building positive and negative pairs from unlabeled image data.

With the positive and negative pairs in place, we can introduce the contrastive loss used in contrastive learning, as opposed to the reconstruction loss in PCA and autoencoders. Let  $\mathbf{z}_{i,k} := \mathbf{g}_\varphi(\mathbf{f}_\theta(\mathbf{x}_{i,k}))$ . Denoting the (cosine) similarity score by  $\text{sim}(\mathbf{z}, \mathbf{z}') := \langle \frac{\mathbf{z}}{\|\mathbf{z}\|}, \frac{\mathbf{z}'}{\|\mathbf{z}'\|} \rangle$ , we train the encoder and projector by minimizing a contrastive loss, e.g., the canonical SimCLR loss (Chen et al., 2020a):

$$\min_{\theta, \varphi} \mathcal{L}(\theta, \varphi) := \underbrace{-\frac{1}{\tau} \sum_i \sum_{k \neq k'} \text{sim}(\mathbf{z}_{i,k}, \mathbf{z}_{i,k'})}_{\text{alignment loss}} + \underbrace{(K-1) \sum_{i,k} \log \left( \sum_{(j,k') \neq (i,k)} e^{\text{sim}(\mathbf{z}_{i,k}, \mathbf{z}_{j,k'})/\tau} \right)}_{\text{uniformity loss}}, \quad (1)$$

where  $\tau > 0$  is known as the *temperature* parameter. Following Wang and Isola (2020), we call the first term the *alignment* loss that promotes feature proximity of positive pairs, and the second term the *uniformity* loss that repels negative pairs.

In fact, the contrastive loss can be viewed as pairwise cross-entropy loss (a.k.a. logistic loss). Let  $s = (i, k)$  denote the index tuple for simplicity. In this case,  $s, s'$  are a positive pair if and only if  $i = i'$ . Then equivalently, one has  $\mathcal{L}(\theta, \varphi) = \sum_s \mathcal{L}_s(\theta, \varphi)$  where

$$\mathcal{L}_s(\theta, \varphi) = - \sum_{s' \neq s} \mathbf{1}\{i = i'\} \log \left( \frac{\exp(\text{sim}(\mathbf{z}_{s'}, \mathbf{z}_s)/\tau)}{\sum_{\bar{s} \neq s} \exp(\text{sim}(\mathbf{z}_{\bar{s}}, \mathbf{z}_s)/\tau)} \right).$$

**Downstream accuracy.** After training, we freeze  $\theta$  and only keep the encoder  $\mathbf{f}_\theta$  (i.e., throwing away the projector). Later, given a downstream task, say a classification problem with labeled data  $(\tilde{y}_i, \tilde{\mathbf{x}}_i)_{i \leq m}$ , one can simply apply logistic regression to learned features  $\{\mathbf{f}_\theta(\tilde{\mathbf{x}}_i)\}$  and labels  $\{\tilde{y}_i\}$ .

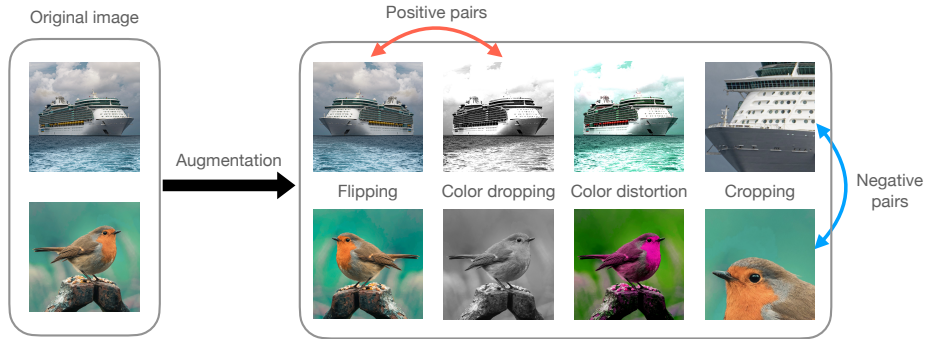


Figure 2: Positive and negative pairs after image augmentation.

### 1.3 Puzzling effect of projectors on the representations

The success of contrastive learning, or more specifically SimCLR (Chen et al., 2020a), can be attributed (at least) to two ingredients: strong data augmentation and the use of projectors. However, the role of projectors is quite puzzling. It is observed in Chen et al. (2020a) that representations learned before projectors outperform those learned after—measured using the downstream linear classification accuracy. This is even true when the projectors themselves are linear; see Figure 8 therein. As a result, the standard practice in contrastive learning is to jointly train the encoder and the projector, and then *remove* the projector completely after training (Balestrierio et al., 2023).

In this paper, we aim to demystify this puzzling phenomenon about projectors in contrastive learning. In particular, we focus on answering the following two questions:

**Q1:** *What geometric structure does contrastive loss minimization induce on the projectors?*

**Q2:** *Why do (linear) projectors affect the generalization properties of learned features?*

For Q1, the contrastive loss is relatively new and thus less understood. It is unclear what properties the optimal solution to contrastive loss minimization possesses. A quantitative characterization will be helpful for understanding contrastive learning.

For Q2, heuristics in the literature are insufficient (see e.g., Section 3.2 in Balestrierio et al. (2023)). One may view a projector as a buffer component: it protects features from distortion due to loss minimization, so its removal after training improves the generalization properties of features. Yet, this argument cannot reconcile with the observation that even linear projectors affect the downstream linear classification accuracy. Therefore, without a thorough investigation, the role of projectors on generalization remains mysterious.

### 1.4 Our contributions

To delineate the effect of projectors, we do not attempt to analyze the encoder or its training dynamics, but rather assume access to a well-trained encoder, which is often achieved at the later stage of training. Through fixing such a good encoder, we make the following empirical and theoretical discoveries that are fundamental to addressing the aforementioned two questions.

1. First, we identify two crucial effects—expansion and shrinkage—induced by the contrastive loss on the projectors. In essence, contrastive loss either expands or shrinks the signal direction in the representations learned by encoders.



2. Secondly, under a simpler projector model, we precisely characterize the downstream linear classification accuracy in the high-dimensional asymptotic limit. Our findings reveal that linear projectors operating in the shrinkage (resp. expansion) regime hinder (resp. improve) the downstream classification accuracy. This provides the first theoretical explanation as to why (linear) projectors impact the downstream performance of learned representations.

We also discuss connections to other empirical phenomena such as dimensional collapse, feature transferability, neural collapse, etc.

## 1.5 Paper organization

In Section 2, we present the empirical findings regarding the effects of the contrastive loss on the projectors, including expansion and shrinkage. In Section 3, we introduce the feature-level Gaussian mixture model and provide empirical evidence as motivation for our modeling approach. Moving on to Section 4, we present a precise approximation of the population contrastive loss under the Gaussian mixture model. Additionally, we theoretically characterize the sharp phase transition between the expansion and shrinkage regimes, and provide the approximation bound along with the finite-sample loss. Section 5 analyzes the impact of the expansion/shrinkage phenomenon of the linear projection head on downstream tasks. We calculate the precise generalization error as a function of the expansion effect. Furthermore, in Section 6, we extend our expansion/shrinkage results from Section 4 to inhomogeneous augmentations. We identify the simultaneous expansion and shrinkage, which aligns with our empirical discoveries using the STL-10 dataset. Simulation details and theoretical proofs are deferred to the appendix.

## 1.6 Notations

$\mathbb{R}_+$  denotes  $\{a \in \mathbb{R} : a \geq 0\}$  and  $\bar{\mathbb{R}}$  denotes  $\mathbb{R} \cup \{\infty\}$ . For a vector  $\mathbf{u}$ , we use  $\|\mathbf{u}\|$  or  $\|\mathbf{u}\|_2$  to denote its Euclidean norm. For vectors  $\mathbf{u}, \mathbf{v}$  of the same length, we use  $\langle \mathbf{u}, \mathbf{v} \rangle := \mathbf{u}^\top \mathbf{v}$  to denote the inner product. For a matrix  $\mathbf{A}$ , we use  $\|\mathbf{A}\|_F$  to denote its Frobenius norm. The identity matrix of size  $p \times p$  is denoted by  $\mathbf{I}_p$  or simply  $\mathbf{I}$ . We use  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  to denote a Gaussian distribution with mean  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ . The notation  $\Phi$  means the cumulative distribution function of the standard Gaussian variable.

For two real-valued sequences  $(a_n)_{n \geq 1}$  and  $(b_n)_{n \geq 1}$ , we use the standard small-o notation:  $a_n = o(b_n)$  means  $\lim_{n \rightarrow \infty} a_n/b_n = 0$ , and sometimes we also write  $a_n \ll b_n$ . For random variable  $X_n$ ,  $X_n = o_{\mathbb{P}}(b_n)$  means  $|X_n|/b_n$  converges in probability to 0 as  $n \rightarrow \infty$ . Moreover, if  $\mathbf{X}_n$  is a random vector, then  $\mathbf{X}_n = o_{\mathbb{P}}(b_n)$  means  $\|\mathbf{X}_n\|/b_n$  converges in probability to 0 as  $n \rightarrow \infty$ .

## 2 Empirical discovery: expansion and shrinkage

We begin with presenting our empirical findings on the two crucial effects—expansion and shrinkage, of contrastive loss on projectors.

**Experimental setup.** We provide a brief description about our experimental setup, and leave the details to the appendix. We freeze the encoder (based on ResNet-18)<sup>1</sup> pretrained on the STL-10 image dataset.<sup>2</sup> Then we apply standard data augmentation techniques such as random cropping

<sup>1</sup>Downloaded from <https://github.com/sthalles/SimCLR>.

<sup>2</sup>Source from <https://cs.stanford.edu/~acoates/stl10/>.




Figure 3: Results with the pretrained encoder and a one-layer linear projector  $\mathbf{W} \in \mathbb{R}^{512 \times 512}$  under the standard SimCLR loss (ResNet-18 and 10-class STL-10 dataset). **Left**: singular values of  $\mathbf{W}$  with varying temperature  $\tau$ . **Middle**: cumulative score plot  $\text{score}_i = \sum_{j \leq i} \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1, c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1, c_2}\|^2$  where  $(c_1, c_2) = (7, 2)$ . **Right**: downstream task accuracy using features  $\mathbf{h}$  versus using embeddings  $\mathbf{z}$  on the test set containing all 10 classes.

and color distortion to generate positive/negative pairs. In the end, we train a linear projector  $\mathbf{W} \in \mathbb{R}^{512 \times 512}$  under different configurations of hyperparameters, including 5 different temperatures.

In what follows, we summarize our empirical findings.

## 2.1 Insights from spectral decomposition

First, for two different class labels  $c_1, c_2$  (e.g.,  $c_1$  denotes airplane while  $c_2$  denotes dog), we use  $\boldsymbol{\mu}_{c_k} := \text{Ave}_{y_i=c_k} \mathbf{h}_i$  ( $k = 1, 2$ ) to denote the average representation after the encoder, where Ave is the average operator. Correspondingly, we denote  $\boldsymbol{\mu}_{c_1, c_2} := \boldsymbol{\mu}_{c_1} - \boldsymbol{\mu}_{c_2} \in \mathbb{R}^p$  to be the difference between class means. It turns out that  $\boldsymbol{\mu}_{c_1, c_2}$  is closely connected to the top/bottom right singular subspaces of the linear projector  $\mathbf{W}$ .

To be more precise, let  $\mathcal{V}_{\text{top}}, \mathcal{V}_{\text{bottom}}$  be the top/bottom singular subspaces containing a few right singular vectors (SVs) of  $\mathbf{W}$  and  $\mathcal{V}_{\text{bulk}}$  be the singular subspace containing the remaining SVs. Clearly one has  $\mathbb{R}^p = \mathcal{V}_{\text{top}} \oplus \mathcal{V}_{\text{bulk}} \oplus \mathcal{V}_{\text{bottom}}$ .

*Our first main empirical finding is that, empirically, the following holds approximately:*

$$\boldsymbol{\mu}_{c_1, c_2} \in \mathcal{V}_{\text{top}} \oplus \mathcal{V}_{\text{bottom}}.$$

In other words, the energy is concentrated on the span of extreme (right) singular vectors. Moreover, on  $\mathcal{V}_{\text{top}}$  and  $\mathcal{V}_{\text{bottom}}$ , the corresponding singular values experience sharp drops. Indeed, in Figure 3, we train  $\mathbf{W}$  on the 10-class STL-10 dataset and calculate the alignment score  $\kappa_{j, c_1, c_2} := \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1, c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1, c_2}\|^2$ , where  $\mathbf{v}_j$  is the  $j$ th right SV of  $\mathbf{W}$ . For each index  $i$ , we report the cumulative score  $\sum_{j \leq i} \kappa_{j, c_1, c_2}$ , which satisfies the normalization  $\sum_{j=1}^p \kappa_{j, c_1, c_2} = 1$ . A wide flat cumulative score suggests orthogonality  $\boldsymbol{\mu}_{c_1, c_2} \perp \mathcal{V}_{\text{bulk}}$ . A geometric interpretation is that  $\mathbf{W}$  is *expanding* vectors in  $\mathcal{V}_{\text{top}}$  and *shrinking* vectors in  $\mathcal{V}_{\text{bottom}}$ .

## 2.2 Expansion/shrinkage affects generalization

*Our second empirical finding is that the expansion/shrinkage effects are highly correlated with downstream accuracy.*

As we observed earlier, when the temperature  $\tau$  increases, expansion gets stronger because the alignment of the top SVs with  $\boldsymbol{\mu}_{7, 2}$  increases. Figure 3 (right) shows that a larger temperature also

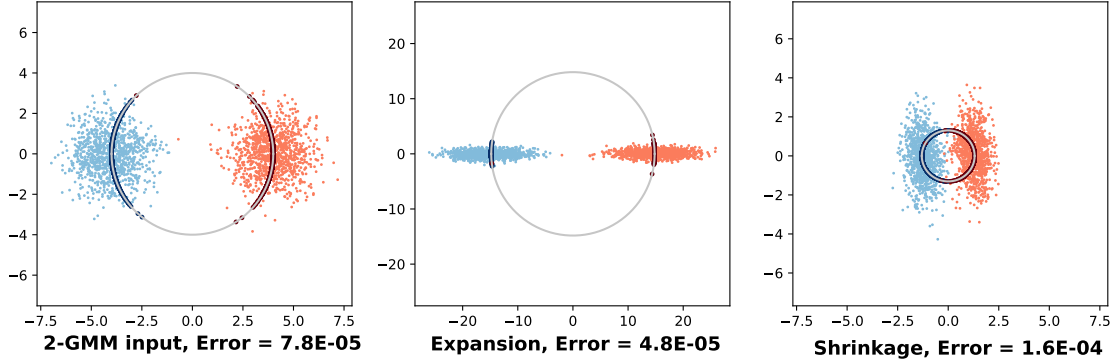


Figure 4: Contrastive loss drives expansion or shrinkage. **Left:** we generate clean 2-component GMM data. **Middle and right:** we use SimCLR loss for training 100-dim data, and then visualize embeddings on a 2D plane under large/small augmentation and large/small temperature (middle/right plot). Also, we show the projection of embeddings onto a circle with varying radii, which is  $z_i/\|z_i\|$  used for calculating cosine similarities. **Bottom text:** expansion decreases test error, while shrinkage reduces signals and increases test error.

leads to an increase in the classification accuracy using the embeddings  $(z_i)_{i \leq n}$ . Figure 3 (right) confirms the surprise that a linear projector can change the generalization performance significantly.

Can we theoretically justify the expansion and shrinkage effects of contrastive loss and their impact on the generalization power of the learned representations?

### 2.3 A simple simulation and heuristic explanations

Here, we present a simple simulation using synthetic input data that recreates the salient characteristics of our empirical findings. We hope to provide some heuristic explanations for the observed phenomena, while leaving the formal proof to later sections.

**A visual illustration.** The main characteristics of the empirical structure of projectors are displayed in a simple clean model. Consider a 2-component Gaussian mixture model  $\frac{1}{2}\mathcal{N}(-\boldsymbol{\mu}, \mathbf{I}_p) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p)$ . We generate data according to this model, and then generate augmented data by adding independent perturbation (Gaussian noise). The data pass through a linear layer  $\mathbf{W} \in \mathbb{R}^{p \times p}$  treated as the projector. We obtain  $\mathbf{W}^*$  by minimizing the SimCLR loss over  $\mathbf{W}$ .

We project and visualize our simulated data in Figure 4 using  $\mathbf{W}^*$ . Under two different hyperparameter settings, the original data (shown in the left plot) are either (i) extended along the separation direction  $\boldsymbol{\mu}$  in the middle plot, or (ii) compressed along  $\boldsymbol{\mu}$  in the right plot.

The simulation result matches earlier empirical findings. (i) In the expansion regime, normalized features (projected to a circle) are more aligned, signaling a large top singular value; and in the shrinkage regime, normalized features are more uniform, signaling a smaller singular value. (ii) The test error decreases in the expansion regime that has a higher temperature, and it increases in the shrinkage regime that has a lower temperature.

**Expansion and shrinkage promote alignment and uniformity.** In Wang and Isola (2020), the contrastive loss is decomposed into two components, which are called the *alignment* loss and the *uniformity* loss—which correspond to the two terms in the RHS of (1). Feature embeddings in the hypersphere are driven by the two opposite forces induced by the two loss components. Our

expansion and shrinkage perspective offers a consistent explanation. If the projector stretches the features along the signal direction thus increasing variance in that direction, then after normalization to the hypersphere the features become more aligned. Conversely, if the projector compressed the features along the signal direction, after normalization features will become more uniform. See Section 4 for a formal analysis.

**Projector as reparametrization changes inductive bias.** Why does the projector, even in the linear case, change the generalization performance on downstream tasks? Reparametrization techniques such as skip connections (He et al., 2016) and batch normalization (Ioffe and Szegedy, 2015) are commonly used in deep learning, yet how they impact generalization is rarely elucidated. Here, we give an explanation from the lens of inductive bias. When the downstream task involves linearly separable data, gradient descent on the commonly used logistic loss produces a sequence of iterates  $\beta^{(1)}, \beta^{(2)}, \dots$  that converge in direction to the max-margin solution (Soudry et al., 2018).

$$\begin{aligned} \max_{\beta} \quad & \min_{i \leq n} y_i \langle \mathbf{z}_i, \beta \rangle \\ \text{subject to} \quad & \|\beta\|_2 \leq 1. \end{aligned}$$

The  $\ell_2$  norm in the constraint plays the role of implicit regularization that is induced by the gradient descent. Applying a linear transformation to features, however, will effectively lead to a different norm. Thus, the generalization properties on downstream tasks are affected by even a linear projector. See Section 5 for the formal analysis.

### 3 Feature-level modeling via GMM

To decouple the encoder and the projector, we assume access to a well-trained encoder and model the output using a well-separated Gaussian mixture model (GMM), where each Gaussian component represents a class. This assumption is also supported by at least two empirical observations: (i) the cluster structure in the feature space, and (ii) convergence speed during training. We defer detailed justifications about this assumption to Section 3.2, before which we detail the model setup using GMMs.

#### 3.1 Model setup

Mathematically, we assume that the outputs of the encoder (i.e., features) are generated from a 2-component GMM (or 2-GMM), that is

$$\mathbf{h}_{0,i} \stackrel{\text{i.i.d.}}{\sim} \frac{1}{2} \mathcal{N}(-\boldsymbol{\mu}, \mathbf{I}_p) + \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p),$$

where  $\boldsymbol{\mu} \in \mathbb{R}^p$  denotes the mean difference between two classes. Conditional on  $\mathbf{h}_{0,i}$ , we construct two augmentations (views) via

$$\mathbf{h}_i, \mathbf{h}_i^+ \mid \mathbf{h}_{0,i} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{h}_{0,i}, \sigma_{\text{aug}}^2 \mathbf{I}_p),$$

where  $\sigma_{\text{aug}}^2 > 0$  represents the augmentation strength. As a result,  $(\mathbf{h}_i, \mathbf{h}_i^+)$  is viewed as a positive pair, and  $(\mathbf{h}_i, \mathbf{h}_j)$ ,  $(\mathbf{h}_i, \mathbf{h}_j^+)$  ( $i \neq j$ ) are considered negative pairs. For the projector, we focus on *linear* projectors, i.e.,  $\mathbf{z} = \mathbf{g}_{\mathbf{W}}(\mathbf{h}) = \mathbf{W}\mathbf{h}$ , where  $\mathbf{W} \in \mathbb{R}^{p \times p}$ . We present a modified SimCLR loss

that is more amenable to theoretical analysis. Given the temperature  $\tau > 0$  and projector outputs  $(\mathbf{z}_i)_{i \leq n}$  where  $\mathbf{z}_i = \mathbf{g}_{\mathbf{W}}(\mathbf{h}_i)$ , we define

$$\mathcal{L}_n(\mathbf{W}) := -\frac{1}{n\tau} \sum_{i=1}^n \mathbb{E}_{\mathbf{z}_i^+, \mathbf{z}_i | \mathbf{h}_{0,i}} [\text{sim}^*(\mathbf{z}_i, \mathbf{z}_i^+)] + \log \left( \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathbf{z}_i^-, \mathbf{z}_i | \mathbf{h}_{0,i}} [e^{\text{sim}^*(\mathbf{z}_i^-, \mathbf{z}_i)/\tau}] \right), \quad (2)$$

$$\text{where } \text{sim}^*(\mathbf{z}_1, \mathbf{z}_2) := \frac{-\|\mathbf{z}_1 - \mathbf{z}_2\|^2/2}{\{\mathbb{E}[\|\mathbf{z}_1\|^2]\}^{1/2} \cdot \{\mathbb{E}[\|\mathbf{z}_2\|^2]\}^{1/2}}.$$

Here,  $\mathbb{E}_{\mathbf{z}_i^+, \mathbf{z}_i | \mathbf{h}_{0,i}}$  denotes expectation over random augmentations *conditioning* on  $\mathbf{h}_{0,i}$  (i.e., using every possible positive pair),  $\mathbf{z}_i^-$  denotes an independent copy of  $\mathbf{z}_i$ , and  $\mathbb{E}_{\mathbf{z}_i^-, \mathbf{z}_i | \mathbf{h}_{0,i}}$  denotes expectation over the negative pair  $(\mathbf{z}_i, \mathbf{z}_i^-)$  conditioning on  $\mathbf{h}_{0,i}$  (i.e., using every possible negative pair).

Our modified SimCLR loss is different from the original SimCLR in the following ways: (i) We consider full batch and training for infinite time (so averages over augmentations are replaced by expectations); (ii) similar to Wang and Isola (2020), we interchange log with the summation; and (iii) we use a variant of similarity score  $\text{sim}^*$  where we replace instance-based normalization with the population level normalization.

*We emphasize that this modified loss is only used for theoretical analysis. In all experiments, we use the original SimCLR contrastive loss.*

### 3.2 Why fixing a well-trained encoder?

Arguably, our theoretical study departs from common practice as we assume a well-trained encoder and fix it to be a simple Gaussian mixture model. However, we would like to argue that this assumption sheds light on the practice.

**Convergence of projectors.** Training an encoder and projector jointly creates complex dynamics that are beyond the scope of this paper. Yet, we observe in experiments that the dynamics are significantly simplified in the later stage of training. Let  $(\boldsymbol{\theta}^{(t)}, \mathbf{W}^{(t)})$  be the parameters at epoch  $t$ , and  $\widetilde{\mathbf{W}}^{(t)}$  be the optimal projector parameters  $\widetilde{\mathbf{W}}^{(t)} := \text{argmin}_{\mathbf{W}} \mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathbf{W})$  while freezing encoder parameters  $\boldsymbol{\theta}^{(t)}$  at epoch  $t$ . It can be seen from Figure 14 in the appendix that when  $t \geq 50$ ,

$$\|\widetilde{\mathbf{W}}^{(t)} - \mathbf{W}^{(t)}\| \ll \|\mathbf{W}^{(t)}\|,$$

which suggests that the projector is close to (conditionally) optimal. This allows us to decouple the dynamics and encoders and projectors into two separate optimization problems:

1. For fixed encoder  $\boldsymbol{\theta}$ , solve  $\min_{\mathbf{W}} \mathcal{L}(\boldsymbol{\theta}, \mathbf{W})$ .
2. Solve  $\min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}, \widetilde{\mathbf{W}}(\boldsymbol{\theta}))$  where  $\widetilde{\mathbf{W}}(\boldsymbol{\theta})$  is the minimizer of the first part.

Our focus is the first optimization problem, assuming that the encoder  $\boldsymbol{\theta}$  is sufficiently good, e.g.,  $\boldsymbol{\theta} = \boldsymbol{\theta}^{(t)}$  for large  $t$ . We remark that the feature-learning process and training dynamics of encoder networks are complicated and elusive for analysis, so it is not uncommon to decouple multiple layers or components by freezing some of them Han et al. (2021); Bietti et al. (2022).

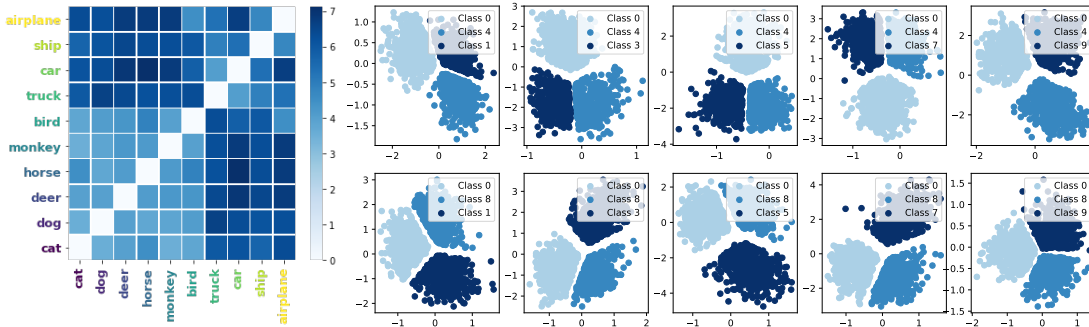


Figure 5: Pretrained STL-10 features from the standard SimCLR model exhibit a strong cluster structure. **Left:** positive pairwise margins of linear classifier imply pairwise linear separability. **Right:** 2D visualization of 3-class subsets of features after 30 epochs. The cluster structure emerges quickly.

**Cluster structure in feature space.** When the encoders are trained well, we observe well-separated clusters formed by pretrained features in contrastive learning, and these clusters contain class information (Böhm et al., 2022). We examine the cluster structure of pretrained features both visually and quantitatively. In Figure 5 (left plot), for every pair of classes, we perform max-margin linear classification on a subset of pretrained features. We find that every pair achieves zero training error with a large margin, which indicates that all classes are linearly separable. In Figure 5 (right plot), we visualize pretrained features by following the same projection technique Müller et al. (2019). For well-clustered pretrained features, the GMM is a natural model, which is the starting point of our analysis.

## 4 Contrastive loss drives expansion/shrinkage

Under the model setup in Section 3.1, we hope to characterize the minimizer of  $\mathcal{L}_n(\mathbf{W})$ . However, the loss  $\mathcal{L}_n$  is very complex, and quantitative results are hard to obtain. It turns out that, under reasonable simplification, the minimization admits an explicit solution. We then quantify the approximation error in the aforementioned simplification.

### 4.1 Quantitative results under simplification

We start our analysis with the infinite-sample case, i.e.,  $n = \infty$ . Let  $\mathcal{L}(\mathbf{W}) := \mathbb{E}_z \mathcal{L}_n(\mathbf{W})$  be the population loss. Denote by  $\mathbf{W} = \sum_{j=1}^p \sigma_j \mathbf{u}_j \mathbf{v}_j^\top$  the SVD of  $\mathbf{W} \in \mathbb{R}^{p \times p}$ , where  $\sigma_1 \geq \dots \geq \sigma_p \geq 0$ ,  $[\mathbf{u}_1, \dots, \mathbf{u}_p]$  and  $[\mathbf{v}_1, \dots, \mathbf{v}_p]$  are orthogonal matrices. Define

$$\alpha := \alpha(\mathbf{W}) = (1 + \sigma_{\text{aug}}^2) \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2.$$

A straightforward calculation yields the following decomposition of the population loss.

**Proposition 4.1.** *We have  $\mathcal{L}(\mathbf{W}) = \mathcal{L}_{\text{align}}(\mathbf{W}) + \mathcal{L}_{\text{unif}}(\mathbf{W})$  with*

$$\begin{aligned} \mathcal{L}_{\text{align}}(\mathbf{W}) &= \frac{\sigma_{\text{aug}}^2}{\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2, \\ \mathcal{L}_{\text{unif}}(\mathbf{W}) &= -\log 2 - \frac{1}{2} \sum_{j=1}^p \log \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right) + \log \left( 1 + \exp \left( - \sum_{j=1}^p \frac{2\sigma_j^2 \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2}{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha} \right) \right) \end{aligned}$$



$$=: -\log 2 + \mathcal{L}_{\text{unif}}^{(1)}(\mathbf{W}) + \mathcal{L}_{\text{unif}}^{(2)}(\mathbf{W}).$$

See Appendix B.1 for the proof.

To understand the two loss components, let us make some observations.

1. The main effect of  $\mathcal{L}_{\text{align}}(\mathbf{W})$  is to align positive pairs. If we only minimize  $\mathcal{L}_{\text{align}}(\mathbf{W})$ , then we need to maximize<sup>3</sup>  $\|\mathbf{W}\boldsymbol{\mu}\|$  for any fixed value of  $\|\mathbf{W}\|$ . In other words,  $\mathbf{W}$  needs to stretch  $\boldsymbol{\mu}$  as much as possible. This is intuitive since positive pairs after such stretching and normalization will be closer to each other; see the middle plot of Figure 4.
2. The main effect of  $\mathcal{L}_{\text{unif}}(\mathbf{W})$  is to expel negative pairs. If we only minimize one critical part  $\mathcal{L}_{\text{unif}}^{(1)}(\mathbf{W})$ , then given fixed  $(\sigma_j)_{j \leq p}$  we need to minimize  $\|\mathbf{W}\boldsymbol{\mu}\|$  (i.e.,  $\mathbf{W}$  compresses  $\boldsymbol{\mu}$  as much as possible), leading to  $|\langle \boldsymbol{\mu}, \mathbf{v}_p \rangle| = 1$ . Embeddings after shrinkage and normalization will be more diversely distributed; see the right plot of Figure 4.

We can gain further insights by considering a *first-order approximation* of the loss: intuitively, if the signal strength  $\|\boldsymbol{\mu}\|^2$  is large, then we may expect  $\alpha \gg \|\mathbf{W}\|_{\text{F}}^2 \geq \max_j \sigma_j^2$ , so we can try a first-order expansion by treating  $\alpha$  as a diverging quantity. This motivates us to define an approximate loss:

$$\begin{aligned} \tilde{\mathcal{L}}(\mathbf{W}) &:= \frac{\sigma_{\text{aug}}^2}{\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2 - \log 2 - \sum_{j=1}^p \frac{(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} + \log \left( 1 + \exp \left( - \sum_{j=1}^p \frac{2\sigma_j^2 \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2}{\tau\alpha} \right) \right) \\ &= -\frac{t}{\tau} + \log \left( \frac{1}{2} + \frac{1}{2} \exp \left( -\frac{2}{\tau} + \frac{2(1 + \sigma_{\text{aug}}^2)}{\tau} t \right) \right), \end{aligned} \quad (3)$$

where  $t := t(\mathbf{W}) = \frac{\|\mathbf{W}\|_{\text{F}}^2}{\alpha(\mathbf{W})}$ . As an approximation, it always holds that  $\mathcal{L}(\mathbf{W}) \geq \tilde{\mathcal{L}}(\mathbf{W})$ . Now that  $\tilde{\mathcal{L}}(\mathbf{W})$  is actually a univariate function, we can state precise quantitative results.

**Theoretical prediction.** Now we are ready to theoretically demonstrate the two key effects—expansion and shrinkage. Let  $F(t; \sigma_{\text{aug}}^2, \tau) := \frac{d}{dt} \tilde{\mathcal{L}}(\mathbf{W})$ .

**Definition 4.2.** A three-parameter configuration  $(\sigma_{\text{aug}}^2, \tau, \|\boldsymbol{\mu}\|^2)$  is said to be in the

- expansion regime if  $F\left(\frac{1}{1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2}; \sigma_{\text{aug}}^2, \tau\right) > 0$ ,
- shrinkage regime if  $F\left(\frac{1}{1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2}; \sigma_{\text{aug}}^2, \tau\right) < 0$ .

The expansion and shrinkage regimes characterize the distinctive behavior of the projector  $\mathbf{W}$  at minimization, as stated in the next theorem.

**Theorem 4.3** (expansion/shrinkage phase transition). *The properties of the minimizer  $\mathbf{W}^*$  of  $\tilde{\mathcal{L}}(\mathbf{W})$  depends on the configuration of  $(\sigma_{\text{aug}}^2, \tau, \|\boldsymbol{\mu}\|^2)$ . Specifically, with the notation  $t^* := t(\mathbf{W}^*)$ ,*

- in the expansion regime,  $\tilde{\mathcal{L}}(\mathbf{W})$  is minimized at  $t^* = 1/(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2)$ , which happens if and only if  $\sigma_2 = \dots = \sigma_p = 0$  and  $\langle \mathbf{v}_1, \boldsymbol{\mu} \rangle^2 = \|\boldsymbol{\mu}\|^2$ ;
- in the shrinkage regime,  $\tilde{\mathcal{L}}(\mathbf{W})$  is minimized at

$$t^* = \frac{1}{1 + \sigma_{\text{aug}}^2} \left( 1 - \frac{\tau}{2} \log(1 + 2\sigma_{\text{aug}}^2) \right) \in \left[ \frac{1}{1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2}, \frac{1}{1 + \sigma_{\text{aug}}^2} \right]$$

that solves  $F(t) = 0$ . Moreover, if  $\sigma_{\text{aug}}^2 \rightarrow 0$ , then  $t^* \rightarrow 1$  and  $\|\mathbf{W}\boldsymbol{\mu}\|/\|\mathbf{W}\|_{\text{F}} \rightarrow 0$ .

<sup>3</sup>Note that  $\alpha$  also depends on  $\mathbf{W}$ .




Figure 6: Expansion/shrinkage phase transition for the configuration  $n = 2000$ ,  $p = 50$ ,  $\|\boldsymbol{\mu}\| = 5$ . **Left:** heatmap of expansion measure  $T^* := T(\mathbf{W}^*)$ ; **green curve:** theoretical phase transition. **Right:**  $T^*$  vs.  $\tau$  with different fixed  $\sigma_{aug}$ .

See Appendix B.3 for the proof.

Several remarks are in order. First, when the configuration  $(\sigma_{aug}^2, \tau, \|\boldsymbol{\mu}\|^2)$  is in the expansion regime, optimizing the contrastive loss leads to a linear projector  $\mathbf{W}^*$  that maximally expands along the signal direction  $\boldsymbol{\mu}$ . Second, in contrast, in the shrinkage regime, the signal direction is compressed in  $\mathbf{W}^*$ , since for fixed value  $\|\mathbf{W}\|_F$ , a vanishing  $\|\mathbf{W}\boldsymbol{\mu}\|$  implies  $\max_j |\sigma_j \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle| \rightarrow 0$ —any singular vector positively correlated with  $\boldsymbol{\mu}$  must have a vanishing singular value.

**Numerical experiments.** Under the 2-GMM model stated in Section 3, in Figure 6 (left), we plot the heatmap of the expansion measure  $T^* := T(\mathbf{W}^*)$  with  $T(\mathbf{W}) := \|\mathbf{W}\boldsymbol{\mu}\|^2 / (\|\mathbf{W}\|_F^2 \|\boldsymbol{\mu}\|^2) \in [0, 1]$  with different pairs of temperature and augmentation strength. The green curve is the transition curve

$$\tau^* = 2\|\boldsymbol{\mu}\|^2 \left\{ (1 + \sigma_{aug}^2 + \|\boldsymbol{\mu}\|^2) \log(1 + 2\sigma_{aug}^2) \right\}^{-1}$$

which separates the expansion and shrinkage regimes. In Figure 6 (right), with a set of  $\sigma_{aug}$ , we plot the expansion measure against varying temperatures which shows that our theoretical prediction is also precise in the regime when  $0 < T^* < 1$ . When  $\tau < \tau^*$  (shrinkage),  $T^*$  is close to zero, indicating a significant compression along  $\boldsymbol{\mu}$ ; when  $\tau > \tau^*$  (expansion),  $T^* \approx 1$ , corresponding to maximal expansion along  $\boldsymbol{\mu}$ .

Despite simplification and approximation, our theoretical prediction does align well with empirical phase transition with the original SimCLR loss.

## 4.2 Extensions to the general case

Recall that  $\tilde{\mathcal{L}}$  is an approximation to the original loss  $\mathcal{L}$ . Therefore it is of interest to quantify the approximation error in the optimal solutions. We also briefly discuss the technical difficulty of extending these analysis to the finite-sample regime.

**Approximation errors.** Compared with  $\tilde{\mathcal{L}}$ , the original loss  $\mathcal{L}$  cannot be expressed as a simple univariate function. Recall that in the approximation, we consider large signal strength  $\|\boldsymbol{\mu}\|$ . Denote  $\mathbf{W}_{simclr}$  as the solution to the original loss  $\mathcal{L}$  and  $t_{simclr} := t(\mathbf{W}_{simclr}) = \|\mathbf{W}_{simclr}\| / ((1 + \sigma_{aug}^2)\|\mathbf{W}_{simclr}\|^2 + \|\mathbf{W}_{simclr}\boldsymbol{\mu}\|^2)$ . Approximation bounds are shown in the next proposition for the expansion and shrinkage regimes, respectively.

**Proposition 4.4.** *The difference between  $t^*$  and  $t_{\text{simclr}}$  has the following bounds:*

- in the expansion regime,

$$0 \leq t_{\text{simclr}} - t^* \leq \left( \frac{2(1 + \sigma_{\text{aug}}^2)}{1 + e^{2/\tau}} - 1 \right)^{-1} \frac{4(1 + \sigma_{\text{aug}}^2)}{\tau e^{1/\tau}} \frac{1}{\|\boldsymbol{\mu}\|^2} + C(\tau, \sigma_{\text{aug}}^2) \|\boldsymbol{\mu}\|^{-4};$$

- in the shrinkage regime,

$$\|t_{\text{simclr}} - t^*\|^2 \leq \frac{e^{2\sigma_{\text{aug}}^2/\tau}}{(1 + \sigma_{\text{aug}}^2)^2} \frac{1}{p-1} + C(\tau) \cdot \left( \frac{\sigma_{\text{aug}}^4}{\|\boldsymbol{\mu}\|^2} \right).$$

where  $C(\tau, \sigma_{\text{aug}}^2), C(\tau) > 0$  are some constants.

See Appendix B.4 for the proof.

We can infer the following properties of the approximation error bounds.

- For the expansion regime, the quality of approximation improves almost linearly in  $\|\boldsymbol{\mu}\|^2$ , and the dominating bound vanishes as  $\tau \rightarrow +\infty$ ;
- For the shrinkage regime, a sufficiently large dimension  $p$  implies a smaller approximation error, and when  $\sigma_{\text{aug}}^2$  is adequately small, the error bound becomes tighter as  $\sigma_{\text{aug}}^2$  increases.

**Finite-sample regime.** Now we comment on the difficulty when dealing with finite samples, as opposed to the infinite-sample regime we have been focusing on thus far. In the finite-sample regime, the loss function becomes

$$\mathcal{L}_n(\mathbf{W}) := -\frac{1}{\tau} \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^+, \mathbf{h} | \mathbf{h}} [\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^+) | \mathbf{h}_0] \right\} + \log \left( \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^-, \mathbf{h} | \mathbf{h}_0} [e^{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)/\tau} | \mathbf{h}_0] \right\} \right),$$

where  $\mathbb{E}_n[a]$  denotes the expectation over the empirical distribution of  $\{\mathbf{h}_{0,i}\}_{i \leq n}$ . To write  $\mathcal{L}_n(\mathbf{W})$  as a more explicit function in  $\mathbf{W}$ , denote  $\mathbf{M} = (\mathbf{I} + (\tau\alpha)^{-1}(1 + \sigma_{\text{aug}}^2)\mathbf{W}^\top \mathbf{W})^{-1}$ . Additionally, define

$$\begin{aligned} S_{\mathbf{h}_0} &:= \exp \left( -\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 - \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 - \boldsymbol{\mu}) \right) \\ &\quad + \exp \left( -\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 + \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 + \boldsymbol{\mu}) \right), \\ \tilde{S}_{\mathbf{h}_0} &:= \exp \left( -\frac{1}{2\tau\alpha} \|\mathbf{W}(\mathbf{h}_0 - \boldsymbol{\mu})\|^2 \right) + \exp \left( -\frac{1}{2\tau\alpha} \|\mathbf{W}(\mathbf{h}_0 + \boldsymbol{\mu})\|^2 \right), \end{aligned}$$

then we have the following proposition.

**Proposition 4.5.** *The finite sample loss  $\mathcal{L}_n(\mathbf{W})$  can be written as*

$$\mathcal{L}_n(\mathbf{W}) = -\log 2 + \frac{\sigma_{\text{aug}}^2}{\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2 + \frac{1}{2} \log \det[\mathbf{M}] + \log(\mathbb{E}_n S_{\mathbf{h}_0}).$$

Further, if  $\alpha \gg \|\mathbf{W}\|_{\text{F}}^2$ ,  $\mathcal{L}_n(\mathbf{W})$  has the following approximation

$$\tilde{\mathcal{L}}_n(\mathbf{W}) = -\log 2 - \frac{1}{2\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2 + \log(\mathbb{E}_n \tilde{S}_{\mathbf{h}_0}).$$

From the proposition, the difficulty in analyzing the finite-sample loss lies in the nonlinearity in  $\log \det[\mathbf{M}]$  as well as the terms  $S_{\mathbf{h}_0}$ . If we assume that  $\alpha \gg \|\mathbf{W}\|_F^2$ , the approximation is based on the fact that  $\mathbf{M} \approx \mathbf{I} - (\tau\alpha)^{-1}(1 + \sigma_{\text{aug}}^2)\mathbf{W}^\top \mathbf{W}$ .

*Remark.* As  $\mathbf{h}_0 \sim \mathcal{N}(y \cdot \boldsymbol{\mu}, \mathbf{I})$ , if we replace  $\mathbb{E}_n$  by  $\mathbb{E}_{\mathbf{h}_0}$  in  $\tilde{\mathcal{L}}_n(\mathbf{W})$ , then the loss can be approximated by

$$-\log 2 - \frac{1}{\tau\alpha} \|\mathbf{W}\|_F^2 + \log \left( 1 + \exp \left( -\frac{2}{\tau\alpha} \|\mathbf{W}\boldsymbol{\mu}\|^2 \right) \right),$$

which is exactly the loss function  $\tilde{\mathcal{L}}(\mathbf{W})$  defined in (3).

## 5 Effect of projectors on generalization

We move on to investigate the effect of projectors on the generalization of learned representations.

### 5.1 One-parameter projector model

To understand the generalization puzzle about projectors, an important empirical observation is that pretrained features typically form *linearly separable* clusters; see Figure 5. Adopting our 2-GMM view, we assume that cluster labels  $\{y_i\}_{i \leq n}$  are revealed to us on downstream tasks, and our data  $\{\mathbf{h}_{0,i}, y_i\}_{i \leq n}$  follow the model

$$\mathbf{h}_{0,i} \mid y_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(y_i \cdot \boldsymbol{\mu}, \mathbf{I}_p), \quad \text{where } y_i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{-1, 1\}).$$

Assuming a linear projector as before, we denote  $\mathbf{z}_i = \mathbf{W}\mathbf{h}_{0,i}$  and  $\{\mathbf{z}_i, y_i\}_{i \leq n}$  are the inputs of the classification problem.

In reality, the minimizer from the standard SimCLR loss (1) is not perfectly characterized by the ideal expansion/shrinkage; see the singular value plot in Figure 3. Nevertheless, to gain insights, we consider a simpler form of linear transforms:

$$\mathcal{W} = \{ \mathbf{W}_\eta = \mathbf{I}_p + \eta \rho^{-1} \boldsymbol{\mu} \boldsymbol{\mu}^\top : \eta > -1 \}, \quad \text{where } \rho := \|\boldsymbol{\mu}\|^2.$$

Since only the right singular vectors of  $\mathbf{W}$  is of interest in terms of expansion/shrinkage, the symmetric projection head is without loss of generality. If we constrain ourselves to this simple one-parameter space  $\mathcal{W}$ , then expansion/shrinkage regimes are solely determined by  $\eta$ . Specifically,  $\eta = 0$  means effectively no projector,  $\eta > 0$  corresponds to the expansion regime, while  $\eta < 0$  corresponds to the shrinkage regime.

$$\begin{cases} \text{If } \eta \in (-1, 0), & \text{then shrinkage regime;} \\ \text{If } \eta \in (0, \infty), & \text{then expansion regime;} \\ \text{If } \eta = 0, & \text{effectively no projector.} \end{cases}$$

### 5.2 Confirming folklore: invariance of test errors in low-dimension regime

In the low-dimensional regime (i.e., when  $p/n \rightarrow 0$ ), the data  $(\mathbf{h}_{0,i}, y_i)_{i \leq n}$  are not linearly separable with probability approaching one. We focus on  $\ell_2$ -regularized logistic regression with an intercept in a fixed dimension  $p$ . Define

$$\ell_n(\gamma, \boldsymbol{\beta}; \lambda_n) = \mathbb{E}_n \left\{ \log \left[ 1 + e^{-y(\gamma + \mathbf{z}^\top \boldsymbol{\beta})} \right] \right\} + \lambda_n \|\boldsymbol{\beta}\|^2, \quad (4)$$

where  $\mathbb{E}_n$  denotes the sample average over  $n$  samples  $i = 1, \dots, n$ . With coefficients  $\gamma, \boldsymbol{\beta}$  and a given  $\mathbf{W} = \mathbf{W}_\eta$ , define the test error as  $\text{Err}(\gamma, \boldsymbol{\beta}; \eta) := \mathbb{P}(\gamma + \tilde{y}\langle \tilde{\mathbf{z}}, \boldsymbol{\beta} \rangle < 0)$ , where  $(\tilde{\mathbf{z}}, \tilde{y})$  is a new independent sample. We further define the expected error  $\overline{\text{Err}}(\eta) = \mathbb{E}[\text{Err}(\hat{\gamma}, \hat{\boldsymbol{\beta}}; \eta)]$ , where  $\hat{\gamma}, \hat{\boldsymbol{\beta}}$  is the logistic regression estimator, and the expectation is taken w.r.t. the training samples. Intuitively, the linear transformation will not change the generalization error. The following result confirms this intuition by proving that the effects of expansion or shrinkage are negligible, unless we add an unreasonably large regularizer; see the second case below.

**Theorem 5.1.** *Consider minimizing the regularized logistic loss function (4).*

1. If  $0 \leq \lambda_n \ll \sqrt{n}$ , then the test error obeys

$$\overline{\text{Err}}(\eta) = \Phi(-\|\boldsymbol{\mu}\|) + o(1).$$

Here, the dominant term (i.e., the first term) in the test error remains the same for varying  $\eta$ .

2. If  $\lambda_n \asymp \sqrt{n}$ , then  $\overline{\text{Err}}(\eta)$  is decreasing in  $\eta$ .

See Appendix C.1 for the proof.

### 5.3 High-dimensional regime: inductive bias matters

In search of the explanation, we then turn to the high-dimensional regime. In the high-dimensional regime, two distinct phenomena arise: first,  $\{\mathbf{h}_{0,i}\}_{i \leq n}$  are linearly separable with high probability. Second, solving logistic regression on linearly separable data via gradient descent results in a max-margin classifier (Rosset et al., 2003; Soudry et al., 2018)—a form of inductive bias (Neyshabur et al., 2014). Recall that  $\mathbf{z}_i = \mathbf{W}\mathbf{h}_{0,i}$ , and the max-margin classifier is given by

$$\max_{\boldsymbol{\beta}} \min_{i \leq n} y_i \langle \mathbf{z}_i, \boldsymbol{\beta} \rangle \quad \text{subject to } \|\boldsymbol{\beta}\| \leq 1. \quad (5)$$

To understand the different inductive biases brought by different linear projectors, let us rewrite the max-margin problem above using the original input data: denoting  $\tilde{\boldsymbol{\beta}} = \mathbf{W}\boldsymbol{\beta}$ , the optimization problem (5) can be reformulated into

$$\begin{aligned} & \max_{\tilde{\boldsymbol{\beta}}} \min_{i \leq n} y_i \langle \mathbf{h}_{0,i}, \tilde{\boldsymbol{\beta}} \rangle \\ & \text{subject to } \|\tilde{\boldsymbol{\beta}}\|_{(\mathbf{W}\mathbf{W}^\top)^{1/2}} \leq 1, \end{aligned}$$

where  $\|\mathbf{u}\|_{\mathbf{A}} := \mathbf{u}^\top \mathbf{A}^{-1} \mathbf{u}$  is a norm associated with the positive definite matrix  $\mathbf{A}$ . Clearly the inductive bias by  $\mathbf{W}$  manifests itself via the induced norm  $\|\cdot\|_{(\mathbf{W}\mathbf{W}^\top)^{1/2}}$ . Therefore, it is reasonable to expect that reparametrization due to expansion or shrinkage affects generalization in high dimensions. The question then boils down to provably demonstrate the effect of  $\mathbf{W}$  on the downstream classification accuracy  $\text{Err}(\boldsymbol{\beta}; \eta) := \mathbb{P}(\tilde{y}\langle \tilde{\mathbf{z}}, \boldsymbol{\beta} \rangle < 0)$  with new independent sample  $(\tilde{\mathbf{z}}, \tilde{y})$ .

### 5.4 Precise asymptotic characterization in high dimension

Before presenting the main result, let us define some useful quantities. The solution is denoted by  $\hat{\boldsymbol{\beta}}$  (unique in the separable case). Second, define  $\mathbf{P}_\mu := \rho^{-1} \boldsymbol{\mu} \boldsymbol{\mu}^\top$ ,  $\mathbf{P}_\mu^\perp := \mathbf{I}_p - \rho^{-1} \boldsymbol{\mu} \boldsymbol{\mu}^\top$ , the ratio  $\hat{u} := \|\mathbf{P}_\mu^\perp \mathbf{W}_\eta \hat{\boldsymbol{\beta}}\| / \|\mathbf{P}_\mu \mathbf{W}_\eta \hat{\boldsymbol{\beta}}\|$ . We have the following theoretical guarantees.

**Theorem 5.2.** *Suppose that  $n/p \rightarrow \delta > 0$ ,  $\mathbf{W}_\eta \in \mathcal{W}$ , and  $\|\boldsymbol{\mu}\|$  is a constant. There exists  $\delta^*(\rho) > 0$  such that the following holds.*

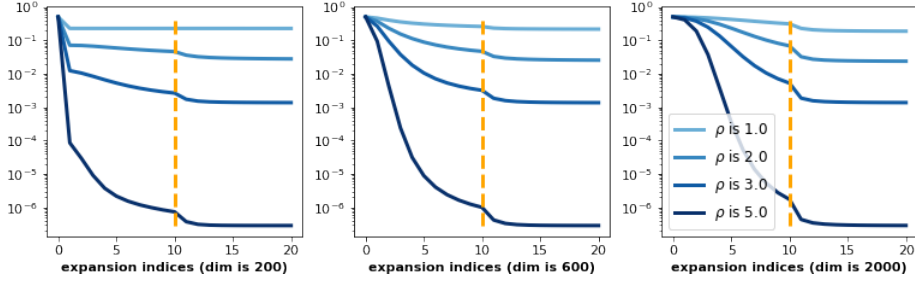


Figure 7: Test error is monotone in  $\eta$  under the one-parameter projector model. We apply a linear transform  $\mathbf{W} \in \mathcal{W}$  on train data from 2-GMM, and calculate the test errors of the max-margin classifiers. Note  $\eta = 0$  (vertical dashed) is equivalent to no transform.

1. (non-separability) If  $\delta > \delta^*(\rho)$ , then with probability approaching one,  $(\mathbf{h}_i, y_i)_{i \leq n}$  is not linearly separable and  $\hat{\beta} = \mathbf{0}$ .
2. (separability) If  $\delta < \delta^*(\rho)$ , then with probability approaching one, there exists a unique solution  $\hat{\beta}$  to (5) with the margin

$$\hat{\kappa} = \min_{i \leq n} y_i \langle \mathbf{z}_i, \hat{\beta} \rangle \xrightarrow{p} \kappa^*(\rho, \eta) > 0.$$

3. (monotonicity of error) If  $\delta < \delta^*(\rho)$ ,  $\hat{u} \xrightarrow{p} u^*(\rho, \eta)$  and  $u^*(\rho, \eta)$  is monotonically decreasing in  $\eta$ . Moreover, the test error obeys

$$\text{Err}(\hat{\beta}; \eta) \xrightarrow{p} \Phi \left( -\rho^{1/2} \sqrt{\frac{1}{1+[u^*(\rho, \eta)]^2}} \right),$$

where  $\Phi$  denotes Gaussian CDF. Thus, the asymptotic test error is decreasing in  $\eta$ .

See Appendix C.2 for the proof, and the precise definitions of  $\kappa^*$  and  $u^*$ .

While separability thresholds (claims 1 and 2) are known in similar models (Deng et al., 2022), the third claim (the most interesting and important one) shows that in the high-dimensional regime, even a linear projector can change the test accuracy of a “linear” max-margin classifier, and the test accuracy increases with the expansion strength  $\eta$ . This partially explains the puzzling effect of projectors on downstream performance. We remark in passing that this result is based on a recent technique known as the convex Gaussian minimax theorem (Gordon, 1988; Thrampoulidis et al., 2015).

**Numerical evidence.** In Figure 7, we use the 2-GMM with separation parameter  $\rho \in \{1, 2, 3, 5\}$  to generate a dataset  $(\mathbf{h}_{0,i})_{i \leq n}$  of size  $n = 1000$  and dimension  $p \in \{200, 600, 2000\}$ . We apply  $\mathbf{W} = \mathbf{I}_p + \eta \rho^{-1} \boldsymbol{\mu} \boldsymbol{\mu}^\top \in \mathcal{W}$  to obtain embeddings  $(\mathbf{z}_i)_{i \leq n}$ . We treat the mixture membership as labels  $(y_i)_{i \leq n}$  and compute the max-margin classifier on  $(\mathbf{z}_i, y_i)_{i \leq n}$ . For varying  $\rho, p, \eta$ , we report the test error of the max-margin classifier and confirm the monotone decreasing property.

## 6 Extensions to inhomogeneous feature augmentation

So far our analysis applies to scenarios where *either* expansion *or* shrinkage appears. However, in practice, it is not uncommon to encounter situations where both expansion and shrinkage appear. This does not render our previous analysis vacuous. As we will demonstrate in this section, we



can extend the previous analysis to the case with inhomogeneous augmentations, under which both expansion and shrinkage can appear. Our treatment throughout this section parallels that in Section 4.

## 6.1 Feature augmentation with spiked covariance

Suppose that features are generated from the same 2-GMM but augmentations are inhomogeneous:

$$\mathbf{h}_{0,i} \stackrel{\text{i.i.d.}}{\sim} \frac{1}{2}\mathcal{N}(-\boldsymbol{\mu}, \mathbf{I}_p) + \frac{1}{2}\mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p), \quad \text{and} \quad \mathbf{h}_i, \mathbf{h}_i^+ | \mathbf{h}_{0,i} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{h}_{0,i}, \sigma_{\text{aug}}^2 \mathbf{A}),$$

where  $\mathbf{A} \succeq \mathbf{I}_p$  is covariance matrix. Throughout this section, we assume  $p > 2$ . This inhomogeneous model is supported by empirical evidence. In the appendix, we show that image-level augmentation (random cropping, color distortion, etc.) does produce inhomogeneous features, which lead to more realistic phenomena.

Now following the same setup as before: we consider linear projectors  $\mathbf{z} = \mathbf{W}\mathbf{h}$ , the modified SimCLR loss (2), the population counterpart  $\mathcal{L}(\mathbf{W}) = \mathbb{E}_{\mathbf{z}} \mathcal{L}_n(\mathbf{W})$  and its approximation  $\tilde{\mathcal{L}}(\mathbf{W})$ . To ease the notation, let  $\tilde{\mathbf{W}} = (\mathbf{I} + \sigma_{\text{aug}}^2)^{1/2} \mathbf{W}$ ,  $\tilde{\boldsymbol{\mu}} = (\mathbf{I} + \sigma_{\text{aug}}^2)^{-1/2} \boldsymbol{\mu}$  and  $\sum_{j=1}^p \tilde{\sigma}_j^2 \tilde{\mathbf{u}}_j^\top \tilde{\mathbf{v}}_j$  be the SVD of  $\tilde{\mathbf{W}}$ . The following proposition gives the explicit formulas of  $\mathcal{L}(\mathbf{W})$  and  $\tilde{\mathcal{L}}(\mathbf{W})$ .

**Proposition 6.1.** *Define*

$$\tilde{\alpha} := \tilde{\alpha}(\mathbf{W}) = \mathbb{E} \|\mathbf{W}\mathbf{h}\|^2 = \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2 + \sigma_{\text{aug}}^2 \text{Tr}(\mathbf{W}^\top \mathbf{A} \mathbf{W}).$$

The loss  $\mathcal{L}(\mathbf{W})$  takes the form  $\mathcal{L}(\mathbf{W}) = \mathcal{L}_{\text{align}}(\mathbf{W}) + \mathcal{L}_{\text{unif}}(\mathbf{W})$ , where

$$\begin{aligned} \mathcal{L}_{\text{align}}(\mathbf{W}) &= -\frac{1}{\tau \tilde{\alpha}} \|\mathbf{W}\|_{\text{F}}^2 \\ \mathcal{L}_{\text{unif}}(\mathbf{W}) &= -\frac{1}{2} \sum_{j=1}^p \log \left( 1 + \frac{2\tilde{\sigma}_j^2}{\tau \tilde{\alpha}} \right) + \log \left( 1 + \exp \left( -\sum_{j=1}^p \frac{2\tilde{\sigma}_j^2 \langle \tilde{\boldsymbol{\mu}}, \tilde{\mathbf{v}}_j \rangle^2}{2\tilde{\sigma}_j^2 + \tau \tilde{\alpha}} \right) \right) - \log 2. \end{aligned}$$

Similar to before, one can consider a first-order approximation  $\tilde{\mathcal{L}}(\mathbf{W})$ :

$$\tilde{\mathcal{L}}(\mathbf{W}) = -\frac{1}{\tau \tilde{\alpha}} \|\mathbf{W}\|_{\text{F}}^2 + \log \left( 1 + \exp \left( -\frac{2\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau \tilde{\alpha}} \right) \right) - \log 2. \quad (6)$$

The loss  $\tilde{\mathcal{L}}(\mathbf{W})$  depends on  $\mathbf{W}$  through  $\|\mathbf{W}\|$ ,  $\|\mathbf{W}\boldsymbol{\mu}\|$  and  $\|\mathbf{W}\mathbf{v}_{\text{aug}}\|$ . We further consider a simple one-spike covariance model

$$\mathbf{A} = \mathbf{I}_p + \rho_{\text{aug}} \mathbf{v}_{\text{aug}} \mathbf{v}_{\text{aug}}^\top, \quad \text{where } \|\mathbf{v}_{\text{aug}}\| = 1. \quad (7)$$

Here  $\rho_{\text{aug}} \geq 0$  quantifies the strength of the spike in augmentations. In particular, setting  $\rho_{\text{aug}} = 0$  recovers the homogeneous case. It is natural to expect that  $\text{span}\{\boldsymbol{\mu}, \mathbf{v}_{\text{aug}}\}$  plays a critical role in its minimization. For that purpose, let us define an orthogonal basis in  $\text{span}\{\boldsymbol{\mu}, \mathbf{v}_{\text{aug}}\}$ . Let  $\tilde{\boldsymbol{\mu}} = \boldsymbol{\mu}/\|\boldsymbol{\mu}\|$ ,  $r = \langle \tilde{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}} \rangle$ , and  $\boldsymbol{\mu}_\perp$  be a vector such that  $\|\boldsymbol{\mu}_\perp\| = 1$ ,  $\langle \boldsymbol{\mu}_\perp, \boldsymbol{\mu} \rangle = 0$  and  $r \langle \boldsymbol{\mu}_\perp, \mathbf{v}_{\text{aug}} \rangle \geq 0$ .

**Theoretical prediction.** On the surface, the loss  $\tilde{\mathcal{L}}(\mathbf{W})$  here is not a univariate function. It turns out, as our proof reveals, that minimization of  $\tilde{\mathcal{L}}(\mathbf{W})$  is equivalent to a univariate minimization problem that depends only on  $T = \|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2/\|\mathbf{W}\|_{\mathbb{F}}^2$ . Define the threshold

$$\tau_1^* = \frac{2(1-r^2)\|\boldsymbol{\mu}\|^2}{\log(1+2\sigma_{\text{aug}}^2)(1+\sigma_{\text{aug}}^2+(1-r^2)\|\boldsymbol{\mu}\|^2)}.$$

We have the following result that characterizes the phase transition of the minimizer  $\mathbf{W}^*$ .

**Theorem 6.2.** (*phase transition under spiked augmentation*) Consider the one-spike inhomogeneous model (7) and the approximate loss (6). Recall the definition  $T := T(\mathbf{W}) = \|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2/\|\mathbf{W}\|_{\mathbb{F}}^2$ . Let  $\mathbf{W}^*$  be a minimizer of  $\tilde{\mathcal{L}}(\mathbf{W})$  and  $T^* = T(\mathbf{W}^*)$ . Then,  $T^*$  is given by the minimization problem

$$\min_{T \in [0,1]} \left( -\frac{1}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 [(r\sqrt{T} - \sqrt{1-r^2}\sqrt{1-T})_+]^2 \right]} + \log \left( 1 + \exp \left( -\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 [(r\sqrt{T} - \sqrt{1-r^2}\sqrt{1-T})_+]^2 \right]} \right) \right) \right).$$

where  $x_+$  denotes  $\max\{x, 0\}$ . Under  $\rho_{\text{aug}} > 0$  and  $0 < |r| < 1$ , we have

- if  $\tau > \tau_1^*$ , then the minimizer is attained at  $T^* \in (1-r^2, 1)$ , which is associated with a rank-one projector  $\mathbf{W}^*$  given by

$$(\mathbf{W}^*)^\top \mathbf{W}^* = (\sqrt{T^*}\bar{\boldsymbol{\mu}} - \sqrt{1-T^*}\boldsymbol{\mu}_\perp)(\sqrt{T^*}\bar{\boldsymbol{\mu}} - \sqrt{1-T^*}\boldsymbol{\mu}_\perp)^\top;$$

- if  $\tau \leq \tau_1^*$ , the minimizer is attained at  $T^* = \frac{\tau(1+\sigma_{\text{aug}}^2)}{2\|\boldsymbol{\mu}\|^2} \log(1+2\sigma_{\text{aug}}^2) \left[1 - \frac{\tau}{2} \log(1+2\sigma_{\text{aug}}^2)\right]^{-1}$  and  $\mathbf{W}^* \mathbf{v}_{\text{aug}} = 0$ .

See Appendix D.2 for the proof.

If  $\rho_{\text{aug}} = 0$  in the optimization problem of Theorem 6.2, then the analysis reduces to the homogeneous case. However, with *any small* spike strength  $\rho_{\text{aug}} > 0$  and assuming nondegeneracy ( $0 < |r| < 1$ ), the phase transition becomes qualitatively very different.

- The phase threshold is smaller as  $\tau_1^* \leq \tau^*$ , which we recall that

$$\tau^* = 2\|\boldsymbol{\mu}\|^2 \left\{ (1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2) \log(1 + 2\sigma_{\text{aug}}^2) \right\}^{-1}.$$

The difference depends on the cosine angle between the signal  $\bar{\boldsymbol{\mu}}$  and spike direction  $\mathbf{v}_{\text{aug}}$ , irrespective of the spike strength  $\rho_{\text{aug}}$ .

- There is no perfect expansion along the signal direction  $\bar{\boldsymbol{\mu}}$ , as  $T^* < 1$  is *always true*.

**Degenerate cases.** It is beneficial to consider two degenerate examples:  $r = 0$  and  $r = 1$ . When  $r = 0$ , the signal direction and spike direction are orthogonal; and when  $r = 1$ , they are perfectly aligned.

Analyzing the optimization problem in Theorem 6.2 for the degenerate examples yields the following characterization.




Figure 8: GMM with inhomogeneous augmentation. (a) **Left**:  $p = 100$  singular values of  $\mathbf{W}$  in descending order with varying  $\tau$  indicated by different colors. (a) **Middle**: Cumulative score for  $\mu$ , i.e.,  $\text{score}_i(\mu) = \sum_{j \leq i} \langle \mu / \|\mu\|, \mathbf{v}_j \rangle^2$ . (a) **Right**: Cumulative score for  $\mathbf{v}_{\text{aug}}$ , i.e.,  $\text{score}_i(\mathbf{v}_{\text{aug}}) = \sum_{j \leq i} \langle \mathbf{v}_{\text{aug}}, \mathbf{v}_j \rangle^2$ . (b) Phase transition showing  $T^*$  against  $\tau$ : theoretical prediction (dashed curve) versus empirical values (circles).




Figure 9: An illustration showing the phase changes in the one-spike inhomogeneous model as we increase  $\tau$ . In the figure ‘E’ is a shorthand for expansion and ‘S’ for shrinkage; see Corollary 6.3.

1. When  $r = 0$  (namely  $\bar{\mu} \perp \mathbf{v}_{\text{aug}}$ ), we have  $\tau_1^* = \tau^*$ , and  $\mathbf{W}^* \mathbf{v}_{\text{aug}} = 0$ , namely pure shrinkage along  $\mathbf{v}_{\text{aug}}$ . Moreover, similar to the homogeneous case, we have shrinkage along  $\bar{\mu}$  if  $\tau \leq \tau^*$  and expansion along  $\bar{\mu}$  if  $\tau > \tau^*$ .
2. When  $|r| = 1$  (namely  $\bar{\mu} \parallel \mathbf{v}_{\text{aug}}$ ), we have expansion along  $\bar{\mu}$  if

$$\tau \cdot \log \left( \frac{2(1 + \sigma_{\text{aug}}^2) \|\mu\|^2}{\|\mu\|^2 + \rho_{\text{aug}} \sigma_{\text{aug}}^2} - 1 \right) > \frac{2 \|\mu\|^2}{1 + \sigma^2 + \|\mu\|^2 + \rho_{\text{aug}} \sigma_{\text{aug}}^2}.$$

In the first example, expansion/shrinkage operates independently along the two directions  $\bar{\mu}$  and  $\mathbf{v}_{\text{aug}}$ . In the second example, the presence of a parallel spike produces a threshold that is nonlinear in  $\rho_{\text{aug}}$  for expansion/shrinkage along  $\bar{\mu}$ .

**Numerical experiments.** Figure 8a shows the singular vectors of  $\mathbf{W}^*$  after training with the SimCLR loss and also the cumulative scores with  $\bar{\mu}$  and  $\mathbf{v}_{\text{aug}}$ , respectively. We vary the temperature  $\tau$ , and fix  $\rho_{\text{aug}} = 5$ ,  $\sigma_{\text{aug}} = 0.5$ , and  $r = \langle \bar{\mu}, \mathbf{v}_{\text{aug}} \rangle = 0.5$ . When  $\tau \leq 2$ , there is consistent shrinkage in the direction of  $\mathbf{v}_{\text{aug}}$ . In addition, for example, when  $\tau = 2$  or 0.795, expansion and shrinkage in the direction of  $\mu$  coexist and  $\mu$  is only spanned by either the top singular vector or the bottom singular vector of  $\mathbf{W}^*$ , which is consistent with Figure 3.

## 6.2 Simultaneous expansion/shrinkage

The signal  $\mu$  and spike  $\mathbf{v}_{\text{aug}}$  exhibit different levels of expansion/shrinkage at the same temperature parameter, which sometimes leads to simultaneous expansion/shrinkage.

Consider the one-spike inhomogeneous model (7) and the approximate loss (6). Assume non-degeneracy  $\rho_{\text{aug}} > 0$  and  $0 < |r| < 1$ . Recall the SVD of  $\mathbf{W}^*$  is  $\sum_{j \leq p} \sigma_j \mathbf{u}_j \mathbf{v}_j^\top$ . Without loss of generality we assume  $\|\mathbf{W}^*\|_F = 1$ .

**Corollary 6.3** (Phase change under varying  $\tau$ ). *As we increase the temperature parameter  $\tau$ , treating other parameters  $\|\boldsymbol{\mu}\|, \sigma_{\text{aug}}^2, \rho_{\text{aug}}, r$  as constants, we experience the following different phases.*

1. When  $\tau \ll 1$ : both shrinkage. We have  $\mathbf{W}^* \mathbf{v}_{\text{aug}} = 0$ , and also  $\|\mathbf{W}^* \bar{\boldsymbol{\mu}}\| \ll 1$ , the latter of which implies  $\max_{1 \leq j \leq p} |\sigma_j \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle| \ll 1$ .
2. When  $\omega(1) \leq \tau \leq \tau_1^*$ : simultaneous expansion and shrinkage. We have  $\mathbf{W}^* \mathbf{v}_{\text{aug}} = 0$ , and also  $\|\mathbf{W}^* \bar{\boldsymbol{\mu}}\| \geq c$  for certain constant  $c > 0$ . There are two jumps in the cumulative score: for certain dimension-free constant  $\epsilon \in (0, 1)$ , we have  $\sum_{j \leq \epsilon^{-1}} \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 \geq \epsilon$  and  $\langle \mathbf{v}_p, \bar{\boldsymbol{\mu}} \rangle^2 \geq \epsilon$ .
3. When  $O(1) \geq \tau > \tau_1^*$ : both expansion.  $\mathbf{W}^*$  is a rank-one matrix, and its right singular vector  $\sqrt{T^*} \bar{\boldsymbol{\mu}} - \sqrt{1 - T^*} \boldsymbol{\mu}_\perp$  has positive cosine angles with both  $\bar{\boldsymbol{\mu}}$  and  $\mathbf{v}_{\text{aug}}$ .
4. When  $\tau \gg 1$ : expansion increasingly aligns with  $\bar{\boldsymbol{\mu}}$ . Note that  $\sqrt{T^*}$  is very close to 1 but always strictly smaller than 1.

See Appendix D.3 for the proof.

Figure 9 gives an illustration of the phase change in the above corollary. It offers a theoretical explanation for the singular value/vector plots for STL-10 data (Figure 3) and simulated GMM data (Figure 8a).

## 7 Discussion and related work

### 7.1 Connections to emerging empirical phenomena

The puzzles about projectors in contrastive learning echo several known phenomena in deep learning.

**Dimensional collapse.** It is often observed that the trained features and embeddings do not span the entire ambient space. To be more precise, the singular values of the feature matrix  $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_n]^\top$  and the trained embedding matrix  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_n]^\top$  (with  $\mathbf{h}_i = \mathbf{f}_\theta(\mathbf{x}_i)$  and  $\mathbf{z}_i = \mathbf{g}_\varphi(\mathbf{h}_i)$ ) contain one or more (approximate) zeros. This phenomenon is known as *dimensional collapse*, which has been repeatedly reported in the literature (Chen et al., 2020a; Jing et al., 2021; Balestrieri and LeCun, 2022).<sup>4</sup> When dimensional collapse occurs in the feature space, we obtain less informative representations. This is generally undesirable according to these papers and requires careful handling due to its adverse effects on generalization. Figure 3 (left) confirms dimensional collapse by showing that the linear projector does not have full rank.

To address dimensional collapse in contrastive learning (and more so in non-contrastive SSL), a line of work proposes to refine loss functions and design structured projectors (Balestrieri et al., 2023), but a systematic treatment is still lacking.

<sup>4</sup>In fact, dimensional collapse is a more salient issue for non-contrastive approaches in SSL (Hua et al., 2021; Tian et al., 2021) due to the lack of negative pairs.

**Transferability of intermediate-layer features.** In supervised learning, it is well observed that trained deep neural networks contain interpretable features that become progressively complex when moving up layers (Zeiler and Fergus, 2014). Therefore, it is natural to use intermediate-layer features pretrained on large datasets for related tasks (Yosinski et al., 2014). At the very top layers, features are believed to be very specific to a classification task, and thus they need to be finetuned on downstream tasks.

Intuitively, projectors bear similarity to those top layers that require finetuning on downstream tasks. For both supervised learning and contrastive learning, minimizing a specific loss seems to reduce the information in features and thus their generality.

**Neural collapse.** In supervised learning, the features in the penultimate layer tend to form a symmetric structure, if we train neural networks many epochs well past the terminal phase where the train error achieves zero. This phenomenon is known as *neural collapse* (Papayan et al., 2020). In short, as training evolves, the penultimate features gradually collapse to their respective class means, which form an equiangular simplex. The highly symmetric and compact cluster structure is observed on the training dataset and rarely on the test dataset (unless the test error is also zero).

Contrastive learning produces weaker cluster structure but the learned features are more general for downstream tasks (Wang and Isola, 2020). The cluster structure resulting from contrastive learning is desirable for generalization, whereas the strong cluster structure resulting from cross-entropy minimization is partly due to the optimization artifact. In fact, a fine-grained intermediate-layer neural collapse suggests that top layers (including the penultimate layer) do not improve and sometimes even harm the generalization properties of features (Galanti et al., 2022).

## 7.2 Related work

**Analysis of contrastive learning.** Contrastive learning has received tremendous attention in the past few years, and a large body of work has been done around this topic. We refer interested readers to the recent overview (Balestriero et al., 2023) for the historical account and for the recent updates. Though empirically successful, contrastive learning also brings various intriguing phenomena including dimensional collapse (Jing et al., 2021; Hua et al., 2021) and behavior of projectors (Chen et al., 2020a,b; Cosentino et al., 2022). These motivate quite a few recent theoretical attempts to explain the success of contrastive learning. Most notable and related is the paper by Wang and Isola (2020), where they discover that contrastive loss promotes both alignment and uniformity of the learned representations. This viewpoint is instrumental in our analysis and understanding, as can be seen from e.g., Proposition 4.1. Our work goes beyond alignment/uniformity by precisely characterizing the effect (namely, expansion and shrinkage) of contrastive loss on the projection head. Jing et al. (2021) also studies the role of projectors by arguing that they prevent dimensional collapse in the representation space. However, no theoretical study of the generalization property is provided. In addition, several recent papers (HaoChen et al., 2021; Wen and Li, 2021; Ji et al., 2021; Lee et al., 2021; Wen and Li, 2022; Von Kügelgen et al., 2021; Saunshi et al., 2022) theoretically study contrastive learning without focusing on the role of projectors.

**Implicit bias and interpolating models.** It is recently discovered that in over-parametrized models (i.e., interpolating models), gradient descent (GD) type algorithms have implicit regularization effects on the model parameters. Relevant to our paper are the results of Soudry et al. (2018) and Gunasekar et al. (2018), where it is proved that for linearly separable data, GD iterates converge in direction to the *max-margin classifier*. It is also important to characterize the generalization error of the solutions with implicit bias. A line of relevant papers include Belkin et al. (2019);

Bartlett et al. (2020); Hastie et al. (2022); Liang and Rakhlin (2020); Bartlett et al. (2020); Hastie et al. (2022); Montanari et al. (2019); Montanari and Zhong (2022); Deng et al. (2022); Mei and Montanari (2022); Liang and Sur (2022); Montanari et al. (2021). Closely related to our Section 5 is the recent paper by Deng et al. (2022), but their main purpose is explaining the double-descent phenomenon rather than studying the effects for expansion/shrinkage.

## References

- Arora, S., Cohen, N., Golowich, N., and Hu, W. (2018). A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*.
- Balestrieri, R., Ibrahim, M., Sobal, V., Morcos, A., Shekhar, S., Goldstein, T., Bordes, F., Bardes, A., Mialon, G., Tian, Y., et al. (2023). A cookbook of self-supervised learning. *arXiv preprint arXiv:2304.12210*.
- Balestrieri, R. and LeCun, Y. (2022). Contrastive and non-contrastive self-supervised learning recover global and local spectral embedding methods. *arXiv preprint arXiv:2205.11508*.
- Bardes, A., Ponce, J., and LeCun, Y. (2021). Vicreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*.
- Bartlett, P. L., Long, P. M., Lugosi, G., and Tsigler, A. (2020). Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070.
- Belkin, M., Hsu, D., Ma, S., and Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854.
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Bietti, A., Bruna, J., Sanford, C., and Song, M. J. (2022). Learning single-index models with shallow neural networks. *arXiv preprint arXiv:2210.15651*.
- Böhm, J. N., Berens, P., and Kobak, D. (2022). Unsupervised visualization of image datasets using contrastive learning. *arXiv preprint arXiv:2210.09879*.
- Bourlard, H. and Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, 59(4):291–294.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020a). A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., and Hinton, G. E. (2020b). Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255.
- Cosentino, R., Sengupta, A., Avestimehr, S., Soltanolkotabi, M., Ortega, A., Willke, T., and Tepper, M. (2022). Toward a geometrical understanding of self-supervised contrastive learning. *arXiv preprint arXiv:2205.06926*.



- Deng, Z., Kammoun, A., and Thrampoulidis, C. (2022). A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 11(2):435–495.
- Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., and Zisserman, A. (2021). With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9568–9577.
- Galanti, T., Galanti, L., and Ben-Shaul, I. (2022). On the implicit bias towards minimal depth of deep neural networks. *arXiv preprint arXiv:2202.09028*.
- Ghojogh, B., Crowley, M., Karray, F., and Ghodsi, A. (2023). *Elements of Dimensionality Reduction and Manifold Learning*. Springer Nature.
- Gordon, Y. (1988). On milman’s inequality and random subspaces which escape through a mesh in  $\mathbb{R}^n$ . In *Geometric aspects of functional analysis*, pages 84–106. Springer.
- Gunasekar, S., Lee, J., Soudry, D., and Srebro, N. (2018). Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pages 1832–1841. PMLR.
- Han, X., Pappas, V., and Donoho, D. L. (2021). Neural collapse under mse loss: Proximity to and dynamics on the central path. *arXiv preprint arXiv:2106.02073*.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. (2021). Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507.
- Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417.
- Hua, T., Wang, W., Xue, Z., Ren, S., Wang, Y., and Zhao, H. (2021). On feature decorrelation in self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9598–9608.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr.
- Jennrich, R. I. (1969). Asymptotic properties of non-linear least squares estimators. *The Annals of Mathematical Statistics*, 40(2):633–643.

- Ji, W., Deng, Z., Nakada, R., Zou, J., and Zhang, L. (2021). The power of contrast for feature learning: A theoretical analysis. *arXiv preprint arXiv:2110.02473*.
- Jing, L., Vincent, P., LeCun, Y., and Tian, Y. (2021). Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348*.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lee, J. D., Lei, Q., Saunshi, N., and Zhuo, J. (2021). Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323.
- Liang, T. and Rakhlin, A. (2020). Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 48(3):1329–1347.
- Liang, T. and Sur, P. (2022). A precise high-dimensional asymptotic theory for boosting and minimum- $\ell_1$ -norm interpolated classifiers. *The Annals of Statistics*, 50:1669–1695.
- Mei, S. and Montanari, A. (2022). The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766.
- Misra, I. and van der Maaten, L. (2019). Self-supervised learning of pretext-invariant representations. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6706–6716.
- Montanari, A., Ruan, F., Sohn, Y., and Yan, J. (2019). The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*.
- Montanari, A. and Zhong, Y. (2022). The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *The Annals of Statistics*, 50(5):2816–2847.
- Montanari, A., Zhong, Y., and Zhou, K. (2021). Tractability from overparametrization: The example of the negative perceptron. *arXiv preprint arXiv:2110.15824*.
- Müller, R., Kornblith, S., and Hinton, G. E. (2019). When does label smoothing help? *Advances in neural information processing systems*, 32.
- Newey, W. K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245.
- Neyshabur, B., Tomioka, R., and Srebro, N. (2014). In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*.
- Papayan, V., Han, X., and Donoho, D. L. (2020). Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663.
- Pearson, K. (1901). Principal components analysis. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 6(2):559.

- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Rosset, S., Zhu, J., and Hastie, T. (2003). Margin maximizing loss functions. In Thrun, S., Saul, L., and Schölkopf, B., editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press.
- Saunshi, N., Ash, J., Goel, S., Misra, D., Zhang, C., Arora, S., Kakade, S., and Krishnamurthy, A. (2022). Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. (2018). The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878.
- Thrapoulidis, C., Oymak, S., and Hassibi, B. (2015). Regularized linear regression: A precise analysis of the estimation error. In Grünwald, P., Hazan, E., and Kale, S., editors, *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 1683–1709, Paris, France. PMLR.
- Tian, Y. (2022). Understanding deep contrastive learning via coordinate-wise optimization. In *Advances in Neural Information Processing Systems*.
- Tian, Y., Chen, X., and Ganguli, S. (2021). Understanding self-supervised learning dynamics without contrastive pairs. *ArXiv*, abs/2102.06810.
- Von Kügelgen, J., Sharma, Y., Gresele, L., Brendel, W., Schölkopf, B., Besserve, M., and Locatello, F. (2021). Self-supervised learning with data augmentations provably isolates content from style. *Advances in neural information processing systems*, 34:16451–16467.
- Wang, T. and Isola, P. (2020). Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR.
- Wen, Z. and Li, Y. (2021). Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pages 11112–11122. PMLR.
- Wen, Z. and Li, Y. (2022). The mechanism of prediction head in non-contrastive self-supervised learning. *arXiv preprint arXiv:2205.06226*.
- Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S., and Pennington, J. (2018). Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, pages 5393–5402. PMLR.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. (2014). How transferable are features in deep neural networks? *Advances in neural information processing systems*, 27.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR.
- Zeiler, M. D. and Fergus, R. (2014). Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer.

# Contents

<b>A Experiments: details and extensions</b>	<b>26</b>
A.1 Experiment setup and details	26
A.2 Additional experiments	28
<b>B Proofs for Section 4</b>	<b>33</b>
B.1 Proof of Proposition 4.1	33
B.2 Justification of approximate loss	35
B.3 Proof of Theorem 4.3	36
B.4 Proof of Proposition 4.4	37
B.4.1 A sandwich formula	37
B.4.2 Approximation bound in expansion regime	38
B.4.3 Approximation bound in shrinkage regime	39
B.5 Proof of Proposition 4.5	42
<b>C Proofs for Section 5</b>	<b>44</b>
C.1 Proof of Theorem 5.1	44
C.1.1 Proof of Lemma C.1	48
C.1.2 Proof of Lemma C.2	50
C.1.3 Proof of Lemma C.3	50
C.2 High-dimensional regime: max-margin classifier	51
C.2.1 Preliminaries	51
C.2.2 Proof of Theorem 5.2 (1).	54
C.2.3 Proof of Theorem 5.2 (2).	55
C.2.4 Proof of Theorem 5.2 (3).	56
C.2.5 Proof of Lemma C.6	57
C.2.6 Proof of Lemma C.7	58
C.2.7 Proof of Lemma C.8	59
C.2.8 Proof of Lemma C.9	60
C.2.9 Proof of Lemma C.10	64
C.2.10 Proof of Lemma C.11	67
<b>D Proof for Section 6</b>	<b>69</b>
D.1 Proof of Proposition 6.1	69
D.2 Proof of Theorem 6.2	70
D.3 Proof of Corollary 6.3	74

## A Experiments: details and extensions

**Reproducibility.** Our code and data are included in the supplemental materials.

### A.1 Experiment setup and details

**Fixed encoder network.** In Figure 3, we freeze the encoder network (ResNet-18) trained and saved in <https://github.com/sthalles/SimCLR> to focus on the behavior of projector under SimCLR loss. The pretrained architecture is trained with the default temperature  $\tau = 0.07$  and the following composite augmentation

```

1 class TransformsSimCLR:
2     """
3     A stochastic data augmentation module that transforms any given data example
4     randomly
5     resulting in two correlated views of the same example,
6     denoted  $x_i$  and  $x_j$ , which we consider as a positive pair.
7     """
8     def __init__(self, size, aug_str=1):
9         color_jitter = torchvision.transforms.ColorJitter(
10            0.8 * aug_str, 0.8 * aug_str, 0.8 * aug_str, 0.2 * aug_str
11        )
12        self.train_transform = torchvision.transforms.Compose(
13            [
14                torchvision.transforms.RandomResizedCrop(size=size),
15                torchvision.transforms.RandomHorizontalFlip(),
16                torchvision.transforms.RandomApply([color_jitter], p=0.8),
17                torchvision.transforms.RandomGrayscale(p=0.2),
18                torchvision.transforms.ToTensor(),
19            ]
20        )
21
22        self.test_transform = torchvision.transforms.Compose(
23            [
24                torchvision.transforms.Resize(size=size),
25                torchvision.transforms.ToTensor(),
26            ]
27        )
28
29    def __call__(self, x):
30        return self.train_transform(x), self.train_transform(x)

```

Listing 1: Image augmentation

**Training details.** For simplicity, we extract and then center pretrained features of 10-class STL-10 images (which conforms to the zero-mean assumption in 2-GMM). We train the projector  $\mathbf{W} \in \mathbb{R}^{512 \times 512}$  for 50 epochs with the batch size of 64.

In all experiments reported in the paper, we use one-layer linear projector  $\mathbf{W} \in \mathbb{R}^{p \times p}$  without the bias term, and this matrix is initialized by a random orthogonal matrix (orthogonal initialization avoids potential optimization artifacts (Xiao et al., 2018; Arora et al., 2018)). We train the linear projector using the standard SimCLR loss function <https://github.com/sthalles/SimCLR>. For downstream accuracy, we use “linear\_model.LogisticRegression” from `sklearn` with very small  $\ell_2$  regularization (choosing  $C = 1000$ ), with the aim to approximate the max-margin classifier when data are linearly separable (Rosset et al., 2003).

To assess the behavior of  $\mathbf{W}$  and the transition between expansion and shrinkage regimes, 10 values of  $\log(\tau)$  are chosen from equi-spaced grids in  $[\log(0.01), \log(10)]$  and two values of the augmentation strength (0.1 and 1.0) are chosen to represent small and moderate augmentation respectively.

**Results for other pairs and deeper encoder.** Since the projector  $\mathbf{W}$  is trained with the 10-class STL-10 dataset, then the plots for singular values of  $\mathbf{W}$  and downstream accuracy are the same as is shown in Figure 3. In Figure 10, the cumulative sums of alignment scores are shown for different pairs which show similar patterns: as  $\tau$  increases, the expansion effect is gradually gained.




Figure 10: Cumulative sums of alignment scores  $\kappa_{j,c_1,c_2} = \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1,c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1,c_2}\|^2$  for different pairs with the pretrained encoder and a one-layer linear projector  $\mathbf{W} \in \mathbb{R}^{512 \times 512}$  under the standard SimCLR loss (ResNet-18 and 10-class STL-10 dataset).

In addition, we also experimented with ResNet-50 as the encoder instead of ResNet-18. Here we present results using ResNet-50 in Figure 11a, which is similar to Figure 3. We also presented results for different pairs in 11b.

We generate  $n = 2000$  data  $(\mathbf{x}_i)_{i \leq n}$  from 2-GMM with  $p = 100$  and  $\boldsymbol{\mu} = 4\mathbf{e}_1$  where  $(\mathbf{e}_k)_{k \leq p}$  forms the canonical basis. The left plot shows the first two coordinates of these data points, which is equivalent to projecting data onto  $\text{span}(\mathbf{e}_1, \mathbf{e}_2)$ .

We add random perturbation  $\mathcal{N}(0, \sigma_{\text{aug}}^2 \mathbf{I}_p)$  to form augmented data, and then we train a linear projector  $\mathbf{W}$  and the standard SimCLR loss on  $(\mathbf{x}_i)_{i \leq n}$ . After training, we calculate embeddings  $\mathbf{z}_i = \mathbf{W}\mathbf{x}_i$ , project the embeddings onto  $\text{span}(\mathbf{W}\mathbf{e}_1, \mathbf{W}\mathbf{e}_2)$  and visualize these 2D projections. We plot the embeddings under an archetypal expansion regime (middle plot,  $\sigma_{\text{aug}} = 0.1, \tau = 0.2$ ) and an archetypal shrinkage regime (right plot,  $\sigma_{\text{aug}} = 1, \tau = 20$ ).

To aid visualization, we add circles in each of the three plots. Note that the normalized embedding  $\mathbf{z}_i / \|\mathbf{z}_i\|$  is used for calculating the cosine similarities in the SimCLR loss. We can interpret the plots using the alignment vs. uniformity perspective (Wang and Isola, 2020): in the expansion regime, the alignment loss is the dominant term and forces concentrations of normalized embeddings, whereas in the shrinkage regime, the uniformity loss is the dominant term and encourages normalized embeddings to be evenly spread.

## A.2 Additional experiments

**Beyond fixed encoders: training SimCLR from scratch.** To explore the expansion/shrinkage phenomenon without freezing the encoder component, we train the entire architecture (e.g., ResNet-50 encoder and a one-layer linear projector) on STL-10 train dataset with 400 epochs and the batch size of 256. We choose 10 values of  $\log(\tau)$  from equi-spaced grids in  $[\log(0.01), \log(10)]$  and we choose the augmentation as the default value 1.0.

As shown in Figure 12, visibly, the dimensional collapse phenomenon is evident in both the feature space and the projector. Our GMM theory does not apply directly to this scenario since in



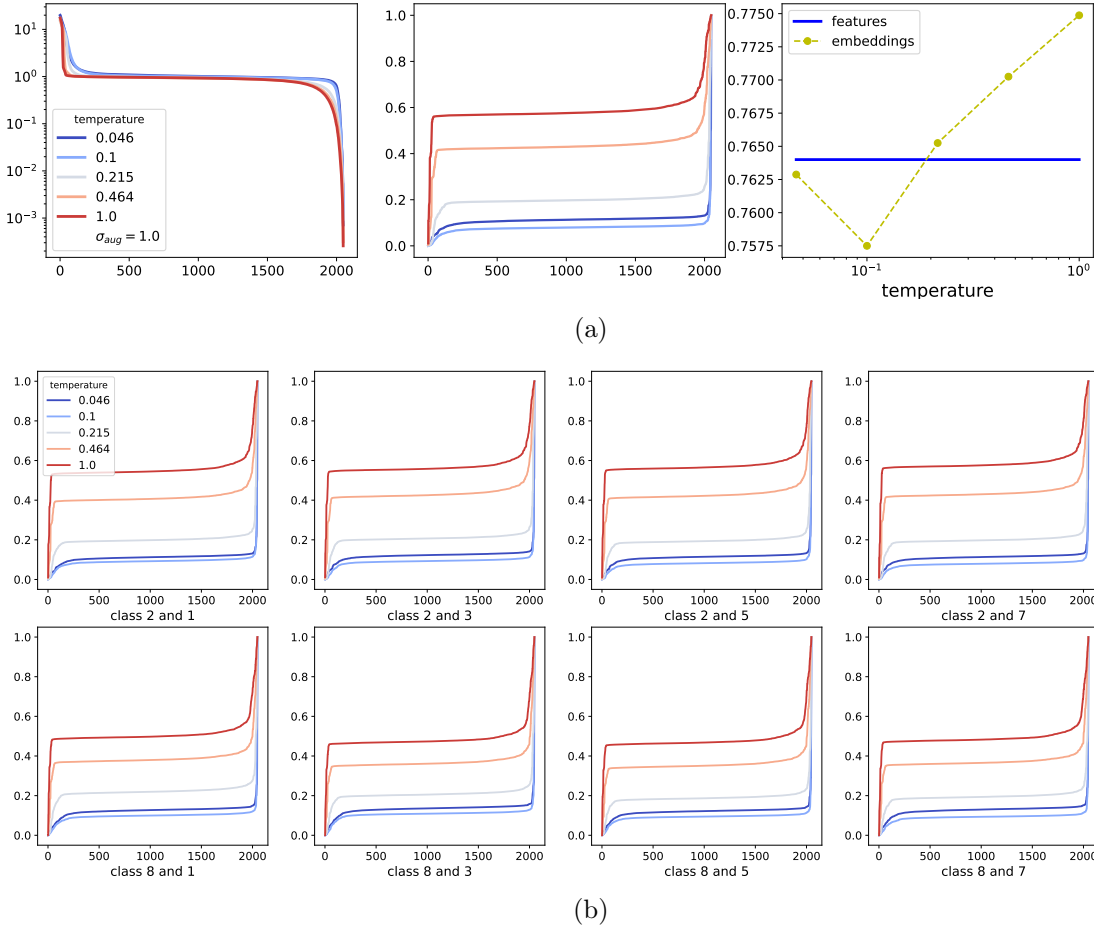


Figure 11: Results with the pretrained encoder and a one-layer linear projector  $\mathbf{W} \in \mathbb{R}^{2048 \times 2048}$  under the standard SimCLR loss (ResNet-50 and 10-class STL-10 dataset). **(a) Left:** singular values of  $\mathbf{W}$  with varying temperature. **(a) Middle:** cumulative sums of alignment scores  $\kappa_{j,c_1,c_2} = \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1,c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1,c_2}\|^2$  for the pair  $(c_1, c_2) = (7, 2)$ . **(a) Right:** comparison between downstream task accuracy with features and embeddings for 10-class evaluation. **(b)** Cumulative sums of alignment scores  $\kappa_{j,c_1,c_2} = \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1,c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1,c_2}\|^2$  for different pairs.

Section 4 we assume that features are generated from a full-dimensional mixture model. Still, our expansion/shrinkage analysis provides partial explanations as summarized below.

1. When we train the encoder and projector simultaneously, their roles and effects are not distinctly separated. Indeed, there are many ways to express  $\mathbf{g}_\varphi \circ \mathbf{f}_\theta$  as function compositions. Thus, the dimensional collapse in the feature space can be interpreted as a shrinkage effect induced by the last few layers in the encoder. Understanding how dimensional collapse emerges progressively across layers is an interesting research direction.
2. Our analysis still provides useful information about downstream accuracy when both the encoder and the projector are trained. For example, when we vary the temperature parameter, the severity of the collapse is correlated with the downstream accuracy; see Figure 12.
3. Our theory matches the empirical results if we restrict the linear transform  $\mathbf{W}$  on the subspace that the feature/embedding vectors span. Figure 13 shows that the singular values/vectors of




Figure 12: Results for SimCLR models trained from scratch. **Left:** comparison between classification accuracy with features (before projection) and embeddings (after projection). **Middle:** singular values of feature matrix  $\mathbf{H}_0 = [\mathbf{h}_{1,0}, \dots, \mathbf{h}_{n,0}]$  with varying temperatures. **Right:** singular values of embedding matrix  $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_n]$  with varying temperatures.

the restricted linear transform. Note that we recover similar cumulative score plots as in the fixed encoder scenario.

Below we provide more detailed explanations for point 2 and 3.

First, from Figure 12 (left), we can see that the downstream accuracy using features is higher than that using embeddings, which validates the practice of using only the features before the projector for classification. When  $\tau \leq 0.129$ , the difference between two curves is decreasing and both achieve the highest value at  $\tau = 0.129$ . However, when  $\tau$  further increases, which disagrees with our previous findings, both accuracy start to decrease. This phenomenon can be explained by the plots of singular values. As we can see from Figure 12 (middle), the features already have dimensional collapse with the one-layer linear projector even when  $\tau$  is small, but the collapse is relatively moderate when  $\tau \leq 0.129$ , which refers to the bunch of curves starting to drop after the index of 1000. When  $\tau > 0.129$ , the collapse becomes much more severe and the effective rank decreases fast below 200 when  $\tau$  goes to 1.0. The singular values of embeddings change accordingly.

The trend in the downstream task accuracy together with the changes in singular values of  $\mathbf{W}$  convey the message that

- when  $\tau$  is moderate, the increase in  $\tau$ , which enhances the expansion of signal (will be shown in the following figures), will improve the downstream task accuracy with embeddings, making it as good as the accuracy with features even when features are undergoing the dimensional collapse;
- when  $\tau$  is large, it poses negative effects in downstream task accuracy in that the features are already low-rank as is shown in Tian (2022), which may lead to the information loss in the data and may further do harm to the training of projector. As a result, the accuracy with either embeddings or features decreases. Also, the benefits from the expansion of signal are surpassed and accuracy with embeddings can be worse than that with features. In contrast, with pretrained model and full-rank features, the accuracy with embeddings can be better than that with features with the benefits from expansion as is shown in Figure 3.

In Figure 13, we first write the SVD of the feature matrix  $\mathbf{H}_0$  as  $\mathbf{H}_0 = \mathbf{U}(\mathbf{H}_0)\mathbf{D}(\mathbf{H}_0)\mathbf{V}^\top(\mathbf{H}_0)$ , where  $\mathbf{D}(\mathbf{H}_0)$  has singular values of  $\mathbf{H}_0$  as diagonal elements:  $\{s_k(\mathbf{H}_0) : 1 \leq k \leq p\}$ . We consider projecting both  $\mathbf{W}$  and  $\boldsymbol{\mu}$  onto the features' top- $r$  right singular subspace, that is we define  $\mathbf{W}^{(r)} = \mathbf{W}\mathbf{V}_r(\mathbf{H}_0)$  and  $\boldsymbol{\mu}^{(r)} = \mathbf{V}_r^\top(\mathbf{H}_0)\boldsymbol{\mu}$  where  $\mathbf{V}_r(\mathbf{H}_0) \in \mathbb{R}^{p \times r}$  is the submatrix consisting of the first




Figure 13: Plots for train-from-scratch experiment with projector  $\mathbf{W}^{(r)} = \mathbf{W}\mathbf{V}_r(\mathbf{H}_0) \in \mathbb{R}^{p \times r}$  and  $\boldsymbol{\mu}^{(r)} = \mathbf{V}_r^\top(\mathbf{H}_0)\boldsymbol{\mu} \in \mathbb{R}^r$  which are both projected onto feature matrix  $\mathbf{H}_0$ 's top- $r$  right singular subspace with  $r = \operatorname{argmax}\{k : s_k(\mathbf{H}_0) \geq 10^{-3}\}$ . **Left:** unsquared alignment scores  $\langle \mathbf{v}_j^{(r)}, \boldsymbol{\mu}_{c_1, c_2}^{(r)} \rangle / \|\boldsymbol{\mu}_{c_1, c_2}^{(r)}\|$  for the pair  $(c_1, c_2) = (7, 2)$ . **Middle:** cumulative sum of alignment scores. **Right:** singular values of  $\mathbf{W}^{(r)}$ .

$r$  columns of  $\mathbf{V}(\mathbf{H}_0)$ . We choose  $r$  by  $r = \operatorname{argmax}\{k : s_k(\mathbf{H}_0) \geq 10^{-3}\}$ . Then, we can instead calculate the SVD of  $\mathbf{W}^{(r)} = \mathbf{U}^{(r)}\mathbf{D}^{(r)}\mathbf{V}^{(r)\top}$ . Denote  $\mathbf{v}_j^{(r)}$  as the  $j$ th right singular vector of  $\mathbf{W}^{(r)}$ .

In the left column, we plot the unsquared alignment score  $\langle \mathbf{v}_j^{(r)}, \boldsymbol{\mu}_{c_1, c_2}^{(r)} \rangle / \|\boldsymbol{\mu}_{c_1, c_2}^{(r)}\|$ , where  $(c_1, c_2) = (7, 2)$ , and the cumulative scores are plotted in the middle column. With the truncated scores, we can see that when  $\tau = 0.077$ , the bottom singular vectors align much better with the  $\boldsymbol{\mu}_{c_1, c_2}^{(r)}$  than top ones (shrinkage regime). As  $\tau$  increases, the expansion and shrinkage effects are comparable to each other with  $\tau = 0.219$ , but when  $\tau$  further increases to 0.599, top singular vectors align better with  $\boldsymbol{\mu}_{c_1, c_2}^{(r)}$  and the expansion effect is dominating. However, expansion effect directly enhances downstream accuracy with full-rank features as is shown in the pretrained experiments and the benefits can be hidden by the dimensional collapse of features as is shown in Figure 12.

**Convergence of the encoder** Recall the decomposition

$$\mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathbf{W}^{(t)}) = \mathcal{L}(\boldsymbol{\theta}^{(t)}, \widetilde{\mathbf{W}}^{(t)}) + \left( \mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathbf{W}^{(t)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)}, \widetilde{\mathbf{W}}^{(t)}) \right),$$

where  $(\boldsymbol{\theta}^{(t)}, \mathbf{W}^{(t)})$  are the parameters at epoch  $t$ , and  $\widetilde{\mathbf{W}}^{(t)} := \operatorname{argmin}_{\mathbf{W}} \mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathbf{W})$ , i.e. the optimal projector with frozen encoder parameters  $\boldsymbol{\theta}^{(t)}$  at epoch  $t$ . We observe that when  $t \geq T_0$  for  $T_0 \approx 50$ , we observe  $\|\widetilde{\mathbf{W}}^{(t)} - \mathbf{W}^{(t)}\| \ll \|\mathbf{W}^{(t)}\|$ .




Figure 14: **Left** the relative change of features and embeddings:  $\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|/\|\boldsymbol{\theta}^{(t)}\|$  and  $\|\mathbf{W}^{(t+1)} - \mathbf{W}^{(t)}\|/\|\mathbf{W}^{(t)}\|$ ; **Right** the standardized distance between  $\mathbf{W}^{(t)}$  and the minimizer  $\widetilde{\mathbf{W}}^{(t)}$ :  $\|\widetilde{\mathbf{W}}^{(t)} - \mathbf{W}^{(t)}\|/\|\mathbf{W}^{(t)}\|$ .




Figure 15: Results with the pretrained encoder and a one-layer linear projector  $\mathbf{W} \in \mathbb{R}^{2048 \times 2048}$  under the standard SimCLR loss (ResNet-50, 10-class STL-10 dataset and feature-level perturbation with  $\sigma_{\text{aug}} = 1.0$ ). **Left**: singular values of  $\mathbf{W}$  with varying temperature. **Middle**: cumulative sums of alignment scores  $\kappa_{j,c_1,c_2} = \langle \mathbf{v}_j, \boldsymbol{\mu}_{c_1,c_2} \rangle^2 / \|\boldsymbol{\mu}_{c_1,c_2}\|^2$  for the pair  $(c_1, c_2) = (7, 2)$ . **Right**: comparison between downstream task accuracy with features and embeddings for 10-class evaluation.

**Feature-level vs. image-level augmentation** Image-level augmentation results in feature-level perturbation in that with the composite augmentation and the fixed encoder  $\mathbf{f}_{\boldsymbol{\theta}_n^*}$ , for each image  $\mathbf{x}_i$ , we have the augmented image  $\mathbf{x}_i^+$  and the mapped features  $\mathbf{h}_i = \mathbf{f}_{\boldsymbol{\theta}_n^*}(\mathbf{x}_i)$ ,  $\mathbf{h}_i^+ = \mathbf{f}_{\boldsymbol{\theta}_n^*}(\mathbf{x}_i^+)$ . Then, the image-level augmentation is associated with the feature-level perturbation  $\mathbf{e}_i = \mathbf{h}_i^+ - \mathbf{h}_i$ . This perturbation is correlated with  $\mathbf{h}_i$  and can be complicated. To make a connection with our 2-GMM theory, we also consider the homogeneous feature-level augmentation and its effect on the expansion/shrinkage phenomenon.

We freeze the encoder network (ResNet-50) as before, then the features are also fixed in this case. For each  $\mathbf{h}_{0,i}$ , we add Gaussian perturbation  $\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \sigma_{\text{aug}}^2 \mathbf{I}_p)$  independently to obtain  $\mathbf{h}_i^+ = \mathbf{h}_{0,i} + \boldsymbol{\epsilon}_i$ . Then, with the same training process, we train the projector  $\mathbf{W} \in \mathbb{R}^{2048 \times 2048}$  without the bias term under the standard SimCLR loss for 50 epochs with the batch size of 64. We choose 10 values of  $\log(\tau)$  from equi-spaced grids in  $[\log(0.01), \log(10)]$  and we choose the augmentation as the default value 1.0.

From Figure 15 (middle), different from the results for image-level augmentation, we can see the dominating expansion effect when  $\tau > 0.215$  where the bottom singular vectors of  $\mathbf{W}$  is merely uncorrelated with  $\boldsymbol{\mu}_{c_1,c_2}$  while the top singular values contribute over 90% of the correlation.

## B Proofs for Section 4

### B.1 Proof of Proposition 4.1

Recall that the population loss  $\mathcal{L}(\mathbf{W}) = \mathbb{E}[\mathcal{L}_n(\mathbf{W})]$  obeys

$$\mathcal{L}(\mathbf{W}) := \mathcal{L}_{\text{align}}(\mathbf{W}) + \mathcal{L}_{\text{unif}}(\mathbf{W}),$$

where we denote

$$\begin{aligned} \mathcal{L}_{\text{align}}(\mathbf{W}) &:= \frac{1}{\tau} \mathbb{E}_{\mathbf{h}, \mathbf{h}^+} \left[ \frac{1}{2} \frac{\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^+\|_2^2}{(\mathbb{E}[\|\mathbf{W}\mathbf{h}\|_2^2] \cdot \mathbb{E}[\|\mathbf{W}\mathbf{h}^+\|_2^2])^{1/2}} \right]; \\ \mathcal{L}_{\text{unif}}(\mathbf{W}) &:= \log \left( \mathbb{E}_{\mathbf{h}^-, \mathbf{h}} \left[ \exp \left( -\frac{1}{2\tau} \frac{\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^-\|_2^2}{(\mathbb{E}[\|\mathbf{W}\mathbf{h}\|_2^2] \cdot \mathbb{E}[\|\mathbf{W}\mathbf{h}^-\|_2^2])^{1/2}} \right) \right] \right). \end{aligned}$$

In what follows, we compute  $\mathcal{L}_{\text{align}}(\mathbf{W})$  and  $\mathcal{L}_{\text{unif}}(\mathbf{W})$ .

**Computing  $\mathcal{L}_{\text{align}}(\mathbf{W})$ .** For the first term, we have

$$\mathbb{E}[\|\mathbf{W}\mathbf{h}\|_2^2] = \mathbb{E}[\|\mathbf{W}\mathbf{h}^+\|_2^2] = (1 + \sigma_{\text{aug}}^2) \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2 = \alpha,$$

where the last relation is the definition of  $\alpha$ . Therefore,  $\mathcal{L}_{\text{align}}$  can be simply written as

$$\mathcal{L}_{\text{align}}(\mathbf{W}) = \frac{1}{2\tau\alpha} \mathbb{E}_{\mathbf{h}, \mathbf{h}^+} [\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^+\|_2^2].$$

In addition, since  $\mathbf{h} = \mathbf{h}_0 + \boldsymbol{\epsilon}$ , and  $\mathbf{h}^+ = \mathbf{h}_0 + \boldsymbol{\epsilon}^+$  with  $\boldsymbol{\epsilon}, \boldsymbol{\epsilon}^+ \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{\text{aug}}^2 \mathbf{I})$ , we obtain

$$\mathbf{h} - \mathbf{h}^+ = \boldsymbol{\epsilon} - \boldsymbol{\epsilon}^+ \sim \mathcal{N}(0, 2\sigma_{\text{aug}}^2 \mathbf{I}),$$

which yields

$$\mathbb{E}_{\mathbf{h}, \mathbf{h}^+} [\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^+\|_2^2] = 2\sigma_{\text{aug}}^2 \|\mathbf{W}\|_{\text{F}}^2.$$

This further leads to the simplified expression for  $\mathcal{L}_{\text{align}}(\mathbf{W})$ :

$$\mathcal{L}_{\text{align}}(\mathbf{W}) = \frac{\sigma_{\text{aug}}^2 \|\mathbf{W}\|_{\text{F}}^2}{\tau\alpha}.$$

**Computing  $\mathcal{L}_{\text{unif}}(\mathbf{W})$ .** For the second term  $\mathcal{L}_{\text{unif}}(\mathbf{W})$ , using the definition of  $\alpha$ , we see that

$$\mathcal{L}_{\text{unif}}(\mathbf{W}) = \log \left( \mathbb{E}_{\mathbf{h}^-, \mathbf{h}} \left[ \exp \left( -\frac{\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^-\|_2^2}{2\tau\alpha} \right) \right] \right). \quad (8)$$

Recall that  $\mathbf{h}$  and  $\mathbf{h}^-$  are two independent draws from the Gaussian mixture model, i.e.,

$$\mathbf{h} = y\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2} \mathbf{g}, \quad \mathbf{h}^- = y^-\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2} \mathbf{g}^-,$$

where  $y$  and  $y^-$  are independent Rademacher random variables,  $\mathbf{g}$  and  $\mathbf{g}^-$  are two independent random vectors from  $\mathcal{N}(0, \mathbf{I})$ . This further yields

$$\mathbf{h} - \mathbf{h}^- \stackrel{\text{d}}{=} (y - y^-)\boldsymbol{\mu} + \sqrt{2(1 + \sigma_{\text{aug}}^2)} \mathbf{g}$$

$$\stackrel{d}{=} \frac{1}{2}\delta_0 + \frac{1}{4}\delta_{2\boldsymbol{\mu}} + \frac{1}{4}\delta_{-2\boldsymbol{\mu}} + \sqrt{2(1 + \sigma_{\text{aug}}^2)}\mathbf{g}.$$

As a result, we obtain

$$\begin{aligned} & \mathbb{E}_{\mathbf{h}^-, \mathbf{h}} \left[ \exp \left( -\frac{\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^-\|_2^2}{2\tau\alpha} \right) \right] \\ &= \mathbb{E}_{\mathbf{g}} \left[ \frac{1}{2} \exp \left( -\frac{(1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) + \frac{1}{2} \exp \left( -\frac{\|\sqrt{2}\mathbf{W}\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) \right]. \end{aligned} \quad (9)$$

Let  $\sum_{j=1}^p \sigma_j \mathbf{u}_j \mathbf{v}_j^\top$  be the singular value decomposition of  $\mathbf{W}$ . Then one has

$$\mathbf{W}\mathbf{g} = \sum_{j=1}^p \sigma_j \mathbf{v}_j^\top \mathbf{g} \mathbf{u}_j, \quad \|\mathbf{W}\mathbf{g}\|_2^2 = \sum_{j=1}^p \sigma_j^2 (\mathbf{v}_j^\top \mathbf{g})^2,$$

and

$$\|\sqrt{2}\mathbf{W}\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{W}\mathbf{g}\|_2^2 = \sum_{j=1}^p \sigma_j^2 (\mathbf{v}_j^\top (\sqrt{2}\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{g}))^2.$$

Taking expectations, we arrive at

$$\begin{aligned} \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{(1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) \right] &= \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\sum_{j=1}^p \frac{(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} (\mathbf{v}_j^\top \mathbf{g})^2 \right) \right] \\ &= \prod_{j=1}^p \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} (\mathbf{v}_j^\top \mathbf{g})^2 \right) \right], \end{aligned}$$

where the last equality uses the independence among  $\{\mathbf{v}_j^\top \mathbf{g}\}_{1 \leq j \leq p}$ . Since  $(\mathbf{v}_j^\top \mathbf{g})^2$  is a  $\chi_1^2$  random variable, we can use the moment generating function of  $\chi_1^2$  to obtain

$$\mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} (\mathbf{v}_j^\top \mathbf{g})^2 \right) \right] = \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right)^{-1/2}.$$

Combining the previous two relations yields

$$\mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{(1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) \right] = \prod_{j=1}^p \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right)^{-1/2}. \quad (10)$$

Using a similar decomposition, we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{\|\sqrt{2}\mathbf{W}\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) \right] \\ &= \prod_{j=1}^p \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{\sigma_j^2}{\tau\alpha} (\sqrt{2}\mathbf{v}_j^\top \boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{v}_j^\top \mathbf{g})^2 \right) \right]. \end{aligned} \quad (11)$$

Elementary calculations tell us that for any  $t > 0$  and any  $a \in \mathbb{R}$ , one has  $\mathbb{E}_{u \sim N(0,1)} [\exp(-t(a+u)^2)] = \frac{1}{\sqrt{1+2t}} \exp\left(-\frac{a^2 t}{2t+1}\right)$ . This allows us to simplify each term in (11) as

$$\begin{aligned} & \mathbb{E}_{u \sim N(0,1)} \left[ \exp \left( -\frac{\sigma_j^2}{\tau\alpha} \left( \sqrt{2} \mathbf{v}_j^\top \boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2} u \right)^2 \right) \right] \\ &= \frac{1}{\sqrt{1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha}}} \exp \left( -\frac{2\sigma_j^2}{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha} (\mathbf{v}_j^\top \boldsymbol{\mu})^2 \right). \end{aligned}$$

The previous two displays taken together lead to

$$\begin{aligned} & \mathbb{E}_{\mathbf{g}} \left[ \exp \left( -\frac{\|\sqrt{2}\mathbf{W}\boldsymbol{\mu} + \sqrt{1 + \sigma_{\text{aug}}^2}\mathbf{W}\mathbf{g}\|_2^2}{\tau\alpha} \right) \right] \\ &= \prod_{j=1}^p \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right)^{-1/2} \exp \left( -\sum_{j=1}^p \frac{2\sigma_j^2}{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha} (\mathbf{v}_j^\top \boldsymbol{\mu})^2 \right). \end{aligned} \quad (12)$$

Substitute (10) and (12) into the identities (9) and (8) to see that

$$\begin{aligned} & \mathbb{E}_{\mathbf{h}^-, \mathbf{h}} \left[ \exp \left( -\frac{\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^-\|_2^2}{2\tau\alpha} \right) \right] \\ &= \frac{1}{2} \prod_{j=1}^p \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right)^{-1/2} \left( 1 + \exp \left( -\sum_{j=1}^p \frac{2\sigma_j^2}{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha} (\mathbf{v}_j^\top \boldsymbol{\mu})^2 \right) \right), \end{aligned}$$

and

$$\begin{aligned} \mathcal{L}_{\text{unif}}(\mathbf{W}) &= -\frac{1}{2} \sum_{j=1}^p \log \left( 1 + \frac{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} \right) \\ &\quad + \log \left( 1 + \exp \left( -\sum_j \frac{2\sigma_j^2}{2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha} (\mathbf{v}_j^\top \boldsymbol{\mu})^2 \right) \right) - \log 2. \end{aligned}$$

These complete the proof of Proposition 4.1.

## B.2 Justification of approximate loss

Here we present justification underlying the approximation (3). In the regime where  $\tau\alpha = \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2 \gg \sigma_j^2$  for each  $j$ , we can use the approximations  $\log(1+x) \approx x$  and  $2(1 + \sigma_{\text{aug}}^2)\sigma_j^2 + \tau\alpha \approx \tau\alpha$  to obtain the approximate loss

$$\begin{aligned} \tilde{\mathcal{L}}_{\text{unif}}(\mathbf{W}) &= -\sum_{j=1}^p \frac{(1 + \sigma_{\text{aug}}^2)\sigma_j^2}{\tau\alpha} + \log \left( 1 + \exp \left( -\sum_j \frac{2\sigma_j^2}{\tau\alpha} (\mathbf{v}_j^\top \boldsymbol{\mu})^2 \right) \right) - \log 2 \\ &= -\frac{(1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\|_{\text{F}}^2}{\tau\alpha} + \log \left( 1 + \exp \left( -\frac{2\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\tau\alpha} \right) \right) - \log 2. \end{aligned}$$

This combined with the alignment term results in

$$\begin{aligned}\tilde{\mathcal{L}}(\mathbf{W}) &= \mathcal{L}_{\text{align}}(\mathbf{W}) + \tilde{\mathcal{L}}_{\text{unif}}(\mathbf{W}) \\ &= -\frac{\|\mathbf{W}\|_{\text{F}}^2}{\tau\alpha} + \log\left(1 + \exp\left(-\frac{2\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\tau\alpha}\right)\right) - \log 2.\end{aligned}$$

This loss function turns out to be a simple univariate function. To see this, we define  $t = \|\mathbf{W}\|_{\text{F}}^2/\alpha$ . As a result, we have  $\|\mathbf{W}\boldsymbol{\mu}\|_2^2 = \alpha - (1 + \sigma_{\text{aug}}^2)\alpha t$ , and hence

$$\begin{aligned}\tilde{\mathcal{L}}(\mathbf{W}) &= -\frac{t}{\tau} + \log\left(1 + \exp\left(-\frac{2\alpha[1 - (1 + \sigma_{\text{aug}}^2)t]}{\tau\alpha}\right)\right) - \log 2 \\ &= -\frac{t}{\tau} + \log\left(1 + \exp\left(\frac{2(1 + \sigma_{\text{aug}}^2)t}{\tau} - \frac{2}{\tau}\right)\right) - \log 2.\end{aligned}$$

### B.3 Proof of Theorem 4.3

Recall the approximate loss function

$$\tilde{\mathcal{L}}(\mathbf{W}) = \ell(t) = -\frac{t}{\tau} + \log\left(\frac{1}{2} + \frac{1}{2}\exp\left(-\frac{2}{\tau} + \frac{2(1 + \sigma_{\text{aug}}^2)}{\tau}t\right)\right),$$

where  $t = \|\mathbf{W}\|_{\text{F}}^2/\alpha \in [1/(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2), 1/(1 + \sigma_{\text{aug}}^2)]$  with  $\alpha = (1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2$ .

Since  $F(t; \sigma_{\text{aug}}^2, \tau) := \frac{d}{dt}\tilde{\mathcal{L}}(\mathbf{W})$ , we have

$$F(t; \sigma_{\text{aug}}^2, \tau) = \frac{1}{\tau} \left\{ 1 + 2\sigma_{\text{aug}}^2 - \frac{2(1 + \sigma_{\text{aug}}^2)}{1 + \exp\left(-\frac{2}{\tau} + \frac{2(1 + \sigma_{\text{aug}}^2)}{\tau}t\right)} \right\},$$

which is a strictly increasing function in  $t$ . In addition,  $F((1 + \sigma_{\text{aug}}^2)^{-1}; \sigma_{\text{aug}}^2, \tau) = \sigma_{\text{aug}}^2/\tau > 0$ . To determine the minimizer of  $\ell(t)$ , it suffices to check the sign of  $F((1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|)^{-1}; \sigma_{\text{aug}}^2, \tau)$ .

**Case 1: expansion regime.** If  $F((1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|)^{-1}; \sigma_{\text{aug}}^2, \tau) > 0$  (i.e., in the *expansion regime*), we have  $F(t; \sigma_{\text{aug}}^2, \tau) > 0$  for any  $t \in [1/(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2), 1/(1 + \sigma_{\text{aug}}^2)]$ . Thus,  $\tilde{\mathcal{L}}(W)$  is strictly increasing in  $t$  and is minimized when  $t = 1/(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2)$ , which implies that

$$\frac{\|\mathbf{W}\|_{\text{F}}^2}{(1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2} = \frac{1}{1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2} \implies \frac{\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\|\mathbf{W}\|_{\text{F}}^2} = \|\boldsymbol{\mu}\|^2. \quad (13)$$

If we write the SVD of  $\mathbf{W}$  as  $\mathbf{W} = \sum_{j=1}^p \sigma_j \mathbf{u}_j \mathbf{v}_j^\top$ , then (13) can be written as

$$\sum_{j=1}^p \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2 = \sum_{j=1}^p \frac{\sigma_j^2}{\sum_{k=1}^p \sigma_k^2} \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2 \leq \frac{\sigma_1^2}{\sum_{k=1}^p \sigma_k^2} \sum_{j=1}^p \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2 \leq \sum_{j=1}^p \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2,$$

where equality holds if and only if  $\sigma_j = 0$  for all  $j \geq 2$  and  $\sigma_1 > 0$ . Then,

$$\|\boldsymbol{\mu}\|^2 = \sum_{j=1}^p \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2 = \langle \mathbf{v}_1, \boldsymbol{\mu} \rangle^2.$$



**Case 2: shrinkage regime.** If  $F((1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|)^{-1}; \sigma_{\text{aug}}^2, \tau) < 0$  (*shrinkage regime*), then  $\tilde{\mathcal{L}}(W)$  is minimized at some  $t^*$  in the interior satisfying

$$F(t^*; \sigma_{\text{aug}}^2, \tau) = \frac{1}{\tau} \left\{ 1 + 2\sigma_{\text{aug}}^2 - \frac{2(1 + \sigma_{\text{aug}}^2)}{1 + \exp\left(-\frac{2}{\tau} + \frac{2(1 + \sigma_{\text{aug}}^2)}{\tau} t^*\right)} \right\} = 0.$$

Solving this provides us with

$$t^* = \frac{1}{1 + \sigma_{\text{aug}}^2} \left\{ 1 - \frac{\tau}{2} \log(1 + 2\sigma_{\text{aug}}^2) \right\}.$$

When  $\sigma_{\text{aug}} \rightarrow 0$ , we have  $t^* \rightarrow 1$ , which implies that

$$\frac{\|\mathbf{W}\|_{\text{F}}^2}{\|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2} \rightarrow 1 \quad \implies \quad \frac{\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\|\mathbf{W}\|_{\text{F}}^2} \rightarrow 0.$$

Note that for any fixed value of  $\|\mathbf{W}\|_{\text{F}}$ ,  $\|\mathbf{W}\boldsymbol{\mu}\|^2 \rightarrow 0$  is equivalent to

$$\sum_{j=1}^p \sigma_j^2 \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle^2 \rightarrow 0$$

from which we have  $\max_{j \leq p} |\sigma_j \langle \mathbf{v}_j, \boldsymbol{\mu} \rangle| \rightarrow 0$ .

## B.4 Proof of Proposition 4.4

### B.4.1 A sandwich formula

Define  $\mathcal{L}_\ell(\mathbf{W}) := \tilde{\mathcal{L}}(\mathbf{W})$ , and  $\mathcal{L}_u(\mathbf{W}) := \tilde{\mathcal{L}}(\mathbf{W}) + \Delta(\mathbf{W})$ , where we denote  $r = 1 + \sigma_{\text{aug}}^2$  and

$$\Delta(\mathbf{W}) := \sum_{j \leq p} \frac{r^2 \sigma_j^4}{\tau^2 \alpha^2} + \exp\left(-\frac{\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau \alpha}\right) \cdot \left[ \exp\left(\sum_{j \leq p} \frac{4r \sigma_j^4 \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2}{\tau^2 \alpha^2}\right) - 1 \right].$$

We have the following sandwich-type result.

**Lemma B.1.** *We have*

$$\mathcal{L}_\ell(\mathbf{W}) \leq \mathcal{L}(\mathbf{W}) \leq \mathcal{L}_u(\mathbf{W}), \quad \text{for all } \mathbf{W} \in \mathbb{R}^{p \times p}.$$

*Proof.* The lower holds by the construction of  $\tilde{\mathcal{L}}(\mathbf{W})$ . Hence we focus on the upper bound.

Since  $\log(1 + x) - x \geq -x^2/2$ , we compare the difference between the first term in  $\mathcal{L}_{\text{unif}}$  and that in  $\tilde{\mathcal{L}}_{\text{unif}}$  to obtain

$$-\frac{1}{2} \sum_{j=1}^p \log\left(1 + \frac{2r\sigma_j^2}{\tau\alpha}\right) + \sum_{j=1}^p \frac{r\sigma_j^2}{\tau\alpha} \leq \sum_{j=1}^p \frac{r^2 \sigma_j^4}{\tau^2 \alpha^2}.$$

Write  $\beta_j = \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2$  for simplicity. We compare the difference between the second terms.

$$\log\left(1 + \exp\left(-\sum_{j=1}^p \frac{2\sigma_j^2 \beta_j}{2r\sigma_j^2 + \tau\alpha}\right)\right) - \log\left(1 + \exp\left(-\sum_{j=1}^p \frac{2\sigma_j^2 \beta_j}{\tau\alpha}\right)\right)$$

$$\begin{aligned}
&= \log \left[ 1 + \frac{\exp \left( - \sum_{j \leq p} 2\sigma_j^2 \beta_j / (2r\sigma_j^2 + \tau\alpha) \right) - \exp \left( - \sum_{j \leq p} 2\sigma_j^2 \beta_j / (\tau\alpha) \right)}{1 + \exp \left( - \sum_{j \leq p} 2\sigma_j^2 \beta_j / (\tau\alpha) \right)} \right] \\
&\stackrel{(i)}{\leq} \exp \left( - \sum_{j \leq p} \frac{2\sigma_j^2 \beta_j}{2r\sigma_j^2 + \tau\alpha} \right) - \exp \left( - \sum_{j \leq p} \frac{2\sigma_j^2 \beta_j}{\tau\alpha} \right) \\
&= \exp \left( - \frac{2\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau\alpha} \right) \cdot \left[ \exp \left( - \sum_{j \leq p} \frac{2\sigma_j^2 \beta_j}{2r\sigma_j^2 + \tau\alpha} + \sum_{j \leq p} \frac{2\sigma_j^2 \beta_j}{\tau\alpha} \right) - 1 \right] \\
&\stackrel{(ii)}{\leq} \exp \left( - \frac{2\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau\alpha} \right) \cdot \left[ \exp \left( \sum_{j \leq p} \frac{4r\sigma_j^4 \beta_j}{\tau^2 \alpha^2} \right) - 1 \right]
\end{aligned}$$

where in (i) we used  $\log(1+x) \leq x$  and  $\exp \left( - \sum_{j \leq p} 2\sigma_j^2 \beta_j / (\tau\alpha) \right) \geq 0$ , and in (ii) we used brute force computation. This completes the proof.  $\square$

We denote the minimizers of  $\mathcal{L}_\ell$ ,  $\mathcal{L}$ ,  $\mathcal{L}_u$  by  $\mathbf{W}_\ell^*$ ,  $\mathbf{W}^*$ ,  $\mathbf{W}_u^*$ , respectively. We also denote the related *expansion measure*

$$t_\ell^* = t(\mathbf{W}_\ell^*), \quad t^* = t(\mathbf{W}^*), \quad t_u^* = t(\mathbf{W}_u^*),$$

where we recall

$$t = t(\mathbf{W}) = \frac{\|\mathbf{W}\|_{\mathbb{F}}^2}{r\|\mathbf{W}\|_{\mathbb{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2}.$$

In what remains, we are mainly interested in showing that  $t^*$  and  $t_\ell^*$  are close. This would suggest that our approximation has small effects on the expansion/shrinkage phenomenon, thus justifying our approximation.

#### B.4.2 Approximation bound in expansion regime

Recall the notation  $r = 1 + \sigma_{\text{aug}}^2$ , and define  $\rho = \|\boldsymbol{\mu}\|$ . We observe from Lemma B.1 that

$$\begin{aligned}
\mathcal{L}_u(\mathbf{W}_\ell^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*) &\geq \mathcal{L}(\mathbf{W}_\ell^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*) \geq \mathcal{L}(\mathbf{W}^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*) \\
&\geq \mathcal{L}_\ell(\mathbf{W}^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*) \\
&\geq \min_{t \in [(r+\rho^2)^{-1}, r^{-1}]} F(t) \cdot (t^* - t_\ell^*).
\end{aligned}$$

Here the last inequality uses the definition of  $F(\cdot)$  as well as the fact that  $\mathcal{L}_\ell$  is convex in  $t$ .

Note that in the expansion regime,  $t_\ell^*$  achieves the smallest value, namely  $t_\ell^* = 1/(r + \rho^2)$ . Therefore we must have  $0 \leq t^* - t_\ell^*$ . In addition, we know that in the expansion regime,  $F(t) > 0$  for all  $t > t_\ell^*$ . As a result, we obtain

$$0 \leq t^* - t_\ell^* \leq \left[ \min_{t \in [(r+\rho^2)^{-1}, r^{-1}]} F(t) \right]^{-1} (\mathcal{L}_u(\mathbf{W}_\ell^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*)). \quad (14)$$

Since  $F(t)$  is increasing in  $t$ , we have

$$\min_{t \in [(r+\rho^2)^{-1}, r^{-1}]} F(t) = F(1/(r + \rho^2)) = -\frac{1}{\tau} + \frac{2r}{\tau} \cdot \left[ 1 + \exp \left( \frac{2}{\tau} - \frac{2r}{\tau(r + \rho^2)} \right) \right]^{-1}.$$

Denoting the following quantity

$$\tilde{\theta}_{\rho,r,\tau} = \frac{1}{1 + e^{2/\tau}} + \frac{2re^{2/\tau}}{(1 + e^{2/\tau})^2 \rho^2 \tau},$$

then as  $\rho \rightarrow \infty$ ,

$$\left| \left[ 1 + \exp\left(\frac{2}{\tau} - \frac{2r}{\tau(r + \rho^2)}\right) \right]^{-1} - \tilde{\theta}_{\rho,r,\tau} \right| = O(\rho^{-4}).$$

Therefore,

$$\begin{aligned} \min_{t \in [(r + \rho^2)^{-1}, r^{-1}]} F(t) &= -\frac{1}{\tau} + \frac{2r}{(1 + e^{2/\tau})\tau} + \frac{4r^2 e^{2/\tau}}{(1 + e^{2/\tau})^2 \rho^2 \tau^2} + O(\rho^{-4}) \\ &= \frac{1}{\tau} \left( \frac{2r}{1 + e^{2/\tau}} - 1 \right) + O(\rho^{-2}). \end{aligned} \quad (15)$$

We view  $\tau$  and  $\sigma_{\text{aug}}^2$  as fixed parameters and let  $\rho \rightarrow \infty$ . Recall that the expansion regime occurs if  $F((r + \rho^2)^{-1}) > 0$  and the shrinkage regime occurs if the reverse inequality is true. Under the asymptotics  $\rho \rightarrow \infty$ , the phase transition boundary simplifies to  $2r > 1 + e^{2/\tau}$  for expansion regime and  $2r < 1 + e^{2/\tau}$  for the shrinkage regime. Thus, the leading term 15 is always positive in this asymptotics.

Since  $\mathbf{W} = \mathbf{W}_\ell^*$  if and only if  $\langle \boldsymbol{\mu}, \mathbf{v}_1 \rangle^2 > 0$ ,  $\langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2 = 0$  for  $j > 1$  and  $\sigma_2 = \dots = \sigma_p = 0$ , we can show that

$$\Delta(\mathbf{W}_\ell^*) = \mathcal{L}_u(\mathbf{W}_\ell^*) - \mathcal{L}_\ell(\mathbf{W}_\ell^*) = \frac{r^2}{\tau^2(r + \rho^2)^2} + \exp\left(-\frac{\rho^2}{\tau(r + \rho^2)}\right) \cdot \left[ \exp\left(\frac{4r\rho^2}{\tau^2(r + \rho^2)^2}\right) - 1 \right],$$

and as  $\rho \rightarrow \infty$ ,

$$\left| \Delta(\mathbf{W}_\ell^*) - \frac{4r}{\tau^2 e^{1/\tau}} \frac{1}{\rho^2} \right| = O(\rho^{-4}). \quad (16)$$

From (15) and (16), we have the control for (14):

$$0 \leq t^* - t_\ell^* \leq \left( \frac{2r}{1 + e^{2/\tau}} - 1 \right)^{-1} \frac{4r}{\tau e^{1/\tau}} \frac{1}{\rho^2} + O(\rho^{-4})$$

This finishes the proof.

### B.4.3 Approximation bound in shrinkage regime

We use the notation  $\mathcal{C}^2$  to denote the space of twice continuously differentiable functions. We first present the general result as follows.

**Lemma B.2.** *Suppose that  $f(\mathbf{x})$ ,  $f_\ell(\mathbf{x})$ ,  $f_u(\mathbf{x})$  are defined for  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^q$  with*

$$f_\ell(\mathbf{x}) \leq f(\mathbf{x}) \leq f_u(\mathbf{x}).$$

*Assume that  $f_\ell(\mathbf{x}) = h(\varphi(\mathbf{x}))$  where  $\varphi : \mathcal{X} \rightarrow \mathcal{Y} \subset \mathbb{R}$  is a surjection satisfying  $\varphi \in \mathcal{C}^2(\mathcal{X})$ , and that  $h : \mathcal{Y} \rightarrow \mathbb{R}$  satisfies  $h \in \mathcal{C}^2(\mathcal{X})$  and is strongly convex. Suppose  $\mathbf{x}_\ell^* \in \mathcal{X}$  and  $y_\ell^*$  satisfy*

$$\varphi(\mathbf{x}_\ell^*) = y_\ell^*, \quad y_\ell^* = \operatorname{argmin}_{y \in \mathcal{Y}} h(y).$$

*Then, the following inequality holds.*

$$\|y_\ell^* - \varphi(\mathbf{x}^*)\|^2 \leq \frac{2(f_u(\mathbf{x}_\ell^*) - f_\ell(\mathbf{x}_\ell^*))}{\inf_{y \in \mathcal{Y}} h''(y)}. \quad (17)$$

*Proof.* By the definitions of  $f_u, f, f_\ell$  and  $\mathbf{x}_\ell^*, \mathbf{x}^*$ , we have

$$f_u(\mathbf{x}_\ell^*) \geq f(\mathbf{x}_\ell^*) \geq f(\mathbf{x}^*) \geq f_\ell(\mathbf{x}^*) = h(\varphi(\mathbf{x}^*)). \quad (18)$$

Since  $h$  is convex and  $\varphi(\mathbf{x}_\ell^*)$  achieves the minimum, we have

$$h(\varphi(\mathbf{x}^*)) - h(\varphi(\mathbf{x}_\ell^*)) \geq \frac{1}{2} \inf_{y \in \mathcal{Y}} h''(y) \cdot \|\varphi(\mathbf{x}^*) - \varphi(\mathbf{x}_\ell^*)\|^2. \quad (19)$$

We combine (18)–(19) and use  $h(\varphi(\mathbf{x}_\ell^*)) = f_\ell(\mathbf{x}_\ell^*)$ ,  $\varphi(\mathbf{x}_\ell^*) = y_\ell^*$ , which leads to the desired inequality (17).  $\square$

To use the above lemma, we set  $f = \mathcal{L}, f_\ell = \mathcal{L}_\ell, f_u = \mathcal{L}_u$ ,

$$h(t) = -\frac{t}{\tau} + \log \left( 1 + \exp \left( -\frac{2}{\tau} + \frac{2rt}{\tau} \right) \right) - \log 2,$$

$$\varphi(\mathbf{W}) = t(\mathbf{W}) = \frac{\|\mathbf{W}\|_{\mathbb{F}}^2}{r\|\mathbf{W}\|_{\mathbb{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2},$$

where we identify  $y$  with  $t$ . We calculate the lower bound on  $h''(t)$  as follows.

$$h'(t) = F(t) = -\frac{1}{\tau} + \frac{2r}{\tau} \left[ 1 + \exp \left( \frac{2}{\tau} - \frac{2rt}{\tau} \right) \right]^{-1},$$

$$h''(t) = F'(t) = \frac{4r^2}{\tau^2} \exp \left( \frac{2}{\tau} - \frac{2rt}{\tau} \right) \cdot \left[ 1 + \exp \left( \frac{2}{\tau} - \frac{2rt}{\tau} \right) \right]^{-2}.$$

We note that

$$\exp \left( \frac{2}{\tau} \right) \geq \exp \left( \frac{2}{\tau} - \frac{2rt}{\tau} \right) \geq \exp \left( \frac{2-2r}{\tau} \right).$$

So making use of the fact that  $\frac{x}{(1+x)^2}$  is decreasing when  $x \geq 1$ , we deduce

$$\begin{aligned} \min_{t \in [(r+\rho)^{-1}, 1]} F'(t) &= \frac{4r^2}{\tau^2} \min \left\{ \frac{e^{2\sigma_{\text{aug}}^2/\tau r}}{(1 + e^{2\sigma_{\text{aug}}^2/\tau r})^2}, \frac{e^{2\rho/\tau(r+\rho)}}{(1 + e^{2\rho/\tau(r+\rho)})^2} \right\} \\ &= \frac{4r^2}{\tau^2} \frac{e^{2\sigma_{\text{aug}}^2/\tau r}}{(1 + e^{2\sigma_{\text{aug}}^2/\tau r})^2} \geq \frac{r^2}{\tau^2} e^{-2\sigma_{\text{aug}}^2/\tau r}. \end{aligned}$$

Note that we have the freedom to choose any  $\mathbf{x}$  (identified as  $\mathbf{W}$ ) as long as  $\varphi(\mathbf{W}) = t_\ell^*$ . This leads to the error bound

$$\|t^* - t_\ell^*\|^2 \leq \frac{\tau^2}{r^2} e^{2\sigma_{\text{aug}}^2/\tau} \cdot \Delta(\mathbf{W}).$$

It remains to control  $\min_{\mathbf{W}: \varphi(\mathbf{W})=t_\ell^*} \Delta(\mathbf{W})$  which will be our target below.

**Characterizing  $\min_{\mathbf{W}: \varphi(\mathbf{W})=t_\ell^*} \Delta(\mathbf{W})$ .** It suffices to provide a feasible solution  $\mathbf{W}$  such that  $\varphi(\mathbf{W}) = t_\ell^*$ . To do so, we construct  $\mathbf{W}$  as follows. As usual, let  $\sum_{j=1}^p \sigma_j \mathbf{u}_j \mathbf{v}_j^\top$  be the singular value decomposition of  $\mathbf{W}$ . Here  $\sigma_j, j \in [p]$ , are not necessarily ordered. Let  $\mathbf{v}_p = \boldsymbol{\mu}/\|\boldsymbol{\mu}\|$ . As a result, we have  $\|\mathbf{W}\|_{\mathbb{F}} = \sum_j \sigma_j$  and  $\|\mathbf{W}\boldsymbol{\mu}\| = \sigma_p^2 \|\boldsymbol{\mu}\|^2 = \sigma_p^2 \rho^2$ . Recall that in the shrinkage regime, we have

$$t_\ell^* = \frac{1}{r} \left( 1 - \frac{\tau}{2} \log(2r-1) \right) = \frac{\|\mathbf{W}\|_{\mathbb{F}}^2}{r\|\mathbf{W}\|_{\mathbb{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2} \in \left[ \frac{1}{r+\rho^2}, \frac{1}{r} \right].$$

Consequently, one can verify that any sequence of singular values  $\boldsymbol{\sigma}$  with  $S_{p-1} = \sum_{j=1}^{p-1} \sigma_j^2$  that obeys

$$\frac{S_{p-1}}{\sigma_p^2} = \frac{t_\ell^*(\rho^2 + r) - 1}{1 - t_\ell^* r} > 0 \quad (20)$$

will be a feasible solution. In particular,

$$\sigma_1^2 = \sigma_2^2 = \dots = \sigma_{p-1}^2 = \frac{S_{p-1}}{p-1}.$$

Now we are ready to compute the quantity of interest

$$\Delta(\mathbf{W}) = \sum_{j \leq p} \frac{r^2 \sigma_j^4}{\tau^2 \alpha^2} + \exp\left(-\frac{\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau \alpha}\right) \cdot \left[ \exp\left(\sum_{j \leq p} \frac{4r \sigma_j^4 \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2}{\tau^2 \alpha^2}\right) - 1 \right].$$

1. First, we have

$$\begin{aligned} \sum_{j \leq p} \frac{r^2 \sigma_j^4}{\tau^2 \alpha^2} &= \sum_{j \leq p} \frac{r^2 \sigma_j^4}{\tau^2 (r \|\boldsymbol{\sigma}\|^2 + \sigma_p^2 \rho^2)^2} = \frac{r^2 (t_\ell^*)^2}{\tau^2} \cdot \frac{\sum_{j \leq p} \sigma_j^4}{\|\boldsymbol{\sigma}\|^4} \\ &= \frac{r^2 (t_\ell^*)^2}{\tau^2} \cdot \frac{(p-1) \frac{S_{p-1}^2}{(p-1)^2} + \sigma_p^4}{\left((p-1) \frac{S_{p-1}}{p-1} + \sigma_p^2\right)^2} \\ &= \frac{r^2 (t_\ell^*)^2}{\tau^2 (p-1)} \cdot \frac{\frac{S_{p-1}^2}{\sigma_p^2} + p - 1}{\left(\frac{S_{p-1}}{\sigma_p^2} + 1\right)^2} \end{aligned}$$

Recall the identity (20) and the fact that  $rt_\ell^* \leq 1$ , we can further write the first term as

$$\begin{aligned} \sum_{j \leq p} \frac{r^2 \sigma_j^4}{\tau^2 \alpha^2} &= \frac{r^2 (t_\ell^*)^2}{\tau^2 (p-1)} \frac{(\rho^2 t_\ell^* - 1 + rt_\ell^*)^2 + (p-1)(1 - rt_\ell^*)^2}{\rho^4 (t_\ell^*)^2} \\ &\leq \frac{r^2}{\tau^2 \rho^4 (p-1)} \cdot (\rho^4 (t_\ell^*)^2 + p(1 - rt_\ell^*)^2) \\ &\leq \frac{r^2}{\tau^2 \rho^4 (p-1)} \cdot \left(\frac{\rho^4}{r^2} + \frac{\tau^2 p}{4} \log^2(2r-1)\right) \\ &= \frac{1}{\tau^2 (p-1)} + \frac{pr^2}{4\rho^4 (p-1)} \log^2(2r-1) = \frac{1}{\tau^2 (p-1)} + \frac{\sigma_{\text{aug}}^4}{\rho^4} + O\left(\frac{\sigma_{\text{aug}}^6}{\rho^4}\right) \quad \sigma_{\text{aug}} \rightarrow 0. \end{aligned}$$

2. Secondly, since  $\exp\left(-\frac{\|\mathbf{W}\boldsymbol{\mu}\|^2}{\tau \alpha}\right) \leq 1$  and  $\sigma_p^2 \rho^2 / \alpha = 1 - rt_\ell^*$ , then

$$\begin{aligned} \exp\left(\sum_{j \leq p} \frac{4r \sigma_j^4 \langle \boldsymbol{\mu}, \mathbf{v}_j \rangle^2}{\tau^2 \alpha^2}\right) - 1 &= \exp\left(\frac{4r \sigma_p^4 \rho^2}{\tau^2 \alpha^2}\right) - 1 = \exp\left(\frac{4r}{\tau^2 \rho^2} (1 - rt_\ell^*)^2\right) - 1 \\ &= \exp\left(\frac{r}{\rho^2} \log^2(2r-1)\right) - 1 = \frac{4\sigma_{\text{aug}}^4}{\rho^2} + O\left(\frac{\sigma_{\text{aug}}^6}{\rho^2}\right) \quad \sigma_{\text{aug}} \rightarrow 0. \end{aligned}$$

Combining pieces above,

$$\Delta(\mathbf{W}) \leq \frac{1}{\tau^2(p-1)} + O\left(\frac{\sigma_{\text{aug}}^4}{\|\boldsymbol{\mu}\|^2}\right).$$

As a result, we can bound the approximation error in the shrinkage regime by

$$\|t^* - t_\ell^*\|^2 \leq \frac{1}{(1 + \sigma_{\text{aug}}^2)^2(p-1)} e^{2\sigma_{\text{aug}}^2/\tau} + O\left(\frac{\sigma_{\text{aug}}^4}{\|\boldsymbol{\mu}\|^2}\right).$$

## B.5 Proof of Proposition 4.5

Here we provide some informal analysis on the finite-sample scenario. We realize that a rigorous analysis is challenging and is left to future research.

Recall the contrastive loss

$$\mathcal{L}_n(\mathbf{W}) := -\frac{1}{\tau} \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^+, \mathbf{h} | \mathbf{h}_0} [\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^+) | \mathbf{h}_0] \right\} + \log \left( \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^-, \mathbf{h} | \mathbf{h}_0} \left[ e^{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)/\tau} | \mathbf{h}_0 \right] \right\} \right),$$

which can be decomposed as two terms

$$\begin{aligned} \mathcal{L}_{n,1}(\mathbf{W}) &= -\frac{1}{\tau} \mathbb{E}_n [\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^+)], \\ \mathcal{L}_{n,2}(\mathbf{W}) &= \log \left( \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^-, \mathbf{h} | \mathbf{h}_0} \left[ e^{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)/\tau} | \mathbf{h}_0 \right] \right\} \right). \end{aligned}$$

Denote  $\alpha = (1 + \sigma_{\text{aug}}^2) \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2$  and  $\mathbf{A} = (2\tau\alpha)^{-1} \mathbf{W}^\top \mathbf{W}$ , with the definition of  $\text{sim}^*$ , we have

$$\mathcal{L}_{n,1}(\mathbf{W}) = \frac{1}{2\tau\alpha} \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^+, \mathbf{h} | \mathbf{h}_0} [\|\mathbf{W}(\mathbf{h} - \mathbf{h}^+)\|_2^2 | \mathbf{h}_0] \right\} = \frac{\sigma_{\text{aug}}^2}{\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2.$$

For  $\mathbf{h}^- \sim \frac{1}{2} \mathcal{N}(-\boldsymbol{\mu}, (1 + \sigma_{\text{aug}}^2)\mathbf{I}) + \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}, (1 + \sigma_{\text{aug}}^2)\mathbf{I})$  and  $\mathbf{h} | \mathbf{h}_0 \sim \mathcal{N}(\mathbf{h}_0, \sigma_{\text{aug}}^2\mathbf{I})$ , we have

$$\begin{aligned} \mathbb{E}_{\mathbf{h}, \mathbf{h}^- | \mathbf{h}_0} \left[ \exp \left( \frac{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)}{\tau} \right) \middle| \mathbf{h}_0 \right] &= \mathbb{E}_{\mathbf{h}, \mathbf{h}^- | \mathbf{h}_0} \left[ \exp \left( -\frac{\|\mathbf{W}(\mathbf{h} - \mathbf{h}^-)\|_2^2}{2\tau\alpha} \right) \middle| \mathbf{h}_0 \right] \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, (1+2\sigma_{\text{aug}}^2)\mathbf{I})} \left[ \exp \left( -\frac{\|\mathbf{W}(\mathbf{h}_0 - \mathbf{z})\|_2^2}{2\tau\alpha} \right) \right] + \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(-\boldsymbol{\mu}, (1+2\sigma_{\text{aug}}^2)\mathbf{I})} \left[ \exp \left( -\frac{\|\mathbf{W}(\mathbf{h}_0 - \mathbf{z})\|_2^2}{2\tau\alpha} \right) \right]. \end{aligned} \quad (21)$$

With  $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, (1 + 2\sigma_{\text{aug}}^2)\mathbf{I})$ , we have

$$\begin{aligned} &\mathbb{E}_{\mathbf{z} | \mathbf{h}_0} \left[ \exp \left( -\frac{\|\mathbf{W}(\mathbf{h}_0 - \mathbf{z})\|_2^2}{2\tau\alpha} \right) \middle| \mathbf{h}_0 \right] \\ &= \left( \frac{1}{2\pi(1 + 2\sigma_{\text{aug}}^2)} \right)^{d/2} \int \exp \left( -\frac{\|\mathbf{W}(\mathbf{h}_0 - \mathbf{z})\|_2^2}{2\tau\alpha} \right) \exp \left( -\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \|\mathbf{z} - \boldsymbol{\mu}\|_2^2 \right) d\mathbf{z} \\ &= \left( \frac{1}{2\pi(1 + 2\sigma_{\text{aug}}^2)} \right)^{d/2} \int \exp \left( -\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ (\mathbf{z} - \boldsymbol{\mu})^\top (\mathbf{z} - \boldsymbol{\mu}) + (\mathbf{h}_0 - \mathbf{z})^\top \tilde{\mathbf{A}} (\mathbf{h}_0 - \mathbf{z}) \right] \right) d\mathbf{z} \end{aligned} \quad (22)$$

where  $\tilde{\mathbf{A}} = 2(1 + 2\sigma_{\text{aug}}^2)\mathbf{A}$ . Rearrange (22), we have

$$(\mathbf{z} - \boldsymbol{\mu})^\top (\mathbf{z} - \boldsymbol{\mu}) + (\mathbf{h}_0 - \mathbf{z})^\top \tilde{\mathbf{A}} (\mathbf{h}_0 - \mathbf{z})$$

$$=(\mathbf{z} - \tilde{\mathbf{a}})^\top \mathbf{M}^{-1}(\mathbf{z} - \tilde{\mathbf{a}}) - (\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0)^\top \mathbf{M}(\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0) + \boldsymbol{\mu}^\top \boldsymbol{\mu} + \mathbf{h}_0^\top \tilde{\mathbf{A}}\mathbf{h}_0.$$

where  $\mathbf{M} = (\mathbf{I} + \tilde{\mathbf{A}})^{-1}$  and  $\tilde{\mathbf{a}} = \mathbf{M}(\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0)$ . Then, the integral in (22) can be simplified as

$$(22) = \sqrt{\det(\mathbf{M})} \exp\left(\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ (\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0)^\top \mathbf{M}(\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0) - \boldsymbol{\mu}^\top \boldsymbol{\mu} - \mathbf{h}_0^\top \tilde{\mathbf{A}}\mathbf{h}_0 \right]\right).$$

Note that  $\mathbf{M}\tilde{\mathbf{A}} = \mathbf{I} - \mathbf{M} = \tilde{\mathbf{A}}\mathbf{M}$ , we can rewrite (21) as

$$\begin{aligned} & \mathbb{E}_{\mathbf{h}^-, \mathbf{h}|\mathbf{h}_0} \left[ \exp\left(\frac{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)}{\tau}\right) \middle| \mathbf{h}_0 \right] \\ &= \frac{1}{2} \sqrt{\det(\mathbf{M})} \exp\left(\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ (\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0)^\top \mathbf{M}(\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0) - \boldsymbol{\mu}^\top \boldsymbol{\mu} - \mathbf{h}_0^\top \tilde{\mathbf{A}}\mathbf{h}_0 \right]\right) \\ & \quad + \frac{1}{2} \sqrt{\det(\mathbf{M})} \exp\left(\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ (-\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0)^\top \mathbf{M}(-\boldsymbol{\mu} + \tilde{\mathbf{A}}\mathbf{h}_0) - \boldsymbol{\mu}^\top \boldsymbol{\mu} - \mathbf{h}_0^\top \tilde{\mathbf{A}}\mathbf{h}_0 \right]\right) \\ &= \frac{1}{2} \sqrt{\det(\mathbf{M})} \exp\left(\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ \boldsymbol{\mu}^\top (\mathbf{M} - \mathbf{I})\boldsymbol{\mu} + \mathbf{h}_0^\top (\mathbf{M} - \mathbf{I})\mathbf{h}_0 - 2\boldsymbol{\mu}^\top (\mathbf{M} - \mathbf{I})\mathbf{h}_0 \right]\right) \\ & \quad \cdot \left( 1 + \exp\left(\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} \left[ 4\boldsymbol{\mu}^\top (\mathbf{M} - \mathbf{I})\mathbf{h}_0 \right]\right) \right). \end{aligned}$$

Further, as  $\mathbf{I} - \mathbf{M} \succcurlyeq 0$ ,

$$\begin{aligned} & \mathbb{E}_n \left\{ \mathbb{E}_{\mathbf{h}^-, \mathbf{h}|\mathbf{h}_0} \left[ \exp\left(\frac{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)}{\tau}\right) \middle| \mathbf{h}_0 \right] \right\} \\ &= \frac{1}{2} \sqrt{\det(\mathbf{M})} \mathbb{E}_n \left\{ \exp\left(-\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 - \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 - \boldsymbol{\mu})\right) \right. \\ & \quad \left. + \exp\left(-\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 + \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 + \boldsymbol{\mu})\right) \right\}. \end{aligned}$$

Taking the logarithm, the uniformity loss can be simplified as

$$\begin{aligned} \mathcal{L}_{n,2}(\mathbf{W}) &= \log \left( \mathbb{E}_n \mathbb{E}_{\mathbf{h}^-, \mathbf{h}|\mathbf{h}_0} \left[ \exp\left(\frac{\text{sim}^*(\mathbf{W}\mathbf{h}, \mathbf{W}\mathbf{h}^-)}{\tau}\right) \right] \right) \\ &= -\log 2 + \frac{1}{2} \log \det[\mathbf{M}] + \log(\mathbb{E}_n S_{\mathbf{h}_0}), \end{aligned}$$

where

$$\begin{aligned} S_{\mathbf{h}_0} &:= \exp\left(-\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 - \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 - \boldsymbol{\mu})\right) \\ & \quad + \exp\left(-\frac{1}{2(1 + 2\sigma_{\text{aug}}^2)} (\mathbf{h}_0 + \boldsymbol{\mu})^\top (\mathbf{I} - \mathbf{M})(\mathbf{h}_0 + \boldsymbol{\mu})\right). \end{aligned}$$

Organizing the terms above, we have

$$\mathcal{L}_n(\mathbf{W}) = -\log 2 + \frac{\sigma_{\text{aug}}^2}{\tau\alpha} \|\mathbf{W}\|_{\text{F}}^2 + \frac{1}{2} \log \det[\mathbf{M}] + \log(\mathbb{E}_n S_{\mathbf{h}_0}).$$



When  $\alpha = (1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\|_{\mathbb{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2 \gg \sigma_j^2$  for any singular value  $\sigma_j$  of  $\mathbf{W}$ , we have  $\mathbf{M} \approx \mathbf{I} - (\tau\alpha)^{-1}(1 + 2\sigma_{\text{aug}}^2)\mathbf{W}^\top\mathbf{W}$  and

$$\log(\det[\mathbf{M}]) = \log\left(\prod_{j=1}^p \left(1 + \frac{1 + 2\sigma_{\text{aug}}^2}{\tau\alpha}\sigma_j^2\right)^{-1}\right) \approx -\frac{1 + 2\sigma_{\text{aug}}^2}{\tau\alpha}\|\mathbf{W}\|_{\mathbb{F}}^2.$$

We can then approximate  $S_n$  by

$$\tilde{S}_{\mathbf{h}_0} := \exp\left(-\frac{1}{2\tau\alpha}\|\mathbf{W}(\mathbf{h}_0 - \boldsymbol{\mu})\|^2\right) + \exp\left(-\frac{1}{2\tau\alpha}\|\mathbf{W}(\mathbf{h}_0 + \boldsymbol{\mu})\|^2\right).$$

As a result, we have the approximation

$$\mathcal{L}_{n,2}(\mathbf{W}) \approx -\log 2 - \frac{1 + 2\sigma_{\text{aug}}^2}{2\tau\alpha}\|\mathbf{W}\|_{\mathbb{F}}^2 + \log\left(\mathbb{E}_n \tilde{S}_{\mathbf{h}_0}\right) =: \tilde{\mathcal{L}}_{n,2}(\mathbf{W}).$$

Therefore, recall that  $\alpha = (1 + \sigma_{\text{aug}}^2)\|\mathbf{W}\|_{\mathbb{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|^2$ , we have the following approximation for  $\mathcal{L}_n(\mathbf{W})$ :

$$\begin{aligned} \tilde{\mathcal{L}}_n(\mathbf{W}) &= \mathcal{L}_{n,1}(\mathbf{W}) + \tilde{\mathcal{L}}_{n,2}(\mathbf{W}) \\ &= -\log 2 - \frac{1}{2\tau\alpha}\|\mathbf{W}\|_{\mathbb{F}}^2 + \log\left(\mathbb{E}_n \tilde{S}_{\mathbf{h}_0}\right). \end{aligned}$$

As  $\mathbf{h}_0 \sim \mathcal{N}(y \cdot \boldsymbol{\mu}, \mathbf{I})$ , then

$$\log\left(\mathbb{E}_{\mathbf{h}_0} \tilde{S}_{\mathbf{h}_0}\right) = -\frac{1}{2}\log\left(1 + \frac{\sigma_j^2}{\tau\alpha}\right) + \log\left(1 + \exp\left(-\frac{2\sigma_j^2}{\sigma_j^2 + \tau\alpha}(\mathbf{v}_j^\top \boldsymbol{\mu})\right)\right).$$

Under the assumption that  $\alpha \gg \|\mathbf{W}\|_{\mathbb{F}}^2$ , we further have

$$\log\left(\mathbb{E}_{\mathbf{h}_0} \tilde{S}_{\mathbf{h}_0}\right) \approx -\frac{1}{2\tau\alpha}\|\mathbf{W}\|_{\mathbb{F}}^2 + \log\left(1 + \exp\left(-\frac{2}{\tau\alpha}\|\mathbf{W}\boldsymbol{\mu}\|^2\right)\right).$$

Accordingly, if we replace  $\mathbb{E}_n$  by  $\mathbb{E}_{\mathbf{h}_0}$  in  $\tilde{\mathcal{L}}_n(\mathbf{W})$ , then the loss can be approximated by

$$-\log 2 - \frac{1}{\tau\alpha}\|\mathbf{W}\|_{\mathbb{F}}^2 + \log\left(1 + \exp\left(-\frac{2}{\tau\alpha}\|\mathbf{W}\boldsymbol{\mu}\|^2\right)\right),$$

which is exactly the loss function  $\tilde{\mathcal{L}}(\mathbf{W})$  defined in (3).

## C Proofs for Section 5

### C.1 Proof of Theorem 5.1

Denote  $\boldsymbol{\theta} = \mathbf{W}\boldsymbol{\beta}$ ,  $\boldsymbol{\Omega} = \mathbf{W}^{-2}$ , and  $\boldsymbol{\xi} = (\gamma, \boldsymbol{\theta}^\top)^\top$ . One can rewrite the logistic loss function as

$$\tilde{\ell}_n(\boldsymbol{\xi}; \mathbf{h}_0) = \ell_{n,\text{logis}}(\boldsymbol{\xi}; \mathbf{h}_0) + \lambda_n \boldsymbol{\theta}^\top \boldsymbol{\Omega} \boldsymbol{\theta} = \mathbb{E}_n \{\ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n)\},$$

where we define

$$\ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n) = \log\left[1 + \exp(-y(\gamma + \mathbf{h}_0^\top \boldsymbol{\theta}))\right] + \lambda_n \boldsymbol{\theta}^\top \boldsymbol{\Omega} \boldsymbol{\theta}$$

and  $\ell_{n,\text{logis}}(\boldsymbol{\xi}; \mathbf{h}_0) = \mathbb{E}_n \log [1 + \exp(-y(\gamma + \mathbf{h}_0^\top \boldsymbol{\theta}))]$ . We also denote

$$\ell_n^*(\boldsymbol{\xi}) = \mathbb{E} \{ \ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n) \},$$

where the loss depends on  $n$  through  $\lambda_n$ .

Since both  $\ell_n^*(\boldsymbol{\xi})$  and  $\tilde{\ell}_n(\boldsymbol{\xi}; \mathbf{h}_0)$  are strictly convex in  $\boldsymbol{\xi}$  and are strongly convex when  $\lambda_n \geq 0$ , we denote

$$\widehat{\boldsymbol{\xi}}_n = (\widehat{\gamma}, \widehat{\boldsymbol{\theta}}^\top)^\top = \operatorname{argmin} \tilde{\ell}_n(\boldsymbol{\xi}; \mathbf{h}_0), \quad \text{and} \quad \boldsymbol{\xi}_n^* = (\gamma_n^*, \boldsymbol{\theta}_n^{*\top})^\top = \operatorname{argmin} \ell_n^*(\boldsymbol{\xi}).$$

**Step 1: characterizing  $\boldsymbol{\xi}_n^*$ .** Under the 2-GMM model, we have the following characterization of the solution  $\boldsymbol{\xi}_n^*$  to the population loss  $\ell_n^*(\boldsymbol{\xi})$ .

**Lemma C.1.** *There exists a constant  $C > 0$  such that the following holds. For any given  $(\lambda_n)_{n \geq 1}$  with  $\lambda_n \geq 0$ , the population loss minimizer  $\boldsymbol{\xi}_n^*$  is unique and has the form*

$$\boldsymbol{\xi}_n^* = (0, \kappa_{\lambda_n} \boldsymbol{\mu}^\top)^\top$$

where  $\kappa_{\lambda_n} \in (0, C)$  is the zero of the function

$$\psi(\kappa) := \mathbb{E} \left\{ \frac{1 + (1 - \kappa) \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})}{(1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G}))^2} \right\} - 2\lambda_n (1 + \eta)^{-2} \kappa.$$

Moreover, if  $\lim_{n \rightarrow \infty} \lambda_n = \infty$ , then  $\lim_{n \rightarrow \infty} \lambda_n \kappa_{\lambda_n} = \frac{1}{4}(1 + \eta)^2$ .

As a corollary, we see that  $\boldsymbol{\xi}_n^*$  has bounded norm.

**Step 2: characterizing  $\widehat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*$ .** Denote the Fisher information matrix

$$\mathbf{I}_n(\boldsymbol{\xi}) = \mathbb{E}_n \left\{ \frac{\bar{\mathbf{h}}_0 \bar{\mathbf{h}}_0^\top \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi})}{(1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}))^2} \right\}, \quad \mathbf{I}(\boldsymbol{\xi}) = \mathbb{E} \left\{ \frac{\bar{\mathbf{h}}_0 \bar{\mathbf{h}}_0^\top \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi})}{(1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}))^2} \right\}.$$

We first have the following lemma for consistency.

**Lemma C.2.** *For any  $(\lambda_n)_{n \geq 1}$  with  $\lambda_n \geq 0$  and any constant  $R > 0$ ,*

$$\sup_{\boldsymbol{\xi} \in \mathbb{R}^{p+1}: \|\boldsymbol{\xi}\| \leq R} \left| \mathbb{E}_n \ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n) - \mathbb{E} \ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n) \right| \xrightarrow{p} 0,$$

$$\sup_{\boldsymbol{\xi} \in \mathbb{R}^{p+1}: \|\boldsymbol{\xi}\| \leq R} \left| \mathbf{I}_n(\boldsymbol{\xi}) - \mathbf{I}(\boldsymbol{\xi}) \right| \xrightarrow{p} 0.$$

As a result, we have

$$\|\widehat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*\| \xrightarrow{p} 0, \quad \text{as } n \rightarrow \infty$$

Based on the consistency result, we consider the score functions for  $\tilde{\ell}_n(\boldsymbol{\xi}; \mathbf{h}_0)$  and  $\ell_n^*(\boldsymbol{\xi})$ . Denoting  $\bar{\boldsymbol{\Omega}} = \operatorname{diag}\{0, \boldsymbol{\Omega}\}$ , we obtain

$$0 = \frac{\partial}{\partial \boldsymbol{\xi}} \tilde{\ell}_n(\widehat{\boldsymbol{\xi}}_n; \mathbf{h}_0) = -\mathbb{E}_n \left\{ \frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y \bar{\mathbf{h}}_0^\top \widehat{\boldsymbol{\xi}}_n)} \right\} + 2\lambda_n \bar{\boldsymbol{\Omega}} \widehat{\boldsymbol{\xi}}_n. \quad (23)$$

$$0 = \frac{\partial}{\partial \boldsymbol{\xi}} \ell_n^*(\boldsymbol{\xi}_n^*) = -\mathbb{E} \left\{ \frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}_n^*)} \right\} + 2\lambda_n \bar{\boldsymbol{\Omega}} \boldsymbol{\xi}_n^*. \quad (24)$$

According to Lemma C.2, we have the first-order approximation

$$\frac{\partial}{\partial \boldsymbol{\xi}} \tilde{\ell}_{n, \text{logis}}(\hat{\boldsymbol{\xi}}_n; \mathbf{h}_0) - \frac{\partial}{\partial \boldsymbol{\xi}} \tilde{\ell}_{n, \text{logis}}(\boldsymbol{\xi}_n^*; \mathbf{h}_0) = \mathbf{I}_n(\boldsymbol{\xi}_n^*)(\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*) + o_{\mathbb{P}}(\|\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*\|).$$

By calculating the difference between (23) and (24), we obtain

$$-\mathbf{I}_n(\boldsymbol{\xi}_n^*)(\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*) - 2\lambda_n \bar{\boldsymbol{\Omega}}(\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*) + (\mathbb{E}_n - \mathbb{E}) \left[ \frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}_n^*)} \right] + o_{\mathbb{P}}(\|\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*\|) = 0.$$

When  $p$  is fixed and  $n \rightarrow +\infty$ , since  $\boldsymbol{\xi}_n^*$  has bounded norm, then by Lemma C.2, we have  $\|(\mathbf{I}_n(\boldsymbol{\xi}_n^*) - \mathbf{I}(\boldsymbol{\xi}_n^*))(\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*)\| = o_{\mathbb{P}}(\|\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*\|)$ . In addition, denote

$$\mathbf{z}_n = \frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}_n^*)},$$

where  $\mathbb{E} \mathbf{z}_n = \mathbf{0}$  and  $\text{cov}(\mathbf{z}_n) = \mathbf{I}(\boldsymbol{\xi}_n^*)$ , then we have the following lemma.

**Lemma C.3.** *We have  $\mathbf{I}(\boldsymbol{\xi}_n^*) \succ \mathbf{0}$  and by the Lindeberg-Feller central limit theorem,*

$$\sqrt{n} (\mathbf{I}(\boldsymbol{\xi}_n^*))^{-1/2} (\mathbb{E}_n - \mathbb{E}) \left[ \frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}_n^*)} \right] \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

Let  $\bar{\boldsymbol{\zeta}} = \sqrt{n}(\hat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^*)$ , then

$$(\mathbf{I}(\boldsymbol{\xi}_n^*))^{-1/2} (\mathbf{I}(\boldsymbol{\xi}_n^*) + 2\lambda_n \bar{\boldsymbol{\Omega}}) \bar{\boldsymbol{\zeta}} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (25)$$

**Step 3: calculating misclassification error.** Conditioning on  $y = 1$ ,

$$\text{Err}_1 = \mathbb{P}(y(\langle \mathbf{h}_0, \hat{\boldsymbol{\theta}} \rangle + \hat{\gamma}) < 0 | y = 1) = \mathbb{P}(\langle \boldsymbol{\mu} + \mathbf{G}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta} \rangle + \hat{\gamma} < 0), \quad (26)$$

where  $\mathbf{G} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , then (26) can be written using the Gaussian cdf as

$$\begin{aligned} (26) &= \Phi \left( -\frac{\hat{\gamma} + \langle \boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta} \rangle}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|} \right) \\ &= \Phi \left( -\|\boldsymbol{\mu}\| \cos(\boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}) - \frac{\hat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|} \right). \end{aligned}$$

By the symmetry of the distribution of  $y$  and the fact that  $y \mathbf{h}_0 \sim \boldsymbol{\mu} + \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ , if we condition on  $y = -1$ ,

$$\begin{aligned} \text{Err}_{-1} &= \mathbb{P}(y(\langle \mathbf{h}_0, \hat{\boldsymbol{\theta}} \rangle + \hat{\gamma}) < 0 | y = -1) = \mathbb{P}(\langle \boldsymbol{\mu} + \mathbf{G}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta} \rangle - \hat{\gamma} < 0) \\ &= \Phi \left( -\|\boldsymbol{\mu}\| \cos(\boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}) + \frac{\hat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|} \right). \end{aligned}$$

Then, we will next analyze the two terms

$$\cos(\boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}) \quad \text{and} \quad \frac{\hat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|}$$

to see how they will affect the misclassification error.

**Step 4: finalizing conclusions.** We consider the following regimes.

**Case 1:**  $\limsup_{n \rightarrow +\infty} \lambda_n \leq \bar{\lambda} < +\infty$ . In this case, by Lemma C.1, we have  $\liminf_{n \rightarrow +\infty} \kappa_{\lambda_n} \geq c_\kappa > 0$ . Then,

$$\begin{aligned} \cos(\boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}) &= 1 + o_{\mathbb{P}}(1) \\ \frac{\widehat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|} &= \frac{\widehat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu}\|} + o_{\mathbb{P}}(1). \end{aligned}$$

Denote  $q(t) = \frac{1}{1+e^t}$ , since  $\mathbf{h}_0 y \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p)$ , one can check that for any  $\boldsymbol{\xi} = (\gamma, \boldsymbol{\theta}^\top)^\top$  and  $2 \leq j \leq p+1$ ,

$$\begin{aligned} [\mathbf{I}(\boldsymbol{\xi})]_{j,1} &= \frac{1}{2} \mathbb{E}[q(\mathbf{h}_0^\top \boldsymbol{\theta} y + \gamma y)(1 - q(\mathbf{h}_0^\top \boldsymbol{\theta} y + \gamma y))(\mathbf{h}_0)_j | y = 1] \\ &\quad + \frac{1}{2} \mathbb{E}[q(\mathbf{h}_0^\top \boldsymbol{\theta} y + \gamma y)(1 - q(\mathbf{h}_0^\top \boldsymbol{\theta} y + \gamma y))(\mathbf{h}_0)_j | y = -1] = 0. \end{aligned}$$

Then, by (25), we have

$$\sqrt{ni(\boldsymbol{\theta}_n^*)} \widehat{\gamma} \xrightarrow{d} \mathcal{N}(0, 1),$$

where  $i(\boldsymbol{\theta}_n^*) := [\mathbf{I}(\boldsymbol{\xi}_n^*)]_{1,1}$  with  $\gamma^* = 0$  and  $[\mathbf{I}(\boldsymbol{\theta}_n^*)]_{1,1}$  denotes the submatrix consisting of the first column and the last row of  $\mathbf{I}(\boldsymbol{\theta}_n^*)$ . Since  $\liminf_{n \rightarrow +\infty} i(\boldsymbol{\theta}_n^*) > 0$ , we have  $\widehat{\gamma} = [ni(\boldsymbol{\theta}_n^*)]^{-1/2} G' + o_{\mathbb{P}}(1)$  with  $G' \sim \mathcal{N}(0, 1)$ , then

$$\text{Err} = \frac{1}{2} \text{Err}_1 + \frac{1}{2} \text{Err}_{-1} = \Phi \left( -\|\boldsymbol{\mu}\| + \frac{[i(\boldsymbol{\theta}_n^*)]^{-1/2} G'}{\sqrt{n} \kappa_{\lambda_n} \|\boldsymbol{\mu}\|} \right) + o_{\mathbb{P}}(1).$$

From the reasoning above, we have  $\kappa_{\lambda_n}^{-1} [ni(\boldsymbol{\theta}_n^*)]^{-1/2} = o(1)$ , then

$$\text{Err} = \Phi(-\|\boldsymbol{\mu}\|) + o_{\mathbb{P}}(1).$$

**Case 2:**  $\limsup_{n \rightarrow +\infty} \lambda_n = +\infty$ . Without loss of generality, consider  $\lambda_n \rightarrow +\infty$  as  $n \rightarrow +\infty$ . By Lemma C.1,  $\lambda_n \kappa_{\lambda_n} = \frac{1}{4}(1 + \eta)^{-2} + o(1)$ , thus we have  $\mathbf{I}(\boldsymbol{\xi}_n^*) = \frac{1}{4} \mathbf{I}_{p+1} + o(1)$ , particularly  $i(\boldsymbol{\theta}_n^*) = \frac{1}{4} + o(1)$ . By (25), we have  $\boldsymbol{\zeta} = O_{\mathbb{P}}(\kappa_{\lambda_n})$ , then similar with before

$$\begin{aligned} \cos(\boldsymbol{\mu}, \kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}) &= 1 + o_{\mathbb{P}}(1) \\ \frac{\widehat{\gamma}}{\|\kappa_{\lambda_n} \boldsymbol{\mu} + \frac{1}{\sqrt{n}} \boldsymbol{\zeta}\|} &= \frac{2}{\sqrt{n} \kappa_{\lambda_n} \|\boldsymbol{\mu}\|} + o_{\mathbb{P}}(1). \end{aligned}$$

1. If  $\sqrt{n}/\lambda_n = o(1)$ , then  $\sqrt{n} \kappa_{\lambda_n} = o(1)$ . We have  $\mathbb{E}[\text{Err}] \rightarrow 1$  as  $\lambda_n \rightarrow +\infty$ .
2. If  $\lambda_n = a \cdot \sqrt{n}$  with a positive constant  $a$ , then  $a\sqrt{n} \kappa_n = \frac{1}{4}(1 + \eta)^2 + o(1)$  and

$$\text{Err} = \Phi \left( -\|\boldsymbol{\mu}\| + \frac{8G'}{a(1 + \eta)^2 \|\boldsymbol{\mu}\|} \right) + o_{\mathbb{P}}(1).$$

We have the following lemma

**Lemma C.4.** Let  $\tau, w > 0$  and  $U \sim \mathcal{N}(0, 1)$ , then

$$\frac{d}{dw} \mathbb{E} \Phi(-\tau + w \cdot U) = \frac{\tau w}{\sqrt{2\pi}} \frac{1}{(1+w^2)^{3/2}} \exp\left(-\frac{\tau^2}{2(1+w^2)}\right) > 0.$$

Then,  $\mathbb{E}[\mathbf{Err}]$  is increasing in  $\frac{8}{a(1+\eta)^2\|\boldsymbol{\mu}\|}$ , thus is decreasing in  $\eta$ .

3. If  $\lambda_n = a \cdot b_n$  with a constant  $a > 0$  and  $b_n/\sqrt{n} = o(1)$ , then  $ab_n\kappa_n = \frac{1}{4}(1+\eta)^2 + o(1)$  and

$$\mathbf{Err} = \Phi\left(-\|\boldsymbol{\mu}\| + \frac{8G'}{a(1+\eta)^2\|\boldsymbol{\mu}\|} \frac{b_n}{\sqrt{n}}\right) + o_{\mathbb{P}}(1) = \Phi(-\|\boldsymbol{\mu}\|) + O_{\mathbb{P}}\left(\frac{b_n}{\sqrt{n}}\right).$$

As a result,  $\mathbb{E}[\mathbf{Err}] = \Phi(-\|\boldsymbol{\mu}\|) + O(b_n/\sqrt{n})$ .

Combining pieces above, low-dimensional logistic regression with  $\ell_2$  regularization has decreasing prediction error in the expansion parameter  $\eta$  only if  $\lambda_n \asymp \sqrt{n}$ . Otherwise,  $\mathbf{Err}$  asymptotically remains constant in  $\eta$ .

### C.1.1 Proof of Lemma C.1

Since the population loss function  $\tilde{\ell}(\boldsymbol{\xi})$  is strongly convex in  $\boldsymbol{\xi}$ , it suffices to show that (i) there exists some  $\kappa := \kappa_{\lambda_n} \in (0, C)$  such that  $\boldsymbol{\xi} = (0, \kappa\boldsymbol{\mu}^\top)^\top$  satisfies the score equation

$$\mathbb{E}\left\{\frac{\bar{\mathbf{h}}_0 y}{1 + \exp(y\bar{\mathbf{h}}_0^\top \boldsymbol{\xi})}\right\} = 2\lambda_n \bar{\boldsymbol{\Omega}} \boldsymbol{\xi}; \quad (27)$$

and that (ii)  $\lim_{n \rightarrow \infty} \lambda_n \kappa_{\lambda_n} = \frac{1}{4}(1+\eta)^{-2}$ . Due to the block structure of  $\bar{\boldsymbol{\Omega}}$ , we separate the verification process into two parts.

**Step 1: check the first coordinate.** Clearly, the first coordinate on the right-hand side of (27) is zero. When it comes to the left-hand side, we have

$$\mathbb{E}\left\{\frac{y}{1 + \exp(y\bar{\mathbf{h}}_0^\top \boldsymbol{\xi})}\right\} = \frac{1}{2} \mathbb{E}\left\{\frac{1}{1 + \exp(\kappa\boldsymbol{\mu}^\top(\boldsymbol{\mu} + \mathbf{G}))}\right\} + \frac{1}{2} \mathbb{E}\left\{\frac{-1}{1 + \exp(\kappa\boldsymbol{\mu}^\top(\boldsymbol{\mu} + \mathbf{G}))}\right\} = 0. \quad (28)$$

Here,  $\mathbf{G} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$ . Thus, (28) holds in the first coordinate for  $\boldsymbol{\xi} = (0, \kappa\boldsymbol{\mu}^\top)^\top$ .

**Step 2: check the remaining coordinates.** Now we move on to the remaining coordinates, i.e., the coordinates corresponding to  $\mathbf{h}_0$ .

The right-hand side of (27) obeys

$$2\lambda_n \bar{\boldsymbol{\Omega}} \boldsymbol{\theta} = 2\lambda_n (1+\eta)^{-2} \kappa \boldsymbol{\mu},$$

by the definition of  $\bar{\boldsymbol{\Omega}}$ . For the left hand side, since  $\mathbf{h}_0 y \stackrel{d}{=} \boldsymbol{\mu} + \mathbf{G} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p)$ , we have

$$\begin{aligned} \mathbb{E}\left\{\frac{\mathbf{h}_0 y}{1 + \exp(y\kappa\mathbf{h}_0^\top \boldsymbol{\mu})}\right\} &= \mathbb{E}\left\{\frac{\boldsymbol{\mu} + \mathbf{G}}{1 + \exp(\kappa\|\boldsymbol{\mu}\|^2 + \kappa\boldsymbol{\mu}^\top \mathbf{G})}\right\} \\ &= \mathbb{E}\left\{\frac{1}{1 + \exp(\kappa\|\boldsymbol{\mu}\|^2 + \kappa\boldsymbol{\mu}^\top \mathbf{G})}\right\} \boldsymbol{\mu} + \mathbb{E}\left\{\frac{\mathbf{G}}{1 + \exp(\kappa\|\boldsymbol{\mu}\|^2 + \kappa\boldsymbol{\mu}^\top \mathbf{G})}\right\}. \end{aligned}$$

Apply Stein's identity to see that

$$\mathbb{E} \left\{ \frac{\mathbf{G}}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})} \right\} = -\kappa \mathbb{E} \left\{ \frac{\exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})}{(1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G}))^2} \right\} \boldsymbol{\mu},$$

which leads to the conclusion

$$\mathbb{E} \left\{ \frac{\mathbf{h}_0 y}{1 + \exp(y \kappa \mathbf{h}_0^\top \boldsymbol{\mu})} \right\} = \mathbb{E} \left\{ \frac{1 + (1 - \kappa) \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})}{(1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G}))^2} \right\} \boldsymbol{\mu}. \quad (29)$$

Then it boils down to showing that there exists a constant  $\kappa > 0$  such that

$$\mathbb{E} \left\{ \frac{1 + (1 - \kappa) \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})}{(1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G}))^2} \right\} = 2\lambda_n (1 + \eta)^{-2} \kappa.$$

Recall the function

$$\psi(\kappa) = \mathbb{E} \left\{ \frac{1 + (1 - \kappa) \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})}{(1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G}))^2} \right\} - 2\lambda_n (1 + \eta)^{-2} \kappa.$$

We see that  $\psi(0) = \frac{1}{2} > 0$ .

(i) First, we claim that there exists  $C > 0$  depending solely on  $\|\boldsymbol{\mu}\|$  such that  $\psi(C) < 0$ . Once this is proved, since  $\psi(\kappa)$  is a continuous function, by the intermediate value theorem, there must exist  $\kappa \in (0, C)$  such that  $\psi(\kappa) = 0$ .

To prove this claim, we use the identity (29) to express  $\psi(\kappa)$  into

$$\begin{aligned} \psi(\kappa) &= \|\boldsymbol{\mu}\|^{-2} \cdot \mathbb{E} \left\{ \frac{\|\boldsymbol{\mu}\|^2 + \mathbf{G}^\top \boldsymbol{\mu}}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \boldsymbol{\mu}^\top \mathbf{G})} \right\} - 2\lambda_n (1 + \eta)^{-2} \kappa \\ &\leq \|\boldsymbol{\mu}\|^{-2} \cdot \mathbb{E} \left\{ \frac{\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)} \right\} \end{aligned}$$

where  $G \sim \mathcal{N}(0, 1)$ . Denote an event  $A = \{\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G < 0\}$ . For any fixed  $\|\boldsymbol{\mu}\| > 0$ , the probability  $\mathbb{P}(A) \in (0, 1)$ . By the dominated convergence theorem,

$$\begin{aligned} \lim_{\kappa \rightarrow \infty} \mathbb{E} \left\{ \frac{(\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G)_+}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)} \right\} &= 0, \\ \lim_{\kappa \rightarrow \infty} \mathbb{E} \left\{ \frac{(\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G)_-}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)} \right\} &= \mathbb{E}[(\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G)_-] > 0. \end{aligned}$$

Combining the two parts, we deduce that  $\limsup_{\kappa \rightarrow \infty} \psi(\kappa) \leq -\mathbb{E}[(\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G)_-] < 0$ .

(ii) Since for any  $a \in \mathbb{R}$ ,  $a < \kappa^{-1} + a \leq \kappa^{-1} \exp(\kappa a)$ , we have

$$\mathbb{E} \left\{ \frac{\|\boldsymbol{\mu}\|^2 + \|\boldsymbol{\mu}\| G}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)} \right\} \leq \kappa^{-1} \mathbb{E} \left\{ \frac{\exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)}{1 + \exp(\kappa \|\boldsymbol{\mu}\|^2 + \kappa \|\boldsymbol{\mu}\| G)} \right\} < \kappa^{-1}.$$

The equation  $\psi(\kappa_{\lambda_n}) = 0$  must imply

$$2\lambda_n (1 + \eta)^{-2} \kappa_{\lambda_n} < \frac{1}{\kappa_{\lambda_n} \|\boldsymbol{\mu}\|^2}.$$

If  $\lambda_n \rightarrow \infty$  as  $n \rightarrow \infty$ , then we must have  $\kappa_{\lambda_n} = o(1)$ . Taking the limit  $\lim_{n \rightarrow \infty} \psi(\kappa_{\lambda_n})$ , we get

$$\frac{1}{2} - 2(1 + \eta)^{-2} \lim_{n \rightarrow \infty} \lambda_n \kappa_{\lambda_n} = 0$$

so  $\lim_{n \rightarrow \infty} \lambda_n \kappa_{\lambda_n} = \frac{1}{4} (1 + \eta)^2$ .

Combining the two steps finishes the proof.

### C.1.2 Proof of Lemma C.2

Recall that for any  $\lambda_n \geq 0$  and any  $\boldsymbol{\xi} = (\gamma, \boldsymbol{\theta}^\top)^\top \in \mathbb{R}^{p+1}$ ,

$$\ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n) = \log \left[ 1 + \exp(-y(\gamma + \mathbf{h}_0^\top \boldsymbol{\theta})) \right] + \lambda_n \boldsymbol{\theta}^\top \boldsymbol{\Omega} \boldsymbol{\theta}.$$

**Step 1: Applying the uniform law of large numbers.** Since the term  $\lambda_n \boldsymbol{\theta}^\top \boldsymbol{\Omega} \boldsymbol{\theta}$  is canceled out, we need to show that for any  $R$  and the compact ball  $B_R = \{\|\boldsymbol{\xi}\| \leq R\}$ ,

$$\sup_{\boldsymbol{\xi} \in B_R} \left| (\mathbb{E}_n - \mathbb{E}) \left\{ \log \left[ 1 + \exp(-y(\gamma + \mathbf{h}_0^\top \boldsymbol{\theta})) \right] \right\} \right| \xrightarrow{p} 0. \quad (30)$$

We see that  $\ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n)$  is continuous in  $\boldsymbol{\xi}$  for any  $\mathbf{h}_0$  and is measurable in  $(\mathbf{h}_0, y)$  for any  $\boldsymbol{\xi}$ . In addition, for any  $R$ ,

$$\log \left[ 1 + \exp(-y(\gamma + \mathbf{h}_0^\top \boldsymbol{\theta})) \right] \leq \log \left[ 1 + \exp(R\|\bar{\mathbf{h}}_0\|) \right] \quad \text{with} \quad \mathbb{E} \left\{ \log \left[ 1 + \exp(R\|\bar{\mathbf{h}}_0\|) \right] \right\} < +\infty.$$

Then, we can apply the uniform law of large numbers (Newey and McFadden, 1994, Lemma 2.4) to show (30).

It is similar for the convergence of Fisher information as for any  $\boldsymbol{\xi} \in \mathbb{R}^{p+1}$ ,

$$\frac{\bar{\mathbf{h}}_0 \bar{\mathbf{h}}_0^\top \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi})}{(1 + \exp(y \bar{\mathbf{h}}_0^\top \boldsymbol{\xi}))^2} \leq \frac{1}{4} \bar{\mathbf{h}}_0 \bar{\mathbf{h}}_0^\top \quad \text{with} \quad \mathbb{E} \left\{ \frac{1}{4} \bar{\mathbf{h}}_0 \bar{\mathbf{h}}_0^\top \right\} = \frac{1}{4} \mathbf{I}_{p+1}.$$

Then, we have

$$\sup_{\boldsymbol{\xi} \in B_R} \left| \mathbf{I}_n(\boldsymbol{\xi}) - \mathbf{I}(\boldsymbol{\xi}) \right| \xrightarrow{p} 0.$$

**Step 2: Proving the convergence of the minimizer.** With  $\widehat{\boldsymbol{\xi}}_n$  and  $\boldsymbol{\xi}_n^*$  defined before, there exists a constant  $R > 0$  such that  $\mathbb{P}(\widehat{\boldsymbol{\xi}}_n \notin B_R) = o(1)$ . Therefore,

$$\ell_n^*(\widehat{\boldsymbol{\xi}}_n) - \ell_n^*(\boldsymbol{\xi}_n^*) = \ell_n^*(\widehat{\boldsymbol{\xi}}_n) - \tilde{\ell}_n(\widehat{\boldsymbol{\xi}}_n; \mathbf{h}_0) + \tilde{\ell}_n(\widehat{\boldsymbol{\xi}}_n; \mathbf{h}_0) - \tilde{\ell}_n(\boldsymbol{\xi}_n^*; \mathbf{h}_0) + \tilde{\ell}_n(\boldsymbol{\xi}_n^*; \mathbf{h}_0) - \ell_n^*(\boldsymbol{\xi}_n^*).$$

Here, we have  $\ell_n^*(\widehat{\boldsymbol{\xi}}_n) - \tilde{\ell}_n(\widehat{\boldsymbol{\xi}}_n; \mathbf{h}_0) \leq \sup_{\boldsymbol{\xi} \in B_R} |(\mathbb{E}_n - \mathbb{E})\ell(\boldsymbol{\xi}; \mathbf{h}_0, \lambda_n)| \xrightarrow{p} 0$ ,  $\tilde{\ell}_n(\widehat{\boldsymbol{\xi}}_n; \mathbf{h}_0) - \tilde{\ell}_n(\boldsymbol{\xi}_n^*; \mathbf{h}_0) \leq 0$  by the optimality of  $\widehat{\boldsymbol{\xi}}_n$ , and  $\tilde{\ell}_n(\boldsymbol{\xi}_n^*; \mathbf{h}_0) - \ell_n^*(\boldsymbol{\xi}_n^*) = o_{\mathbb{P}}(1)$  by the law of large numbers. Consequently, we have  $\ell_n^*(\widehat{\boldsymbol{\xi}}_n) - \ell_n^*(\boldsymbol{\xi}_n^*) \leq o_{\mathbb{P}}(1)$ . Since  $\ell_n^*(\boldsymbol{\xi})$  is strictly convex and by definition  $\boldsymbol{\xi}_n^* = \operatorname{argmin}_{\boldsymbol{\xi} \in \mathbb{R}^{p+1}} \ell_n^*(\boldsymbol{\xi})$ , we further have  $\widehat{\boldsymbol{\xi}}_n - \boldsymbol{\xi}_n^* = o_{\mathbb{P}}(1)$ .

### C.1.3 Proof of Lemma C.3

Denoting

$$\mathbf{z}_{n,i} = \frac{\bar{\mathbf{h}}_{0,i} y_i}{1 + \exp(y_i \bar{\mathbf{h}}_{0,i}^\top \boldsymbol{\xi}_n^*)},$$

then it suffices to show that for any  $\varepsilon > 0$ ,

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\{ \|\mathbf{z}_{n,i}\|^2 \mathbb{1}\{\|\mathbf{z}_{n,i}\| \geq \varepsilon \sqrt{n}\} \right\} = \mathbb{E} \left\{ \|\mathbf{z}_{n,1}\|^2 \mathbb{1}\{\|\mathbf{z}_{n,1}\| \geq \varepsilon \sqrt{n}\} \right\} \xrightarrow{p} 0.$$



Let  $Z_{n,i} = \|\mathbf{z}_{n,i}\|^2$ , then  $\mathbb{E}Z_{n,i} = \text{Tr}(\mathbb{E}(\mathbf{z}_{n,i}\mathbf{z}_{n,i}^\top)) = \text{Tr}(\mathbf{I}(\boldsymbol{\xi}_n^*))$ . Let  $W_{n,i} = Z_{n,i}\mathbb{1}\{Z_{n,i} \geq \epsilon^2 n\}$ . We can verify that  $0 \leq W_{n,i} \leq Z_{n,i}$  and for any  $\delta > 0$ , by Markov's inequality,

$$\mathbb{P}(W_{n,i} \geq \delta) = \mathbb{P}(Z_{n,i} \geq \epsilon^2 n) \leq \frac{\text{Tr}(\mathbf{I}(\boldsymbol{\xi}_n^*))}{n\epsilon^2} \rightarrow 0, \quad \text{as } n \rightarrow \infty,$$

which shows that  $W_{n,i} \xrightarrow{p} 0$ . Then, by the dominated convergence theorem, we have

$$\lim_{n \rightarrow +\infty} \mathbb{E}W_{n,i} = \mathbb{E} \lim_{n \rightarrow +\infty} W_{n,i} = 0,$$

which verifies Lindeberg's condition. By the Lindeberg-Feller theorem for multivariate random variables, we have

$$\sqrt{n} [\mathbf{I}(\boldsymbol{\xi}_n^*)]^{-1/2} (\mathbb{E}_n \mathbf{z}_{n,i} - \mathbb{E} \mathbf{z}_{n,i}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

## C.2 High-dimensional regime: max-margin classifier

Written in terms of the input data  $\{\mathbf{h}_{0,i}\}$ , the max-margin classifier (5) is equivalent to

$$\begin{aligned} & \max_{\boldsymbol{\beta} \in \mathbb{R}^p} \min_{i \leq n} y_i \langle \boldsymbol{\beta}, \mathbf{h}_{0,i} \rangle \\ & \text{subject to } \|\boldsymbol{\beta}\|_{\boldsymbol{\Omega}} = \sqrt{\boldsymbol{\beta}^\top \boldsymbol{\Omega} \boldsymbol{\beta}} \leq 1, \end{aligned} \tag{31}$$

where  $\boldsymbol{\Omega} := (\mathbf{I} + \eta \boldsymbol{\mu}_0 \boldsymbol{\mu}_0^\top)^{-2} = \mathbf{I} - \left(1 - \frac{1}{(1+\eta)^2}\right) \boldsymbol{\mu}_0 \boldsymbol{\mu}_0^\top$ , and  $\boldsymbol{\mu}_0 = \frac{1}{\|\boldsymbol{\mu}\|} \boldsymbol{\mu}$ .

### C.2.1 Preliminaries

Before presenting the proof of Theorem 5.2, we first introduce important quantities and their equivalence.

Define a compact set  $\Theta = \{\boldsymbol{\beta} \in \mathbb{R}^p : \boldsymbol{\beta}^\top \boldsymbol{\Omega} \boldsymbol{\beta} \leq 1\}$ . Similar to the development in Montanari et al. (2019), for any positive margin  $\kappa > 0$ , we have the equivalence

$$\begin{aligned} \mathcal{E}_{n,p,\kappa} &= \{\text{there exists } \boldsymbol{\beta} \in \Theta : y_i \langle \mathbf{h}_{0,i}, \boldsymbol{\beta} \rangle \geq \kappa \text{ for all } i \leq n\} \\ &= \left\{ \min_{\boldsymbol{\beta} \in \Theta} \max_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{H}_0 \boldsymbol{\beta}) = 0 \right\}. \end{aligned}$$

To simplify the notation, for any set  $\mathbf{B}_p \subseteq \mathbb{R}^p$ , we define the quantity

$$\xi_{n,p,\kappa}(\mathbf{B}_p) = \min_{\boldsymbol{\beta} \in \mathbf{B}_p} \max_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{H}_0 \boldsymbol{\beta}).$$

In particular, we set  $\xi_{n,p,\kappa} := \xi_{n,p,\kappa}(\Theta)$ , which obeys

$$\{\xi_{n,p,\kappa} > 0\} \iff \mathcal{E}_{n,p,\kappa}^c \quad \text{and} \quad \{\xi_{n,p,\kappa} = 0\} \iff \mathcal{E}_{n,p,\kappa}.$$

As a result, our goal is to analyze whether  $\xi_{n,p,\kappa}$  is positive.

**Step 1: Applying CGMT.** Recall the Gaussian mixture model where  $\mathbf{H}_0 = \mathbf{y}\boldsymbol{\mu}^\top + \mathbf{G}$  and each entry of  $\mathbf{G} \in \mathbb{R}^{n \times p}$  is independently drawn from  $\mathcal{N}(0, 1)$ . As a result, we have  $\boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{H}_0 \boldsymbol{\beta}) = \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \mathbf{y} - \mathbf{G} \boldsymbol{\beta})$ , which is a bilinear form of the Gaussian random matrix  $\mathbf{G}$ . We plan to use Gordon's comparison inequality to simplify the calculation of  $\xi_{n,p,\kappa}(\mathbf{B}_p)$ . To do so, we introduce another quantity for an arbitrary set  $\mathbf{B}_p$ :

$$\xi_{n,p,\kappa}^{(1)}(\mathbf{B}_p) = \min_{\boldsymbol{\beta} \in \mathbf{B}_p} \max_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \left\{ \|\boldsymbol{\lambda}\| \mathbf{g}^\top \boldsymbol{\beta} + \|\boldsymbol{\beta}\| \mathbf{h}^\top \boldsymbol{\lambda} + \kappa \boldsymbol{\lambda}^\top \mathbf{y} - \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \boldsymbol{\lambda}^\top \mathbf{y} \right\},$$

where  $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$  and  $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$  are two independent Gaussian vectors. Similar to before, we denote  $\xi_{n,p,\kappa}^{(1)} := \xi_{n,p,\kappa}^{(1)}(\Theta)$ . The following lemma connects  $\xi_{n,p,\kappa}(\mathbf{B}_p)$  with  $\xi_{n,p,\kappa}^{(1)}(\mathbf{B}_p)$ , which is a simple corollary of Gordon's comparison inequality (Gordon, 1988; Thrampoulidis et al., 2015; Montanari et al., 2019).

**Lemma C.5.** *For any  $t \in \mathbb{R}$  and any compact set  $\mathbf{B}_p$ ,*

$$\begin{aligned} \mathbb{P}(\xi_{n,p,\kappa} \leq t) &\leq 2\mathbb{P}\left(\xi_{n,p,\kappa}^{(1)} \leq t\right) \quad \text{and} \quad \mathbb{P}(\xi_{n,p,\kappa} \geq t) \leq 2\mathbb{P}\left(\xi_{n,p,\kappa}^{(1)} \geq t\right). \\ \mathbb{P}(\xi_{n,p,\kappa}(\mathbf{B}_p) \leq t) &\leq 2\mathbb{P}\left(\xi_{n,p,\kappa}^{(1)}(\mathbf{B}_p) \leq t\right). \end{aligned}$$

*If in addition  $\mathbf{B}_p$  is convex, then*

$$\mathbb{P}(\xi_{n,p,\kappa}(\mathbf{B}_p) \geq t) \leq 2\mathbb{P}\left(\xi_{n,p,\kappa}^{(1)}(\mathbf{B}_p) \geq t\right).$$

**Step 2: Connecting  $\xi_{n,p,\kappa}^{(1)}$  with  $\tilde{\xi}_{n,p,\kappa}^{(1)}$ .** We then move on to characterizing  $\xi_{n,p,\kappa}^{(1)}$  by translating the min-max problem into a single minimization problem. For any set  $\mathbf{B} \subseteq \mathbb{R} \times \mathbb{R}_+$ , we define a quantity

$$\tilde{\xi}_{n,p,\kappa}^{(1)}(\mathbf{B}) = \min_{(\gamma, z) \in \mathbf{B}} \frac{1}{\sqrt{p}} \left\{ -z \|\mathbf{g}\| + \left\| \left( \sqrt{\gamma^2 + z^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}.$$

In particular, we set  $\tilde{\xi}_{n,p,\kappa}^{(1)} := \tilde{\xi}_{n,p,\kappa}^{(1)}(\tilde{\Theta}_\eta)$ , where  $\tilde{\Theta}_\eta := \{(\gamma, z) \in \mathbb{R} \times \mathbb{R}_+ : (1 + \eta)^{-2} \gamma^2 + z^2 \leq 1\}$ . We have the following relation between  $\tilde{\xi}_{n,p,\kappa}^{(1)}$  and  $\xi_{n,p,\kappa}^{(1)}$ .

**Lemma C.6.** *With the definition above, we have*

$$\left| \left( \tilde{\xi}_{n,p,\kappa}^{(1)} \right)_+ - \xi_{n,p,\kappa}^{(1)} \right| \xrightarrow{p} 0.$$

See Appendix C.2.5 for the proof.

**Step 3: Connecting  $\tilde{\xi}_{n,p,\kappa}^{(1)}$  with  $\xi_{n,p,\kappa}^{(2)}$ .** It turns out that  $\tilde{\xi}_{n,p,\kappa}^{(1)}$  can be further simplified for analytical purposes. For any set  $\mathbf{B} \subseteq \mathbb{R} \times \mathbb{R}_+$ , we define a new quantity

$$\xi_{n,p,\kappa}^{(2)}(\mathbf{B}) = \min_{(\gamma, z) \in \mathbf{B}} \left\{ -z + \sqrt{\delta} \sqrt{\mathbb{E} \left( \sqrt{\gamma^2 + z^2} G + \kappa - \gamma \|\boldsymbol{\mu}\| \right)_+^2} \right\}.$$

Similar to before, we simply denote  $\xi_{n,p,\kappa}^{(2)} := \xi_{n,p,\kappa}^{(2)}(\tilde{\Theta}_\eta)$ . The two quantities of interest can be related via the uniform law of large numbers as is shown in the following lemma.

**Lemma C.7.** *With the definition of  $\xi_{n,p,\kappa}^{(2)}$ , we have*

$$\left| \xi_{n,p,\kappa}^{(1)} - \left( \xi_{n,p,\kappa}^{(2)} \right)_+ \right| \xrightarrow{p} 0.$$

See Appendix C.2.6 for the proof.

Following the chain of equivalence, we switch our goal to study  $\xi_{n,p,\kappa}^{(2)}$ .

**Step 4: Analyzing  $\xi_{n,p,\kappa}^{(2)}$ .** In view of the chain of arguments above, analyzing the margin boils down to analyzing the positivity of  $\xi_{n,p,\kappa}^{(2)}$ , which is equivalent to the positivity of the following function

$$F_\delta(\gamma, z, \kappa) := -z^2 + \delta \mathbb{E} \left( \sqrt{\gamma^2 + z^2} G + \kappa - \gamma \|\boldsymbol{\mu}\| \right)_+^2. \quad (32)$$

To do so, the following lemma is useful in which  $\mathbf{B}_v := \{\boldsymbol{\beta} \in \mathbb{R}^p : \|\mathbf{P}_\mu^\perp \boldsymbol{\beta}\| \leq v \|\mathbf{P}_\mu \boldsymbol{\beta}\|\}$  is a convex cone for each choice of  $v \geq 0$ .

**Lemma C.8.** *Define the event*

$$\mathcal{E}_{n,p,\kappa,v} := \{ \text{there exists } \boldsymbol{\beta} \in \Theta \cap \mathbf{B}_v : y_i \langle \mathbf{h}_{0,i}, \boldsymbol{\beta} \rangle \geq \kappa \text{ for all } i \leq n \}.$$

1. *If there exists  $(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq v\gamma\}$  such that  $F_\delta(\gamma, z, \kappa) < 0$ , then*

$$\mathbb{P}(\mathcal{E}_{n,p,\kappa,v}) = 1 - o(1).$$

2. *If for any  $(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq v\gamma\}$ , we have  $F_\delta(\gamma, z, \kappa) > 0$ , then*

$$\mathbb{P}(\mathcal{E}_{n,p,\kappa,v}^c) = 1 - o(1).$$

See Appendix C.2.7 for the proof.

*Remark.* This proposition still holds when  $v = +\infty$ . This will be useful for proving Theorem 5.2 (1).

**Step 5: Analyzing  $F_\delta(\gamma, z, \kappa)$  via an equivalent form.** Lemma C.8 motivates us to focus on the positivity of function  $F_\delta$  (32). To this end, we define a helpful function  $f_\delta : \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}_+ \mapsto \mathbb{R}$

$$f_\delta(u, \kappa, c) = -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2} G + \kappa \sqrt{u^2 + c} - \|\boldsymbol{\mu}\| \right)_+^2 \right].$$

Fix any  $c > 0$ , we consider the optimization problem

$$\begin{aligned} & \sup_{\kappa \geq 0} \kappa \\ & \text{subject to } \inf_{u \geq 0} f_\delta(u, \kappa, c) \leq 0. \end{aligned} \quad (33)$$

In fact,  $c = \frac{1}{(1+\eta)^2}$  is just a reparametrization of  $\eta$ . Denote by  $\kappa^* = \kappa^*(c)$  the corresponding maximizer and  $u^* = u^*(c)$ , the smallest minimizer of  $\inf_{u \geq 0} f_\delta(u, \kappa^*(c), c)$ . The following lemma demonstrates that both are well defined.

**Lemma C.9.** *For every  $c > 0$ , one has  $0 \leq \kappa^*(c) < +\infty$ , and  $0 < u^*(c) < +\infty$ . And they attain the corresponding maximum and minimum.*

See Appendix C.2.8 for the proof.

Both  $\kappa^*$  and  $u^*$  play an essential role in our later developments.

### C.2.2 Proof of Theorem 5.2 (1).

We plan to invoke Lemma C.8 part 2 to prove this claim. More precisely, we aim to show that for any fixed  $\kappa > 0$ , when  $\delta > \delta^*(\rho)$  with<sup>5</sup>

$$\delta^*(\rho) = \left( \inf_{u>0} \frac{1}{u^2} \mathbb{E} \left\{ \left( \sqrt{1+u^2}G - \|\boldsymbol{\mu}\| \right)_+^2 \right\} \right)^{-1},$$

one has  $\inf_{(\gamma,z) \in \tilde{\Theta}_\eta} F_\delta(\gamma, z, \kappa) > 0$ .

**Case 1:**  $\gamma = 0$ . In this case, we have  $F_\delta(0, z, \kappa) = -z^2 + \delta \mathbb{E} \left[ (zG + \kappa)_+^2 \right]$ . If in addition  $z = 0$ , one has  $F_\delta(0, 0, \kappa) > 0$ . Otherwise for any  $0 < z \leq 1$ , one has

$$F_\delta(0, z, \kappa) = z^2 \left( -1 + \delta \mathbb{E} \left[ \left( G + \frac{\kappa}{z} \right)_+^2 \right] \right).$$

Note that since  $z \leq 1$ , we have

$$\begin{aligned} -1 + \delta \mathbb{E} \left[ \left( G + \frac{\kappa}{z} \right)_+^2 \right] &\geq -1 + \delta \mathbb{E} \left[ (G + \kappa)_+^2 \right] \\ &> -1 + \delta \mathbb{E} \left[ (G)_+^2 \right] \quad \text{since } \kappa > 0 \\ &> -1 + \delta^*(\rho) \frac{1}{2} \geq 0, \end{aligned}$$

where the last relation uses the fact that  $\delta^*(\rho) \geq 2$ .

**Case 2:**  $\gamma > 0$ . In this case, we have

$$F_\delta(\gamma, z, \kappa) = \gamma^2 \left\{ -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G + \frac{\kappa}{\gamma} - \|\boldsymbol{\mu}\| \right)_+^2 \right] \right\},$$

where we denote  $u = z/\gamma$ . It is easy to see that when  $u = 0$ , one has  $F_\delta(\gamma, z, \kappa) > 0$ . Therefore we focus on the case when  $u > 0$ .

By the definition of  $\delta^*(\rho)$ , we know that for any  $u > 0$ ,

$$\frac{1}{\delta} < \frac{1}{\delta^*(\rho)} \leq \frac{1}{u^2} \mathbb{E} \left[ \left( \sqrt{1+u^2}G - \|\boldsymbol{\mu}\| \right)_+^2 \right].$$

As a result, for any  $u > 0$ , one has

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G - \|\boldsymbol{\mu}\| \right)_+^2 \right] > 0,$$

---

<sup>5</sup>To see that  $\delta^*(\rho)$  is well-defined, let  $\alpha = 1/u$ . Then we need to check if

$$\inf_{\alpha \geq 0} \mathbb{E} \left\{ \left( \sqrt{1+\alpha^2}G - \|\boldsymbol{\mu}\| \alpha \right)_+^2 \right\} > 0.$$

At  $\alpha = 0$ ,  $\mathbb{E}[(G)_+^2] > 0$  indicates that the infimum is strictly positive in the neighborhood of 0; for  $\alpha \geq C_\alpha$ ,  $\inf_{\alpha \geq 0} \mathbb{E} \left\{ \left( \sqrt{1+\alpha^2}G - \|\boldsymbol{\mu}\| \alpha \right)_+^2 \right\} \geq \inf_{\alpha \geq 0} \alpha^2 \mathbb{E} \left\{ (G - \|\boldsymbol{\mu}\|)_+^2 \right\} \geq C_\alpha^2 \mathbb{E} \left\{ (G - \|\boldsymbol{\mu}\|)_+^2 \right\} > 0$ . Thus,  $\delta^*(\rho) < +\infty$  is well-defined. As a final remark, when  $\|\boldsymbol{\mu}\| = 0$  (i.e., the pure noise case), we have  $\delta^*(0) = 2$ . As a result,  $\delta^*(\rho) \geq 2$  for  $\rho \geq 0$ .

which further implies

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G + \frac{\kappa}{\gamma} - \|\mu\| \right)_+^2 \right] > 0.$$

This proves that  $F_\delta(\gamma, z, \kappa) > 0$  for all  $\gamma, z$ .

### C.2.3 Proof of Theorem 5.2 (2).

In view of Part 1 of Lemma C.8, we only need to show that for some  $\kappa > 0$ , there exist some  $(\gamma, z)$  such that  $F_\delta(\gamma, z, \kappa) < 0$ . Since  $\delta < \delta^*(\rho)$ , we have

$$\frac{1}{\delta} > \frac{1}{u^2} \mathbb{E} \left[ \left( \sqrt{1+u^2}G - \|\mu\| \right)_+^2 \right]$$

for some  $u > 0$ . As a result, we have

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G - \|\mu\| \right)_+^2 \right] < 0.$$

Due to continuity, we know that there exists some  $L > 0$ , such that

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G + \kappa\sqrt{u^2+c} - \|\mu\| \right)_+^2 \right] < 0, \quad \text{for all } \kappa < L.$$

Let  $\gamma = \frac{1}{\sqrt{u^2+c}}$ , and  $z = u \cdot \gamma$ . We then have

$$F_\delta(\gamma, z, \kappa) = \gamma^2 \left\{ -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1+u^2}G + \frac{L}{2}\sqrt{u^2+c} - \|\mu\| \right)_+^2 \right] \right\} < 0.$$

This finishes the proof.

We then turn to the claim regarding the margin. Recall that we work in the regime where  $\delta < \delta^*(\rho)$ . Fix any  $\varepsilon > 0$ . We aim to prove that  $\mathcal{E}_{n,p,\kappa^*+\varepsilon}$  does not hold and that  $\mathcal{E}_{n,p,\kappa^*-\varepsilon}$  holds.

**Step 1: Quantifying  $\mathcal{E}_{n,p,\kappa^*+\varepsilon}$ .** Similar to the proof of Theorem 5.2 (1), it suffices to prove that for all  $\gamma, z$ ,  $F_\delta(\gamma, z, \kappa^* + \varepsilon) > 0$ .

**Case 1:  $\gamma = 0$ .** It is simple to see that  $F_\delta(0, 0, \kappa^* + \varepsilon) > 0$ . Therefore we focus on the case for any  $0 < z \leq 1$ . Using similar arguments as in Theorem 5.2 (1), we have

$$F_\delta(0, z, \kappa^* + \varepsilon) = z^2 \left( -1 + \delta \mathbb{E} \left[ \left( G + \frac{\kappa^* + \varepsilon}{z} \right)_+^2 \right] \right).$$

Note that since  $z \leq 1$ , we have

$$-1 + \delta \mathbb{E} \left[ \left( G + \frac{\kappa^* + \varepsilon}{z} \right)_+^2 \right] \geq -1 + \delta \mathbb{E} \left[ (G + \kappa^*)_+^2 \right].$$

When  $2 \leq \delta < \delta^*(\rho)$ , we have  $\kappa^* \geq 0$ , and hence  $F_\delta(0, z, \kappa^* + \varepsilon) > 0$ . When  $\delta < 2$ , since  $\kappa^* \geq \kappa(\delta)$ , where  $\kappa(\delta)$  is the unique solution to

$$\frac{1}{\delta} = \mathbb{E} \left[ (G + \kappa)_+^2 \right],$$

we again have  $F_\delta(0, z, \kappa^* + \varepsilon) > 0$ .

**Case 2:**  $\gamma > 0$ . In this case, we have

$$F_\delta(\gamma, z, \kappa^* + \varepsilon) = \gamma^2 \left\{ -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2 G} + \frac{\kappa^* + \varepsilon}{\gamma} - \|\mu\| \right)_+^2 \right] \right\},$$

where we denote  $u = z/\gamma$ . By the definition of  $\kappa^*$ , we know that for all  $u \geq 0$ ,

$$f_\delta(u, \kappa^* + \varepsilon, c) = -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2 G} + (\kappa^* + \varepsilon) \sqrt{u^2 + c} - \|\mu\| \right)_+^2 \right] > 0.$$

For all  $0 < \gamma \leq \frac{1}{\sqrt{u^2 + c}}$ , we further have

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2 G} + \frac{\kappa^* + \varepsilon}{\gamma} - \|\mu\| \right)_+^2 \right] > -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2 G} + (\kappa^* + \varepsilon) \sqrt{u^2 + c} - \|\mu\| \right)_+^2 \right] > 0.$$

This finishes the claim.

**Step 2: Quantifying  $\mathcal{E}_{n,p,\kappa^*-\varepsilon}$ .** Similarly, for  $\kappa^* - \varepsilon$ , since  $f_\delta(u, \kappa, c)$  is increasing in  $\kappa$  and  $\inf_{u \geq 0} f_\delta(u, \kappa^*, c) = 0$ , there exists some  $u > 0$ , such that

$$-u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2 G} + (\kappa^* - \varepsilon) \sqrt{u^2 + c} - \|\mu\| \right)_+^2 \right] < 0.$$

Going through the same argument as in part 1 completes the proof of this part.

Combine the arguments above, for any  $\varepsilon > 0$ , we see that with probability approaching 1,  $\kappa^* - \varepsilon < \hat{\kappa} < \kappa^* + \varepsilon$ , which proves that  $\hat{\kappa} \xrightarrow{P} \kappa^*$ .

#### C.2.4 Proof of Theorem 5.2 (3).

We start with Lemma C.10 showing that  $u^* = u^*(c) = u^*(\rho, \eta)$ , defined previously, is indeed the limit of  $\hat{u} := \frac{\|\mathbf{P}_\mu^\perp \hat{\beta}\|}{\|\mathbf{P}_\mu \hat{\beta}\|}$ .

**Lemma C.10.** Assume that  $\delta < \delta^*(\rho)$ . Fix any  $c > 0$  and any  $\varepsilon > 0$ ,

$$\mathcal{E}_{n,p,\kappa^*,u^*+\varepsilon} \quad \text{and} \quad \mathcal{E}_{n,p,\kappa^*,u^*-\varepsilon}^c \quad \text{hold.}$$

As a result, we have

$$\hat{u} := \frac{\|\mathbf{P}_\mu^\perp \hat{\beta}\|}{\|\mathbf{P}_\mu \hat{\beta}\|} \xrightarrow{P} u^*.$$

See Appendix C.2.9 for the proof.

Based on the lemma, recall the definition  $\rho = \|\mu\|^2$ ,  $u^*(\rho, \eta) = u^*(c)$ , and

$$\text{Err}(\hat{\beta}; \eta) = \mathbb{P} \left( y' \langle \mathbf{h}'_0, \hat{\beta} \rangle < 0 \right) = \Phi \left( -\langle \mu, \frac{\hat{\beta}}{\|\hat{\beta}\|} \rangle \right), \quad \text{Err}^*(\eta) = \Phi \left( -\|\mu\| \sqrt{\frac{1}{1 + u^{*2}(\rho, \eta)}} \right).$$

From Lemma C.10, we have

$$\left| \langle \mu, \frac{\hat{\beta}}{\|\hat{\beta}\|} \rangle - \|\mu\| \sqrt{\frac{1}{1 + u^{*2}(\rho, \eta)}} \right| = o_{\mathbb{P}}(1),$$

which implies that

$$\text{Err}(\widehat{\boldsymbol{\beta}}; \eta) \xrightarrow{p} \text{Err}^*(\eta) = \Phi \left( -\|\boldsymbol{\mu}\| \sqrt{\frac{1}{1 + u^{*2}(\rho, \eta)}} \right).$$

Since  $\text{Err}^*(\eta)$  is monotonically decreasing in  $u^*(\rho, \eta)$ , then asymptotically, the classification error of the original max-margin solution is monotonically decreasing in  $u^*(\rho, \eta)$ . It remains to show that  $u^*(\rho, \eta)$  is monotonically decreasing in  $\eta$ , which is equivalent to showing that  $u^*(c)$  is monotonically increasing in  $c$ . This is provided in the following lemma.

**Lemma C.11.** *With the definition of  $(u^*, \kappa^*, c)$ , for any  $c > 0$ ,*

1.  $u^* = u^*(c)$  is a continuous function in  $c$ ;
2.  $u^* = u^*(c)$  is monotonically increasing in  $c$ .

See Appendix C.2.10 for the proof.

These finish the proof of Theorem 5.2.

### C.2.5 Proof of Lemma C.6

Recall that

$$\xi_{n,p,\kappa}^{(1)}(\mathbf{B}_p) = \min_{\boldsymbol{\beta} \in \mathbf{B}_p} \max_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \left\{ \|\boldsymbol{\lambda}\| \mathbf{g}^\top \boldsymbol{\beta} + \|\boldsymbol{\beta}\| \mathbf{h}^\top \boldsymbol{\lambda} + \kappa \boldsymbol{\lambda}^\top \mathbf{y} - \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \boldsymbol{\lambda}^\top \mathbf{y} \right\}.$$

Denote

$$\ell(r_\lambda, \boldsymbol{\lambda}_0) := \frac{1}{\sqrt{p}} \left\{ \|\boldsymbol{\lambda}\| \mathbf{g}^\top \boldsymbol{\beta} + \|\boldsymbol{\beta}\| \mathbf{h}^\top \boldsymbol{\lambda} + \kappa \boldsymbol{\lambda}^\top \mathbf{y} - \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle \boldsymbol{\lambda}^\top \mathbf{y} \right\}.$$

**Step 1: inner maximization.** We focus on the inner maximization problem, which is equivalent to

$$\max_{r_\lambda \in [0,1], \|\boldsymbol{\lambda}_0\|=1, \mathbf{y} \odot \boldsymbol{\lambda}_0 \geq 0} \ell(r_\lambda, \boldsymbol{\lambda}_0) = \frac{1}{\sqrt{p}} r_\lambda \mathbf{g}^\top \boldsymbol{\beta} + \frac{1}{\sqrt{p}} r_\lambda \left[ \|\boldsymbol{\beta}\| \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right]^\top (\mathbf{y} \odot \boldsymbol{\lambda}_0),$$

where  $\gamma = \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle / \|\boldsymbol{\mu}\|$ .

When  $r_\lambda = 0$ , we have  $\ell(r_\lambda, \boldsymbol{\lambda}_0) = 0$ . When  $r_\lambda > 0$ , since  $\|\mathbf{y} \odot \boldsymbol{\lambda}_0\| = \|\boldsymbol{\lambda}_0\| = 1$  and  $\mathbf{y} \odot \boldsymbol{\lambda}_0 \geq 0$ , we have for any vector  $\boldsymbol{\nu}$  with at least one nonnegative coordinate,

$$\max_{\|\boldsymbol{\lambda}_0\|=1, \mathbf{y} \odot \boldsymbol{\lambda}_0 \geq 0} \boldsymbol{\nu}^\top (\mathbf{y} \odot \boldsymbol{\lambda}_0) = \|\boldsymbol{\nu}_+\|. \quad (34)$$

In fact, letting  $\boldsymbol{\nu} = \|\boldsymbol{\beta}\| \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n$ , we see that  $\boldsymbol{\nu}$  has i.i.d. coordinates, and for each coordinate, the probability for it to be negative is given by

$$\mathbb{P}(\|\boldsymbol{\beta}\| g + \kappa - \gamma \|\boldsymbol{\mu}\| < 0) = \Phi \left( -\frac{\kappa}{\|\boldsymbol{\beta}\|} + \frac{\gamma \|\boldsymbol{\mu}\|}{\|\boldsymbol{\beta}\|} \right) \leq \Phi \left( \frac{\gamma \|\boldsymbol{\mu}\|}{\|\boldsymbol{\beta}\|} \right) \leq \Phi(\|\boldsymbol{\mu}\|).$$

As a result, with probability at least  $1 - \Phi(\|\boldsymbol{\mu}\|)^n$ ,  $\boldsymbol{\nu}$  has at least one nonnegative entry, and hence (34) holds.

Combine the two cases to arrive at the conclusion that: with probability at least  $1 - \Phi(\|\boldsymbol{\mu}\|)^n$ ,

$$\max_{r_\lambda \in [0,1], \|\boldsymbol{\lambda}_0\|=1, \mathbf{y} \odot \boldsymbol{\lambda}_0 \geq 0} \ell(r_\lambda, \boldsymbol{\lambda}_0) = \left( \frac{1}{\sqrt{p}} \left\{ \mathbf{g}^\top \boldsymbol{\beta} + \left\| (\|\boldsymbol{\beta}\| \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n)_+ \right\| \right\} \right)_+,$$

where we recall that  $\gamma = \langle \boldsymbol{\mu}, \boldsymbol{\beta} \rangle / \|\boldsymbol{\mu}\|$ .



**Step 2: outer minimization.** Now we turn to the outer minimization over  $\boldsymbol{\beta} \in \Theta = \{\boldsymbol{\beta} \in \mathbb{R}^p : \boldsymbol{\beta}^\top \boldsymbol{\Omega} \boldsymbol{\beta} \leq 1\}$ . Using a reparametrization  $\boldsymbol{\beta} = \gamma \boldsymbol{\mu}_0 + \boldsymbol{\beta}_\perp$ , where  $\boldsymbol{\beta}_\perp = r \cdot \boldsymbol{w}$  with  $z = \|\boldsymbol{\beta}_\perp\|$  and  $\boldsymbol{w} \perp \boldsymbol{\mu}$ ,  $\|\boldsymbol{w}\| = 1$ , we see that with high probability  $\xi_{n,p,\kappa}^{(1)}$  is equal to

$$\bar{\xi}_{n,p,\kappa}^{(1)} = \frac{1}{\sqrt{p}} \min_{(\gamma,z) \in \tilde{\Theta}_\eta} \min_{\boldsymbol{w} \perp \boldsymbol{\mu}, \|\boldsymbol{w}\|=1} \left\{ \gamma \boldsymbol{g}^\top \boldsymbol{\mu}_0 + z \boldsymbol{g}^\top \boldsymbol{w} + \left\| \left( \sqrt{\gamma^2 + z^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}_+.$$

Here we recall that  $\tilde{\Theta}_\eta = \{(\gamma, z) \in \mathbb{R} \times \mathbb{R}_+ : (1 + \eta)^{-2} \gamma^2 + z^2 \leq 1\}$ .

We can solve the inner minimization, that is

$$\begin{aligned} & \min_{\boldsymbol{w} \perp \boldsymbol{\mu}, \|\boldsymbol{w}\|=1} \left\{ \gamma \boldsymbol{g}^\top \boldsymbol{\mu}_0 + z \boldsymbol{g}^\top \boldsymbol{w} + \left\| \left( \sqrt{\gamma^2 + z^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}_+ \\ &= \left\{ \gamma \boldsymbol{g}^\top \boldsymbol{\mu}_0 - z \|\mathbf{P}_{\boldsymbol{\mu}}^\perp \boldsymbol{g}\| + \left\| \left( \sqrt{\gamma^2 + z^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}_+ \end{aligned}$$

Note that  $\gamma \boldsymbol{g}^\top \boldsymbol{\mu}_0 / \sqrt{p} = o_{\mathbb{P}}(1)$ . This together with the boundedness of  $\tilde{\Theta}_\eta$  implies that

$$\bar{\xi}_{n,p,\kappa}^{(1)} - \min_{(\gamma,z) \in \tilde{\Theta}_\eta} \frac{1}{\sqrt{p}} \left\{ -r \|\mathbf{g}\| + \left\| \left( \sqrt{\gamma^2 + r^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}_+ = o_{\mathbb{P}}(1).$$

Note that  $\min_{(\gamma,r) \in \tilde{\Theta}_\eta} \frac{1}{\sqrt{p}} \left\{ -z \|\mathbf{g}\| + \left\| \left( \sqrt{\gamma^2 + z^2} \mathbf{h} \odot \mathbf{y} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\| \right\}_+ = \left( \tilde{\xi}_{n,p,\kappa}^{(1)} \right)_+$ .

Combine all the pieces together to complete the proof.

### C.2.6 Proof of Lemma C.7

Define two functions

$$\begin{aligned} \phi_{n,\boldsymbol{\mu},\delta,\kappa}^{(1)}(\gamma, z) &= \left[ \mathbb{E} \left( \sqrt{\gamma^2 + z^2} G + \kappa - \gamma \|\boldsymbol{\mu}\| \right)_+^2 \right]^{1/2}, \\ \phi_{n,\boldsymbol{\mu},\delta,\kappa}^{(2)}(\gamma, z; \mathbf{h}) &= \frac{1}{\sqrt{n}} \left\| \left( \sqrt{\gamma^2 + z^2} \tilde{\mathbf{h}} + (\kappa - \gamma \|\boldsymbol{\mu}\|) \mathbf{1}_n \right)_+ \right\|, \end{aligned}$$

where  $\tilde{\mathbf{h}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$  and  $G \sim \mathcal{N}(0, 1)$ . By the uniform law of large numbers (Jennrich, 1969) (see also (Montanari et al., 2019) (Lemma 6.2)), we have

$$\sup_{(\gamma,z) \in \tilde{\Theta}_\eta} \left| \phi_{n,\boldsymbol{\mu},\delta,\kappa}^{(1)}(\gamma, z) - \phi_{n,\boldsymbol{\mu},\delta,\kappa}^{(2)}(\gamma, z; \mathbf{h}) \right| = o_{\mathbb{P}}(1),$$

In addition, since  $p^{-1/2} \|\mathbf{g}\| = 1 + o_{\mathbb{P}}(1)$ , we have

$$\left| \tilde{\xi}_{n,p,\kappa}^{(1)} - \xi_{n,p,\kappa}^{(2)} \right| \xrightarrow{p} 0,$$

which implies the desired claim.

### C.2.7 Proof of Lemma C.8

**Step 1: establish equivalence.** We notice that for the convex set  $\Theta \cap \mathbf{B}_v$ ,

$$\mathcal{E}_{n,p,\kappa,v} = \mathcal{E}_{n,p,\kappa}(\Theta \cap \mathbf{B}_v).$$

Denote the quantity

$$\xi_{n,p,\kappa,v} = \xi_{n,p,\kappa}(\mathbf{B}_v \cap \Theta) = \min_{\beta \in \mathbf{B}_v \cap \Theta} \max_{\|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{H}_0 \boldsymbol{\beta}),$$

for which we have

$$\{\xi_{n,p,\kappa,v} > 0\} \iff \mathcal{E}_{n,p,\kappa,v}^c \quad \text{and} \quad \{\xi_{n,p,\kappa,v} = 0\} \iff \mathcal{E}_{n,p,\kappa,v}.$$

Following the equivalence established in Lemma C.5, C.6 and C.7, denote

$$\begin{aligned} \xi_{n,p,\kappa,v}^{(1)} &= \xi_{n,p,\kappa}^{(1)}(\mathbf{B}_v \cap \Theta) \\ \xi_{n,p,\kappa,v}^{(2)} &= \xi_{n,p,\kappa}^{(2)}(\tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq v\gamma\}). \end{aligned}$$

Then, it is directly implied by Lemma C.5 that for any  $t \in \mathbb{R}$ ,

$$\mathbb{P}(\xi_{n,p,\kappa,v} \leq t) \leq 2\mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \leq t), \quad \mathbb{P}(\xi_{n,p,\kappa,v} \geq t) \leq 2\mathbb{P}(\xi_{n,p,\kappa,v}^{(2)} \geq t). \quad (35)$$

By Lemma C.6 and C.7, we further have

$$\left| \xi_{n,p,\kappa,v}^{(1)} - \left( \xi_{n,p,\kappa,v}^{(2)} \right)_+ \right| \xrightarrow{p} 0. \quad (36)$$

**Step 2: analyze  $F_\delta(\gamma, z, \kappa)$ .** We consider the following two cases.

1. By the definition of  $F_\delta(\gamma, z, \kappa)$ , if there exists  $(\gamma, z) \in \tilde{\Theta} \cap \{(\gamma, z) : z \leq v\gamma\}$  such that  $F_\delta(\gamma, z, \kappa) < 0$ , then

$$\xi_{n,p,\kappa,v}^{(2)} = \min_{(\gamma,z) \in \tilde{\Theta} \cap \{(\gamma,z): z \leq v\gamma\}} \left\{ -z + \sqrt{\delta} \sqrt{\mathbb{E} \left( \sqrt{\gamma^2 + z^2} G + \kappa - \gamma \|\boldsymbol{\mu}\| \right)_+^2} \right\} < 0,$$

which, according to (36), implies that  $\mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} > 0) = o(1)$  and for any  $\varepsilon > 0$ ,

$$\mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \geq \varepsilon) = o(1).$$

As  $\mathbb{P}(\xi_{n,p,\kappa,v} > t)$  is right-continuous in  $t$ , then

$$\mathbb{P}(\xi_{n,p,\kappa,v} > 0) = \lim_{t \rightarrow 0^+} \mathbb{P}(\xi_{n,p,\kappa,v} > t).$$

Combining (35) with the definition of  $\mathcal{E}_{n,p,\kappa,v}^c$ , we have

$$\mathbb{P}(\mathcal{E}_{n,p,\kappa,v}^c) = \mathbb{P}(\xi_{n,p,\kappa,v} > 0) = \lim_{\varepsilon \rightarrow 0^+} \mathbb{P}(\xi_{n,p,\kappa,v} > \varepsilon) \leq \lim_{\varepsilon \rightarrow 0^+} \mathbb{P}(\xi_{n,p,\kappa,v} \geq \varepsilon) \leq 2 \lim_{\varepsilon \rightarrow 0^+} \mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \geq \varepsilon).$$

In addition, we verify the validity of changing the order of  $\lim_{n \rightarrow +\infty, n/p \rightarrow \delta}$  and  $\lim_{\varepsilon \rightarrow 0^+}$ . Let  $q_{n,p}(t) = \mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \geq t)$  and  $\tilde{q}_{n,p}(t) = \lim_{s \rightarrow t^+} q_{n,p}(s) = \mathbb{P}(\xi_{n,p,\kappa,v}^{(2)} > t)$ , which satisfies that  $q_{n,p}(t)$  is left-continuous,  $\tilde{q}_{n,p}(t)$  is right-continuous, and for any  $t > 0$

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} q_{n,p}(t) = \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \tilde{q}_{n,p}(t) = 0.$$

Further, we have

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \tilde{q}_{n,p}(0) = \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} > 0) = 0.$$

By the fact that  $\tilde{q}_{n,p}(t)$  is non-increasing in  $t$ , we have

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \sup_{t \geq 0} |\tilde{q}_{n,p}(t) - 0|,$$

which indicates that  $\tilde{q}_{n,p}(t)$  uniformly converges to 0. Therefore, by Moore-Osgood theorem,

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \lim_{\varepsilon \rightarrow 0+} q_{n,p}(\varepsilon) = \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \lim_{\varepsilon \rightarrow 0} \tilde{q}_{n,p}(\varepsilon) = \lim_{\varepsilon \rightarrow 0} \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \tilde{q}_{n,p}(\varepsilon) = 0,$$

thus, we have

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \lim_{\varepsilon \rightarrow 0+} \mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \geq \varepsilon) = 0 \quad \implies \quad \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \mathbb{P}(\mathcal{E}_{n,p,\kappa,v}^c) = 0.$$

2. If, instead, for any  $(\gamma, z) \in \tilde{\Theta} \cap \{(\gamma, z) : z \leq v\gamma\}$ , we have  $F_\delta(\gamma, z, \kappa) > 0$ . Since the set  $(\gamma, z) \in \tilde{\Theta} \cap \{(\gamma, z) : z \leq v\gamma\}$  is convex and compact, then

$$\xi_{n,p,\kappa,v}^{(2)} = \min_{(\gamma,z) \in \tilde{\Theta} \cap \{(\gamma,z): z \leq v\gamma\}} \left\{ -z + \sqrt{\delta} \sqrt{\mathbb{E} \left( \sqrt{\gamma^2 + z^2 G + \kappa - \gamma \|\boldsymbol{\mu}\|} \right)_+^2} \right\} > 0,$$

which implies that

$$\mathbb{P}(\mathcal{E}_{n,p,\kappa,v}) = \mathbb{P}(\xi_{n,p,\kappa,v} = 0) \leq 2\mathbb{P}(\xi_{n,p,\kappa,v}^{(1)} \leq 0) = o(1).$$

### C.2.8 Proof of Lemma C.9

In this section, we will first show the relationship between  $F_\delta(\gamma, z, \kappa)$  and  $f_\delta(u, \kappa, c)$ . In the next step, we show the boundedness of  $\kappa^*(c)$  and  $u^*(c)$ .

**Step 1: Relationship between  $F_\delta(\gamma, z, \kappa)$  and  $f_\delta(u, \kappa, c)$ .** Consider the properties of  $F_\delta(\gamma, z, \kappa)$ . Since  $F_\delta(\gamma, z, \kappa) < F_\delta(-\gamma, z, \kappa)$  for  $\gamma > 0$ , in the following we only consider  $\gamma \in \mathbb{R}_+$ . Denote  $u = z/\gamma$ , then  $F_\delta(\gamma, z, \kappa)$  can be written as a function of  $(\gamma, u, \kappa)$ :

$$\tilde{F}_\delta(\gamma, u, \kappa) = \gamma^2 \left( -u^2 + \delta \mathbb{E} \left\{ \left( \sqrt{1 + u^2 G + \frac{\kappa}{\gamma} - \|\boldsymbol{\mu}\|} \right)_+^2 \right\} \right),$$

where  $(\gamma, u) \in \tilde{\Omega} = \{(\gamma, u) \in \mathbb{R}_+ \times \bar{\mathbb{R}}_+ : \gamma^2 (u^2 + c) \leq 1\}$  with  $c = (1 + \eta)^{-2}$  and  $\bar{\mathbb{R}}_+ = \mathbb{R}_+ \cup \{\infty\}$ . We have the following problem

$$\begin{cases} \sup & \kappa \\ \text{subject to} & \exists (\gamma, u) \in \tilde{\Omega}, \tilde{F}_\delta(\gamma, u, \kappa) < 0. \end{cases} \quad (37)$$

Note that when  $\kappa$  approaches the supremum, for any  $(\gamma', u) \notin \partial \tilde{\Omega}$  such that  $\tilde{F}_\delta(\gamma', u, \kappa) < 0$ , there exists  $(\gamma, u) \in \partial \tilde{\Omega}$  with  $\gamma > \gamma'$  at which  $\tilde{F}_\delta(\gamma, u, \kappa) < \tilde{F}_\delta(\gamma', u, \kappa)$ . The reason is that when  $\gamma^2 > (\gamma')^2$  and  $\gamma^{-2} \tilde{F}_\delta(\gamma, u, \kappa) < (\gamma')^{-2} \tilde{F}_\delta(\gamma', u, \kappa) < 0$ , we have  $\tilde{F}_\delta(\gamma, u, \kappa) < \tilde{F}_\delta(\gamma', u, \kappa)$ . Then, the

above shows that the minimizer  $(\gamma, u)$  must be located on the boundary  $\partial\tilde{\Omega}$ , that is  $\gamma^{-2} = u^2 + c$ . If we also view  $c$  as a variable, recall the definition  $f_\delta(u, \kappa, c)$  with

$$g_\delta(u, \kappa, c) = \tilde{F}_\delta((u^2 + c)^{-1/2}, u, \kappa) \\ = \frac{1}{u^2 + c} \cdot \left\{ -u^2 + \delta \mathbb{E} \left\{ \left( \sqrt{1 + u^2} G + \kappa \sqrt{u^2 + c} - \|\boldsymbol{\mu}\| \right)_+^2 \right\} \right\}$$

and the optimization problem in (33)

$$\begin{aligned} & \max \quad \kappa \\ & \text{subject to} \quad \inf_{u \geq 0} f_\delta(u, \kappa, c) \leq 0, \end{aligned}$$

where  $\kappa^* = \kappa^*(c)$  is the corresponding maximizer and  $u^* = u^*(c)$ , short for  $u^*(\rho, \eta)$ , is the smallest minimizer of  $\min_{u \geq 0} f_\delta(u, \kappa, c)$ .

We have the following lemma.

**Lemma C.12.** *If  $\tilde{\kappa}$  is the maximizer to (37) and  $(\tilde{u}, \tilde{\gamma})$  is the minimizer to  $F_\delta(\gamma, u, \tilde{\kappa})$  such that  $\tilde{u}$  is the smallest minimizer for  $u$ , then*

$$\tilde{\kappa} = \kappa^*, \quad \tilde{u} = u^*, \quad \tilde{\gamma} = \frac{1}{\sqrt{u^{*2} + c}}.$$

*Proof.* Fix  $c \in \mathbb{R}_+$  and for any  $\kappa \in \mathbb{R}_+$ , we obtain from the reasoning above that for any  $(\gamma, u)$  such that  $\tilde{F}_\delta(\gamma, u, \kappa) \leq 0$ ,

$$\tilde{F}_\delta(\gamma, u, \kappa) \geq \tilde{F}_\delta((u^2 + c)^{-1/2}, u, \kappa) = g_\delta(u, \kappa, c).$$

As a result, when  $\inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa) \leq 0$ ,

$$\inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa) \geq \inf_{u \geq 0} g_\delta(u, \kappa, c).$$

In addition, for the other direction,

$$\inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa) \leq \inf_{(\gamma, u) \in \tilde{\Omega}, \gamma = (u^2 + c)^{-2}} \tilde{F}_\delta(\gamma, u, \kappa) = \inf_{u \geq 0} g_\delta(u, \kappa, c).$$

Combining the two inequalities, we have

$$\inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa) = \inf_{u \geq 0} g_\delta(u, \kappa, c),$$

which implies that the following feasible sets are the same

$$\left\{ \kappa > 0 : \inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa) < 0 \right\} = \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0 \right\}.$$

Then, it is equivalent to show that (33) has the same solution with the following problem

$$\begin{cases} \sup & \kappa \\ \text{subject to} & \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0. \end{cases}$$

We denote  $\tilde{u}$  and  $\tilde{\kappa}$  as the solution to (C.2.8) such that  $\tilde{u}$  is the smallest one if the minimizer is not unique. Since  $\bar{\mathbb{R}}_+$  is a closed set, then  $\inf_{u \geq 0} g_\delta(u, \kappa, c) = \min_{u \geq 0} g_\delta(u, \kappa, c)$ . We further have

$$\left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0 \right\} \subseteq \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) \leq 0 \right\},$$

indicating that

$$\kappa^* = \sup \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) \leq 0 \right\} \geq \sup \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0 \right\} = \tilde{\kappa}.$$

Recall that the solution  $(u^*, \kappa^*)$  for (33) has the property that  $u^*$  is the smallest minimizer. Since  $g_\delta(u, \kappa, c)$  is strictly increasing in  $\kappa$ , then for any  $\varepsilon > 0$ , there exists  $\kappa_\varepsilon \geq \kappa^* - \varepsilon$  such that

$$\inf_{u \geq 0} g_\delta(u, \kappa_\varepsilon, c) < \inf_{u \geq 0} g_\delta(u, \kappa^*, c) \leq 0,$$

which means

$$\kappa_\varepsilon \in \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0 \right\}. \quad (38)$$

If  $\tilde{\kappa} < \kappa^*$ , let  $\epsilon = (\kappa^* - \tilde{\kappa})/2$ . Then,  $\kappa_\epsilon > \tilde{\kappa}$  and (38) shows that there is a feasible  $\kappa$  for that is larger than  $\tilde{\kappa}$ . This draws the contradiction and we have

$$\tilde{\kappa} = \sup \left\{ \kappa > 0 : \inf_{u \geq 0} g_\delta(u, \kappa, c) < 0 \right\} = \kappa^*.$$

By definition,  $(\tilde{u}, \tilde{\gamma})$  is the minimizer to  $F_\delta(\gamma, u\gamma, \kappa^*)$  where  $\tilde{u}$  is the smallest one if the minimizer is not unique, then by the reasoning above that the minimizer locates on the boundary of  $\tilde{\Theta}_\eta$ , it is equivalent that  $\tilde{u}$  is the minimizer to  $g_\delta(u, \kappa^*, c)$ . By the uniqueness of the smallest minimizer to  $g_\delta(u, \kappa^*, c)$ , we have  $\tilde{u} = u^*$ . By the boundary equation, we have  $\tilde{\gamma} = 1/\sqrt{u^{*2} + c}$ .  $\square$

**Step 2: Proving  $\kappa^*(c) < +\infty$ .** It suffices to provide some  $\bar{\kappa} < +\infty$  such that  $\inf_{u \geq 0} f_\delta(u, \bar{\kappa}, c) > 0$ . Recall that

$$\begin{aligned} f_\delta(u, \kappa, c) &= -u^2 + \delta \mathbb{E} \left[ \left( \sqrt{1 + u^2} G + \kappa \sqrt{u^2 + c} - \|\boldsymbol{\mu}\| \right)_+^2 \right] \\ &= u^2 \left\{ -1 + \delta \frac{1 + u^2}{u^2} \mathbb{E} \left[ \left( G + \kappa \sqrt{\frac{u^2 + c}{u^2 + 1}} - \sqrt{\frac{1}{1 + u^2}} \|\boldsymbol{\mu}\| \right)_+^2 \right] \right\} \\ &\geq u^2 \left\{ -1 + \delta \mathbb{E} \left[ \left( G + \kappa \sqrt{\frac{u^2 + c}{u^2 + 1}} - \sqrt{\frac{1}{1 + u^2}} \|\boldsymbol{\mu}\| \right)_+^2 \right] \right\}. \end{aligned}$$

Since  $\mathbb{E}[(G + a)_+]^2$  is increasing in  $a$  and

$$\kappa \sqrt{\frac{u^2 + c}{u^2 + 1}} - \sqrt{\frac{1}{1 + u^2}} \|\boldsymbol{\mu}\| \geq (c \wedge 1) \kappa - \sqrt{\frac{1}{1 + u^2}} \|\boldsymbol{\mu}\|.$$

Then, it suffices to show that there exists some  $0 \leq \bar{\kappa} < +\infty$  such that

$$\inf_{u \geq 0} \left\{ -1 + \delta \mathbb{E} \left[ \left( G + (c \wedge 1) \bar{\kappa} - \sqrt{\frac{1}{1 + u^2}} \|\boldsymbol{\mu}\| \right)_+^2 \right] \right\} = -1 + \delta \mathbb{E} \left[ (G + (c \wedge 1) \bar{\kappa} - \|\boldsymbol{\mu}\|)_+^2 \right] > 0.$$

Since  $c \wedge 1 > 0$ , then  $\mathbb{E}[(G + (c \wedge 1)\kappa - \|\boldsymbol{\mu}\|_+^2)]$  is an increasing and continuous function in  $\kappa$ , and

$$\lim_{\kappa \rightarrow +\infty} \{-1 + \delta \mathbb{E}[(G + (c \wedge 1)\kappa - \|\boldsymbol{\mu}\|_+^2)]\} = +\infty.$$

Consequently, there exists  $\bar{\kappa} < +\infty$  such that  $-1 + \delta \mathbb{E}[(G + (c \wedge 1)\bar{\kappa} - \|\boldsymbol{\mu}\|_+^2)] > 0$ , which proves that  $\kappa^*(c) < +\infty$ .

**Step 3: Proving  $u^*(c) < +\infty$ .** For a given  $c$ , denote

$$\mathcal{A} = \{(u, \kappa) \in \mathbb{R}_+ \times \mathbb{R}_+ : f_\delta(u, \kappa, c) \leq 0\}, \quad \mathcal{A}_\kappa = \{u \in \mathbb{R}_+ : (u, \kappa) \in \mathcal{A}\}.$$

Then, to show that  $u^* < +\infty$ , it is equivalent to show that  $\mathcal{A}_{\kappa^*} \setminus \{+\infty\} \neq \emptyset$ .

(1)  $2 < \delta < \delta^*(\|\boldsymbol{\mu}\|)$ . We first consider  $\delta \in (2, \delta^*(\rho))$ . As  $u \rightarrow +\infty$ , we have for any  $\kappa \geq 0$ ,

$$\lim_{u \rightarrow +\infty} f_\delta(u, \kappa, c) = \lim_{u \rightarrow +\infty} \left\{ \left( \frac{\delta}{2} - 1 \right) u^2 \right\} = +\infty,$$

which implies that  $+\infty \notin \mathcal{A}_\kappa$  for any  $\kappa \geq 0$ . Since  $\mathcal{A}_{\kappa^*} \neq \emptyset$  by the existence of the max-margin solution, then  $\mathcal{A}_{\kappa^*} \setminus \{+\infty\} \neq \emptyset$ .

(2)  $\delta \leq 2$ . When  $\delta \leq 2$ , let  $\kappa_0$  be the unique solution to

$$\frac{1}{\delta} = \mathbb{E} \left[ (G + \kappa)_+^2 \right].$$

Since  $f_\delta(u, \kappa, c)$  is increasing in  $\kappa$ , then for  $\kappa > \kappa'$ ,  $\mathcal{A}_\kappa \subseteq \mathcal{A}_{\kappa'}$ . Also, for  $\kappa < \kappa_0$ , we have  $\delta \mathbb{E} \left[ (G + \kappa)_+^2 \right] < 1$ ,

$$\lim_{u \rightarrow +\infty} f_\delta(u, \kappa, c) = \lim_{u \rightarrow +\infty} \left\{ \left( \delta \mathbb{E} \left[ (G + \kappa)_+^2 \right] - 1 \right) u^2 \right\} = -\infty,$$

which implies that  $\mathcal{A}_\kappa \neq \emptyset$  and  $+\infty \in \mathcal{A}_\kappa$  for  $\kappa < \kappa_0$ .

Further, if  $\kappa > \kappa_0$ , we have

$$\delta \mathbb{E} \left[ (G + \kappa)_+^2 \right] > 1,$$

which implies that

$$\lim_{u \rightarrow +\infty} f_\delta(u, \kappa, c) = \lim_{u \rightarrow +\infty} \left\{ -u^2 + \left( \delta \mathbb{E} \left[ (G + \kappa)_+^2 \right] \right) u^2 \right\} = +\infty.$$

Thus,  $+\infty \notin \mathcal{A}_\kappa$  for  $\kappa > \kappa_0$ . Then, to complete the proof that  $u^* < +\infty$ , we need to show that  $\mathcal{A}_{\kappa_0} \neq \emptyset$ . Consider

$$\sqrt{1 + u^2}G + \kappa_0 \sqrt{u^2 + c} = Gu + \kappa_0 u + \frac{G}{\sqrt{1 + u^2 + u}} + \frac{c}{\sqrt{u^2 + c + u}},$$

there exists  $C > 0$  such that for  $u \geq C$ , we have

$$\frac{c}{\sqrt{u^2 + c + u}} < \frac{1}{2} \|\boldsymbol{\mu}\|.$$

Then, we have

$$\begin{aligned}
& \mathbb{E} \left\{ \left( \sqrt{1+u^2}G + \kappa_0 \sqrt{u^2+c} - \|\boldsymbol{\mu}\| \right)_+^2 \right\} \leq \mathbb{E} \left\{ \left( Gu + \kappa_0 u + \frac{G}{\sqrt{1+u^2+u}} - \frac{1}{2} \|\boldsymbol{\mu}\| \right)_+^2 \right\} \\
& \leq \mathbb{E} \left\{ \left( Gu + \kappa_0 u - \frac{1}{4} \|\boldsymbol{\mu}\| \right)_+^2 \mathbb{1} \left\{ \frac{G}{\sqrt{1+u^2+u}} < \frac{1}{4} \|\boldsymbol{\mu}\| \right\} \right\} \\
& \quad + \mathbb{E} \left\{ \left( Gu + \kappa_0 u + \frac{G}{\sqrt{1+u^2+u}} \right)_+^2 \mathbb{1} \left\{ \frac{G}{\sqrt{1+u^2+u}} > \frac{1}{4} \|\boldsymbol{\mu}\| \right\} \right\} \\
& \leq \mathbb{E} \left\{ \left( Gu + \kappa_0 u - \frac{1}{4} \|\boldsymbol{\mu}\| \right)_+^2 \right\} + \mathbb{E} \left\{ \left( Gu + \kappa_0 u + \frac{G}{\sqrt{1+u^2+u}} \right)_+^2 \mathbb{1} \left\{ \frac{G}{\sqrt{1+u^2+u}} > \frac{1}{4} \|\boldsymbol{\mu}\| \right\} \right\} \\
& := I_1(u) + I_2(u),
\end{aligned}$$

where  $\lim_{u \rightarrow +\infty} I_2(u) = 0$ . For  $I_1(u)$ , consider the function

$$q(t) = \delta \mathbb{E} \left\{ (G + \kappa_0 - t)_+^2 \right\} - 1,$$

then  $q(0) = 1$  and  $q'(t) = -2\delta \mathbb{E} \left\{ (G + \kappa_0 - t)_+ \right\} < 0$ . At  $t \approx 0$ , we have the approximation  $q(t) = q'(0)t + o(t)$ . If we take  $t = \|\boldsymbol{\mu}\|/(4u)$ ,  $-u^2 + \delta I_1(u)$  can be approximated by

$$-u^2 + \delta I_1(u) = \left\{ \delta \mathbb{E} \left\{ \left( G + \kappa_0 - \frac{\|\boldsymbol{\mu}\|}{4u} \right)_+^2 \right\} - 1 \right\} u^2 = \frac{\|\boldsymbol{\mu}\|}{4} u q'(0) + o(u),$$

thus there exists sufficiently large  $u \geq C$  such that  $-u^2 + \delta(I_1(u) + I_2(u)) < 0$ , which implies that  $\mathcal{A}_{\kappa_0} \neq \emptyset$ .

We consider the constrained problem with arbitrarily small  $\varepsilon > 0$ ,

$$\begin{cases} \max & \kappa \\ \text{subject to} & \inf_{u \geq 0, \kappa \geq \kappa_0 + \varepsilon} f_\delta(u, \kappa, c) \leq 0, \end{cases} \quad (39)$$

Let  $(\tilde{u}, \tilde{\kappa})$  be any solution to (39), where the existence is obtained by the continuity of  $g_\delta(u, \kappa, c)$  in  $\kappa$ . On one hand, since (39) is a constrained version of (33), then we have  $\tilde{\kappa} \leq \kappa^*$ ; on the other hand, when  $\tilde{\kappa} < \kappa^*$ , then  $(u^*, \kappa^*)$  is a feasible point for the constrained problem, thus  $\kappa^* \leq \tilde{\kappa}$ . Then, we have  $\kappa^* = \tilde{\kappa} > \kappa_0$ , indicating that  $0 < u^* < +\infty$  since  $+\infty \notin \mathcal{A}_\kappa$  for  $\kappa > \kappa_0$ .

Combining pieces above, we have shown that in either regime  $\delta \in (2, \delta^*(\rho))$  or  $\delta \leq 2$ , the solution to (33) satisfies that  $0 < u^* < +\infty$ , which completes the proof.

### C.2.9 Proof of Lemma C.10

Recall the notation

$$\begin{aligned}
\mathcal{E}_{n,p,\kappa,v} & := \{ \text{there exists } \boldsymbol{\beta} \in \Theta \cap \mathbf{B}_v : y_i \langle \mathbf{h}_{0,i}, \boldsymbol{\beta} \rangle \geq \kappa \text{ for all } i \leq n \}, \\
\xi_{n,p,\kappa,v} & := \xi_{n,p,\kappa}(\mathbf{B}_v \cap \Theta) = \min_{\boldsymbol{\beta} \in \mathbf{B}_v \cap \Theta} \max_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{p}} \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{H}_0 \boldsymbol{\beta}).
\end{aligned}$$

where  $\mathbf{B}_v = \{ \boldsymbol{\beta} \in \mathbb{R}^p : \|\mathbf{P}_\mu^\perp \boldsymbol{\beta}\| \leq v \|\mathbf{P}_\mu \boldsymbol{\beta}\| \}$  is a convex cone for each choice of  $v \in \bar{\mathbb{R}}_+$ . Accordingly, we denote

$$\xi_{n,p,\kappa,v}^{(2)} := \xi_{n,p,\kappa}^{(2)}(\tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq v\gamma\})$$

$$= \min_{(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq v\gamma\}} \left\{ -z + \sqrt{\delta} \sqrt{\mathbb{E} \left( \sqrt{\gamma^2 + z^2 G + \kappa - \gamma \|\boldsymbol{\mu}\|} \right)_+^2} \right\}$$

Then, by definition,

$$\xi_{n,p,\kappa} = \min_{v \in \bar{\mathbb{R}}_+} \xi_{n,p,\kappa,v}, \quad \xi_{n,p,\kappa}^{(2)} = \min_{v \in \bar{\mathbb{R}}_+} \xi_{n,p,\kappa,v}^{(2)}$$

Recall the definition of  $\hat{\kappa}$  and  $\hat{u}$  in Theorem 5.2, then

$$\hat{\kappa} = \sup_{\kappa > 0} \{\xi_{n,p,\kappa} = 0\}, \quad \hat{u} \in \{v : \xi_{n,p,\hat{\kappa},v} = 0\}.$$

**Step 1: Proving  $\mathcal{E}_{n,p,\kappa^*,u^*+\varepsilon}$  and  $\mathcal{E}_{n,p,\kappa^*,u^*-\varepsilon}^c$ .** For any  $\varepsilon > 0$ , we first show that

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \mathbb{P}(\xi_{n,p,\kappa^*,u^*-\varepsilon} > 0) = 1.$$

By the result in Lemma C.8, we need to show that

$$\inf_{(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq (u^* - \varepsilon)\gamma\}} F_\delta(\gamma, z, \kappa^*) > 0,$$

up to the change of variables, which is equivalent to

$$\inf_{(\gamma, u) \in \tilde{\Omega} \cap \{(\gamma, u) : z \leq u^* - \varepsilon\}} \tilde{F}_\delta(\gamma, u, \kappa^*) > 0.$$

Recall the definition of  $u^*$  as the smallest minimizer of  $f_\delta(u, \kappa^*, c)$  in (33). By Section C.2.8 and Lemma C.12, we have  $\tilde{F}_\delta((u^{*2} + c)^{-1/2}, u^*, \kappa^*) = 0$ . Assume  $\inf_{(\gamma, u) \in \tilde{\Omega} \cap \{(\gamma, u) : u \leq u^* - \varepsilon\}} \tilde{F}_\delta(\gamma, u, \kappa^*) \leq 0$ , which implies that there exists  $v^* \leq u^* - \varepsilon < u^*$  and  $\gamma_1$  such that

$$\tilde{F}_\delta(\gamma_1, v^*, \kappa^*) \leq \tilde{F}_\delta((u^{*2} + c)^{-1/2}, u^*, \kappa^*) = \inf_{(\gamma, u) \in \tilde{\Omega}} \tilde{F}_\delta(\gamma, u, \kappa^*),$$

then  $(\gamma_1, v^*)$  is also the minimizer to the unconstrained problem. This is a contradiction to the definition that  $u^*$  is the smallest minimizer, thus

$$\inf_{(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq (u^* - \varepsilon)\gamma\}} F_\delta(\gamma, z, \kappa^*) > 0 \quad \implies \quad \lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \mathbb{P}(\xi_{n,p,\kappa^*,u^*-\varepsilon} > 0) = 1.$$

In addition, since  $\tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq u^*\gamma\} \subseteq \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq (u^* + \varepsilon)\gamma\}$ ,

$$\begin{aligned} \inf_{(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq (u^* + \varepsilon)\gamma\}} F_\delta(\gamma, z, \kappa^*) &\leq \inf_{(\gamma, z) \in \tilde{\Theta}_\eta \cap \{(\gamma, z) : z \leq u^*\gamma\}} F_\delta(\gamma, z, \kappa^*) \leq 0 \\ &= \inf_{(\gamma, u) \in \tilde{\Omega} \cap \{(\gamma, u) : u \leq u^*\}} \tilde{F}_\delta(\gamma, u, \kappa^*) \\ &= \inf_{u \in [0, u^*]} \tilde{F}_\delta((u^{*2} + c)^{-1/2}, u^*, \kappa^*) \leq 0. \end{aligned}$$

Then, for any  $\varepsilon > 0$ , we have

$$\mathcal{E}_{n,p,\kappa^*,u^*+\varepsilon} \quad \text{and} \quad \mathcal{E}_{n,p,\kappa^*,u^*-\varepsilon}^c \quad \text{hold.}$$



**Step 2: Proving  $\hat{u} \geq u^*$ .** Suppose the claim is wrong, then there exists  $\varepsilon > 0$  such that  $\hat{u} \leq u^* - \varepsilon < u^*$ , which indicates that

$$\hat{\boldsymbol{\beta}} \in \Theta \cap \mathbf{B}_{u^* - \varepsilon}.$$

Then, with  $\hat{\boldsymbol{\beta}}$ ,

$$y_i \langle \mathbf{h}_{0,i}, \hat{\boldsymbol{\beta}} \rangle \geq \hat{\kappa}, \quad \text{for all } i \leq n.$$

By Lemma C.8, we have

$$\inf_{0 \leq u \leq u^* - \varepsilon} f_\delta(u, \hat{\kappa}, c) \leq 0.$$

As  $\hat{\kappa} \xrightarrow{p} \kappa^*$  and  $f_\delta(u, \kappa, c)$ , we have

$$\inf_{0 \leq u \leq u^* - \varepsilon} f_\delta(u, \kappa^*, c) \leq 0,$$

which implies that  $\mathcal{E}_{n,p,\kappa^*,u^*-\varepsilon}$  holds. This draws the contradiction and proves that  $\hat{u} \geq u^*$ .

**Step 3: Proving  $\hat{u} \leq u^* + 2\varepsilon$ .** For any  $\varepsilon > 0$ , as  $\mathbf{B}_{u^*} \subseteq \mathbf{B}_{u^* + \varepsilon}$ , we already have

$$\lim_{n \rightarrow +\infty, n/p \rightarrow \delta} \mathbb{P}(\xi_{n,p,\kappa^*,u^*+\varepsilon} = 0) = 1.$$

Denote by  $\hat{\boldsymbol{\beta}}^\varepsilon$  any maximizer to the optimization problem (31) with an additional constraint  $\boldsymbol{\beta} \in \mathbf{B}_{u^* + \varepsilon}$ , and let  $\hat{u}^\varepsilon = \|\mathbf{P}_\mu^\perp \hat{\boldsymbol{\beta}}^\varepsilon\| / \|\mathbf{P}_\mu \hat{\boldsymbol{\beta}}^\varepsilon\|$ . Our next goal is to show that with probability approaching one,  $\hat{\boldsymbol{\beta}}$  satisfies  $\|\mathbf{P}_\mu^\perp \hat{\boldsymbol{\beta}}\| \leq (u^* + 2\varepsilon) \|\mathbf{P}_\mu \hat{\boldsymbol{\beta}}\|$ . Once it is proved, then  $\hat{u} \xrightarrow{p} u^*$  since  $\varepsilon$  is arbitrary.

Consider any  $\boldsymbol{\beta}$  with  $\|\boldsymbol{\beta}\|_\Omega \leq 1$  and  $\boldsymbol{\beta} \notin \mathbf{B}_{u^* + 2\varepsilon}$  where we recall

$$\mathbf{B}_{u^* + 2\varepsilon} = \{\mathbf{x} \in \mathbb{R}^p : \|\mathbf{P}_\mu^\perp \mathbf{x}\| \leq (u^* + 2\varepsilon) \|\mathbf{P}_\mu \mathbf{x}\|\}.$$

Since the region  $\Theta = \{\mathbf{x} : \mathbf{x}^\top \Omega \mathbf{x} \leq 1\}$  is an ellipsoid, it must be strictly convex. For any  $\boldsymbol{\beta} \in \Theta$ , the interpolant

$$\boldsymbol{\beta}_\lambda := \lambda \hat{\boldsymbol{\beta}}^\varepsilon + (1 - \lambda) \boldsymbol{\beta}, \quad \text{where } \lambda \in (0, 1)$$

lies in  $\Theta$ . Since  $\|\hat{\boldsymbol{\beta}}^\varepsilon - \boldsymbol{\beta}\| \geq c_\varepsilon > 0$  for certain constant  $c_\varepsilon$ , we can find (sufficiently small) constants  $\lambda \in (0, 1)$  and  $\alpha > 0$  such that

$$\boldsymbol{\beta}_\lambda \in \mathbf{B}_{u^* + \varepsilon}, \quad \text{and} \quad (1 + \alpha) \boldsymbol{\beta}_\lambda \in \Theta.$$

Let us use  $\kappa(\boldsymbol{\beta}) = \min_{i \leq n} y_i \langle \mathbf{h}_{0,i}, \boldsymbol{\beta} \rangle$  to denote the margin associated with  $\boldsymbol{\beta}$ . By the definition of  $\hat{\boldsymbol{\beta}}^\varepsilon$  and the fact that  $(1 + \alpha) \boldsymbol{\beta}_\lambda \in \mathbf{B}_{u^* + \varepsilon}$ , we have

$$\kappa((1 + \alpha) \boldsymbol{\beta}_\lambda) \leq \kappa(\hat{\boldsymbol{\beta}}^\varepsilon). \quad (40)$$

By linearity, we also have

$$\kappa((1 + \alpha) \boldsymbol{\beta}_\lambda) \geq (1 + \alpha) \lambda \kappa(\hat{\boldsymbol{\beta}}^\varepsilon) + (1 + \alpha)(1 - \lambda) \kappa(\boldsymbol{\beta}). \quad (41)$$

Combining (40) and (41), we have

$$\kappa(\boldsymbol{\beta}) \leq \frac{1 - (1 + \alpha)\lambda}{(1 + \alpha)(1 - \lambda)} \kappa(\hat{\boldsymbol{\beta}}^\varepsilon) < \kappa(\hat{\boldsymbol{\beta}}^\varepsilon).$$

This proves that if  $\boldsymbol{\beta} = \hat{\boldsymbol{\beta}} \notin \mathbf{B}_{u^* + 2\varepsilon}$ , then  $\hat{\boldsymbol{\beta}} < \kappa(\hat{\boldsymbol{\beta}}^\varepsilon)$ , which is a contradiction. Then, we have shown that  $\hat{\boldsymbol{\beta}} \in \mathbf{B}_{u^* + 2\varepsilon}$ . Consequently, we have  $\hat{u} \in [u^*, u^* + 2\varepsilon)$  for any  $\varepsilon > 0$ , thus

$$\hat{u} = \frac{\|\mathbf{P}_\mu^\perp \hat{\boldsymbol{\beta}}\|}{\|\mathbf{P}_\mu \hat{\boldsymbol{\beta}}\|} \xrightarrow{p} u^*.$$

### C.2.10 Proof of Lemma C.11

By the definition of  $\kappa^*(c)$  and  $u^*(c)$ , it is easy to check that

$$\begin{cases} f_\delta(u^*(c), \kappa^*(c), c) &= 0, \\ \partial_u f_\delta(u^*(c), \kappa^*(c), c) &= 0. \end{cases}$$

In other words, for any  $c$ ,  $(u^*(c), \kappa^*(c), c)$  is the zero point of the systems of equations

$$\begin{cases} f(u, \kappa, c) &= 0, \\ \partial_u f(u, \kappa, c) &= 0, \end{cases} \quad (42)$$

where we identify  $f_\delta$  with  $f$  to simply the notation.

Fix any  $c_0$ , and its corresponding  $u^*(c_0), \kappa^*(c_0)$ . We can compute that

$$\partial_c \partial_u f - \frac{\partial_c f}{\partial_\kappa f} \partial_\kappa \partial_u f \Big|_{u^*(c_0), \kappa^*(c_0), c_0} < 0 \quad (43)$$

As a result, the implicit function theorem tells us that in an open neighborhood around  $u^*(c_0)$ , the zero point of (42) can be written as a continuously differentiable function  $u \rightarrow \kappa^*(u), c^*(u)$ . Moreover, implicit function theorem also tells us that

$$\left( \partial_c \partial_u f - \frac{\partial_c f}{\partial_\kappa f} \partial_\kappa \partial_u f \right) \Big|_{u^*(c_0), \kappa^*(c_0), c_0} \cdot c'(u^*(c_0)) + \partial_u^2 f \Big|_{u^*(c_0), \kappa^*(c_0), c_0} = 0$$

As  $u^*(c_0)$  is a minimizer for  $f(u, \kappa^*(c_0), c_0)$ , we have  $\partial_u^2 f \Big|_{u^*(c_0), \kappa^*(c_0), c_0} \geq 0$ . In all, we conclude that  $c'(u^*(c_0)) \geq 0$ .

We have the following lemma which shows the continuity of  $u^*(c)$ .

**Lemma C.13.** *Both  $\kappa^*(c)$  and  $u^*(c)$  are continuous in  $c$ .*

See the end of this section for the proof.

According to Lemma C.13, in an open neighborhood  $\mathcal{O}$  of  $(u_0^*, \kappa_0^*, c_0) = (u^*(c_0), \kappa^*(c_0), c_0)$ , for any  $c_1 > c_0$  such that  $c_1 - c_0$  is sufficiently small, denote  $(u_1^*, \kappa_1^*)$  as the solution to (33) with  $c = c_1$ . Since both  $u^*(c)$  and  $\kappa^*(c)$  are continuous in  $c$ , we have  $(u_0^*, \kappa_0^*, c_0), (u_1^*, \kappa_1^*, c_1) \in \mathcal{O}$ . By the implicit function theorem, in the neighborhood  $\mathcal{O}$ ,  $c$  can be written as a differentiable function  $c(u^*)$  of  $u^*$  with  $0 \leq c'(u^*) < +\infty$ . Suppose we do not have  $u_0^* < u_1^*$ . If  $u_0^* > u_1^*$ , by the mean value theorem, then there exists  $u_2^* \in (u_1^*, u_0^*)$  such that  $c'(u_2^*) < 0$ , which draws the contradiction. If  $u_0^* = u_1^*$ , then  $u_0^*$  is mapped to two distinct values and the function  $c(u)$  is not well-defined. Therefore, we have  $u_0^* < u_1^*$  and  $u^*(c)$  is increasing in  $c$ .

**Proof of (43).** Denote  $Z(u) = \sqrt{1+u^2}G + \kappa\sqrt{u^2+c} - \|\boldsymbol{\mu}\|$ . By direct calculations, we have

1.  $\partial_\kappa f = 2\delta\sqrt{u^2+c}\mathbb{E}[Z_+(u)],$
2.  $\partial_c f = \frac{\kappa\delta}{\sqrt{u^2+c}}\mathbb{E}[Z_+(u)],$
3.  $\partial_u f = -2u + 2u\delta\mathbb{P}(Z(u) \geq 0) + 2\delta\frac{\kappa u}{\sqrt{u^2+c}}\mathbb{E}[Z_+(u)],$
4.  $\partial_\kappa \partial_u f = 2u\delta\frac{\sqrt{u^2+c}}{\sqrt{u^2+1}}\phi\left(\frac{\sqrt{u^2+c}}{\sqrt{u^2+1}}\kappa - \frac{\|\boldsymbol{\mu}\|}{\sqrt{u^2+1}}\right) + 2u\frac{\delta}{\sqrt{u^2+c}}\mathbb{E}[Z_+(u)] + 2u\delta\kappa\mathbb{P}(Z(u) \geq 0),$

$$5. \partial_c \partial_u f = u \delta \frac{\kappa}{\sqrt{u^2+1}} \frac{1}{\sqrt{u^2+c}} \phi \left( \frac{\sqrt{u^2+c}}{\sqrt{u^2+1}} \kappa - \frac{\|\boldsymbol{\mu}\|}{\sqrt{u^2+1}} \right) - u \delta \kappa \frac{1}{(u^2+c)^{3/2}} \mathbb{E}[Z_+(u)] + \delta \frac{\kappa^2}{u^2+c} \mathbb{P}(Z(u) \geq 0).$$

Denoting  $\nu = \frac{\sqrt{u^2+c}}{\sqrt{u^2+1}} \kappa - \frac{\|\boldsymbol{\mu}\|}{\sqrt{u^2+1}}$ , one then has

$$\frac{\partial_c f}{\partial_\kappa f} \partial_\kappa \partial_u f = \delta u \kappa \sqrt{\frac{1}{(u^2+1)(u^2+c)}} \phi(\nu) + \frac{\delta u \kappa}{(u^2+c)^{3/2}} \mathbb{E}[Z_+(u)] + \frac{\delta u \kappa^2}{u^2+c} \mathbb{P}(Z(u) \geq 0).$$

As the result, for any  $0 < u < +\infty$ , we have

$$\partial_c \partial_u f - \frac{\partial_c f}{\partial_\kappa f} \partial_\kappa \partial_u f = -\frac{2\delta u \kappa}{(u^2+c)^{3/2}} \mathbb{E}[Z_+(u)] < 0.$$

**Proof of Lemma C.13.** Denote  $\boldsymbol{w} = (\gamma, u)$  and

$$\kappa(\boldsymbol{w}) = \sup_{\kappa \geq 0} \{ \tilde{F}_\delta(\gamma, u, \kappa) \leq 0 \},$$

where  $\kappa(\boldsymbol{w}) = 0$  if  $\{ \kappa \geq 0 : \tilde{F}_\delta(\gamma, u, \kappa) \leq 0 \} = \emptyset$ , then  $\kappa(\boldsymbol{w})$  is continuous in  $\boldsymbol{w}$ . Rewrite  $U(c) = \{ (\gamma, u) \in \mathbb{R} \times \mathbb{R}_+ : \gamma^2(u^2+c) \leq 1 \}$ , which satisfies that  $U(c_1) \subseteq U(c_2)$  for  $c_1 \geq c_2$ , and denote

$$\kappa^*(c) = \sup_{\boldsymbol{w} \in U(c)} \kappa(\boldsymbol{w})$$

Consider any  $0 < c < +\infty$  and  $\boldsymbol{w}^*(c)$  is the associated solution to (37).

**Step 1: Continuity of  $\kappa^*(c)$ .** We first show that  $\kappa^*(c)$  is continuous in  $c$ . Otherwise, without loss of generality, there exists  $\varepsilon > 0$  such that there exists an increasing sequence  $\{c_m\}_{m \geq 1}$  with  $c/2 < c_m < c$  and  $c - c_m \leq \frac{1}{m}$  which satisfies that  $|\kappa^*(c_m) - \kappa^*(c)| \geq \varepsilon$ , which forms a non-increasing sequence  $\{\kappa^*(c_m)\}_{m=1}^\infty$  and  $\{\boldsymbol{w}^*(c_m)\}_{m=1}^\infty \subseteq U(c/2)$ . Recall that  $\kappa^*(c)$  is non-increasing in  $c$ , then  $\kappa^*(c_m) \geq \kappa^*(c) + \varepsilon$  and  $\kappa^*(c_m) \in [\kappa^*(c), \kappa^*(c/2)]$ . Then, for the sequence in a compact set, there is a subsequence  $\{c_{m_k}\}_{k=1}^\infty \subseteq \{c_m\}_{m=1}^\infty$  such that

$$\lim_{k \rightarrow +\infty} c_{m_k} = c, \quad \lim_{k \rightarrow +\infty} \kappa^*(c_{m_k}) = \tilde{\kappa}, \quad \lim_{k \rightarrow +\infty} \boldsymbol{w}^*(c_{m_k}) = \tilde{\boldsymbol{w}} \in \bigcap_{k=1}^\infty U(c_{m_k}) = U(c).$$

Thus, there exists  $K \in \mathbb{N}$  such that for any  $k \geq K$ ,

$$\kappa(\tilde{\boldsymbol{w}}) \geq \kappa(\boldsymbol{w}^*(c_{m_k})) - \frac{\varepsilon}{2} = \kappa^*(c_{m_k}) - \frac{\varepsilon}{2} \geq \kappa^*(c) + \frac{\varepsilon}{2} > \kappa^*(c),$$

which violates the fact that  $\boldsymbol{w}^*(c) = \operatorname{argmax}_{\boldsymbol{w} \in U(c)} \kappa(\boldsymbol{w})$ . This proves the left-continuity of  $\kappa^*(c)$  and the right-continuity follows the same technique.

**Step 2: Continuity of  $\boldsymbol{w}^*(c)$ .** Suppose  $\boldsymbol{w}^*(c) = (\gamma^*(c), u^*(c))$  is not continuous in  $c$ , particularly for the second coordinate, then there exists  $\varepsilon > 0$  such that for any  $m \in \mathbb{N}$ , there exists an increasing sequence  $\{c_m\}_{m \geq 1}$  with  $c/2 < c_m < c$  and  $c - c_m \leq \frac{1}{m}$  such that  $|u^*(c_m) - u^*(c)| \geq \varepsilon$ , which forms a sequence  $\{\boldsymbol{w}^*(c_m)\}_{m=1}^\infty$ . Also, we have  $\boldsymbol{w}^*(c_m) = (\gamma^*(c_m), u^*(c_m)) \in U(c_m) \subseteq U(c/2)$  for each  $m$ . As  $U(c/2)$  is compact, then there is a subsequence  $\{c_{m_k}\}_{k=1}^\infty \subseteq \{c_m\}_{m=1}^\infty$  such that

$$\lim_{k \rightarrow +\infty} c_{m_k} = c, \quad \lim_{k \rightarrow +\infty} \boldsymbol{w}^*(c_{m_k}) = \tilde{\boldsymbol{w}} = (\tilde{\gamma}, \tilde{u}) \in \bigcap_{k=1}^\infty U(c_{m_k}) = U(c).$$

Since  $\tilde{F}_\delta(\mathbf{w}, \kappa)$  is continuous in  $(\mathbf{w}, \kappa)$  and for any  $k$ ,  $\tilde{F}_\delta(\mathbf{w}^*(c_{m_k}), \kappa^*(c_{m_k})) = 0$ , then  $\tilde{F}_\delta(\tilde{\mathbf{w}}, \tilde{\kappa}) = 0$  with  $|\tilde{u} - u^*(c)| \geq \varepsilon$ . Denote

$$\mathcal{B}_k = \left\{ \mathbf{w} : \tilde{F}_\delta(\mathbf{w}, \kappa^*(c_{m_k})) \leq 0 \right\}, \quad \mathcal{B}_\infty = \left\{ \mathbf{w} : \tilde{F}_\delta(\mathbf{w}, \kappa^*(c)) \leq 0 \right\}$$

Since  $\kappa^*(c_{m_k})$  is non-increasing as  $k$  increases with  $\lim_{k \rightarrow +\infty} \kappa^*(c_{m_k}) = \kappa^*(c)$ , and  $\tilde{F}_\delta(\mathbf{w}, \kappa)$  is continuous and increasing in  $\kappa$ , we have

$$\mathcal{B}_k \subseteq \mathcal{B}_{k+1} \subseteq \mathcal{B}_\infty, \quad \bigcup_{k \geq 1} \mathcal{B}_k = \mathcal{B}_\infty.$$

By definition,  $u^*(c_{m_k}) = \inf\{u : (\gamma, u) \in \mathcal{B}_k\}$ , then as  $u^*(c_{m_k}) \rightarrow \tilde{u}$ ,

$$\tilde{u} \leq \inf\{u : (\gamma, u) \in \mathcal{B}_k\}, \quad \text{for any } k.$$

Consequently, we have

$$\tilde{u} \leq \inf_{k \geq 1} \inf\{u : (\gamma, u) \in \mathcal{B}_k\} = \inf \left\{ u : (\gamma, u) \in \bigcup_{k \geq 1} \mathcal{B}_k \right\} = \inf\{u : (\gamma, u) \in \mathcal{B}_\infty\}.$$

In addition, by definition,  $u^*(c) = \inf\{u : (\gamma, u) \in \mathcal{B}_\infty\}$ , we have  $\tilde{u} \leq u^*(c)$  and  $\tilde{u}, u^*(c) \in \{u : (\gamma, u) \in \mathcal{B}_\infty\}$  such that  $|\tilde{u} - u^*(c)| \geq \varepsilon > 0$ , which draws the contradiction. Thus,  $u^*(c)$  is continuous in  $c$ .

## D Proof for Section 6

### D.1 Proof of Proposition 6.1

Recall that we assumed features and augmentations with inhomogeneous covariance take the following form.

$$\begin{aligned} \mathbf{h}_{0,i} &\stackrel{\text{i.i.d.}}{\sim} \frac{1}{2} \mathcal{N}(-\boldsymbol{\mu}, \mathbf{I}_p) + \frac{1}{2} \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_p) \\ \mathbf{h}_i, \mathbf{h}_i^+ | \mathbf{h}_{0,i} &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{h}_{0,i}, \sigma_{\text{aug}}^2 \mathbf{A}). \end{aligned}$$

where  $\mathbf{A} \succeq \mathbf{I}_p$ . The special case where  $\mathbf{A} = \mathbf{I}_p$  is already analyzed before. Here, we can write

$$\begin{aligned} \mathbf{h}_i &\stackrel{d}{=} \xi_i \boldsymbol{\mu} + (\mathbf{I}_p + \sigma_{\text{aug}}^2 \mathbf{A})^{1/2} \mathbf{g}_i, \quad \text{where} \\ \xi_i \perp \mathbf{g}_i, \quad \xi_i &\sim \text{Unif}(\{\pm 1\}), \quad \mathbf{g}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p). \end{aligned}$$

We have

$$\begin{aligned} \mathbb{E}[\|\mathbf{W}\mathbf{h}_i\|^2] &= \|\mathbf{W}\boldsymbol{\mu}\|^2 + \text{Tr}(\mathbf{W}(\mathbf{I}_p + \sigma_{\text{aug}}^2 \mathbf{A})\mathbf{W}^\top) \\ &= \|\mathbf{W}\boldsymbol{\mu}\|^2 + \|\mathbf{W}\|_{\text{F}}^2 + \sigma_{\text{aug}}^2 \text{Tr}(\mathbf{W}\mathbf{A}\mathbf{W}^\top) =: \tilde{\alpha}, \\ \mathbb{E}[\|\mathbf{W}\mathbf{h}_i^+ - \mathbf{W}\mathbf{h}_i\|^2] &= 2\sigma_{\text{aug}}^2 \text{Tr}(\mathbf{W}\mathbf{A}\mathbf{W}^\top) \\ \|\mathbf{W}\mathbf{h}_i - \mathbf{W}\mathbf{h}_i^-\|^2 &= \|\mathbf{W}(\xi_i - \xi_i^-)\boldsymbol{\mu} + \mathbf{W}(\mathbf{I}_p + \sigma_{\text{aug}}^2 \mathbf{A})^{1/2}(\mathbf{g}_i - \mathbf{g}_i^-)\|^2. \end{aligned}$$

Consider the same infinite-sample loss  $\mathcal{L}(\mathbf{W})$  in Section 4. The alignment loss can be written as

$$\mathcal{L}_{\text{align}}(\mathbf{W}) = \frac{1}{2\tau\tilde{\alpha}} \mathbb{E}_{\mathbf{h}, \mathbf{h}^+} [\|\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^+\|^2] = \frac{\sigma_{\text{aug}}^2}{\tau\tilde{\alpha}} \text{Tr}(\mathbf{W}\mathbf{A}\mathbf{W}^\top).$$

Further, the difference between the negative pair has the distribution

$$\mathbf{h} - \mathbf{h}^- \stackrel{d}{=} \frac{1}{2}\delta_0 + \frac{1}{4}\delta_{2\boldsymbol{\mu}} + \frac{1}{4}\delta_{-2\boldsymbol{\mu}} + \sqrt{2}(\mathbf{I} + \sigma_{\text{aug}}^2 \mathbf{A})^{1/2} \mathbf{g},$$

where  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})$ . Let  $\widetilde{\mathbf{W}} = \mathbf{W}(\mathbf{I} + \sigma_{\text{aug}}^2 \mathbf{A})^{1/2}$  and  $\widetilde{\boldsymbol{\mu}} = (\mathbf{I} + \sigma_{\text{aug}}^2 \mathbf{A})^{-1/2} \boldsymbol{\mu}$ . Then, we have

$$\mathbf{W}\mathbf{h} - \mathbf{W}\mathbf{h}^- \stackrel{d}{=} \frac{1}{2}\delta_0 + \frac{1}{4}\delta_{2\widetilde{\mathbf{W}}\widetilde{\boldsymbol{\mu}}} + \frac{1}{4}\delta_{-2\widetilde{\mathbf{W}}\widetilde{\boldsymbol{\mu}}} + \sqrt{2}\widetilde{\mathbf{W}}\mathbf{g},$$

which has the same form as in Section B.1. Therefore, with the SVD of  $\widetilde{\mathbf{W}} = \sum_{j=1}^p \widetilde{\sigma}_j \widetilde{\mathbf{u}}_j \widetilde{\mathbf{v}}_j^\top$ , we can calculate the following expectation

$$\begin{aligned} & \mathbb{E}_{\mathbf{h}^-, \mathbf{h}} \left[ \exp \left( -\frac{\|\widetilde{\mathbf{W}}\mathbf{h} - \widetilde{\mathbf{W}}\mathbf{h}^-\|_2^2}{2\tau\widetilde{\alpha}} \right) \right] \\ &= \frac{1}{2} \prod_{j=1}^p \left( 1 + \frac{2\widetilde{\sigma}_j^2}{\tau\widetilde{\alpha}} \right)^{-1/2} \left( 1 + \exp \left( -\sum_{j=1}^p \frac{2\widetilde{\sigma}_j^2}{2\widetilde{\sigma}_j^2 + \tau\widetilde{\alpha}} (\widetilde{\mathbf{v}}_j^\top \widetilde{\boldsymbol{\mu}})^2 \right) \right). \end{aligned}$$

As a result, we have decomposition  $\mathcal{L}(\mathbf{W}) = \mathcal{L}_{\text{align}}(\mathbf{W}) + \mathcal{L}_{\text{unif}}(\mathbf{W})$ , where

$$\begin{aligned} \mathcal{L}_{\text{align}}(\mathbf{W}) &= \frac{\sigma_{\text{aug}}^2}{\tau\widetilde{\alpha}} \text{Tr}(\mathbf{W}\mathbf{A}\mathbf{W}^\top) \\ \mathcal{L}_{\text{unif}}(\mathbf{W}) &= -\frac{1}{2} \sum_{j=1}^p \log \left( 1 + \frac{2\widetilde{\sigma}_j^2}{\tau\widetilde{\alpha}} \right) + \log \left( 1 + \exp \left( -\sum_{j=1}^p \frac{2\widetilde{\sigma}_j^2 \langle \widetilde{\boldsymbol{\mu}}, \widetilde{\mathbf{v}}_j \rangle^2}{2\widetilde{\sigma}_j^2 + \tau\widetilde{\alpha}} \right) \right) - \log 2. \end{aligned}$$

Further, if we assume that  $\widetilde{\alpha} \gg \|\mathbf{W}\|_{\text{F}}^2$ , then  $\mathcal{L}_{\text{unif}}(\mathbf{W})$  can be approximated by

$$\begin{aligned} \widetilde{\mathcal{L}}_{\text{unif}}(\mathbf{W}) &= -\sum_{j=1}^p \frac{\widetilde{\sigma}_j^2}{\tau\widetilde{\alpha}} + \log \left( 1 + \exp \left( -\sum_{j=1}^p \frac{2\widetilde{\sigma}_j^2 \langle \widetilde{\boldsymbol{\mu}}, \widetilde{\mathbf{v}}_j \rangle^2}{\tau\widetilde{\alpha}} \right) \right) - \log 2 \\ &= -\frac{\|\widetilde{\mathbf{W}}\|_{\text{F}}^2}{\tau\widetilde{\alpha}} + \log \left( 1 + \exp \left( -\frac{2\|\widetilde{\mathbf{W}}\widetilde{\boldsymbol{\mu}}\|_2^2}{\tau\widetilde{\alpha}} \right) \right) - \log 2 \\ &= -\frac{1}{\tau\widetilde{\alpha}} \left( \|\mathbf{W}\|_{\text{F}}^2 + \sigma_{\text{aug}}^2 \text{Tr}(\mathbf{W}\mathbf{A}\mathbf{W}^\top) \right) + \log \left( 1 + \exp \left( -\frac{2\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\tau\widetilde{\alpha}} \right) \right) - \log 2. \end{aligned}$$

We can then introduce an approximate loss  $\widetilde{\mathcal{L}}(\mathbf{W}) = \mathcal{L}_{\text{align}}(\mathbf{W}) + \widetilde{\mathcal{L}}_{\text{unif}}(\mathbf{W})$ , where

$$\widetilde{\mathcal{L}}(\mathbf{W}) = -\frac{1}{\tau\widetilde{\alpha}} \|\mathbf{W}\|_{\text{F}}^2 + \log \left( 1 + \exp \left( -\frac{2\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\tau\widetilde{\alpha}} \right) \right).$$

## D.2 Proof of Theorem 6.2

Recall that the minimization problem is  $\min_{\mathbf{W}} \widetilde{\mathcal{L}}(\mathbf{W})$ , where

$$\begin{aligned} \widetilde{\mathcal{L}}(\mathbf{W}) &= -\frac{1}{\tau\widetilde{\alpha}} \|\mathbf{W}\|_{\text{F}}^2 + \log \left( 1 + \exp \left( -\frac{2\|\mathbf{W}\boldsymbol{\mu}\|_2^2}{\tau\widetilde{\alpha}} \right) \right), \\ \widetilde{\alpha} &= (1 + \sigma_{\text{aug}}^2) \|\mathbf{W}\|_{\text{F}}^2 + \|\mathbf{W}\boldsymbol{\mu}\|_2^2 + \sigma_{\text{aug}}^2 \rho_{\text{aug}} \|\mathbf{W}\mathbf{v}_{\text{aug}}\|_2^2. \end{aligned}$$

First, notice that the loss  $\tilde{\mathcal{L}}(\mathbf{W})$  is scale-invariant. Hence without loss of generality, we can assume  $\|\mathbf{W}\|_{\mathbb{F}}^2 = 1$ . Then  $\tilde{\mathcal{L}}(\mathbf{W})$  only depends on the values of

$$T(\mathbf{W}) = \|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2, \quad \text{and} \quad \|\mathbf{W}\mathbf{v}_{\text{aug}}\|^2.$$

Second, we can rewrite the optimization problem into a nested one:

$$\min_{T \in [0,1]} \left[ \min_{\mathbf{W}: T(\mathbf{W})=T} \tilde{\mathcal{L}}(\mathbf{W}) \right]. \quad (44)$$

It is easy to see that given any fixed value of  $T(\mathbf{W})$ , the loss  $\tilde{\mathcal{L}}(\mathbf{W})$  is decreasing as  $\|\mathbf{W}\mathbf{v}_{\text{aug}}\|^2$  decreases. Thus, the inner minimization problem is easy to solve: we only need to determine the smallest possible value of  $\|\mathbf{W}\mathbf{v}_{\text{aug}}\|^2$ , which will be the focus below.

Consider the Gram–Schmidt orthogonalization for  $\bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}}$ : Let  $\mathbf{v}_{\perp} \in \text{span}(\bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}})$  be a unit vector orthogonal to  $\bar{\boldsymbol{\mu}}$ ; define  $a \in [-1, 1]$  to be coefficient in the orthogonal decomposition

$$\bar{\boldsymbol{\mu}} = r\mathbf{v}_{\text{aug}} + a\mathbf{v}_{\perp},$$

where we recall that  $r = \langle \bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}} \rangle \in [-1, 1]$ . It is also clear that  $a^2 + r^2 = 1$ .

We split the inner minimization into two cases.

**Case 1:**  $0 \leq T \leq a^2$ . We claim that there exists some  $\mathbf{W}$  such that  $\|\mathbf{W}\|_{\mathbb{F}}^2 = 1$ ,  $\|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2 = T$  and  $\|\mathbf{W}\mathbf{v}_{\text{aug}}\| = 0$ . To prove this claim, we will construct a  $\mathbf{W}$  of the form

$$\mathbf{W} = \sum_{j=1}^p s_j \bar{\mathbf{u}}_j \bar{\mathbf{v}}_j^{\top},$$

such that the required equalities hold. Let us set

$$\begin{aligned} \bar{\mathbf{v}}_p &= \mathbf{v}_{\text{aug}}, & s_p &= 0, \\ s_1^2 &= \frac{T}{a^2} \leq 1, & \bar{\mathbf{v}}_1 &= \mathbf{v}_{\perp}. \end{aligned}$$

Choose  $(\bar{\mathbf{v}}_j)_{2 \leq j \leq p-1}$  be any orthogonal basis in the orthogonal complement of  $\text{span}(\bar{\mathbf{v}}_1, \bar{\mathbf{v}}_p)$ , and  $\bar{\mathbf{u}}_1, \dots, \bar{\mathbf{u}}_p$  be any orthogonal basis in  $\mathbb{R}^p$ . Also, choose any  $s_2, \dots, s_{p-1} \geq 0$  such that  $s_2^2 + \dots + s_{p-1}^2 = 1 - T/a^2$ .

Our construction in fact gives an SVD of  $\mathbf{W}$ , with  $(s_p, \bar{\mathbf{v}}_p)$  being the bottom singular pair. So we have  $\mathbf{W}\mathbf{v}_{\text{aug}} = s_p \mathbf{v}_{\text{aug}} = 0$ . Moreover,

$$\begin{aligned} \|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2 &= \sum_{j=1}^p s_j^2 \langle \bar{\mathbf{v}}_j, \bar{\boldsymbol{\mu}} \rangle^2 \\ &= \sum_{j=1}^{p-1} s_j^2 \langle \bar{\mathbf{v}}_j, \bar{\boldsymbol{\mu}} \rangle^2 \quad \text{since } s_p = 0 \\ &= s_1^2 \langle \bar{\mathbf{v}}_1, \bar{\boldsymbol{\mu}} \rangle^2 \quad \text{by orthogonality } \boldsymbol{\mu} \perp (\bar{\mathbf{v}}_j)_{2 \leq j \leq p-1} \\ &= s_1^2 a^2 = T. \end{aligned}$$

Thus the inner minimization of (44) is solved at  $\|\mathbf{W}\mathbf{v}_{\text{aug}}\| = 0$ .

In all, under the constraint  $T \leq a^2$ , the problem of minimizing  $\tilde{\mathcal{L}}(\mathbf{W})$  boils down to

$$\min_{T \leq a^2} -\frac{1}{\tau [(1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T]} + \log \left( 1 + \exp \left( -\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau [(1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T]} \right) \right).$$

This objective is in fact the same as that in Section 4 via change of variables. Denote  $t = [(1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T]^{-1} \in [(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2 a^2)^{-1}, (1 + \sigma_{\text{aug}}^2)^{-1}]$ . Then the above objective function can be written as

$$h(t) = -\frac{t}{\tau} + \log \left( 1 + \exp \left( -\frac{2}{\tau} + 2\frac{1 + \sigma_{\text{aug}}^2}{\tau} t \right) \right),$$

for which

$$h'(t) = \frac{1 + 2\sigma_{\text{aug}}^2}{\tau} - 2\frac{1 + \sigma_{\text{aug}}^2}{\tau} \frac{1}{1 + \exp \left( -\frac{2}{\tau} + 2\frac{1 + \sigma_{\text{aug}}^2}{\tau} t \right)}.$$

As in the homogeneous case, note that  $h'(t)$  is increasing in  $t$  and  $h'((1 + \sigma_{\text{aug}}^2)^{-1}) = \sigma_{\text{aug}}^2/\tau > 0$ . Further, notice that

$$h'((1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2 a^2)^{-1}) > 0 \iff \tau > \frac{2a^2\|\boldsymbol{\mu}\|^2}{\log(1 + 2\sigma_{\text{aug}}^2)(1 + \sigma_{\text{aug}}^2 + a^2\|\boldsymbol{\mu}\|^2)} = \tau_1^*.$$

Then, if  $\tau > \tau_1^*$ , then  $h'((1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2 a^2)^{-1}) > 0$  and  $h(t)$  is increasing in  $t \in [(1 + \sigma_{\text{aug}}^2 + \|\boldsymbol{\mu}\|^2 a^2)^{-1}, (1 + \sigma_{\text{aug}}^2)^{-1}]$ , so the minimum of  $h(t)$  is achieved when  $t$  is the smallest in the interval, or equivalently  $T = a^2$ . If  $\tau \leq \tau_1^*$ , the minimizer in the interval  $[0, a^2]$  is determined by  $h'(t) = 0$ , which yields

$$t^* = \frac{1}{1 + \sigma_{\text{aug}}^2} \left( 1 - \frac{\tau}{2} \log(1 + 2\sigma_{\text{aug}}^2) \right).$$

As a result, we have

$$T^* = \frac{1}{\|\boldsymbol{\mu}\|^2} \left( \frac{1}{t^*} - (1 + \sigma_{\text{aug}}^2) \right) = \frac{\tau(1 + \sigma_{\text{aug}}^2)}{2\|\boldsymbol{\mu}\|^2} \log(1 + 2\sigma_{\text{aug}}^2) \left[ 1 - \frac{\tau}{2} \log(1 + 2\sigma_{\text{aug}}^2) \right]^{-1}. \quad (45)$$

**Case 2:**  $a^2 < T \leq 1$ . Note that we implicitly assume  $|a| < 1$  since  $|a| = 1$  is always covered in Case 1. In order to solve the inner minimization problem of (44), it suffices to consider the semidefinite program

$$\begin{aligned} \min \quad & \mathbf{v}_{\text{aug}}^\top \mathbf{A} \mathbf{v}_{\text{aug}} \\ \text{subject to} \quad & \text{Tr}(\mathbf{A}) = 1, \quad \bar{\boldsymbol{\mu}}^\top \mathbf{A} \bar{\boldsymbol{\mu}} = T, \quad \mathbf{A} \succeq 0, \end{aligned} \quad (46)$$

where we define  $\mathbf{A} = \mathbf{W}^\top \mathbf{W} \in \mathbb{R}^{p \times p}$ . The square-root matrix of the minimizer gives  $\mathbf{W}$  that obeys  $\|\mathbf{W} \bar{\boldsymbol{\mu}}\|^2 = T$  and achieves the minimum of the inner minimization problem.

The dual problem of (46) is given by

$$\begin{aligned} \max_{\lambda_1, \lambda_2 \in \mathbb{R}} \quad & -\lambda_1 - \lambda_2 T \\ \text{subject to} \quad & \mathbf{v}_{\text{aug}} \mathbf{v}_{\text{aug}}^\top + \lambda_1 \mathbf{I}_p + \lambda_2 \bar{\boldsymbol{\mu}} \bar{\boldsymbol{\mu}}^\top \succeq 0. \end{aligned} \quad (47)$$

By duality (in particular, complementary slackness), we know that the optimal  $\mathbf{A}$  and  $\lambda_1, \lambda_2$  obey

$$\left\langle \mathbf{v}_{\text{aug}} \mathbf{v}_{\text{aug}}^\top + \lambda_1 \mathbf{I} + \lambda_2 \bar{\boldsymbol{\mu}} \bar{\boldsymbol{\mu}}^\top, \mathbf{A} \right\rangle = 0. \quad (48)$$

For now, we will assume that  $a > 0$ ; the case  $a = 0$  (i.e.,  $|\langle \bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}} \rangle| = 1$ ) will be discussed later. We consider two separate cases: (1) the optimal  $\lambda_1 = 0$ , and (2) the optimal  $\lambda_1 > 0$ .

First, when  $\lambda_1 = 0$ , it is easy to check that the optimal choice of  $\lambda_2$  is also 0, which yields the optimal objective 0. Second, when  $\lambda_1 > 0$ , we claim that the optimal  $\mathbf{A}$  is a rank-one matrix, and it is given by

$$\mathbf{A} = (\sqrt{T} \bar{\boldsymbol{\mu}} - \sqrt{1-T} \bar{\boldsymbol{\mu}}_{\perp})(\sqrt{T} \bar{\boldsymbol{\mu}} - \sqrt{1-T} \bar{\boldsymbol{\mu}}_{\perp})^{\top}, \quad (49)$$

where  $\bar{\boldsymbol{\mu}}_{\perp} = a\mathbf{v}_{\text{aug}} - r\mathbf{v}_{\perp}$ . To see this, we note that  $\mathbf{v}_{\text{aug}}\mathbf{v}_{\text{aug}}^{\top} + \lambda_1\mathbf{I}_p + \lambda_2\bar{\boldsymbol{\mu}}\bar{\boldsymbol{\mu}}^{\top}$  has rank at least  $p-2$  since it is two rank-one updates of  $\lambda_1\mathbf{I}_p$ . In fact, we will show that this matrix has rank exactly  $p-1$ . Otherwise, we must have

$$\begin{aligned} \mathbf{v}_{\text{aug}}^{\top} (\mathbf{v}_{\text{aug}}\mathbf{v}_{\text{aug}}^{\top} + \lambda_1\mathbf{I}_p + \lambda_2\bar{\boldsymbol{\mu}}\bar{\boldsymbol{\mu}}^{\top}) \mathbf{v}_{\text{aug}} &= 0, \\ \bar{\boldsymbol{\mu}}^{\top} (\mathbf{v}_{\text{aug}}\mathbf{v}_{\text{aug}}^{\top} + \lambda_1\mathbf{I}_p + \lambda_2\bar{\boldsymbol{\mu}}\bar{\boldsymbol{\mu}}^{\top}) \bar{\boldsymbol{\mu}} &= 0. \end{aligned}$$

These are equivalent to

$$\begin{aligned} 1 + \lambda_1 + \lambda_2 \langle \bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}} \rangle^2 &= 0, \\ \lambda_1 + \lambda_2 + \langle \bar{\boldsymbol{\mu}}, \mathbf{v}_{\text{aug}} \rangle^2 &= 0, \end{aligned}$$

which implies  $a^2 \cdot (1 - \lambda_2) = 0$ . This further implies  $\lambda_2 = 1$  since  $a > 0$ . However, we know that when  $\lambda_1 > 0$ , the optimal  $\lambda_2$  would not be 1, as  $\lambda_2 = 0$  is strictly better. Therefore, we have shown that  $\mathbf{v}_{\text{aug}}\mathbf{v}_{\text{aug}}^{\top} + \lambda_1\mathbf{I}_p + \lambda_2\bar{\boldsymbol{\mu}}\bar{\boldsymbol{\mu}}^{\top}$  has rank  $p-1$ . By (48) primal optimal  $\mathbf{A}$  is a rank-one matrix, which together with the primal constraint  $\bar{\boldsymbol{\mu}}^{\top} \mathbf{A} \bar{\boldsymbol{\mu}} = T$  yields the conclusion that the optimal  $\mathbf{A}$  is uniquely determined by (49).

As a result, the optimal objective value of (46) is

$$\|\mathbf{W}\mathbf{v}_{\text{aug}}\|^2 = \langle \sqrt{T} \bar{\boldsymbol{\mu}} - \sqrt{1-T} \bar{\boldsymbol{\mu}}_{\perp}, \mathbf{v}_{\text{aug}} \rangle^2 = (r\sqrt{T} - a\sqrt{1-T})^2.$$

As this is always than 0 (the optimal objective when  $\lambda_1 = 0$ ), we conclude that the maximal objective of (47) is  $\langle \sqrt{T} \bar{\boldsymbol{\mu}} - \sqrt{1-T} \bar{\boldsymbol{\mu}}_{\perp}, \mathbf{v}_{\text{aug}} \rangle^2 = (r\sqrt{T} - a\sqrt{1-T})^2$ .

In the corner case  $a = 0$ , trivially, the objective is a constant:  $\mathbf{v}_{\text{aug}}^{\top} \mathbf{A} \mathbf{v}_{\text{aug}} = \bar{\boldsymbol{\mu}}^{\top} \mathbf{A} \bar{\boldsymbol{\mu}} = T$ , which is the same as  $(r\sqrt{T} - a\sqrt{1-T})^2$  in this case.

In all, under the constraint  $a^2 < T \leq 1$ , the original minimization problem (44) becomes

$$\begin{aligned} \min_{1 \geq T > a^2} & - \frac{1}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 (r\sqrt{T} - a\sqrt{1-T})^2 \right]} \\ & + \log \left( 1 + \exp \left( - \frac{2\|\boldsymbol{\mu}\|^2 T}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 (r\sqrt{T} - a\sqrt{1-T})^2 \right]} \right) \right). \end{aligned}$$

**Step 3: combining two cases.** Recall the notation  $x_+ = \max\{x, 0\}$ . Introduce the unified loss  $\ell(T)$  that subsumes both cases:

$$\begin{aligned} \ell(T) &= - \frac{1}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 [(r\sqrt{T} - a\sqrt{1-T})_+]^2 \right]} \\ & + \log \left( 1 + \exp \left( - \frac{2\|\boldsymbol{\mu}\|^2 T}{\tau \left[ (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 [(r\sqrt{T} - a\sqrt{1-T})_+]^2 \right]} \right) \right). \end{aligned}$$



Notice that  $\ell(T) \geq \ell^*(T)$  for each  $T \in [0, 1]$  where

$$\ell^*(T) := -\frac{1}{\tau [(1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T]} + \log \left( 1 + \exp \left( -\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau [(1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T]} \right) \right),$$

and  $\ell(T) = \ell^*(T)$  for  $T \in [0, a^2]$ . In addition, the two functions satisfy that

$$\ell^*(a^2) = \ell(a^2), \quad (\ell^*)'(a^2) = \ell'(a^2).$$

If  $\tau \leq \tau_1^*$ , we have seen in Step 1 that the minimizer  $T^*$  of  $\ell^*(T)$  is given by (45) and  $T^* \in [0, a^2]$ . Since  $\ell(T) \geq \ell^*(T)$  and equality holds if  $T \leq a^2$ , it verifies that  $T^*$  is also the minimizer of  $\ell(T)$ .

If  $\tau > \tau_1^*$ , the analysis in Step 1 shows that for  $T \in [0, a^2]$ ,  $(\ell^*)(T)$  (and thus  $\ell(T)$ ) is monotonically decreasing. Further, denote

$$h(T) = (1 + \sigma_{\text{aug}}^2) + \|\boldsymbol{\mu}\|^2 T + \rho_{\text{aug}} \sigma_{\text{aug}}^2 (r\sqrt{T} - a\sqrt{1-T})^2,$$

which is an increasing function of  $T$  when  $T \geq a^2$ . The derivative of  $\ell(T)$  takes the following form:

$$\ell'(T) = \frac{1}{\tau} \frac{h'(T)}{h^2(T)} - \frac{2\|\boldsymbol{\mu}\|^2}{\tau} \frac{\exp\left(-\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau h(T)}\right)}{1 + \exp\left(-\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau h(T)}\right)} \frac{h(T) - Th'(T)}{h^2(T)}.$$

If  $a \neq 0$  and  $\rho_{\text{aug}} > 0$ , we have  $\lim_{T \rightarrow 1^-} h'(T) = +\infty$ , then  $\ell'(1-) > 0$ , which indicates that the minimizer is attained in the interval  $(a^2, 1)$ . (Note that  $T^* = 1$  is impossible!)

**Step 4: degenerate cases.** If  $\rho_{\text{aug}} = 0$ , our analysis is reduced to the homogeneous case, where  $T^* = 1$  if  $\tau > \tau^*$ .

If  $a = 0$  and  $\rho_{\text{aug}} > 0$ , we have  $\bar{\boldsymbol{\mu}} = \pm \mathbf{v}_{\text{aug}}$  and

$$\ell(T) = -\frac{1}{\tau [1 + \sigma_{\text{aug}}^2 + LT]} + \log \left( 1 + \exp \left( -\frac{2\|\boldsymbol{\mu}\|^2 T}{\tau [1 + \sigma_{\text{aug}}^2 + LT]} \right) \right)$$

where  $L = \|\boldsymbol{\mu}\|^2 + \rho_{\text{aug}} \sigma_{\text{aug}}^2$ . The analysis of the  $\ell(T)$  is similar to the homogeneous case after a change of variables.

### D.3 Proof of Corollary 6.3

As a consequence of Theorem 6.2, most of the statements in Corollary 6.3 are straightforward. Below we will focus on Item 2 and 3. Let  $S_k = \sum_{j \leq k} \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2$  for the  $k$ -th cumulative score for  $\boldsymbol{\mu}$ , where  $1 \leq k \leq p-1$ .

First, by the ordering of singular values, we have

$$(k+1)\sigma_{k+1}^2 \leq \|\mathbf{W}\|_F^2 = 1 \quad \Rightarrow \quad \sigma_{k+1} \leq \frac{1}{\sqrt{k+1}}.$$

Therefore, we have

$$\sum_{j > k} \sigma_j^2 \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 \leq \sigma_{k+1}^2 \cdot \sum_{j > k} \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 \leq \frac{1}{k+1} (1 - S_k).$$

We thus obtain an upper bound on  $\|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2$ :

$$\begin{aligned}\|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2 &= \sum_{j \leq k} \sigma_j^2 \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 + \sum_{j > k} \sigma_j^2 \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 \\ &\leq \sum_{j \leq k} \langle \mathbf{v}_j, \bar{\boldsymbol{\mu}} \rangle^2 + \frac{1}{k+1} (1 - S_k) \\ &\leq \frac{k}{k+1} S_k + \frac{1}{k+1}.\end{aligned}$$

If  $\|\mathbf{W}\bar{\boldsymbol{\mu}}\|^2 \geq c$  for certain constant  $c > 0$ , we can choose  $k \geq 1$  to the smallest integer with  $k \geq (2c)^{-1} - 1$ . Combining the upper bound and lower bound, we obtain  $S_k \geq c/4$ . This proves the first jump of the cumulative score. The second jump is because  $S_p - S_{p-1} = \langle \mathbf{v}_p, \bar{\boldsymbol{\mu}} \rangle^2 = r^2 > 0$  due to nondegeneracy.

Moreover, if  $\tau > \tau_1^*$ , the top right singular vector of  $\mathbf{W}$  is  $\mathbf{v}_1 = \sqrt{T^*} \bar{\boldsymbol{\mu}} - \sqrt{1 - T^*} \boldsymbol{\mu}_\perp$  where  $T^* > \sqrt{1 - r^2}$ . Since

$$\langle \mathbf{v}_1, \mathbf{v}_{\text{aug}} \rangle = r\sqrt{T^*} - \sqrt{1 - r^2} \sqrt{1 - T^*}$$

is monotone increasing in  $T^*$ . The inner product must be positive if  $\tau > \tau_1^*$ . Also,  $\langle \mathbf{v}_1, \bar{\boldsymbol{\mu}} \rangle = \sqrt{T^*} > 0$ , so this inner product is also positive.