

The Power of Contrast for Feature Learning: A Theoretical Analysis

Wenlong Ji

*Department of Statistics
Stanford University
Stanford, CA 94305, USA*

JWL2000@STANFORD.EDU

Zhun Deng

*Department of Computer Science
Columbia University
New York, NY 10027, USA*

ZHUNDENG@G.HARVARD.EDU

Ryumei Nakada

*Department of Statistics
Rutgers University
Piscataway, NJ 08854, USA*

RYUMEI.N@RUTGERS.EDU

James Zou

*Department of Biomedical Data Science
Stanford University
Stanford, CA 94305, USA*

JAMESZ@STANFORD.EDU

Linjun Zhang

*Department of Statistics
Rutgers University
Piscataway, NJ 08854, USA*

LINJUN.ZHANG@RUTGERS.EDU

Editor: Aapo Hyvarinen

Abstract

Contrastive learning has achieved state-of-the-art performance in various self-supervised learning tasks and even outperforms its supervised counterpart. Despite its empirical success, theoretical understanding of the superiority of contrastive learning is still limited. In this paper, under linear representation settings, (i) we provably show that contrastive learning outperforms the standard autoencoders and generative adversarial networks, two classical generative unsupervised learning methods, for both feature recovery and in-domain downstream tasks; (ii) we also illustrate the impact of labeled data in supervised contrastive learning. This provides theoretical support for recent findings that contrastive learning with labels improves the performance of learned representations in the in-domain downstream task, but it can harm the performance in transfer learning. We verify our theory with numerical experiments.

Keywords: Self-Supervised Learning, Contrastive Learning, Principal Component Analysis, Spiked Covariance Model, Supervised Contrastive Learning

1. Introduction

Deep supervised learning has achieved great success in various applications, including computer vision (Krizhevsky et al., 2012), natural language processing (Vaswani et al., 2017), and scientific computing (Han et al., 2018). However, its dependence on manually assigned labels, which is usually difficult and costly, has motivated research into alternative approaches to exploit unlabeled data. Self-supervised learning is a promising approach that leverages the unlabeled data itself as supervision and learns representations that are beneficial to potential in-domain downstream tasks.

At a high level, there are two common approaches for feature extraction in self-supervised learning: generative and contrastive (Liu et al., 2021; Jaiswal et al., 2021).

Both approaches aim to learn latent representations of the original data, while the difference is that the generative approach focused on minimizing the reconstruction error from latent representations, and the contrastive approach targets to decrease the similarity between the representations of contrastive pairs constructed by data augmentation. Recent works have shown the benefits of contrastive learning in practice (Chen et al., 2020a; He et al., 2020; Chen et al., 2020b,c). However, these works did not explain the popularity of contrastive learning — *what is the advantage of contrastive learning and where does it come from?*

Additionally, recent works aim to further improve contrastive learning by introducing label information. Specifically, Khosla et al. (2020) proposed the supervised contrastive learning, where the contrasting procedures are performed across different classes rather than different instances. With the help of label information, their proposed method outperforms self-supervised contrastive learning and classical cross-entropy-based supervised learning. However, despite this improvement in in-domain downstream tasks, Islam et al. (2021) found that such improvement in transfer learning is limited and even negative for such supervised contrastive learning. This phenomenon motivates us to rethink *the impact of labeled data* in the contrastive learning framework.

In this paper, we first establish a theoretical framework to study contrastive learning under the linear representation setting. Under this framework, we provide a theoretical analysis of the feature learning performance of the contrastive learning on the spiked covariance model (Bai and Yao, 2012; Yao et al., 2015; Zhang et al., 2018) and theoretically justify why contrastive learning outperforms standard autoencoders and generative adversarial networks (GANs) (Goodfellow et al., 2014) —contrastive learning is able to remove more noise by constructing contrastive samples via augmentations. Moreover, we investigate the impact of label information in the contrastive learning framework and provide a theoretical justification of why labeled data help to gain accuracy in in-domain regression and classification while can hurt multi-task transfer learning.

1.1 Related Works

The idea of contrastive learning was firstly proposed in Hadsell et al. (2006) as an effective method to perform dimensional reduction. Following this line of research, Dosovitskiy et al. (2014) proposed to perform instance discrimination by creating surrogate classes for each instance and Wu et al. (2018) further proposed to preserve a memory bank as a dictionary of negative samples. Other extensions based on this memory bank approach include He

et al. (2020); Misra and Maaten (2020); Tian et al. (2020); Chen et al. (2020c). Rather than keeping a costly memory bank, another line of work exploits the benefit of mini-batch training where different samples are treated as negative to each other (Ye et al., 2019; Chen et al., 2020a). Moreover, Khosla et al. (2020) explores the supervised version of contrastive learning where pairs are generated based on label information.

Despite its success in practice, the theoretical understanding of contrastive learning is still limited. Previous works provide provable guarantees for contrastive learning under conditional independence assumption (or its variants) (Arora et al., 2019; Lee et al., 2021; Tosh et al., 2021; Tsai et al., 2020). Specifically, they assume the two contrastive views are independent conditioned on the label and show that contrastive learning can provably learn representations beneficial for in-domain downstream tasks. In addition to this line of research, there exist several alternative perspectives for studying the theoretical properties of contrastive learning. To name a few, Wang and Isola (2020); Graf et al. (2021) explored the representation geometry, HaoChen et al. (2021) analyzed the augmentation graph, Tian (2022) proposed a two-player game theory framework, Zimmermann et al. (2021) demonstrated the connection between contrastive learning and nonlinear Independent Component Analysis (Hyvärinen et al., 2009), Saunshi et al. (2022) showed that the importance of inductive bias in contrastive learning, and Jing et al. (2021) investigated the dimensional collapse phenomenon. Furthermore, Tian et al. (2021); Wang et al. (2021) have also explored the ability of self-supervised learning to learn features even without contrastive pairs, specifically in the context of linear representation settings.

More relevant to this paper, Wen and Li (2021) considered representation learning under the sparse coding model and studied the optimization properties in shallow ReLU neural networks. However, the assumptions that features are extremely sparse and signals follow Gaussian distribution seem strong for real data. Garg and Liang (2020) studied the combination of supervised learning and self-supervised learning. They derived sample complexity bounds in a PAC-learning style for various settings. Specifically, the authors assume that there is a ground-truth representation such that it can keep both self-supervised loss and supervised loss at a very low threshold. However, as the authors admit, it is hard to determine such a threshold in practical settings. For example, since the unlabeled data and labeled data come from different domains, such as Image-Net and CIFAR-10, domain-specific features may have a much lower loss compared with domain-transferable features.

While the aforementioned previous works aim to demonstrate that contrastive learning is capable of learning meaningful representations, it was left untouched why contrastive learning *outperforms* other representation learning methods. We also shed light on the impact of labeled data in a contrastive learning framework, which is underexplored in prior works. A detailed comparison with existing literature is deferred to Appendix A.1.

1.2 Outline

This paper is organized as follows. Section 2 provides the setup for the data-generating process and the loss function. In Section 3, we review the connection between PCA and autoencoders/GANs. We also establish a theoretical framework to study contrastive learning in the linear representation setting. Under this framework, we evaluate the feature recovery performance and in-domain downstream task performance of contrastive learning

and autoencoders. In Section 4, we analyze the supervised contrastive learning. In Section 5, we verify our theoretical results given in Sections 3 and 4. Finally, we summarize our analysis and provide future directions in Section 6.

1.3 Notations

In this paper, we use O, Ω, Θ to hide universal constants and we write $a_k \lesssim b_k$ for two sequences of positive numbers $\{a_k\}$ and $\{b_k\}$ if and only if there exists a universal constant $C > 0$ such that $a_k < Cb_k$ for any k . We write $a_k \asymp b_k$ when $a_k \lesssim b_k$ and $a_k \gtrsim b_k$ holds simultaneously. We use $\|\cdot\|, \|\cdot\|_2, \|\cdot\|_F$ to represent the ℓ_2 norm of vectors, the spectral norm of matrices, and Frobenius norm of matrices respectively. Let $\mathbb{O}_{d,r}$ be a set of $d \times r$ orthogonal matrices. Namely, $\mathbb{O}_{d,r} \triangleq \{U \in \mathbb{R}^{d \times r} : U^\top U = I_r\}$. We write $n \gg d$ when there exists a sufficiently small constant c depending on the constant and independent of n, d and r such that $d/n < c$ holds. $d \gg r$ is defined similarly. We use $|A|$ to denote the cardinality of a set A . For any $n \in \mathbb{N}^+$, let $[n] = \{1, 2, \dots, n\}$. We use $\|\sin \Theta(U_1, U_2)\|_F$ to refer to the sine distance between two orthogonal matrices $U_1, U_2 \in \mathbb{O}_{d,r}$, which is defined by: $\|\sin \Theta(U_1, U_2)\|_F \triangleq \|U_{1\perp}^\top U_2\|_F$, where $U_{1\perp} \in \mathbb{O}_{d-r,r}$ is any orthogonal complement of U_1 . More properties of sine distance can be found in Section A.3. We use $\{e_i\}_{i=1}^d$ to denote the canonical basis in d -dimensional Euclidean space \mathbb{R}^d , that is, e_i is the vector whose i -th coordinate is 1 and all the other coordinates are 0. Let $\mathbb{I}\{A\}$ be an indicator function that takes 1 when A is true, otherwise takes 0. We write $a \vee b$ and $a \wedge b$ to denote $\max(a, b)$ and $\min(a, b)$, respectively.

2. Setup

Here we introduce loss and data-generative models that will be used for the theoretical analysis later.

2.1 Linear Representation Settings for Contrastive Learning

Given an input $x \in \mathbb{R}^d$, contrastive learning aims to learn a low-dimensional representation $h = f(x; \theta) \in \mathbb{R}^r$ by contrasting different samples, that is, maximizing the agreement between positive pairs, and minimizing the agreement between negative pairs. Suppose we have n data points $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{d \times n}$ from the population distribution \mathcal{D} . The contrastive learning task can be formulated to the following optimization problem:

$$\min_{\theta} \mathcal{L}(\theta) = \min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(x_i, \mathcal{B}_i^{Pos}, \mathcal{B}_i^{Neg}; f(\cdot; \theta)) + \lambda R(\theta), \quad (1)$$

where $\ell(\cdot)$ is a contrastive loss and $\lambda R(\theta)$ is a regularization term; $\mathcal{B}_i^{Pos}, \mathcal{B}_i^{Neg}$ are the sets of positive samples and negative samples corresponding to x_i , the details of which are described below.

Linear Representation and Regularization Term We consider the linear representation function $f(x; W) = Wx$, where the parameter θ is a matrix $W \in \mathbb{R}^{r \times d}$. This linear representation setting has been widely adopted in other theory papers to understand self-supervised contrastive learning (Jing et al., 2021; Wang et al., 2021; Tian et al., 2021) and

shed light upon other complex machine learning phenomena such as in Tripuraneni et al. (2021). Moreover, since regularization techniques have been widely adopted in contrastive learning practice (Chen et al., 2020a; He et al., 2020; Grill et al., 2020), we further consider penalizing the representation by a regularization term $R(W) = \|WW^\top\|_F^2/2$ to encourage the orthogonality of W and therefore promote the diversity of w_i to learn different representations. The reason we use such quadratic regularization instead of a standard ℓ_2 regularization is to encourage a diverse representation in the linear representation setting by penalizing on the similarity $\langle w_i, w_j \rangle^2$, we defer a formal discussion and numerical experiments about this regularization in the Appendix A.2.

Linear Contrastive Loss The contrastive loss is set to be the average similarity (measured by the inner product) between positive pairs minus that between negative pairs:

$$\ell(x, \mathcal{B}_x^{Pos}, \mathcal{B}_x^{Neg}, f(\cdot; \theta)) = - \sum_{x^{Pos} \in \mathcal{B}_x^{Pos}} \frac{\langle f(x, \theta), f(x^{Pos}, \theta) \rangle}{|\mathcal{B}_x^{Pos}|} + \sum_{x^{Neg} \in \mathcal{B}_x^{Neg}} \frac{\langle f(x, \theta), f(x^{Neg}, \theta) \rangle}{|\mathcal{B}_x^{Neg}|}, \quad (2)$$

where $\mathcal{B}_x^{Pos}, \mathcal{B}_x^{Neg}$ are sets of positive samples and negative samples corresponding to x . This loss function has been commonly used in contrastive learning (Hadsell et al., 2006) and metric learning (Schroff et al., 2015; He et al., 2018). In Khosla et al. (2020), the authors show that the inner-product based linear loss (2) is an approximation of the NT-Xent contrastive loss when one positive and one negative are used, which has been highlighted in recent contrastive learning practice (Sohn, 2016; Wu et al., 2018; Oord et al., 2018; Chen et al., 2020a). In Li et al. (2021), the authors proposed the SSL-HSIC contrastive loss, which can be reduced to this linear loss when the kernel $k(\cdot, \cdot)$ is chosen to be a simple inner product. Following Li et al. (2021), we provide the results in Table 1, which shows that linear contrastive loss can also work well with some additional training techniques.

Testing Accuracy	InfoNCE	Linear contrastive loss
CIFAR10	65.11 ± 0.51	66.07 ± 0.46
STL10	71.02 ± 0.47	70.30 ± 0.31

Table 1: **InfoNCE loss v.s. Linear contrastive loss.** We train a ResNet-18 encoder on CIFAR-10 and STL-10 datasets with different contrastive loss functions. To train the linear contrastive loss, we follow the HSIC regularization techniques used in Li et al. (2021), which helps linear contrastive loss yield comparable performance to standard InfoNCE. We repeat each experiment for ten runs and report the mean and standard deviation of accuracy. Detailed experimental settings can be found in Section 5.2.

2.2 Generation of Positive and Negative Pairs

There are two common approaches to generating positive and negative pairs, depending on whether or not label information is available. When the label information is not available, the typical strategy is to generate different views of the original data via augmentation (Hadsell et al., 2006; Chen et al., 2020a). Two views of the same data point serve as the positive pair for each other, while those of different data serve as negative pairs.

Definition 1 (Augmented Pairs Generation in the Self-supervised Setting) *Given two augmentation functions $g_1, g_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and n training samples $\mathcal{B} = \{x_i\}_{i \in [n]}$, the augmented views are given by: $\{(g_1(x_i), g_2(x_i))\}_{i \in [n]}$. Then for each view $g_v(x_i)$, $v = 1, 2$, the corresponding positive samples and negative samples are defined by: $\mathcal{B}_{i,v}^{Pos} = \{g_s(x_i) : s \in [2] \setminus \{v\}\}$ and $\mathcal{B}_{i,v}^{Neg} = \{g_s(x_j) : s \in [2], j \in [n] \setminus \{i\}\}$.*

The loss function of the self-supervised contrastive learning problem can then be written as:

$$\mathcal{L}_{\text{SelfCon}}(W) = -\frac{1}{2n} \sum_{i=1}^n \sum_{v=1}^2 \left[\langle W g_v(x_i), W g_{[2] \setminus \{v\}}(x_i) \rangle - \sum_{j \neq i} \sum_{s=1}^2 \frac{\langle W g_v(x_i), W g_s(x_j) \rangle}{2n-2} \right] + \frac{\lambda}{2} \|WW^\top\|_F^2. \quad (3)$$

In particular, we adopt the following augmentation in our analysis.

Definition 2 (Random Masking Augmentation) *The two views of the original data are generated by randomly dividing its dimensions into two sets, that is, $g_1(x_i) = Ax_i$, and $g_2(x_i) = (I - A)x_i$, where $A = \text{diag}(a_1, \dots, a_d) \in \mathbb{R}^{d \times d}$ is the diagonal masking matrix with $\{a_i\}_{i=1}^d$ being i.i.d. random variables sampled from a Bernoulli distribution with mean $1/2$.*

Remark 3 *In this paper, we focus on random masking augmentation, which has also been used in other works on the theoretical understanding of contrastive learning, eg. Wen and Li (2021). However, our primary interest lies in comparing the performance of contrastive learning with autoencoders and analyzing the impact of labeled data, while their work focuses on understanding the training process of neural networks in contrastive learning. Random masking augmentation is an analog of the random cropping augmentation used in practice. As shown in Chen et al. (2020a), cropping augmentation achieves overwhelming performance on linear evaluation (ImageNet top-1 accuracy) compared with other augmentation methods, please see Figure 5 in Chen et al. (2020a) for details.*

When the label information is available, Khosla et al. (2020) proposed the following approach to generate positive and negative pairs.

Definition 4 (Pairs Generation in the Supervised Setting) *In a K -class classification problem, given n_k samples for each class $k \in [K]$: $\{x_i^k : i \in [n_k]\}_{k=1}^K$ and let $n = \sum_{k=1}^K n_k$, the corresponding positive samples and negative samples for x_i^k are defined by $\mathcal{B}_{i,k}^{Pos} = \{x_j^k : j \in [n_k] \setminus \{i\}\}$ and $\mathcal{B}_{i,k}^{Neg} = \{x_j^s : s \in [K] \setminus k, j \in [n_s]\}$. That is, the positive samples are the remaining ones in the same class with x_i^k and the negative samples are the samples from different classes.*

Correspondingly, the loss function of the supervised contrastive learning problem can be written as:

$$\mathcal{L}_{\text{SupCon}}(W) = -\frac{1}{nK} \sum_{k=1}^K \sum_{i=1}^{n_k} \left[\sum_{j \neq i} \frac{\langle W x_i^k, W x_j^k \rangle}{n-1} - \sum_{j=1}^n \sum_{s \neq k} \frac{\langle W x_i^k, W x_j^s \rangle}{n(K-1)} \right] + \frac{\lambda}{2} \|WW^\top\|_F^2. \quad (4)$$

2.3 Data Generating Process

In real-world scenarios, data often comprises both signal (relevant information) and noise (irrelevant distractions). For instance, in image classification, the signal might be the primary subject of interest, while the noise could represent background elements. Self-supervised learning methods, without predefined tasks, aim to extract generalized patterns from data, ideally capturing as much of the signal as possible. It is commonly understood that signals tend to exhibit specific low-complexity structures, often being low-rank and showing higher correlations across coordinates. In contrast, background noise might lack a distinct structure, potentially being dense (or full rank) with lower coordinate correlations. To delve into this structural difference more rigorously, we consider an additive data-generating model. Here, the observed data emerges as a combination of a low-rank signal and dense noise.

$$x = U^*z + \xi, \quad \text{Cov}(z) = \nu^2 I_r, \quad \text{Cov}(\xi) = \Sigma, \quad (5)$$

where $z \in \mathbb{R}^r$ and $\xi \in \mathbb{R}^d$ are both zero mean sub-Gaussian independent random variables, and $\nu \in \mathbb{R}$ is a constant represents the signal strength. In particular, $U^* \in \mathbb{O}_{d,r}$ and $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$. The first term U^*z represents the signal of interest residing in a low-dimensional subspace spanned by the columns of U^* . The second term ξ is the dense noise with heteroskedastic noise. Given that, the ideal low-dimensional representation is to compress the observed x into a low-dimensional representation spanned by the columns of U^* . This model is known as the spiked covariance model (Johnstone, 2001; Bai and Yao, 2012; Yao et al., 2015; Zhang et al., 2018). It was proposed from the empirical observation that the eigenvalues of the sample covariance matrix of phoneme data have few "spikes", which corresponds to the low-dimensional structure of data generation. The model has been used in the literature of PCA (Johnstone, 2001; Deshpande and Montanari, 2014; Zhang et al., 2018) and Contrastive Learning (Wen and Li, 2021).

In this paper, we aim to *learn a good projection* $W \in \mathbb{R}^{r \times d}$ onto a lower-dimensional subspace from the observation x . Since the information of W is invariant with the transformation $W \leftarrow OW$ for any $O \in \mathbb{O}_{r,r}$, the essential information of W is contained in the right eigenvector of W . Thus, we quantify the goodness of the representation W using the sine distance $\|\sin \Theta(U, U^*)\|_F$, where U is the top- r right eigenspace of W . It is notable that we only assume that noise and signal follow a sub-Gaussian distribution. This includes bounded noise/signals such as images, sound data, or text data.

3. Comparison of Self-Supervised Contrastive Learning and Autoencoders/GANs

Generative and contrastive learning are two popular approaches of self-supervised learning. Recent experiments have highlighted the improved performance of contrastive learning compared with the generative approach. For example, in Figure 1 of Chen et al. (2020a) and Figure 7 of Liu et al. (2021), it is observed that state-of-the-art contrastive self-supervised learning has more than 10 percent improvement over state-of-the-art generative self-supervised learning, with the same number of parameters. In this section, we rigorously demonstrate the advantage of contrastive learning over autoencoders/GANs, the representative methods in generative self-supervised learning, by investigating the linear representation settings

under the spiked covariance model (5). The investigation is conducted for both feature recovery and in-domain downstream tasks.

Hereafter, we focus on the linear representation settings. This section is organized as follows: in Section 3.1 we first review the connection between principal component analysis (PCA) and autoencoders/GANs, which are two representative methods in generative approaches in self-supervised learning, under linear representation settings. Then we establish the connection between contrastive learning and PCA in Section 3.2. Based on these connections, we make the comparison between contrastive learning and autoencoder on feature recovery ability (Section 3.3) and in-domain downstream performance (Section 3.4).

3.1 Autoencoders, GANs and PCA

Autoencoders are popular unsupervised learning methods to perform dimensional reduction. Autoencoders learn two functions: encoder $f : \mathbb{R}^d \rightarrow \mathbb{R}^r$ and decoder $g : \mathbb{R}^r \rightarrow \mathbb{R}^d$. While the encoder f compresses the original data into low-dimensional features, and the decoder g recovers the original data from those features. It can be formulated to be the following optimization problem for samples $\{x_i\}_{i=1}^n$ (Ballard, 1987; Fan et al., 2019):

$$\min_{f,g} \mathbb{E}_x \mathcal{L}(x, g(f(x))). \quad (6)$$

By minimizing this loss, autoencoders try to preserve the essential features to recover the original data in the low-dimensional representation. In our setting, we consider the class of linear functions for f and g . The loss function is set as the mean squared error. Write $f(x) = W_{\text{AE}}x$ and $g(x) = W_{\text{DE}}x$. Namely, we consider the following problem.

$$\min_{W_{\text{AE}}, W_{\text{DE}}} \frac{1}{n} \|X - W_{\text{DE}}W_{\text{AE}}X\|_F^2.$$

Let $X = (x_1, \dots, x_n) \in \mathbb{R}^{d \times n}$. By Theorem 2.4.8 in Golub and Loan (1996), the optimal solution is given by the eigenspace of XX^\top , which exactly corresponds to the result of PCA. Thus, in linear representation settings, autoencoders are equivalent to PCA, which is also often known as *undercomplete linear autoencoders* (Bourlard and Kamp, 1988; Plaut, 2018; Fan et al., 2019). We write the obtained low-rank representation by autoencoders as

$$W_{\text{AE}} = (U_{\text{AE}}\Sigma_{\text{AE}}V_{\text{AE}}^\top)^\top, \quad (7)$$

where U_{AE} is the top- r eigenvectors of matrix XX^\top , Σ_{AE} is a diagonal matrix of spectral values and $V_{\text{AE}} = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

We also note that GANs (Goodfellow et al., 2014) is related to PCA. Namely, Feizi et al. (2020) showed that the global solution for GANs recovers the empirical PCA solution as the generative model.

To see this, let \mathcal{W}_2 be the second-order Wasserstein distance. Also let \mathcal{G} be the set of linear generator functions from $\mathbb{R}^r \rightarrow \mathbb{R}^d$. Consider the following \mathcal{W}_2 GAN optimization problem:

$$\min_{g \in \mathcal{G}} \mathcal{W}_2^2(\mathbb{P}_n, \mathbb{P}_{g(Z)}), \quad (8)$$

where \mathbb{P}_n denotes the empirical distribution of i.i.d. data $x_1, \dots, x_n \in \mathbb{R}^d$ and $\mathbb{P}_{g(Z)}$ is the generated distribution with generator g and $Z \sim N(0, I_r)$. Note that the optimization problem Equation (8) can be written as $\min_{\mathbb{P}_n, Z} \min_{g \in \mathcal{G}} \mathbb{E}[\|X - g(Z)\|^2]$, where the first minimization is over probability distributions which have marginals \mathbb{P}_n and \mathbb{P}_Z . By Theorem 2 in Feizi et al. (2020), the optimizer of problem Equation (8) is obtained as $\hat{g} : Z \mapsto \hat{G}Z$, where \hat{G} satisfies $\hat{G}\hat{G}^\top = U_{\text{AE}}\Sigma_{\text{AE}}^2U_{\text{AE}}^\top$. This implies $W_{\text{AE}}^\top : \mathbb{R}^r \rightarrow \mathbb{R}^d$ is also a solution to the optimization problem Equation (8). Hence GANs learn the PCA solution as a generator.

From this equivalence among ordinary PCA, autoencoders, and GANs, we **only focus on autoencoders hereafter** for brevity.

3.2 Contrastive Learning and Diagonal-Deletion PCA

Here we bridge PCA and contrastive learning with certain augmentations under the linear representation setting. Recall that the optimization problem for self-supervised contrastive learning is formulated as:

$$\min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}_{\text{SelfCon}}(W) := -\frac{1}{2n} \sum_{i=1}^n \sum_{v=1}^2 \left[\langle Wg_v(x_i), Wg_{[2] \setminus \{v\}}(x_i) \rangle - \sum_{j \neq i} \sum_{s=1}^2 \frac{\langle Wg_v(x_i), Wg_s(x_j) \rangle}{2n-2} \right] + \frac{\lambda}{2} \|WW^\top\|_F^2. \quad (9)$$

To compare contrastive learning with autoencoders, we now derive the solution of the optimization problem (9). We start with the general result for self-supervised contrastive learning with augmented pairs generation in Definition 1, and then turn to the special case of random masking augmentation (Definition 2).

Proposition 5 *For two fixed augmentation functions $g_1, g_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$, denote the augmented data matrices as $X_1 = [g_1(x_1), \dots, g_1(x_n)] \in \mathbb{R}^{d \times n}$ and $X_2 = [g_2(x_1), \dots, g_2(x_n)] \in \mathbb{R}^{d \times n}$, when the augmented pairs are generated as in Definition 1, all the optimal solutions of contrastive learning problem (9) are given by:*

$$W_{CL} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$X_1 X_2^\top + X_2 X_1^\top - \frac{1}{2(n-1)} (X_1 + X_2)(1_n 1_n^\top - I_n)(X_1 + X_2)^\top, \quad (10)$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

The proof is given in Appendix B.1.

Proposition 5 is a general result for augmented pairs generation with fixed and deterministic augmentation functions. The result itself only depends on the augmented data matrices, thus it is straightforward to generalize to the case where different augmentation functions are applied to different samples, we omit it here for the simplicity of notations. Moreover, when the augmentation is sampled from a stochastic distribution, we can also

characterize the optimal solution of the expected loss in the same way. Specifically, if we apply the random masking augmentation (2), we can further obtain a result to characterize the optimal solution. For any square matrix $A \in \mathbb{R}^{d \times d}$, we denote $D(A)$ to be A with all off-diagonal entries set to be zero and $\Delta(A) = A - D(A)$ to be A with all diagonal entries set to be zero. Then we have the following corollary for random masking augmentation.

Corollary 6 *Under the same conditions as in Proposition 5, if we use random masking (Definition 2) as our augmentation function, then the minimizer of the expected loss function of contrastive learning problem (9) over the distribution of random augmentations (i.e., $\mathbb{E}_{g_1, g_2} \mathcal{L}_{SelfCon}(W)$) is given by:*

$$W_{CL} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n)X^\top, \quad (11)$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

The proof is given in Appendix B.2.

With Proposition 5 and Corollary 6 established, we can find that the self-supervised contrastive learning equipped with augmented pairs generation and random masking augmentation can eliminate the effect of random noise on the diagonal entries of the observed covariance matrix. Since $\text{Cov}(\xi) = \Sigma$ is a diagonal matrix, when the diagonal entries $\text{Cov}(U^*z) = \nu^2 U^* U^{*\top}$ only take a small proportion of the total Frobenius norm, the contrasting augmented pairs will preserve the core features while eliminating most of the random noise and give a more accurate estimation of core features.

3.3 Feature Recovery from Noisy Data

After bridging both autoencoder and contrastive learning with PCA, now we can perform the analysis of feature recovery ability to understand the benefit of contrastive learning over autoencoders. As mentioned above, our target is to recover the subspace spanned by the columns of U^* , which can further help us obtain information on the unobserved z that is important for in-domain downstream tasks. However, the observed data has a covariance matrix of $\nu^2 U^* U^{*\top} + \Sigma$ rather than the desired $\nu^2 U^* U^{*\top}$, which brings difficulty to representation learning. We demonstrate that contrastive learning can better exploit the structure of core features and obtain better estimation than autoencoders in this setting.

We start with autoencoders. In the noiseless case, the covariance matrix is $\nu^2 U^* U^{*\top}$ and autoencoders can *perfectly* recover the core features. However, in noisy cases, the random noises sometimes perturb the core features, which makes autoencoders fail to learn the core features. Such noisy cases are widespread in real applications such as measurement errors and backgrounds in images such as grasses and sky. Interestingly, we will later show that contrastive learning can better recover U^* despite the presence of large noise.

To provide rigorous analysis, we first introduce the incoherent constant (Candès and Recht, 2009).

Definition 7 (Incoherent Constant) *We define the incoherence constant of $U \in \mathbb{O}_{d,r}$ as*

$$I(U) = \max_{i \in [d]} \left\| e_i^\top U \right\|^2. \quad (12)$$

Intuitively, the incoherent constant measures the degree of the incoherence of the distribution of entries among different coordinates, or loosely speaking, the similarity between U and canonical basis $\{e_i\}_{i=1}^d$. For uncorrelated random noise, the covariance matrix is diagonal and its eigenspace is exactly spanned by the canonical basis $\{e_i\}_{i=1}^d$ (if the diagonal entries in Σ are all different), which attains the maximum value of the incoherent constant. On the contrary, the core features usually exhibit certain correlation structures and the corresponding eigenspace of the covariance matrix is expected to have a lower incoherent constant.

We then introduce a few assumptions which our theoretical results are built on. Recall that in the spiked covariance model (5), $x = U^*z + \xi$, $\text{Cov}(z) = \nu^2 I_r$ and $\text{Cov}(\xi) = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$.

Assumption 8 (Regular Covariance Condition) *The condition number of covariance matrix $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ satisfies $\kappa := \sigma_{(1)}^2 / \sigma_{(d)}^2 < C$, where $\sigma_{(j)}^2$ represents the j -th largest number among $\sigma_1^2, \dots, \sigma_d^2$ and $C > 0$ is a universal constant.*

Assumption 9 (Signal to noise ratio condition) *Define the signal-to-noise ratio $\rho := \nu / \sigma_{(1)}$, we assume $\rho = \Theta(1)$, implying that the covariance of noise is of the same order as that of the core features.*

Assumption 10 (Incoherent Condition) *The incoherent constant of the core feature matrix $U^* \in \mathbb{O}_{d,r}$ satisfies $I(U^*) = O(r \log d/d)$.*

The incoherent constant often appears in the literature of matrix completion (Candès and Recht, 2009) and PCA (Zhang et al., 2018). The order of $I(U^*)$ can be arbitrary as long as it decreases to 0 as $d \rightarrow \infty$. One can directly adapt the later results to this setting. If U is distributed uniformly on $\mathbb{O}_{d,r}$, then the expectation of incoherent constant is of order $r \log d/d$.

Lemma 11 (Expectation of incoherent constant over a uniform distribution)

$$\mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} I(U) = O\left(\frac{r}{d} \log d\right). \quad (13)$$

Thus, we set $I(U^*)$ to the order $r \log d/d$ for simplicity. The proof is given in Appendix B.6.

Here we provide a remark on the implication of our assumptions, and we defer a further discussion on how to generalize our main results under weaker assumptions in Remark 15.

Remark 12 *The three assumptions above can be explained as follows: Assumption 8 implies that the variances of all dimensions are of the same order. For Assumption 9, we focus on a large noise regime where the noise may hurt the estimation significantly. Here we assume the ratio lies in a constant range, but our theory can easily adapt to the case where ρ has a decreasing order. Specifically, for Theorems 13, 14, 17 and 18 presented below, we derive an explicit dependence on ρ of each result in the appendix. One can check Equations (46), (57), (59), (60), (61) and (62) for details. Assumption 10 implies a stronger correlation among the coordinates of core features, which is the essential property to distinguish them from random noise.*

Now we are ready to present our first result, showing that the autoencoders are unable to recover the core features in the large-noise regime. Due to the equivalence among PCA, autoencoders, and GANs we presented in Section 3.1, for brevity, we only focus on autoencoders hereafter.

Theorem 13 (Recovery Ability of Autoencoders, Lower Bound) *Consider the spiked covariance model (5), under Assumptions 8-10 and $n > d \gg r$, let W_{AE} be the learned representation of autoencoders with singular value decomposition $W_{AE} = (U_{AE}\Sigma_{AE}V_{AE}^\top)^\top$ (as in Equation (7)). If we further assume $\{\sigma_i^2\}_{i=1}^d$ are different from each other and $\sigma_{(1)}^2/(\sigma_{(r)}^2 - \sigma_{(r+1)}^2) < C_\sigma$ for some universal constant C_σ . Then there exist two universal constants $C_\rho > 0$, $c \in (0, 1)$, such that when $\rho < C_\rho$, we have*

$$\mathbb{E} \|\sin \Theta(U^*, U_{AE})\|_F \geq c\sqrt{r}. \quad (14)$$

The proof is given in Appendix B.7. The condition $d \gg r$ means that there exists a sufficiently small constant $c > 0$ independent of d and r such that $r/d < c$ holds. The additional assumptions $\{\sigma_i^2\}_{i=1}^d$ are different from each other and $\sigma_{(1)}^2/(\sigma_{(r)}^2 - \sigma_{(r+1)}^2) < C_\sigma$ for some universal constant C_σ are made to ensure the identifiability of top- r eigenspace. We need these conditions to guarantee the uniqueness of U_{AE} . As an extreme example, the top- r eigenspace of the identity matrix can be any r -dimensional subspace and thus not unique. To avoid discussing such arbitrariness of the output, we make these assumptions to guarantee the separability of the eigenspace.

Then we investigate the feature recovery ability of the self-supervised contrastive learning approach.

Theorem 14 (Recovery Ability of Contrastive Learning, Upper Bound) *Under the spiked covariance model (5), random masking augmentation in Definition 2, Assumptions 8-10 and $n > d \gg r$, let W_{CL} be any solution that minimizes Equation (3), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have*

$$\mathbb{E} \|\sin \Theta(U^*, U_{CL})\|_F \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}. \quad (15)$$

The proof is given in Appendix B.9. The two terms in equation (15) can be explained as follows: the first term is due to the shift between the distributions of the augmented data and the original data. Specifically, the random masking augmentation generates two views with disjoint nonzero coordinates and thus can mitigate the influence of random noise on

the diagonal entries in the covariance matrix. However, such augmentation slightly hurts the estimation of core features. This bias, appearing as the first term in Equation (15), is measured by the incoherent constant defined in Equation (12). The second term corresponds to the estimation error of the population covariance matrix.

Theorems 13 and 14 characterize the difference in feature recovery ability between autoencoders and contrastive learning. The autoencoders fail to recover most of the core features in the large-noise regime since $\|\sin \Theta(U, U^*)\|_F$ has a trivial upper bound \sqrt{r} . In contrast, with the help of data augmentation, the contrastive learning approach mitigates the corruption of random noise while preserving core features. As n and d increase, it yields a consistent estimator of core features and further leads to better performance in the in-domain downstream tasks, as shown in the next section.

Remark 15 *Here we discuss the potential generalization of our results to the setting with weaker assumptions. Intuitively speaking, the random masking augmentation exploits the prior knowledge that the core features in the original signal are more structural across different coordinates compared with the random noise. Thus the essential requirements are*

1. *Noise is less correlated between different coordinates compared with core features.*
2. *Core features and noise are very different, i.e., $\|\sin \Theta(U^*, U_\Sigma)\|_F$ is large.*

Those two requirements correspond to the diagonal assumption on Σ and incoherent assumption on U^ (Assumption 10). In particular, the latter is to give a lower bound for $\|\sin \Theta(U^*, U_\Sigma)\|_F$ when Σ is diagonal and heteroskedasticity. For more general Σ and U^* , it suffices to assume $\frac{\|\Delta(\Sigma)\|_2}{\nu^2} = o(1)$, and $\|\sin \Theta(U^*, U_\Sigma)\|_F = \Omega(\sqrt{r})$, and we can still draw a similar comparison under these assumptions. Notice that by Lemma 11, when U^* is randomly chosen we will immediately have $\mathbb{E}\|\sin \Theta(U^*, U_\Sigma)\|_F = (1 - o(1))\sqrt{r}$. Similar arguments also apply to all of the later results in this paper and we omit them for simplicity.*

Remark 16 *Similar random masking augmentation (as in Definition 2) can also apply to autoencoders. Although directly applying this augmentation would not work as well since it will not affect the optimal solution (see discussion in Appendix C), an alternative strategy is to reconstruct the whole data from the masked one. This method was originally proposed as denoising autoencoders (DAEs) for general augmentation Vincent et al. (2008), and was proven to be powerful with masking augmentation in a recently proposed representation learning method, masked autoencoders (MAEs) (He et al., 2021). DAEs are a variant of autoencoders that are trained to reconstruct the original image from randomly masked patches. It has been found that DAEs (especially MAEs) outperform other self-supervised methods like MoCo v3, DINO, and BEiT after fine-tuning (He et al., 2021). More specifically, under the same setup described in Section 2, let A be the random masking augmentation defined in Definition 2. We adopt the symmetric linear encoders and decoders. Given samples $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$, we formally define the loss minimization problem¹ of DAEs as*

$$\min_{W \in \mathbb{R}^{r \times d}: WW^\top = 2I_r} \frac{1}{n} \mathbb{E}_A \left[\|W^\top W A X - X\|_F^2 \right]. \quad (16)$$

1. In (He et al., 2021), the loss function is computed on masked coordinates only, but as the authors noted, “This choice is purely result-driven: computing the loss on all pixels leads to a slight decrease in accuracy (e.g., 0.5%).” Hence we will analyze the loss with respect to all coordinates for simplicity.

Notice that the DAEs may not preserve the norm of the input since $\mathbb{E}_A[\|Ax_i\|^2] = (1/2)\|x_i\|^2$. As a result, we optimize the loss under the scaled constraint $WW^\top = 2I_r$.

Then, we claim that under the same conditions as in Theorem 14, DAEs behave similarly to contrastive learning: Let W_{DAE} be any solution that minimizes equation (16), and denote its singular value decomposition as $W_{DAE} = (U_{DAE}\Sigma_{DAE}V_{DAE}^\top)^\top$, then we have (the proof is given in Appendix C.)

$$\mathbb{E} \|\sin \Theta(U^*, U_{DAE})\|_F \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}. \quad (17)$$

From Theorem 13, we know that under high dimensional settings with large sample sizes, DAEs (or masked autoencoders) significantly outperform classic autoencoders. Moreover, compared to Theorem 14, the upper bounds of DAEs are the same as contrastive learning with random masking augmentations. We also provide experimental results on synthetic datasets to verify this result in Appendix C. Although they have similar performance in our linear representation framework because both of them exploit the masking views to eliminate noise, the difference could arise from other aspects such as network architecture and training algorithms, for example, He et al. (2021) used a vision Transformer (Dosovitskiy et al., 2020) while Chen et al. (2020a) used a ResNet (He et al., 2016).

3.4 Performance on In-Domain Downstream Tasks

In the previous section, we have seen that contrastive learning can recover the core feature effectively. In practice, we are interested in using the learned features on in-domain downstream tasks. He et al. (2020) experimentally showed the overwhelming performance of linear classifiers trained on representations learned with contrastive learning against several supervised learning methods in those in-domain downstream tasks.

Following the recent success, here we evaluate the in-domain downstream performance of simple predictors, which take a linear transformation of the representation as an input. Let W_{CL} and W_{AE} be the learned representations based on train data $X \in \mathbb{R}^{n \times d}$. We observe a new signal $\tilde{x} = U^* \tilde{z} + \tilde{\xi}$ independent of X following the spiked covariance model (5). For simplicity, assume \tilde{z} follows $N(0, \nu^2 I_r)$ and is independent of $\tilde{\xi}$. We consider two major types of in-domain downstream tasks: classification and regression. For the binary classification task, we observe a new supervised sample \tilde{y} following the binary response model:

$$\tilde{y} | \tilde{z} \sim \text{Ber}(F(\langle \tilde{z}, w^* \rangle / \nu)), \quad (18)$$

where $F: \mathbb{R} \rightarrow [0, 1]$ is a known monotone increasing function satisfying $1 - F(u) = F(-u)$ for any $u \in \mathbb{R}$, and $w^* \in \mathbb{R}^r$ is a unit vector of coefficients. Notice that our model (18) includes a logistic model (when $F(u) = 1/(1 + e^{-u})$) and probit models (when $F(u) = \Phi(u)$, where Φ is the cumulative distribution function of the standard normal distribution.) We can also interpret model (18) as a shallow neural network model with width r for binary classification. For the regression task, we observe a new supervised sample \tilde{y} following the linear regression model:

$$\tilde{y} = \langle \tilde{z}, w^* \rangle / \nu + \tilde{c}, \quad (19)$$

where $\check{\epsilon} \sim (0, \sigma_\epsilon^2)$ is independent of \check{z} , and $w^* \in \mathbb{R}^r$ is a unit vector of coefficients as before. We can interpret this model as a principal component regression model (PCR) (Jolliffe, 1982) under standard error-in-variables settings², where we assume that the coefficients lie in a low-dimensional subspace spanned by column vectors of U^* . We either estimate or predict the signal based on the observed samples contaminated by the measurement error $\check{\xi}$. For details of PCR in error-in-variables settings, see, for example, Čevič et al. (2020); Agarwal et al. (2020); Bing et al. (2021).

In classification setting, we specify 0-1 loss, that is, $\ell_c(\delta) \triangleq \mathbb{I}\{\check{y} \neq \delta(\check{x})\}$ for some predictor δ taking values in $\{0, 1\}$. For regression task, we employ the squared error loss $\ell_r(\delta) \triangleq (\check{y} - \delta(\check{x}))^2$. Based on some learned representation W , we consider a class of linear predictors. Namely, $\delta_{W,w}(\check{x}) \triangleq \mathbb{I}\{F(w^\top W \check{x}) \geq 1/2\}$ for classification task and $\delta_{W,w}(\check{x}) \triangleq w^\top W \check{x}$ for regression task, where $w \in \mathbb{R}^r$ is a weight vector $w \in \mathbb{R}^r$. Note that the learned representation depends only on unsupervised samples X . Let $\mathbb{E}_{\mathcal{D}}[\cdot]$ and $\mathbb{E}_{\mathcal{E}}[\cdot]$ the expectations with respect to (X, Z) and $(\check{y}, \check{x}, \check{z})$, respectively.

Our goal as stated above is to bound the prediction risk of predictors $\{\delta_{W,w} : w \in \mathbb{R}^r\}$ constructed upon the learned representations W_{CL} and W_{AE} , that is, the quantity $\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell(\delta_{W_{\text{CL}},w})]$ and $\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell(\delta_{W_{\text{AE}},w})]$.

Now we state our results on the performance of the in-domain downstream prediction task.

Theorem 17 (Excess Risk for In-Domain Downstream Task: Upper Bound) *Suppose the conditions in Theorem 14 hold. Then, for the regression task, we have*

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{W_{\text{CL}},w})]] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U^{*\top},w})] \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}},$$

and for the classification task,

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{W_{\text{CL}},w})]] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U^{*\top},w})] = O\left(\frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}\right) \wedge 1.$$

The proofs are given in Appendix B.11.

This result shows that the price of estimating U^* by contrastive learning on an in-domain downstream prediction task can be made small in the case where the core feature lies in a relatively low-dimensional subspace, and the number of samples is relatively large compared to the ostensible dimension of data.

However, the in-domain downstream performance of autoencoders is not as good as contrastive learning. We obtain the following lower bound for the in-domain downstream prediction risk with the autoencoders.

Theorem 18 (Excess Risk for In-Domain Downstream Task: Lower Bound) *Suppose the conditions in Theorem 13 hold. Assume $r \leq r_c$ holds for some constant $r_c > 0$. Additionally assume that $\rho = \Theta(1)$ is sufficiently small and $n \gg d \gg r$. Then, For the regression*

2. In error-in-variables settings, the bias term of the measurement error appears in prediction and estimation risk. Since our focus lies in proving a better performance of contrastive learning against autoencoders, we ignore the unavoidable bias term here by considering the excess risk.

task,

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U_{AE}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U^*, w})] \geq c'_c,$$

and for classification task, if F is differentiable at 0 and $F'(0) > 0$, then

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U_{AE}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U^*, w})] \geq c'_r,$$

where $c'_r > 0$ and $c'_c > 0$ are constants independent of n and d .

The proof is given in Appendix B.11. The condition $n \gg d \gg r$ means that there exists a sufficiently small constant $c > 0$ independent of n , d and r such that $d/n \vee r/d < c$ holds.

The constant c' appearing in Theorem 18 is a constant term independent of d and n . Thus, when d is sufficiently large compared to r and d/n is small, the upper bound of in-domain downstream task performance via contrastive learning in Theorem 17 is smaller than the lower bound of in-domain downstream task performance via autoencoders. The assumption of $r \leq r_c$ in Theorem 18 is assumed for clarity of presentation. Using the same techniques in the proof of Theorem 18, one can obtain a constant lower bound for autoencoders with slightly stronger assumptions, for example, $\rho^2 = O(1/\log d)$ with $n \gg dr$, without assuming $r \leq r_c$. Our theory can be adapted to both of these assumptions.

4. The Impact of Labeled Data in Supervised Contrastive Learning

Recent works have explored adding label information to improve contrastive learning (Khosla et al., 2020). Empirical results show that label information can significantly improve the accuracy of the in-domain downstream tasks. However, when domain shift is considered, the label information hardly improves and even hurts transferability (Islam et al., 2021). For example, in Table 2 of Khosla et al. (2020) and the first column in Table 4 of Islam et al. (2021), supervised contrastive learning shows significant improvement with 7%-8% accuracy increase on in-domain downstream classification on ImageNet and Mini-ImageNet. On the contrary, in Table 4 of Khosla et al. (2020) and Table 4 of Islam et al. (2021), supervised contrastive learning hardly increases the predictive accuracy compared to the self-supervised contrastive learning (the difference of mean accuracy is less than 1%) and can harm significantly on some datasets (e.g. 5.5% lower for SUN 397 in Table 4 of Khosla et al. (2020)). These results indicate that some mechanisms in supervised contrastive learning hurt model transferability while the improvement in source tasks is significant. Moreover, in Table 4 of Islam et al. (2021), it is observed that combining supervised learning and self-supervised contrastive learning together achieves the best transfer learning performance compared to each of them individually. Motivated by those empirical observations, in this section, we aim to investigate the impact of labeled data in contrastive learning and provide a theoretical foundation for these phenomena.

4.1 Feature Mining in Multi-Class Classification

We first demonstrate the impact of labels in contrastive learning under the standard single-sourced (i.e. no transfer learning) setting. Suppose our samples are drawn from $r + 1$

different classes with probability p_k for class $k \in [r + 1]$, and $\sum_{k=1}^{r+1} p_k = 1$. For each class, samples are generated from a class-specific Gaussian distribution:

$$x^k = \mu^k + \xi^k, \quad \xi^k \sim \mathcal{N}(0, \Sigma^k), \quad \forall k = 1, 2, \dots, r + 1. \quad (20)$$

We assume the norms of $\mu^k, \forall k \in [r + 1]$ are in the same order, that is, denote $\nu = \|\mu^1\|/\sqrt{r}$, we have $\|\mu^k\| = O(\sqrt{r}\nu), \forall k \in [r + 1]$. We further assume $\Sigma^k = \text{diag}(\sigma_{1,k}^2, \dots, \sigma_{d,k}^2)$, denote $\sigma_{(1)}^2 = \max_{1 \leq i \leq d, 1 \leq j \leq r+1} \sigma_{i,j}^2$ and assume $\sum_{k=1}^{r+1} p_k \mu^k = 0$, where the last assumption is added to ensure identifiability since the classification problem (20) is invariant under translation. Denote $\Lambda = \sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top}$, we assume $\text{rank}(\Lambda) = r$ and $C_1 \nu^2 < \lambda_{(r)}(\Lambda) < \lambda_{(1)}(\Lambda) < C_2 \nu^2$ for two universal constants C_1 and C_2 . We remark that this model is a labeled version of the spiked covariance model (5) since the core features and random noise are both sub-Gaussian. We use $r + 1$ classes to ensure that μ^k 's span an r -dimensional space, and denote its orthonormal basis as U^* . Recall that our target is to recover U^* .

Remark 19 *Here we focus on explaining the impact of label information in the SupCon algorithm (Khosla et al., 2020). SupCon is designed for multi-class classification tasks and it requires using the class label to find positive samples. In that case, labels from a linear function of the latent features can not be used as class labels in a multi-class classification setting. Hence we proposed to use the Gaussian Mixture Model (20) to generate class labels while keeping the most consistency with models used earlier (i.e., the spiked covariance model (5)).*

As introduced in Definition 4, the supervised contrastive learning introduced by Khosla et al. (2020) allows us to generate contrastive pairs using labeled information and discriminate instances across classes. When we have both labeled data and unlabeled data, we can perform contrastive learning based on pairs that are generated separately for the two types of data.

Data Generating Process Formally, let us consider the case where we draw n samples as unlabeled data $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$ from the Gaussian mixture model (20) with $p_1 = p_2 = \dots = p_{r+1}$. For the labeled data, we draw $(r + 1)m$ samples; m samples for each of the $r + 1$ classes in the Gaussian mixture model, and denote them as $\hat{X} = [\hat{x}_1, \dots, \hat{x}_{(r+1)m}] \in \mathbb{R}^{d \times (r+1)m}$. We discuss the above case for simplicity. More general versions that allow different sample sizes for each class are considered in Theorem D.2 (in the appendix). We study the following hybrid loss to illustrate how the label information helps promote performance over self-supervised contrastive learning:

$$\min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}(W) := \min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}_{\text{SelfCon}}(W) + \alpha \mathcal{L}_{\text{SupCon}}(W), \quad (21)$$

where $\alpha > 0$ is the ratio between supervised loss and self-supervised contrastive loss. Here we consider this generalized hybrid loss to show the benefit of exploiting additional unlabeled data. If we choose $\alpha \rightarrow \infty$ it will correspond to the original SupCon loss.

We first provide a high-level explanation of why label information can help learn core features. When the label information is unavailable, no matter how much (unlabeled) data we have, we can only take them (and their augmented views) as positive samples. In such a scenario, performing augmentation leads to an unavoidable trade-off between estimation

bias and accuracy. However, if we have additional class information, we can contrast between data in the same class to extract more beneficial features that help distinguish a particular class from others and therefore reduce the bias.

Theorem 20 *Suppose the labeled and unlabeled samples are generated as the process mentioned above. If Assumptions 8-10 hold, $n > d \gg r$ and let W_{CL} be any solution that minimizes the supervised contrastive learning problem in Equation (21), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have*

$$\mathbb{E}\|\sin \Theta(U_{CL}, U^*)\|_F \lesssim \frac{1}{1+\alpha} \left(\frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}} \right) + \frac{\alpha}{1+\alpha} \sqrt{\frac{dr}{m}}.$$

The proof is given in Appendix D.2

Corollary 21 *From Theorem 20, it directly follows that when we have m labeled data for each class and no unlabeled data ($\alpha \rightarrow \infty$),*

$$\mathbb{E}\|\sin \Theta(U_{CL}, U)\|_F \lesssim \sqrt{\frac{dr}{m}}.$$

The first bound in Theorem 20 demonstrates how the effect of labeled data changes with the ratio α in the hybrid loss in Equation (21). In addition, compared with Theorem 14, when we only have labeled data ($\alpha \rightarrow \infty$), the second bound in Theorem 20 indicates that with labeled data being available, the supervised contrastive learning can yield consistent estimation as $m \rightarrow \infty$ while the self-supervised contrastive learning consists of an irreducible bias term $O(r^{3/2} \log d/d)$. At a high level, label information can help gain accuracy by creating more positive samples for a single anchor and therefore extract more decisive features. One should notice a caveat that when labeled data is extremely rare compared to unlabeled data, the estimation of supervised contrastive learning suffers from high variance. In comparison, self-supervised contrastive learning, which can exploit a much larger number of samples, may outperform it.

4.2 Information Filtering in Multi-Task Transfer Learning

In this section, we show that the theoretical tools developed in this paper can be used to illustrate the role of label information when using contrastive learning in the transfer learning setting. Label information can tell us the beneficial information for the in-domain downstream task, and learning with labeled data will filter out useless information and preserve the decisive parts of core features. However, in transfer learning, the label information is sometimes rather found to hurt the performance of contrastive learning. For example, in Table 4 of Islam et al. (2021), while supervised contrastive learning gains 8% improvement in source tasks by incorporating label information, it improves only 1% on generalizing to new datasets on average and can even hurt on some datasets. Such observation implies that label information in contrastive learning has very different roles for generalization on source tasks and new tasks. In this section, we consider two regimes of transfer learning – tasks are insufficient/abundant. In both regimes, we provide theories to support the empirical observations and further demonstrate how to wisely combine supervised and self-supervised

contrastive learning to avoid those harms and achieve better performance. Specifically, we consider a transfer learning problem with **regression setting and binary classification setting**. Suppose we have T source tasks which share a common data generative model (5). For the t -th task, the labels are generated by $y^t = \langle w_t, z \rangle / \nu$ in a regression setting while $y^t = \text{sign}(\langle w_t, z \rangle)$ in a binary classification setting, where $w_t \in \mathbb{R}^r$ is a unit vector varying across tasks. These two settings share the same distribution for x only differ in the way to generate the labels y^t .

To incorporate label information, we maximize the Hilbert-Schmidt Independence Criteria (HSIC) (Gretton et al., 2005; Barshan et al., 2011), which has been widely used in literature (Song et al., 2007a,b,c; Barshan et al., 2011).

4.2.1 HILBERT-SCHMIDT INDEPENDENT CRITERIA

Gretton et al. (2005) proposed the Hilbert Schmidt Independent Criteria (HSIC) to measure the dependence between two random variables. It computes the Hilbert-Schmidt norm of the cross-covariance operator associated with their Reproducing Kernel Hilbert Space (RKHS). Such measurement has been widely used as a supervised loss function in feature selection (Song et al., 2007c), feature extraction (Song et al., 2007a), clustering (Song et al., 2007b) and supervised PCA (Barshan et al., 2011).

The basic idea behind HSIC is that two random variables are independent if and only if any bounded continuous functions of the two random variables are uncorrelated. Let \mathcal{F} be a separable RKHS containing all continuous bounded real-valued functions mapping from \mathcal{X} to \mathbb{R} and \mathcal{G} be that for maps from \mathcal{Y} to \mathbb{R} . For each point $x \in \mathcal{X}$, there exists a corresponding element $\phi \in \mathcal{F}$ such that $\langle \phi(x), \phi(x') \rangle_{\mathcal{F}} = k(x, x')$, where $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a unique positive definite kernel. Likewise, define the kernel $l(\cdot, \cdot)$ and feature map ψ for \mathcal{G} . The empirical HSIC is defined as follows.

Definition 22 (Empirical HSIC (Gretton et al., 2005)) *Let $Z := \{(x_1, y_1), \dots, (x_m, y_m)\} \subseteq \mathcal{X} \times \mathcal{Y}$ be a series of m independent and identically distributed observations. An estimator of HSIC, written as $\text{HSIC}(Z, \mathcal{F}, \mathcal{G})$, is given by*

$$\text{HSIC}(Z, \mathcal{F}, \mathcal{G}) := (m - 1)^{-2} \text{tr}(KHLH),$$

where $H, K, L \in \mathbb{R}^{m \times m}$, $K_{ij} := k(x_i, x_j)$, $L_{ij} := l(y_i, y_j)$ and $H := I_m - (1/m)\mathbf{1}_m\mathbf{1}_m^\top$.

In our setting, we aim to maximize the dependency between learned features $WX \in \mathbb{R}^{r \times n}$ and label $y \in \mathbb{R}^n$ via HSIC. Substituting $K \leftarrow X^\top W^\top WX$ and $L \leftarrow yy^\top$, we obtain our supervised loss for the representation matrix W :

$$\text{HSIC}(X, y; W) = \frac{1}{(n - 1)^2} \text{tr}\left(X^\top W^\top WXHy y^\top H\right). \quad (22)$$

A more commonly used supervised loss in the regression task is the mean squared error. Here we explain the equivalence of maximizing HSIC with penalty $\|WW^\top\|_F^2$ and minimizing the mean squared error in the regression task.

Recall that in the contrastive learning framework, we first learn the representation via a linear transformation and then perform linear regression to learn a predictor with the learned representation. Consider the mean squared error $\mathcal{L}(\delta) = (1/n) \sum_{i=1}^n (\delta(x_i) - y_i)^2$, where

$\delta(x_i)$ is a predictor of y_i . Also consider a linear class of predictors $\delta_{W,w}(x_i) = w^\top W x_i$ with parameter $w \in \mathbb{R}^r$. Assume that both X and y are centered. For any fixed representation W , the minimum mean squared error is given by

$$\begin{aligned} \min_{w \in \mathbb{R}^r} \mathcal{L}(\delta_{W,w}) &= \frac{1}{n} \|(WX)^\top w^* - y\|^2 \\ &= \frac{1}{n} \left(y^\top y - \text{tr} \left(X^\top W^\top (W X X^\top W^\top)^{-1} W X y y^\top \right) \right), \end{aligned}$$

where $w^* = (W X X^\top W^\top)^{-1} W X y$. Ignoring the constant term $y^\top y$, it can be seen that the only essential difference between the minimization problem $\min_{w \in \mathbb{R}^r} \mathcal{L}(\delta_{W,w})$ and maximizing HSIC in Equation (22) is the normalization term $(W X X^\top W^\top)^{-1/2}$ for W . Thus, minimizing the mean squared error is equivalent to maximizing HSIC with regularization term $\|W W^\top\|_F^2$. Since L_{SelfCon} contains the regularization term $\|W W^\top\|_F^2$, we can jointly use L_{SelfCon} and HSIC as a surrogate for the standard regression error to avoid the singularity of learned representation W .

4.2.2 MAIN RESULTS

First, we consider the regression setting. Before stating our results, we prepare some notations. Suppose we have n unlabeled data $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$ and m labeled data for each source task $\hat{X}^t = [\hat{x}_1^t, \dots, \hat{x}_m^t]$, $y^t = [y_1^t, \dots, y_m^t]$, $\forall t = 1, \dots, T$ where x_i and \hat{x}_j^t are independently drawn from the spiked covariance model (5), we learn the linear representation via the joint optimization:

$$\min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}(W) := \min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}_{\text{SelfCon}}(W) - \alpha \sum_{t=1}^T \text{HSIC}(\hat{X}^t, y^t; W), \quad (23)$$

where $\alpha > 0$ is a pre-specified ratio between the self-supervised contrastive loss and HSIC. A more general setting, where the ratio α and the number of labeled data for each source task are allowed to depend on t , is considered in the appendix, see Section D.2 for details. We now present a theorem showing the recoverability of W by minimizing the hybrid loss function (23).

Theorem 23 *In the regression setting where $y^t = \langle w_t, z \rangle / \nu$, suppose Assumptions 8-10 hold for the spiked covariance model (5) and $n > d \gg r$, if we further assume that $\alpha > C$ for some constant C , $T < r$ and w_t 's are orthogonal to each other, and let W_{CL} be any solution that optimizes the problem in Equation (23), and denote its singular value decomposition as $W_{CL} = (U_{CL} \Sigma_{CL} V_{CL}^\top)^\top$, then we have:*

$$\mathbb{E} \|\sin \Theta(U_{CL}, U^*)\|_F \lesssim \sqrt{r-T} \left(\frac{r \log d}{d} + \sqrt{\frac{d}{n}} + \alpha T \sqrt{\frac{d}{m}} \wedge 1 \right) + \sqrt{T} \left(\frac{r \log d}{\alpha d} + \frac{1}{\alpha} \sqrt{\frac{d}{n}} + T \sqrt{\frac{d}{m}} \right). \quad (24)$$

The proof is given in Appendix D.6.

Similar to Section 3.4, we can obtain an in-domain downstream task risk in a supervised contrastive learning setting. Consider a new test task where a label is generated by $\check{y} =$

$\langle w^*, \tilde{z} \rangle / \nu$ with $\tilde{x} = U^* \tilde{z} + \tilde{\xi}$. Recall that the loss in the in-domain downstream task is measured by the squared error: $\ell_r(\delta) := (\tilde{y} - \delta(\tilde{x}))^2$. We obtain the following result.

Theorem 24 *Suppose the conditions in Theorem 23 hold. Then,*

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{W_{CL}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{U^* \tau, w})] \right] \\ & \lesssim \sqrt{r-T} \left(\frac{r \log d}{d} + \sqrt{\frac{d}{n}} + \alpha T \sqrt{\frac{d}{m}} \wedge 1 \right) + \sqrt{T} \left(\frac{r \log d}{\alpha d} + \frac{1}{\alpha} \sqrt{\frac{d}{n}} + T \sqrt{\frac{d}{m}} \right). \end{aligned} \quad (25)$$

The proof is given in Appendix D.7.

In Theorem 23 and Theorem 24, as α goes to infinity (corresponding to the case where we only use the supervised loss), the upper bounds in Equations (24) and (25) are reduced to $\sqrt{r-T} + T^{3/2} \sqrt{d/m}$, which is worse than the $r^{3/2} \log d/d$ rate obtained by self-supervised contrastive learning (Theorem 14). This implies that when the model focuses mainly on the supervised loss, the algorithm will extract the information only beneficial for the source tasks and fail to estimate other parts of core features. As a result, when the target task has a very different distribution, labeled data will bring extra bias and therefore hurt the transferability. Additionally, one can minimize the right-hand side of Equation (24) to obtain a sharper rate. Specifically, we can choose an appropriate α such that the upper bound becomes $\sqrt{r^2(r-T)} \log d/d$ (when $n, m \rightarrow \infty$), obtaining a smaller rate than that of the self-supervised contrastive learning. These facts provide theoretical foundations for the recent empirical observations that smartly combining supervised and self-supervised contrastive learning achieves significant improvement in transferability compared with performing each of them individually (Islam et al., 2021).

Remark 25 *A heuristic intuition of this surprising fact is that when tasks are not diverse enough, supervised training will only focus on the features that are helpful to predict the labels of source tasks and ignore other features. For example, we have unlabeled images which contain cats or dogs and the background can be sandland or forest. If the source task focuses on classifying the background, supervised learning will not learn features associated with cats and dogs, while self-supervised learning can learn these features since they are helpful to discriminate different images. As a result, although supervised learning can help to classify sandland and forest, it can hurt performance on the classification of dogs and cats and we should incorporate self-supervised contrastive learning to learn these features.*

When the tasks are abundant enough then estimation via labeled data can recover core features completely. Similar to Theorem 23 and Theorem 24, we have the following results.

Theorem 26 *In the regression setting where $y^t = \langle w_t, z \rangle / \nu$, suppose Assumptions 8-10 hold for the spiked covariance model (5) and $n > d \gg r$, if we further assume that $T > r$ and $\lambda_{(r)}(\sum_{i=1}^T w_i w_i^\top) > c$ for some constant $c > 0$, suppose W_{CL} is the optimal solution of optimization problem Equation (23), and denote its singular value decomposition as $W_{CL} = (U_{CL} \Sigma_{CL} V_{CL}^\top)^\top$, then we have:*

$$\mathbb{E} \|\sin \Theta(U_{CL}, U^*)\|_F \lesssim \frac{\sqrt{r}}{\alpha + 1} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + T \sqrt{\frac{dr}{m}}. \quad (26)$$

The proof is given in Appendix D.8.

Theorem 27 *Suppose the conditions in Theorem 26 hold. Then,*

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{W_{CL}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U^{\star\tau}, w})] \lesssim \frac{\sqrt{r}}{\alpha + 1} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + T \sqrt{\frac{dr}{m}}. \quad (27)$$

The proof is given in Appendix D.9.

Theorem 26 and Theorem 27 show that in the case where tasks are abundant, as α goes to infinity (corresponding to the case where we use the supervised loss only), the upper bounds in Equations (26) and (27) are reduced to $T\sqrt{rd/m}$. This rate can be worse than the $\sqrt{r^3} \log d/d + \sqrt{rd/n}$ rate obtained by self-supervised contrastive learning when m is small. Recall that when the number of tasks is small, labeled data introduce extra bias term $\sqrt{r-T}$ (Theorem 23 and Theorem 24). We note that when the tasks are abundant enough, the harm of labeled data is mainly due to the variance brought by the labeled data. When m is sufficiently large, supervised learning on source tasks can yield a consistent estimation of core features, whereas self-supervised contrastive learning can not.

In extending our results from regression to the binary classification setting, the only difference is in the label generation process, and we can obtain similar results with some modification of the proofs. The corresponding feature recovery bounds of Theorem 23 (where the tasks are insufficient) and Theorem 26 (where the tasks are abundant) are stated as follows:

Theorem 28 *In the classification setting where $y^t = \text{sign}(\langle w_t, z \rangle / \nu)$, suppose the conditions in Theorem 23 hold and z in the spiked covariance model (5) follows a Gaussian distribution, then we have:*

$$\mathbb{E} \|\sin \Theta(U_{CL}, U^*)\|_F \lesssim \sqrt{r-T} \left(\frac{r \log d}{d} + \sqrt{\frac{d}{n}} + \alpha T \sqrt{\frac{d}{m}} \wedge 1 \right) + \sqrt{T} \left(\frac{r \log d}{\alpha d} + \frac{1}{\alpha} \sqrt{\frac{d}{n}} + T \sqrt{\frac{d}{m}} \right). \quad (28)$$

Theorem 29 *In the classification setting where $y^t = \text{sign}(\langle w_t, z \rangle / \nu)$, suppose the conditions in Theorem 26 hold and z in the spiked covariance model (5) follows a Gaussian distribution, then we have:*

$$\mathbb{E} \|\sin \Theta(U_{CL}, U^*)\|_F \lesssim \frac{\sqrt{r}}{\alpha + 1} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + T \sqrt{\frac{dr}{m}}. \quad (29)$$

As the feature recovery bounds remain to be the same, the counterpart of the in-domain downstream tasks results, Theorem 24 and Theorem 27, in this classification setting follows immediately. For the sake of space, we defer the generalized version and proofs in Theorem D.10 and Theorem D.11 in the appendix.

5. Numerical Experiments

5.1 Linear Model with Synthetic Data

To verify our theory, we conducted numerical experiments on the spiked covariance model (5) under a linear representation setting and contrastive loss functions defined in (3) and

$\log_e(\alpha)$	-5	-4	-3	-2	-1	0	1	2	3	4	5
$T = 8, r = 10$	0.0242	0.0231	0.0199	0.0141	0.0122	0.0125	0.0184	0.0345	0.0499	0.0535	0.0587
$T = 20, r = 10$	0.0223	0.0163	0.0156	0.0096	0.0079	0.0055	0.0064	0.0064	0.0067	0.0070	0.0079

Table 2: In-domain downstream performance in transfer learning against the penalty parameter α . T is the number of source tasks.

$\log_e(\alpha)$	-5	-4	-3	-2	-1	0	1	2	3	4	5
$T = 8, r = 10$	2.0373	2.0371	2.0228	1.9908	2.0021	2.0055	2.0010	2.0362	2.0699	2.0705	2.0813
$T = 20, r = 10$	2.0352	2.0292	2.0030	1.9871	1.9740	1.9690	1.9766	1.9702	1.9790	1.9714	1.9672

Table 3: Feature recovery performance in transfer learning against the penalty parameter α . T is the number of source tasks.

(23). As we have explicitly formulated the loss function and derived its equivalent form in the main body and appendix, we simply minimize the corresponding loss by gradient descent to find the optimal linear representation W . For self-supervised contrastive learning with random masking augmentation, we independently draw the augmentation function by Definition 2 and apply them to all samples in each iteration. To ensure convergence, we set the maximum number of iterations for it (typically 10000 or 50000 depending on dimension d).

We report two criteria to evaluate the quality of the representation, in-domain downstream error, and sine distance. To obtain the sine distance for a learned representation W , we perform singular value decomposition to get $W = (U\Sigma V^\top)^\top$ and then compute $\|\sin \Theta(U, U^*)\|_F$. To obtain the in-domain downstream task performance, in the comparison between autoencoders and contrastive learning, we first draw n labeled data from spiked covariance model (5) with labels generated as in Section 3.4, then we train the model by using the data without labels to obtain the linear representation W , and learn a linear predictor w using the data with labels and compute the regression error. In the transfer learning setting, we draw some labeled data from the source tasks and additional unlabeled data. The number of labeled data is set to be $m = 1000$ and the number of unlabeled data is set to be $n = 1000$. Then train with them to obtain the linear representation W , and draw labeled data from a new source task to learn a linear predictor w to compute the regression error. In particular, we subtract the optimal regression error obtained by the best representation $U^{*\top}$ for each regression error and report the difference, or more precisely, the excess risk as in-domain downstream performance.

The results are reported in Fig. 1 and 2 and Table 2 and 3. As predicted by Theorems 13 and 14, the feature recovery error and in-domain downstream task risk of contrastive learning decrease as d increases (Fig. 1: **Left**) and as n increases (Fig. 1: **Center**) while that of autoencoders is insensible to the changes in d and n . Consistent with our theory, in Fig. 1: **Right**, it is observed that when tasks are not abundant, the transfer performance exhibit a U -shaped curve, and the best result is achieved by choosing an appropriate α . When tasks are abundant and labeled data are sufficient, the error remains small when we take large α .

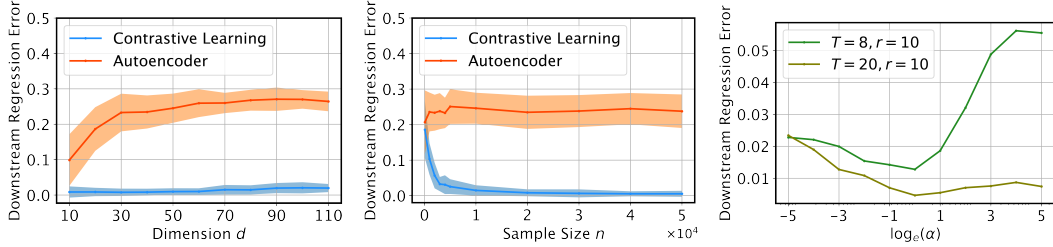


Figure 1: The vertical axes indicate the downstream regression error. We subtract the regression error of the ground truth features to measure the excess error. **Left:** Comparison of in-domain downstream task performance between contrastive learning and autoencoders the dimension d . The sample size n is set as $n = 20000$. **Center:** Comparison of in-domain downstream task performance between contrastive learning and autoencoders the dimension n . The dimension d is set as $d = 40$. **Right:** In-domain downstream task performance in transfer learning against penalty parameter α in log scale. T is the number of source tasks and r is the dimension of the representation function. We set the number of labeled data and unlabeled data as $m = 1000$ and $n = 1000$ respectively.

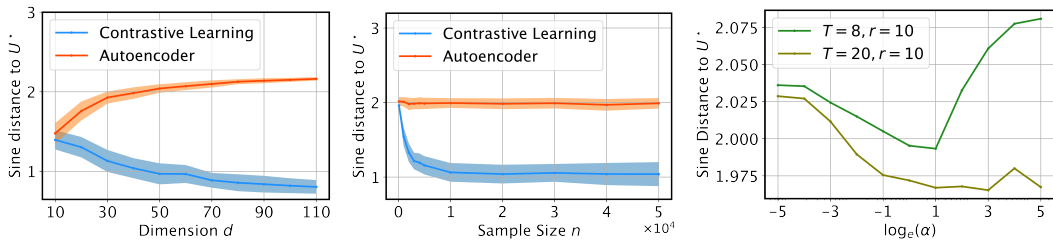


Figure 2: **Left:** Comparison of learned feature between contrastive learning and autoencoders against the dimension d . The sample size n is set as $n = 20000$. **Center:** Comparison of feature recovery performance between contrastive learning and autoencoders against the dimension n . The dimension d is set as $d = 40$. **Right:** Feature recovery performance in transfer learning against penalty parameter α in log scale. T is the number of source tasks. We set the number of labeled data and unlabeled data as $m = 1000$ and $n = 1000$ respectively.

5.2 Neural Nets with Real-World Dataset

In this section, we provide experimental results in real-world datasets to support our theoretical results. Although our model settings and assumptions might be violated under this scenario, as we shall see, our findings still remain valid in practice.

We conduct the experiments using the datasets STL-10 (Coates et al., 2011) and CIFAR-10 (Krizhevsky, 2009) with the neural nets architecture ResNet-18 (He et al., 2016). Our experiments are carried out based on linear evaluation following SimCLR (Chen et al., 2020a), where we first train a ResNet-18 encoder and a two-layer MLP projector with unlabeled augmented data. We then freeze the encoder, train a logistic regression on top of it with labeled data, and lastly evaluate the performance on the test data. Following Chen et al. (2020a), we apply augmentations including resized cropping, horizontal flipping, color distortion, and Gaussian blurring to generate the augmented data, and use the InfoNCE loss function to train the network. All training is carried out with the Adam optimizer (Kingma and Ba, 2015), batch size 256, learning rate 3×10^{-4} , weight decay 10^{-4} , and a cosine annealing learning rate scheduler for 100 epochs. Our codes are implemented in Pytorch and run on an NVIDIA V100 GPU.

Contrastive learning v.s. standard autoencoders: Here we provide real-world evidence for our theoretical findings in Section 3.3 by comparing the performance of contrastive learning versus a standard autoencoder. The architecture of encoders is the same for these two methods, and we use an inversed ResNet-18 as the decoder. During the training time, we use the encoder-decoder architecture and mean squared error loss to reconstruct the input, and then we train a linear classifier on the features learned by the encoder. The results are listed in Table 4, we can find that contrastive learning demonstrates superior performance over the standard autoencoder.

Testing Accuracy	Contrastive Learning	Standard Autoencoder
CIFAR10	65.11 ± 0.51	44.76 ± 0.16
STL10	71.02 ± 0.47	39.00 ± 0.58

Table 4: Comparison of linear evaluation performance of contrastive learning and autoencoders.

The impact of labeled data in transfer learning Now we provide real-world evidence for our theoretical findings in Section 4.2. Following the joint optimization formulation in equation 23, we combine the InfoNCE loss function for unlabeled data and cross-entropy loss for labeled data with a ratio α . For both STL-10 and CIFAR-10 datasets, we divide the test data into two sets, one consists of the first five classes and the other one consists of the remaining five classes. During training, we use the training data as unlabeled data and the first set of test data as the labeled data to train the model jointly, and then train a linear classifier with the second set of test data on features learned by the encoder. As predicted by Theorem 23, when α is small, introducing labeled data from the first five classes would be beneficial to learn better representations and improve the performance on the last five classes; when α is large, labeled data from the first five classes will make the model only focus on features that are useful to discriminate the first five classes and ignore

other features, thus introducing labeled data could be harmful to the performance on the last five classes. Testing accuracy on the last five classes with different α are listed in table 5, it is observed that the accuracy first increases and then decreases as α grows, which is consistent with our theoretical results.

Testing Accuracy	$\alpha = 0.0$	$\alpha = 0.1$	$\alpha = 0.2$	$\alpha = 0.3$	$\alpha = 0.4$	$\alpha = 0.5$	$\alpha = 0.6$	$\alpha = 0.7$	$\alpha = 0.8$	$\alpha = 0.9$	$\alpha = 1.0$
CIFAR10	74.27	75.52	74.86	75.31	75.21	74.86	74.46	73.85	72.20	69.17	51.31
STL10	82.56	83.54	83.27	83.24	83.21	83.03	82.34	82.11	80.94	76.88	52.37

Table 5: **Transfer learning performance with different α .** For each experiment, we report the average accuracy for three independent runs.

6. Conclusion

In this work, we establish a theoretical framework to study contrastive learning under the linear representation setting. We theoretically prove that contrastive learning, compared with autoencoders and GANs, can obtain a better low-rank representation under the spiked covariance model, which further leads to better performance in in-domain downstream tasks. We also highlight the impact of labeled data in supervised contrastive learning and multi-task transfer learning: labeled data can reduce the domain shift bias in contrastive learning, but it harms the learned representation in transfer learning. To our knowledge, our result is the first theoretical result to guarantee the success of contrastive learning by comparing it with existing representation learning methods. However, to get a tractable analysis, like many other theoretical works in representation learning (Du et al., 2020; Lee et al., 2021; Tripuraneni et al., 2021), our work starts with linear representations, which still provides important insights. Recently, Wen and Li (2021) and Refinetti and Goldt (2022) studied the training dynamics of autoencoders and contrastive learning with nonlinear shallow neural networks. Extending our results to these more complex models is an interesting direction for future work.

Acknowledgement

L.Z. is supported by National Science Foundation DMS 2015378. J.Z. is supported by the National Science Foundation (CCF 1763191 and CAREER 1942926), the US National Institutes of Health (P30AG059307 and U01MH098953) and grants from the Silicon Valley Foundation and the Chan-Zuckerberg Initiative.

Appendix A. Background and omitted discussion

A.1 Comparison with other works

Here we compare the results in this paper with some closely related works.

To rigorously analyze contrastive learning, we consider the random masking augmentation strategy which is also analyzed in Wen and Li (2021). In Wen and Li (2021), the authors aim to understand the training dynamics of contrastive learning in a shallow non-linear neural network and focus more on dealing with nonlinearity. In comparison, our work focus on the comparison between contrastive learning and autoencoders and the role of label information in contrastive learning. To make the problem mathematically tractable, we adopt a linear model, which is simple but enough to shed light on many mysterious phenomena in practice. Moreover, while they assume a sparse coding model, where the features are extremely sparse, and Gaussianity of signals and noise, our analysis only requires that the features are sub-Gaussian (5). Furthermore, our technique allows the signal-to-noise ratio to have different orders, as long as it decreases slowly, while their analysis is restricted to a particular signal-to-noise ratio. Tian (2022) studied the relationship between contrastive learning and PCA from a game-theoretic point of view. Specifically, the authors decompose the gradient descent on the contrastive loss into two dynamics, namely the max-player and min-player. It is proven that in deep linear networks, the max-player is equivalent to PCA and the landscape has no spurious minimum. While the results on max-player can be applied to a family of contrastive loss, it is still difficult to analyze the min-player in a general setting. In our paper, we use a linear contrastive loss (2) to explicitly obtain the features learned by contrastive learning. Moreover, our results can be directly extended to a deep linear network setting by the equivalence of a single linear transformation and a deep linear network. The major difference is the non-convexity of the loss landscape.

Garg and Liang (2020) studied the combination of supervised learning and self-supervised learning. They viewed training with unlabeled data as functional regularization on learning the representation function, and obtained sample complexity bounds in a PAC-learning style for various settings. In particular, they found that such functional regularization can help to reduce the amount of labeled data needed, and showed autoencoders and masked self-supervision as two concrete examples. Apart from Garg and Liang (2020), this paper focuses on a regime in combining self-supervised learning and supervised learning, where a trade-off between labeled data and unlabeled data exists.

Specifically, Theorem 3 in Garg and Liang (2020) assumes that a ground truth representation exists such that it can keep both self-supervised loss and supervised loss at a very low threshold. However, as the authors admit, it is hard to determine such a threshold in practical settings. For example, since the unlabeled data and labeled data come from different domains, such as Image-Net and CIFAR-10, domain-specific features may have a much lower loss compared with domain-transferable features. In our paper, we first study the regime where tasks are not diverse enough in Theorem 23 (which corresponds to the case where ground truth does not exist) and show the trade-off between supervised loss and self-supervised loss. Then in Theorem 26 we show that when tasks are abundant (which corresponds to the case where ground truth exists), labeled data helps to achieve better error bounds, which is similar to the result of Garg and Liang (2020). Our result of Theorem

23 provides novel insight into the regime where tasks are not diverse, which has been left untouched in the literature.

In Li et al. (2021), the authors proposed a novel self-supervised loss function based on HSIC and discussed the relationship between InfoNCE and the proposed SSL-HSIC loss. The SSL-HSIC loss measures the dependence between the output features and one-hot encoded labels (which serve as the indicators of positive samples) and minimizing the SSL-HSIC loss encourages the network to discriminate augmented views from different samples. In comparison to this self-supervised loss, we use HSIC in Section 4.2 as a supervised loss to measure the dependence between output features and the true labels, which is a common usage of HSIC in previous works (Barshan et al., 2011; Song et al., 2007c). Moreover, we want to point out that the proposed estimator of SSL-HSIC (see Equation (11) in Li et al. (2021)) can be reduced to the linear loss we use in this paper when the kernel $k(\cdot, \cdot)$ is chosen to be a simple inner product. The authors argued that the standard InfoNCE loss may yield meaningless features in some cases thus the proposed HSIC-based loss could be a better alternative, and provided empirical results illustrating the comparable performance of SSL-HSIC. It remains to be further explored the benefits of such an HSIC-based method.

Lee et al. (2021) studied self-supervised learning under a conditional independence assumption, and showed that with the optimal representation learned in pretext tasks, the in-domain downstream risk is guaranteed to be small. In the contrastive learning context, for example, an image classification downstream task, such an assumption implies that the augmented views generated from the same picture are roughly independent conditional on its ground-truth class, which could be too strong since two views are usually strongly correlated. Such an assumption would be closer to supervised contrastive learning as we have discussed in Section 4.1 where we contrast two independent samples from the same class, such sample pairs are independent conditional on the true label, but it requires label information and thus not applied to self-supervised contrastive learning. Compared with this work, we studied a specific data-generating model where the two views are obtained by practical augmentation and thus could be more close to a real-world setting. It is also remarkable that in Lee et al. (2021), their analysis can be adapted to the nonlinear representation setting in the sense that they directly assumed that the optimal representation is obtained. However, the representation learning in the contrastive learning context, even under a linear representation setting, could be non-convex. Thus our analysis starts from a simple setting to obtain a deep understanding of what could be learned in the pretext tasks.

Saunshi et al. (2022) proposed a novel perspective that theoretical analysis of contrastive learning must take the inductive bias into account. It is shown that without considering the function class, it is possible that the learned features totally fail in downstream classification tasks. Furthermore, it is shown that within the linear representation class, contrastive self-supervised learning is guaranteed to learn meaningful features under certain conditions. In our paper, we have restricted ourselves to a similar linear representation setting and avoided the collapse raised by complicated models. Compared with their analysis in a linear representation setting, our bounds provide an exact order while their results still need to quantify the expressivity and inconsistency measure, which is very difficult without very strong assumptions. Moreover, their analysis requires a finite input space and augmentation set, which is much stronger than our data-generating model.

Later than our paper, Fu et al. (2022) considered a similar issue of transfer learning performance of the supervised contrastive method. The authors argue that directly minimizing the supervised contrastive method will lead to feature collapse, which implies that within each class, all data points will have the same embedding and thus the supervised contrastive method loses information that could be useful for transfer learning. This intuition is similar to our setting in Section 4.2, and based on it the authors proposed a new loss function that combines the supervised contrastive loss and within-class self-supervised loss together. At a high level, this new loss function is quite similar to equation 23 since they are both linear interpolations of self-supervised contrastive loss and supervised loss, and the common motivation is to encourage the model to learn more background features. And our analysis in Section 4.2 can provide a theoretical foundation for when could such interpolation work and how to choose the ratio α .

A.2 Disucssion about the regularization term

In this paper we use a quadratic regularization term $R_1(W) = \|WW^\top\|_F^2$ instead of a standard ℓ_2 regularization term $R_2(W) = \|W\|_F^2$. Denote $W^T = [w_1, \dots, w_r]$, then we can write these two terms as:

$$R_1(W) = \sum_i^r \|w_i\|^4 + \sum_{i,j} \langle w_i, w_j \rangle^2, \quad R_2(W) = \sum_i^r \|w_i\|^2$$

The main difference between these two terms is that except for penalizing the norm of representation w_i , it also penalizes the similarity between different representations. In particular, in the linear representation setting, where we will deal with optimization problems like

$$\min_{W \in \mathbb{R}^{r \times d}} \text{tr}(WAW^\top) + \frac{\lambda}{2} R(W)$$

where A is a symmetric matrix determined by data and augmentation, we can easily find that the ℓ_2 regularization would fail. To see this, we can rewrite the loss function as

$$\text{tr}(WAW^\top) + \frac{\lambda}{2} \|W\|_F^2 = \sum_{i=1}^r (w_i^\top A w_i + \frac{\lambda}{2} \|w_i\|^2) = \sum_{i=1}^r w_i^\top (A + \frac{\lambda}{2} I) w_i,$$

it is easy to find that the optimal solution of each w_i would be at infinity. Moreover, even if we add constraints like $\|w_i\| < C, \forall i \in [r]$, the optimal solution of each w_i would all be the eigenvector corresponding to the smallest eigenvalue of A thus the model would only learn a single representation from the data. In contrast, the quadratic regularization term encourages the diversity of representation by penalizing the similarity between representations, i.e., $\langle w_i, w_j \rangle^2$. In this situation, we have:

$$\text{tr}(WAW^\top) + \frac{\lambda}{2} \|WW^\top\|_F^2 = \frac{\lambda}{2} \|W^\top W\|_F + \frac{1}{\lambda} \|A\|_F^2 - \frac{1}{2\lambda} \|A\|_F^2$$

and it is easy to find that the optimal solution of w_i would be finite and each w_i corresponds to different eigenvectors of A and is orthogonal to each other, which implies that they would learn totally different representations. As a result, the quadratic regularization term would

Testing Accuracy	weight decay=0.0001	quadratic regularization=0.0001
CIFAR10	65.11 \pm 0.51	65.54 \pm 0.26
STL10	71.02 \pm 0.47	71.39 \pm 0.39

Table 6: **Quadratic regularization v.s. weight decay.** We compare the top-1 accuracy of linear classifiers trained on features learned by SimCLR with the ResNet-18 encoder and different regularization methods. We repeat each experiment for 5 runs and report the mean and standard deviation. More details are provided in Section 5.2

be more helpful to self-supervised learning with linear representation. In real-world practice, ℓ_2 regularization can still work well with the help of non-linearity and normalization techniques, but we also provide an empirical observation that applying quadratic regularization would be helpful to improve the performance compared with using standard weight decay. We would also like to point out that in the linear regime, the choice of $\lambda \in \mathbb{R}_+$ would only affect the norm of w_i and makes no difference to the direction of w_i and the quality of representation, thus we do not specify this value in our analysis. Similar regularization techniques are also used in Liu et al. (2021) for theoretical analysis in the linear representation setting.

We verify this conjecture in neural networks, and the results are provided in Table 6. The first column corresponds to training with standard weight decay, and in the setting of the second column, we do not apply weight decay for each weight matrix of fully connected layers in the encoder, and add an additional regularization term $\lambda \|W^T W\|_F^2$ on the loss function instead. Note that the quadratic regularization does not apply to convolution layers and bias terms, thus we keep the weight decay on these parameters. We search the regularization parameter in each settings from $\lambda = 0.1, 0.01, 0.001, 0.0001, 0.00001, 0.000001$ and find that $\lambda = 0.0001$ yields the best performance in each settings and datasets.

It is observed that quadratic regularization slightly improves the performance of contrastive learning, which is consistent with our intuition.

A.3 Background on distance between subspaces

In this section, we will provide some basic properties of sine distance between subspaces. Recall the definition:

$$\|\sin \Theta(U_1, U_2)\|_F \triangleq \left\| U_{1\perp}^\top U_2 \right\|_F = \left\| U_{2\perp}^\top U_1 \right\|_F. \quad (30)$$

where $U_1, U_2 \in \mathbb{O}_{d,r}$ are two orthogonal matrices. Similarly, we can also define:

$$\|\sin \Theta(U_1, U_2)\|_2 \triangleq \left\| U_{1\perp}^\top U_2 \right\|_2 = \left\| U_{2\perp}^\top U_1 \right\|_2.$$

We first give two equivalent definitions of this distance:

Proposition A.1

$$\|\sin \Theta(U_1, U_2)\|_F^2 = r - \left\| U_1^\top U_2 \right\|_F^2$$

Proof Write $U = [U_1, U_{1\perp}] \in \mathbb{O}_{d,d}$. We have

$$r = \|U_2\|_F^2 = \|U^\top U_2\|_F^2 = \left\| U_{1\perp}^\top U_2 \right\|_F^2 + \left\| U_1^\top U_2 \right\|_F^2,$$

then by definition of sine distance, we can obtain the desired equation. \blacksquare

Proposition A.2

$$\|\sin \Theta(U_1, U_2)\|_F^2 = \frac{1}{2} \|U_1 U_1^\top - U_2 U_2^\top\|_F^2$$

Proof Expand the right hand and use Proposition A.1 we have:

$$\begin{aligned} \frac{1}{2} \|U_1 U_1^\top - U_2 U_2^\top\|_F^2 &= \frac{1}{2} (\|U_1 U_1^\top\|_F^2 + \|U_2 U_2^\top\|_F^2 - 2 \operatorname{tr}(U_1 U_1^\top U_2 U_2^\top)) \\ &= \frac{1}{2} (r + r - 2 \operatorname{tr}(U_1^\top U_2 U_2^\top U_1)) \\ &= r - \|U_1^\top U_2\|_F^2 = \|\sin \Theta(U_1, U_2)\|_F^2. \end{aligned}$$

With Propositions A.1 and A.2, it is easy to verify its properties to be a distance function. Obviously, we have $0 \leq \|\sin \Theta(U_1, U_2)\|_F \leq \sqrt{r}$ and $\|\sin \Theta(U_1, U_2)\|_F = \|\sin \Theta(U_2, U_1)\|_F$ by definition. Moreover, we have the following results:

Lemma A.3 (Lemma 1 in Cai and Zhang (2018)) For any $U, V \in \mathbb{O}_{d,r}$,

$$\|\sin \Theta(U, V)\|_2 \leq \inf_{O \in \mathbb{O}_{r,r}} \|UO - V\|_2 \leq \sqrt{2} \|\sin \Theta(U, V)\|_2, \quad (31)$$

and

$$\|\sin \Theta(U, V)\|_F \leq \inf_{O \in \mathbb{O}_{r,r}} \|UO - V\|_F \leq \sqrt{2} \|\sin \Theta(U, V)\|_F. \quad (32)$$

Proposition A.4 (Identity of indiscernibles)

$$\|\sin \Theta(U_1, U_2)\|_F = 0 \Leftrightarrow \exists O \in \mathbb{O}^{r \times r}, \text{ s.t. } U_1 O = U_2$$

Proof It is a straightforward corollary by definition:

$$\begin{aligned} \|\sin \Theta(U_1, U_2)\|_F = 0 &\Leftrightarrow \left\| U_{1\perp}^\top U_2 \right\|_F = 0 \Leftrightarrow U_{2\perp} \perp U_1 \\ &\Leftrightarrow \exists O \in \mathbb{O}^{r \times r}, \text{ s.t. } U_1 O = U_2. \end{aligned}$$

Proposition A.5 (Triangular inequality)

$$\|\sin \Theta(U_1, U_2)\|_F \leq \|\sin \Theta(U_1, U_3)\|_F + \|\sin \Theta(U_2, U_3)\|_F$$

Proof By the triangular inequality for Frobenius norm we have:

$$\|U_1 U_1^\top - U_2 U_2^\top\|_F \leq \|U_1 U_1^\top - U_3 U_3^\top\|_F + \|U_2 U_2^\top - U_3 U_3^\top\|_F,$$

then apply Proposition A.2 to replace the Frobenius norm with sine distance we can finish the proof. \blacksquare

Appendix B. Omitted proofs for Section 3

B.1 Proofs for Section 3.1 and Section 3.2

In this section, we will provide the proof of Proposition 5 and Corollary 6, the restatement of them and the detailed proof can be found in Proposition B.1 and Corollary B.2.

Proposition B.1 (Restatement of Proposition 5) *For two fixed augmentation functions $g_1, g_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$, denote the augmented data matrices as $X_1 = [g_1(x_1), \dots, g_1(x_n)] \in \mathbb{R}^{d \times n}$ and $X_2 = [g_2(x_1), \dots, g_2(x_n)] \in \mathbb{R}^{d \times n}$, when the augmented pairs are generated as in Definition 1, the optimal solution of contrastive learning problem (9) is given by:*

$$W_{CL} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$X_1 X_2^\top + X_2 X_1^\top - \frac{1}{2(n-1)} (X_1 + X_2)(1_r 1_r^\top - I_r)(X_1 + X_2)^\top, \quad (33)$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

Proof [Proof of Proposition 5] When augmented pairs generation in Definition 1 is applied, the contrastive loss can be written as:

$$\begin{aligned} \mathcal{L}_{\text{SelfCon}}(W) &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{n} \sum_{i=1}^n [\langle Wg_1(x_i), Wg_2(x_i) \rangle \\ &\quad - \frac{1}{4(n-1)} \sum_{j \neq i} \langle Wg_1(x_i) + Wg_2(x_i), Wg_1(x_j) + Wg_2(x_j) \rangle] \\ &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{n} \sum_{i=1}^n \langle Wg_1(x_i), Wg_2(x_i) \rangle \\ &\quad + \frac{1}{4n(n-1)} \sum_{i=1}^n \sum_{j \neq i} \langle Wg_1(x_i) + Wg_2(x_i), Wg_1(x_j) + Wg_2(x_j) \rangle \\ &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{2n} \text{tr}(X_1^\top W^\top W X_2 + X_2^\top W^\top W X_1) \\ &\quad + \frac{1}{4n(n-1)} \text{tr}((1_n 1_n^\top - I_n)(X_1 + X_2)^\top W^\top W (X_1 + X_2)) \\ &= \frac{\lambda}{2} \|WW^\top\|_F^2 \\ &\quad - \frac{1}{2n} \text{tr} \left((X_2 X_1^\top + X_1 X_2^\top - \frac{1}{2(n-1)} (X_1 + X_2)(1_n 1_n^\top - I_n)(X_1 + X_2)^\top) W^\top W \right) \\ &= \frac{1}{2} \left\| \lambda W^\top W - \frac{1}{2n\lambda} \left(X_2 X_1^\top + X_1 X_2^\top - \frac{1}{2(n-1)} (X_1 + X_2)(1_n 1_n^\top - I_n)(X_1 + X_2)^\top \right) \right\|_F^2 \\ &\quad - \left\| \frac{1}{2n\lambda} \left(X_2 X_1^\top + X_1 X_2^\top - \frac{1}{2(n-1)} (X_1 + X_2)(1_n 1_n^\top - I_n)(X_1 + X_2)^\top \right) \right\|_F^2. \end{aligned}$$

Note that the last term only depends on X , and the first term implies that when W_{CL} is the optimal solution, $\lambda W_{\text{CL}} W_{\text{CL}}^\top$ is the best rank- r approximation of $\frac{1}{(n-1)\lambda} X H X^\top$, where $H := \mathbf{1}_n \mathbf{1}_n^\top - I_n$. Applying Lemma E.4 to the first term, we can conclude that W_{CL} satisfies the desired conditions. \blacksquare

Corollary B.2 (Restatement of Corollary 6) *Under the same conditions as in Proposition 5, if we use random masking (Definition 2) as our augmentation function, then in expectation over the data augmentation, the optimal solution of contrastive learning problem (9) is given by:*

$$W_{\text{CL}} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$\Delta(X X^\top) - \frac{1}{n-1} X (\mathbf{1}_n \mathbf{1}_n^\top - I_n) X^\top, \quad (34)$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

Proof [Proof of Corollary 6] Following the proof of Proposition 5, now we only need to compute the expectation over the augmentation distribution defined in Definition 2:

$$\begin{aligned} \mathcal{L}_{\text{SelfCon}}(W) &= \frac{\lambda}{2} \|W W^\top\|_F^2 - \mathbb{E}_{(g_1, g_2)} \left[\frac{1}{n} \sum_{i=1}^n \langle W g_1(x_i), W g_2(x_i) \rangle \right. \\ &\quad \left. - \frac{1}{4(n-1)} \sum_{j \neq i} \langle W g_1(x_i) + W g_2(x_i), W g_1(x_j) + W g_2(x_j) \rangle \right] \\ &= \frac{\lambda}{2} \|W W^\top\|_F^2 - \mathbb{E}_{(g_1, g_2)} \left[\frac{1}{2n} \text{tr}((X_2 X_1^\top + X_1 X_2^\top \right. \\ &\quad \left. - \frac{1}{2(n-1)} (X_1 + X_2)(\mathbf{1}_n \mathbf{1}_n^\top - I_n)(X_1 + X_2)^\top) W^\top W \right]. \end{aligned} \quad (35)$$

Note that by the definition of random masking augmentation, we have $X_1 = AX, X_2 = (I - A)X$, which implies $X_1 + X_2 = X$. On the other hand, X_1 and X_2 have no common nonzero entries, hence the matrix $X_1 X_2^\top + X_2 X_1^\top$ only consists of off-diagonal entries and each of the off-diagonal entry denoted as x_{ij} appears if and only if $a_i + a_j = 1$. Moreover, if it appears, we must have x_{ij} equals to the (i, j) -th element of $X X^\top$. With this result, we

can then compute the expectation in Equation (35):

$$\begin{aligned}
 \mathcal{L}_{\text{SelfCon}}(W) &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \mathbb{E}_{(g_1, g_2)} \left[\frac{1}{2n} \text{tr}((X_2 X_1^\top + X_1 X_2^\top) \right. \\
 &\quad \left. - \frac{1}{2(n-1)} (X_1 + X_2)(1_n 1_n^\top - I_n)(X_1 + X_2)^\top) W^\top W \right] \\
 &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{2n} \text{tr} \left(\left(\frac{1}{2} \Delta(X X^\top) - \frac{1}{2(n-1)} X(1_n 1_n^\top - I_n) X^\top \right) W^\top W \right) \\
 &= \frac{1}{2} \left\| \lambda W^\top W - \frac{1}{4n\lambda} \left(\Delta(X X^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \right\|_F^2 \\
 &\quad - \left\| \frac{1}{4n\lambda} \left(\Delta(X X^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \right\|_F^2.
 \end{aligned}$$

By a similar argument as in the proof of Proposition 5, we can conclude that W_{CL} satisfies the desired conditions. \blacksquare

Remark B.3 *Note that the two views generated by random masking augmentation have disjoint non-zero dimensions, hence contrasting such positive pairs yields a correlation between different dimensions only. That is why the first term in equation (34) appears to be $\Delta(X X^\top)$ where the diagonal entries are eliminated.*

B.2 Proofs for Section 3.3

In this section, we will prove Lemma 11, Theorems 13 and 14 in Section 3.3. The restatement and proof of them can be found in Lemma B.6, Theorem B.7, and Theorem B.9.

Before starting the proof, we give two technical lemmas to help the proof.

Lemma B.4 (Uniform distribution on the unit sphere (Marsaglia, 1972)) *If x_1, x_2, \dots, x_n i.i.d. $\sim \mathcal{N}(0, 1)$, then $(x_1/\sqrt{\sum_{i=1}^n x_i^2}, \dots, x_n/\sqrt{\sum_{i=1}^n x_i^2})$ is uniformly distributed on the unit sphere $\mathbb{S}^d = \{(x_1, \dots, x_n) \in \mathbb{R}^n : \sum_{i=1}^n x_i^2 = 1\}$.*

Lemma B.5 *If x_1, x_2, \dots, x_n i.i.d. $\sim \mathcal{N}(0, 1)$, then:*

$$\mathbb{E} \max_{1 \leq i \leq n} x_i^2 \leq 2 \log(n).$$

Proof Denote $Y = \max_{1 \leq i \leq n} x_i^2$, then we have:

$$\exp(t\mathbb{E}Y) \leq \mathbb{E} \exp(tY) \leq \mathbb{E} \sum_{i=1}^n \exp(tx_i^2) = n\mathbb{E} \exp(tx_i^2).$$

Note that the moment-generating function of chi-square distribution with v degrees of freedom is:

$$M_X(t) = (1 - 2t)^{-v/2}.$$

Then combine this fact with Equation (B.2) we have:

$$\exp(t\mathbb{E}Y) \leq n(1 - 2t)^{-\frac{1}{2}},$$

which implies:

$$\mathbb{E}Y \leq \frac{\log(n)}{t} - \frac{1 - 2t}{2t}, \quad \forall t < \frac{1}{2}.$$

In particular, take $t \rightarrow \frac{1}{2}$ yields:

$$\mathbb{E}Y \leq 2 \log(n)$$

as desired. ■

Lemma B.6 (Restatement of Lemma 11)

$$\mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} I(U^*) = O\left(\frac{r}{d} \log d\right). \quad (36)$$

Proof [Proof of Lemma 11] Denote the columns of U as $U = [u_1, \dots, u_r] \in \mathbb{O}_{d,r}$, we have:

$$\begin{aligned} \mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} I(U) &= \mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} \max_{i \in [d]} \sum_{j=1}^r |e_i^\top u_j|^2 \\ &\leq \mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} \sum_{j=1}^r \max_{i \in [d]} |e_i^\top u_j|^2 \\ &= r \mathbb{E}_{u \sim \text{Uniform}(\mathbb{S}^d)} \max_{i \in [d]} |e_i^\top u|^2. \end{aligned}$$

By Lemma B.4 we can transform this expectation on the uniform sphere distribution into normalized multivariate Gaussian variables:

$$\mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} I(U) = r \mathbb{E}_{x_1, \dots, x_d} \frac{\max_{i \in [d]} x_i^2}{\sum_{j=1}^d x_j^2}. \quad (37)$$

where x_1, x_2, \dots, x_d are i.i.d. standard normal random variables. Apply Chebyshev's inequality we know that:

$$\mathbb{P}\left(\left|\frac{1}{d} \sum_{j=1}^d x_j^2 - 1\right| > \epsilon\right) \leq \frac{2}{d\epsilon^2}.$$

In particular, take $\epsilon = 1$ we have:

$$\mathbb{P}\left(\sum_{j=1}^d x_j^2 < \frac{d}{2}\right) \leq \frac{8}{d}.$$

Then take it back into Equation (37) and apply Lemma B.5 we obtain:

$$\begin{aligned}
 \mathbb{E}_{U \sim \text{Uniform}(\mathbb{O}_{d,r})} I(U) &= r \mathbb{E}_{x_1, \dots, x_d} \frac{\max_{i \in [d]} x_i^2}{\sum_{j=1}^d x_j^2} \mathbb{I} \left\{ \sum_{i=1}^d x_j^2 < \frac{d}{2} \right\} \\
 &\quad + r \mathbb{E}_{x_1, \dots, x_d} \frac{\max_{i \in [d]} x_i^2}{\sum_{j=1}^d x_j^2} \mathbb{I} \left\{ \sum_{i=1}^d x_j^2 \geq \frac{d}{2} \right\} \\
 &\leq r \mathbb{P} \left(\sum_{i=1}^d x_j^2 < \frac{d}{2} \right) + \frac{2r}{d} \mathbb{E}_{x_1, \dots, x_d} \max_{i \in [d]} x_i^2 \\
 &\leq \frac{8r}{d} + \frac{4r \log d}{d}
 \end{aligned}$$

as desired. ■

Now we start proving our main results. Note that U_{AE} is the top- r left eigenspace of the observed covariance matrix and U^* is that of the core feature covariance matrix, and by Assumption 9 the observed covariance matrix is dominated by the covariance of random noise. The Davis-Kahan theorem provides a technique to estimate the eigenspace distance via estimating the difference between target matrices. We will adopt this technique to prove the lower bound of the feature recovery ability of autoencoders in Theorem 13.

Theorem B.7 (Restatement of Theorem 13) *Consider the spiked covariance model Eq. (5), under Assumptions 8-10 and $n > d \gg r$, let W_{AE} be the learned representation of autoencoder with singular value decomposition $W_{AE} = (U_{AE} \Sigma_{AE} V_{AE}^\top)^\top$ (as in Eq. (7)). If we further assume $\{\sigma_i^2\}_{i=1}^d$ are different from each other and $\sigma_{(1)}^2 / (\sigma_{(r)}^2 - \sigma_{(r+1)}^2) < C_\sigma$ for some universal constant C_σ . Then there exist two universal constants $C_\rho > 0, c \in (0, 1)$, such that when $\rho < C_\rho$, we have*

$$\mathbb{E} \|\sin \Theta(U^*, U_{AE})\|_F \geq c\sqrt{r}. \quad (38)$$

Proof [Proof of Theorem 13] Denote $M = \nu^2 U^* U^{*\top}$ to be the target matrix, $x_i = U^* z_i + \xi_i$, $i = 1, 2, \dots, n$ to be the samples generated from model 5 and let $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$, $Z = [z_1, \dots, z_n] \in \mathbb{R}^{r \times n}$, $E = [\xi_1, \dots, \xi_n] \in \mathbb{R}^{d \times n}$ to be the corresponding matrices. In addition, we write the column mean matrix $\bar{X} \in \mathbb{R}^{n \times d}$ of a matrix $X \in \mathbb{R}^{n \times d}$ to be $\bar{X} = \frac{1}{n} X 1_n 1_n^\top$, that is, each column of \bar{X} is the column mean of X . We denote the sum of variance σ_i^2 as $\sigma_{\text{sum}}^2 = \sum_{i=1}^d \sigma_i^2$. As shown in Equation (7), autoencoders find the top- r eigenspace of the following matrix:

$$\hat{M}_1 = \frac{1}{n} X (I_n - \frac{1}{n} 1_n 1_n^\top) X^\top = \frac{1}{n} (U^* Z + E)(U^* Z + E)^\top - \frac{1}{n} (U^* \bar{Z} + \bar{E})(U^* \bar{Z} + \bar{E})^\top.$$

The rest of the proof is divided into three steps for the sake of presentation.

Step 1. Bound the difference between \hat{M}_1 and Σ In this step, we aim to show that the data recovery of autoencoders is dominated by the random noise term. Note that $\Sigma = \text{Cov}(\xi) = \mathbb{E}\xi\xi^\top$, we just need to bound the norm of the following matrix:

$$\hat{M}_1 - \Sigma = \frac{1}{n}U^*ZZ^\top U^{*\top} + \frac{1}{n}(U^*ZE^\top + EZ^\top U^{*\top}) + \left(\frac{1}{n}EE^\top - \Sigma\right) - \frac{1}{n}(U^*\bar{Z} + \bar{E})(U^*\bar{Z} + \bar{E})^\top, \quad (39)$$

and we will deal with these four terms separately.

1. For the first term, note that $\mathbb{E}zz^\top = \nu^2 I_r$, the first term can then be divided into two terms

$$\frac{1}{n}U^*ZZ^\top U^{*\top} = M + U^*\left(\frac{1}{n}ZZ^\top - \mathbb{E}zz^\top\right)U^{*\top}. \quad (40)$$

Then apply the concentration inequality of Wishart-type matrices (Lemma E.3) we have:

$$\mathbb{E}\left\|\frac{1}{n}ZZ^\top - \mathbb{E}zz^\top\right\|_2 \leq \left(\sqrt{\frac{r}{n}} + \frac{r}{n}\right)\nu^2.$$

Plug it back into (40) we obtain the bound for the first term:

$$\left\|\frac{1}{n}UZZ^\top U^\top\right\|_2 \leq \|M\|_2 + \|U\|_2 \left\|\frac{1}{n}ZZ^\top - \mathbb{E}zz^\top\right\|_2 \|U\|_2 \leq \left(1 + \sqrt{\frac{r}{n}} + \frac{r}{n}\right)\nu^2. \quad (41)$$

2. For the second term, since Z and E are independent, we must have $\mathbb{E}U^*ZE^\top = 0$, so apply Lemma E.2 twice we have:

$$\begin{aligned} \frac{1}{n}\mathbb{E}\|EZ^\top U^*\|_2 &= \frac{1}{n}\mathbb{E}_Z[\mathbb{E}_E[\|EZ^\top U^*\|_2|Z]] \\ &\lesssim \frac{1}{n}\mathbb{E}_Z[\|Z\|_2(\sigma_{\text{sum}} + r^{1/4}\sqrt{\sigma_{\text{sum}}\sigma_{(1)}} + \sqrt{r}\sigma_{(1)})] \\ &\lesssim \frac{1}{n}\mathbb{E}_Z[\|Z\|_2]\sqrt{d}\sigma_{(1)} \\ &\lesssim \frac{1}{n}\sqrt{d}\sigma_{(1)}(r^{1/2}\nu + (nr)^{1/4}\nu + n^{1/2}\nu) \\ &\lesssim \frac{\sqrt{d}}{\sqrt{n}}\sigma_{(1)}\nu. \end{aligned} \quad (42)$$

3. For the third term, apply Lemma E.3 again yields:

$$\mathbb{E}\left\|\frac{1}{n}EE^\top - \Sigma\right\|_2 \leq \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right)\sigma_{(1)}^2. \quad (43)$$

4. For the last term, note that each column of \bar{Z} and \bar{E} are the same, so we can rewrite it as:

$$\frac{1}{n}(U^*\bar{Z} + \bar{E})(U^*\bar{Z} + \bar{E})^\top = (U^*\bar{z} + \bar{\xi})(U^*\bar{z} + \bar{\xi})^\top,$$

where $\bar{z} = \frac{1}{n} \sum_{i=1}^n z_i$ and $\bar{\xi} = \frac{1}{n} \sum_{i=1}^n \xi_i$. Since z and ξ are independent zero mean sub-Gaussian random variables and $\text{Cov}(z) = \nu^2 I_r$, $\text{Cov}(\xi) = \Sigma$, we can conclude that:

$$\begin{aligned} \mathbb{E} \left\| \frac{1}{n} (U^* \bar{Z} + \bar{E})(U^* \bar{Z} + \bar{E})^\top \right\|_2 &\leq \mathbb{E} \|\bar{z}\bar{z}^\top\|_2 + 2\mathbb{E} \|\bar{z}\bar{\xi}^\top\|_2 + \mathbb{E} \|\bar{\xi}\bar{\xi}^\top\|_2 \\ &\lesssim \frac{r\nu^2}{n} + \frac{\sqrt{d}}{\sqrt{n}} \sigma_{(1)} \nu + \frac{d\sigma_{(1)}^2}{n}. \end{aligned} \quad (44)$$

To sum up, combine equations (41)(42)(43)(44) together we obtain the upper bound for the 2 norm expectation of matrix $\hat{M} - \Sigma$:

$$\mathbb{E} \|\hat{M}_1 - \Sigma\|_2 \lesssim \nu^2 \left(1 + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sqrt{\frac{d}{n}} \sigma_{(1)} \nu. \quad (45)$$

Step 2. Bound the sine distance between eigenspaces As we have shown in step 1, the target matrix of autoencoders is close to the covariance matrix of random noise, that is, Σ . Note that Σ is assumed to be a diagonal matrix with different elements, hence its eigenspace only consists of canonical basis e_i . Denote U_Σ to be the top- r eigenspace of Σ and $\{e_i\}_{i \in C}$ to be its corresponding basis vectors, apply the Davis-Kahan Theorem E.1 we can conclude that:

$$\begin{aligned} \mathbb{E} \|\sin \Theta(U_{\text{AE}}, U_\Sigma)\|_F &\leq \frac{2\sqrt{r} \mathbb{E} \|\hat{M}_1 - \Sigma\|_2}{\sigma_{(r)}^2 - \sigma_{(r+1)}^2} \\ &\lesssim \sqrt{r} \frac{1}{\sigma_{(1)}^2} \left(\nu^2 \left(1 + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sqrt{\frac{d}{n}} \sigma_{(1)} \nu \right) \\ &\lesssim \sqrt{r} \left(\rho^2 + \sqrt{\frac{d}{n}} + \rho \sqrt{\frac{d}{n}} \right). \end{aligned}$$

Step 3. Obtain the final result by triangular inequality By Assumption 10 we know that the distance between canonical basis and the eigenspace of core features can be large:

$$\begin{aligned} \|\sin \Theta(U^*, U_\Sigma)\|_F^2 &= \|U_{\Sigma^\perp}^\top U^*\|_F^2 = \sum_{i \in [d] \setminus C} \|e_i^\top U^*\|^2 = \|U^*\|_F^2 - \sum_{i \in C} \|e_i^\top U^*\|^2 \\ &\geq r - rI(U^*) = r - O\left(\frac{r^2}{d} \log d\right). \end{aligned}$$

Then apply the triangular inequality of sine distance (Proposition A.5) we can obtain the lower bound of autoencoders.

$$\begin{aligned} \mathbb{E} \|\sin \Theta(U_{\text{AE}}, U^*)\|_F &\geq \mathbb{E} \|\sin \Theta(U^*, U_\Sigma)\|_F - \mathbb{E} \|\sin \Theta(U_{\text{AE}}, U_\Sigma)\|_F \\ &\geq \sqrt{r} - O\left(\frac{r}{\sqrt{d}} \sqrt{\log d}\right) - O\left(\sqrt{r} \left(\rho^2 + \sqrt{\frac{d}{n}} + \rho \sqrt{\frac{d}{n}}\right)\right). \end{aligned} \quad (46)$$

By Assumption 9, it implies that when n and d are sufficiently large and ρ is sufficiently small (smaller than a given constant $C_\rho > 0$), there exists a universal constant $c \in (0, 1)$ such that:

$$\mathbb{E} \|\sin \Theta(U_{\text{AE}}, U^*)\|_F \geq c\sqrt{r}. \quad \blacksquare$$

To start the proof, we introduce a technical lemma first.

Lemma B.8 (Lemma 4 in Zhang et al. (2018)) *If $M \in \mathbb{R}^{p \times p}$ is any square matrix and $\Delta(M)$ is the matrix M with diagonal entries set to 0, then*

$$\|\Delta(M)\|_2 \leq 2\|M\|_2.$$

Here, factor " 2 " in the statement above cannot be improved.

Theorem B.9 (Restatement of Theorem 14) *Under the spiked covariance model Eq.(5), random masking augmentation in Definition 2, Assumptions 8-10 and $n > d \gg r$, let W_{CL} be any solution that minimizes Eq.(3), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have*

$$\mathbb{E} \|\sin \Theta(U^*, U_{CL})\|_F \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}. \quad (47)$$

Proof [Proof of Theorem 14] The proof strategy is quite similar to that of Theorem 13 and we follow the notation defined in the first paragraph of that proof. As we have shown in Corollary 6, under our linear representation setting, the contrastive learning algorithm finds the top- r eigenspace of the following matrix:

$$\begin{aligned} \hat{M}_2 &= \frac{1}{n} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \\ &= \frac{1}{n} \Delta((U^*Z + E)(U^*Z + E)^\top) - \frac{1}{n-1} (U^*\bar{Z} + \bar{E})(U^*\bar{Z} + \bar{E})^\top \\ &\quad + \frac{1}{n(n-1)} (U^*Z + E)(U^*Z + E)^\top. \end{aligned}$$

To prove the theorem, first we need to bound the difference between \hat{M}_2 and M . We aim to show that the contrastive learning algorithm is dominated by the core feature term. Note that $\Sigma = \mathbb{E}Uzz^\top U^\top$, we just need to bound the norm of the following matrix:

$$\begin{aligned} \hat{M}_2 - M &= \left(\frac{1}{n} \Delta(U^*ZZ^\top U^{*\top}) - M \right) + \frac{1}{n} \Delta(U^*ZE^\top + EZ^\top U^{*\top}) + \frac{1}{n} \Delta(EE^\top) \\ &\quad - \frac{1}{n-1} (U^*\bar{Z} + \bar{E})(U^*\bar{Z} + \bar{E})^\top + \frac{1}{n(n-1)} (U^*Z + E)(U^*Z + E)^\top. \end{aligned} \quad (48)$$

and we will also deal with these five terms separately.

1. For the first term, we can divide it into two parts:

$$\frac{1}{n}\Delta(U^*ZZ^\top U^{*\top}) - M = \Delta\left(\frac{1}{n}U^*ZZ^\top U^{*T} - M\right) + \Delta(M) - M. \quad (49)$$

Then apply Lemma B.8 and Lemma E.3 we have:

$$\mathbb{E}\|\Delta\left(\frac{1}{n}U^*ZZ^\top U^{*\top} - M\right)\|_2 \leq 2\mathbb{E}\|\frac{1}{n}U^*ZZ^\top U^{*\top} - M\|_2 \leq 2\left(\sqrt{\frac{r}{n}} + \frac{r}{n}\right)\nu^2.$$

Using the incoherent condition $I(U) = O\left(\frac{r}{d}\log d\right)$, we know that:

$$\|M - \Delta(M)\|_2 \leq \nu^2 \max_{i \in [d]} \|e_i^\top U^*\|_2^2 = \nu^2 I(U^*) \lesssim \frac{r}{d} \log d \nu^2.$$

Combine the two equations above together we obtain the bound for the first term:

$$\mathbb{E}\|\frac{1}{n}\Delta(U^*ZZ^\top U^{*\top}) - M\|_2 \leq \mathbb{E}\|\Delta\left(\frac{1}{n}U^*ZZ^\top U^{*\top} - M\right)\|_2 + \|M - \Delta(M)\|_2 \quad (50)$$

$$\lesssim \nu^2 \left(\frac{r}{d} \log d + \frac{r}{n} + \sqrt{\frac{r}{n}}\right). \quad (51)$$

2. For the second term, apply equation (42) yields:

$$\frac{1}{n}\mathbb{E}\|\Delta(U^*ZE^\top + EZ^\top U^{*\top})\|_2 \leq \frac{4}{n}\mathbb{E}\|EZ^\top U^{*\top}\|_2 \lesssim \frac{\sqrt{d}}{\sqrt{n}}\sigma_{(1)}\nu. \quad (52)$$

3. For the third term, apply equation (43) yields:

$$\mathbb{E}\|\frac{1}{n}\Delta(EE^\top)\|_2 = \mathbb{E}\|\Delta\left(\frac{1}{n}EE^\top - \Sigma\right)\|_2 \leq 2\|\frac{1}{n}EE^\top - \Sigma\|_2 \lesssim \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right)\sigma_{(1)}^2. \quad (53)$$

4. For the fourth term, apply equation (44) yields:

$$\begin{aligned} \mathbb{E}\|\frac{1}{n-1}(U^*\bar{Z} + \bar{E})(U\bar{Z} + \bar{E})^\top\|_2 &\lesssim \mathbb{E}\|\frac{1}{n}(U\bar{Z} + \bar{E})(U\bar{Z} + \bar{E})^\top\|_2 \\ &\lesssim \frac{r\nu^2}{n} + \frac{\sqrt{d}}{\sqrt{n}}\sigma_{(1)}\nu + \frac{d\sigma_{(1)}^2}{n}. \end{aligned} \quad (54)$$

5. For the last term, by equations (41)(42)(43) we know:

$$\begin{aligned} &\mathbb{E}\|\frac{1}{n}(U^*Z + E)(U^*Z + E)^\top\|_2 \\ &\lesssim \|\Sigma\|_2 + \left(1 + \sqrt{\frac{r}{n}} + \frac{r}{n}\right)\nu^2 + \sqrt{\frac{d}{n}}\sigma_{(1)}\nu + \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right)\sigma_{(1)}^2. \end{aligned}$$

Thus we can conclude that:

$$\mathbb{E}\|\frac{1}{n(n-1)}(U^*Z + E)(U^*Z + E)^\top\|_2 \lesssim \frac{d}{n}\sigma_{(1)}^2 + \frac{r}{n}\nu^2. \quad (55)$$

To sum up, combine equations (50)(52)(53)(54)(55) together we obtain the upper bound for the 2 norm expectation of matrix $\hat{M}_2 - M$:

$$\mathbb{E}\|\hat{M}_2 - M\|_2 \lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sigma_{(1)} \nu \sqrt{\frac{d}{n}}. \quad (56)$$

With the upper bound for $\|\hat{M}_2 - M\|_2$, simply apply Lemma E.1 we can obtain the desired bound for sine distance:

$$\begin{aligned} \mathbb{E}\|\sin \Theta(U_{\text{CL}}, U^*)\|_F &\leq \frac{2\sqrt{r}\mathbb{E}\|\hat{M}_2 - M\|_2}{\nu^2} \\ &\lesssim \sqrt{r} \frac{1}{\nu^2} \left(\nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sigma_{(1)} \nu \sqrt{\frac{d}{n}} \right) \\ &= \sqrt{r} \left(\left(\frac{r}{d} \log d + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \rho^{-2} \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \rho^{-1} \sqrt{\frac{d}{n}} \right) \\ &\lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}. \end{aligned} \quad (57)$$

Moreover, there exists an orthogonal matrix $\hat{O} \in \mathbb{O}^{r \times r}$ depending on U_{CL} such that:

$$\mathbb{E}\|U^\top U_{\text{CL}} \hat{O} - I_r\|_F = \mathbb{E}\|U_{\text{CL}} \hat{O} - U\|_F \leq \frac{2\sqrt{r}\mathbb{E}\|\hat{M}_2 - M\|_2}{\nu^2} \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}.$$

which finishes the proof. \blacksquare

B.3 Proofs for Section 3.4

In this section, we will provide the proof of Theorems 17 and 18 with both regression and classification settings. The corresponding statement and proof can be found in Theorems B.10 and B.11.

For notation simplicity, define the prediction risk of predictor δ for classification and regression tasks as $\mathcal{R}_c(\delta) := \mathbb{E}_{\mathcal{D}}[\ell_c(\delta)]$ and $\mathcal{R}_r(\delta) := \mathbb{E}_{\mathcal{D}}[\ell_r(\delta)]$, respectively. Define $\Sigma_x := \nu^2 U^* U^{*\top} + \Sigma$. We write $\delta_{U,w}$ for $\delta_{U^\top, w}$ with a slight abuse of notation.

For two matrices A and B of the same order, we define $A \succeq B$ when $A - B$ is positive semi-definite.

Theorem B.10 (Restatement of Theorem 17) *Suppose the conditions in Theorem 14 hold. Then, for the classification task, we have*

$$\mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{W_{\text{CL}}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U^{*\top}, w})] \right] = O \left(\frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}} \right) \wedge 1,$$

and for regression tasks,

$$\mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{W_{\text{CL}}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U^{*\top}, w})] \right] \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}.$$

Theorem B.11 (Restatement of Theorem 18) *Suppose the conditions in Theorem 13 hold. Assume $r \leq r_c$ holds for some constant $r_c > 0$. Additionally assume that $\rho = \Theta(1)$ is sufficiently small and $n \gg d \gg r$. Then, For the regression task,*

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U_{AE}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_r(\delta_{U^*, w})] \geq c'_c,$$

and for classification task, if F is differentiable at 0 and $F'(0) > 0$, then

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U_{AE}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}}[\ell_c(\delta_{U^*, w})] \geq c'_r,$$

where $c'_r > 0$ and $c'_c > 0$ are constants independent of n and d .

The proofs of Theorem B.10 and B.11 relies on Lemma B.15, B.16, B.17, B.18 and B.20 which are proved later in this section.

Proof [Proof of Theorem B.10: Classification Task Part] Lemma B.17 gives for any $U \in \mathbb{O}_{d,r}$,

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*, w})] \tag{58}$$

$$\leq ((\kappa(1 + \rho^2))^3 + \kappa\rho^2(1 + \rho^{-2})^2 + (\kappa\rho^2 \vee 1)^{-1})\mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2]. \tag{59}$$

Substituting $U \leftarrow U_{AE}$ combined with Assumption 9 and $\kappa = O(1)$ concludes the proof. \blacksquare

Proof [Proof of Theorem B.10: Regression Part] Note that under Assumption 9 and $\kappa = O(1)$, $(1 + \rho^{-2})/(1 + \kappa^{-1}\rho^{-2})^2 = O(1)$. Lemma B.20 gives for any $U \in \mathbb{O}_{d,r}$,

$$\mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*, w})] = O((1 + \rho^{-2})\mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2]\|w^*\|^2). \tag{60}$$

Theorem 14 with substitution $U \leftarrow U_{AE}$ gives the desired result. \blacksquare

Proof [Proof of Theorem B.11: Classification Part] Lemma B.16 gives that for $c_1 := 1 - 1/(2\kappa r_c) \in (0, 1)$, we can take $n \gg d \gg r$ and sufficiently small $\rho > 0$ so that $\mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U_{AE}, U^*)\|_F^2] \geq c_1 r$ holds. By Lemma B.18,

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U_{AE}, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*, w})] \\ & \gtrsim \frac{(1 + \rho^2)^{3/2}}{(1 + \kappa\rho^2)^{3/2}}\rho^2 \left(\frac{1}{1 + \rho^2} - \kappa(r - \|\sin \Theta(U_{AE}, U^*)\|_F^2) \right) \\ & \geq \frac{(1 + \rho^2)^{3/2}}{(1 + \kappa\rho^2)^{3/2}}\rho^2 \left(\frac{1}{1 + \rho^2} - \kappa(1 - c_1)r \right) \\ & \geq \frac{(1 + \rho^2)^{3/2}}{(1 + \kappa\rho^2)^{3/2}}\rho^2 \left(\frac{1}{1 + \rho^2} - \frac{1}{2} \right), \end{aligned} \tag{61}$$

where the last inequality follows since $r \leq r_c$. If we further take $\rho = \Theta(1) < 1/2$, the right hand becomes a positive constant. This concludes the proof. \blacksquare

Proof [Proof of Theorem B.11: Regression Part] From proposition B.19, we have

$$\begin{aligned} & \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U_{AE}, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*, w}) \\ &= w^{\star \top} \left((I + (1/\nu^2)U^{\star \top} \Sigma U^{\star})^{-1} \right. \\ & \quad \left. - U^{\star \top} U_{AE} (U_{AE}^{\top} U^{\star} U^{\star \top} U_{AE} + (1/\nu^2)U_{AE}^{\top} \Sigma U_{AE})^{-1} U_{AE}^{\top} U^{\star} \right) w^{\star}. \end{aligned}$$

Thus from Lemma B.15,

$$\begin{aligned} & \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U_{AE}, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*, w}) \\ & \geq \left(\frac{1}{1 + \rho^{-2}} + \rho^2 \kappa(\|\sin \Theta(U_{AE}, U^{\star})\|_F^2 - r) \right) \|w^{\star}\|^2. \end{aligned} \quad (62)$$

Using Lemma B.16 and by the same argument in the proof of Theorem 18: Classification Part, we conclude the proof. \blacksquare

Lemma B.12 For any $U \in \mathbb{O}_{d,r}$,

$$\lambda_{\min}(\nu^2 U^{\star \top} U (U^{\top} \Sigma_x U)^{-1} U^{\top} U^{\star}) \geq \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} (1 - \|\sin \Theta(U, U^{\star})\|_2^2).$$

Proof Since $\lambda_{\min}(AC) \geq \lambda_{\min}(A)\lambda_{\min}(C)$ for symmetric positive semi-definite matrices A and C ,

$$\begin{aligned} & \lambda_{\min}(\nu^2 U^{\star \top} U (U^{\top} \Sigma_x U)^{-1} U^{\top} U^{\star}) \\ & \geq \lambda_{\min}(U^{\top} U^{\star} U^{\star \top} U) \lambda_{\min}(\nu^2 (U^{\top} \Sigma_x U)^{-1}) \\ & \geq \lambda_{\min}(I - (I - U^{\top} U^{\star} U^{\star \top} U)) \frac{\nu^2}{\lambda_{\max}(\nu^2 U^{\top} U^{\star} U^{\star \top} U + U^{\top} \Sigma U)} \\ & \geq \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} (1 - \|\sin \Theta(U, U^{\star})\|_2^2), \end{aligned}$$

where we used Weyl's inequality $\lambda_{\min}(A + C) \geq \lambda_{\min}(A) - \|C\|_2$ in the second inequality. \blacksquare

Lemma B.13 For any $U \in \mathbb{O}_{d,r}$,

$$\lambda_{\max}(\nu^2 U^{\star \top} U (U^{\top} \Sigma_x U)^{-1} U^{\top} U^{\star}) \leq \frac{\nu^2}{\nu^2 (1 - \|\sin \Theta(U, U^{\star})\|_2) + \sigma_{(d)}^2}.$$

Proof Since $\|AC\|_2 \leq \|A\|_2\|C\|_2$,

$$\begin{aligned}
 \lambda_{\max}(\nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star) &\leq \lambda_{\max}(\nu^2 (U^\top \Sigma_x U)^{-1}) \\
 &\leq \frac{\nu^2}{\lambda_{\min}(\nu^2 U^\top U^\star U^{\star\top} U + U^\top \Sigma U)} \\
 &\leq \frac{\nu^2}{\lambda_{\min}(\nu^2 I - \nu^2 (I - U^\top U^\star U^{\star\top} U) + U^\top \Sigma U)} \\
 &\leq \frac{\nu^2}{\nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2) + \sigma_{(d)}^2},
 \end{aligned}$$

where we used Weyl's inequality $\lambda_{\min}(A + C) \geq \lambda_{\min}(A) - \|C\|_2$ and $\lambda_{\min}(\nu^2 I + U^\top \Sigma U) \geq \nu^2 + \sigma_{(d)}^2$. \blacksquare

Lemma B.14 For any $U \in \mathbb{O}_{d,r}$,

$$\begin{aligned}
 &\|\nu^2 (U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star\|_2 \\
 &= O\left(\frac{1}{1 - \|\sin \Theta(U, U^\star)\|_2^2 + \kappa^{-1} \rho^{-2}} \frac{1 + \rho^{-2}}{1 + \kappa^{-1} \rho^{-2}} \|\sin \Theta(U, U^\star)\|_2\right).
 \end{aligned}$$

Proof Observe that

$$\begin{aligned}
 &\|(U^{\star\top} \Sigma_x U^\star)^{-1} - U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star\|_2 \\
 &\leq \|(U^{\star\top} \Sigma_x U^\star)^{-1} - (U^\top \Sigma_x U)^{-1}\|_2 + \|(U^\top \Sigma_x U)^{-1} - U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star\|_2 \\
 &:= (T1) + (T2).
 \end{aligned}$$

For the term (T1),

$$\begin{aligned}
 (T1) &= \|(U^\top \Sigma_x U)^{-1} (U^\top \Sigma_x U) (U^{\star\top} \Sigma_x U^\star)^{-1} - (U^\top \Sigma_x U)^{-1} (U^{\star\top} \Sigma_x U^\star) (U^{\star\top} \Sigma_x U^\star)^{-1}\|_2 \\
 &\leq \|(U^\top \Sigma_x U)^{-1}\|_2 \|U^\top \Sigma_x U - U^{\star\top} \Sigma_x U^\star\|_2 \|(U^{\star\top} \Sigma_x U^\star)^{-1}\|_2.
 \end{aligned}$$

Note

$$\begin{aligned}
 \|U^\top \Sigma_x U - U^{\star\top} \Sigma_x U^\star\|_2 &= \|\nu^2 U^\top U^\star U^{\star\top} U - \nu^2 I + U^\top \Sigma U - U^{\star\top} \Sigma U^\star\|_2 \\
 &\leq \nu^2 \|\sin \Theta(U, U^\star)\|_2^2 + \|U^\top \Sigma (U - U^\star) + (U - U^\star)^\top \Sigma U^\star\|_2 \\
 &\leq \nu^2 \|\sin \Theta(U, U^\star)\|_2^2 + 2\sigma_{(1)}^2 \|U - U^\star\|_2.
 \end{aligned}$$

Also we have $\lambda_{\min}(U^\top \Sigma_x U) \geq \nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2^2) + \sigma_{(d)}^2$ from the proof of Lemma B.13 and $\lambda_{\min}(U^{\star\top} \Sigma_x U^\star) \geq \nu^2 + \sigma_{(d)}^2$. Therefore

$$(T1) \leq \frac{1}{(\nu^2 + \sigma_{(d)}^2)(\nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2^2) + \sigma_{(d)}^2)} (\nu^2 \|\sin \Theta(U, U^\star)\|_2^2 + 2\sigma_{(1)}^2 \|U - U^\star\|_2).$$

For the term (T2),

$$\begin{aligned}
 (T2) &= \|(U^\top \Sigma_x U)^{-1} - U^{\star\top} (U^\star + (U - U^\star)) (U^\top \Sigma_x U)^{-1} (U^\star + (U - U^\star))^\top U^\star\|_2 \\
 &= \|-U^{\star\top} (U - U^\star) (U^\top \Sigma_x U)^{-1} - (U^\top \Sigma_x U)^{-1} (U - U^\star)^\top U^\star \\
 &\quad - U^{\star\top} (U - U^\star) (U^\top \Sigma_x U)^{-1} (U - U^\star)^\top U^\star\|_2 \\
 &\leq \frac{1}{\nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2^2) + \sigma_{(d)}^2} (2\|U - U^\star\|_2 + \|U - U^\star\|_2^2).
 \end{aligned}$$

From Lemma A.3, $\|\sin \Theta(U, U^\star)\|_2 \leq \|U - U^\star\|_2$. Finally from these results and $\|U - U^\star\|_2^2 \leq 2\|U - U^\star\|_2$,

$$\begin{aligned}
 &\|\nu^2 (U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star\|_2 \\
 &= O\left(\frac{\nu^2}{\nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2^2) + \sigma_{(d)}^2} \frac{\nu^2 + \sigma_{(1)}^2}{\nu^2 + \sigma_{(d)}^2} \|U - U^\star\|_2\right).
 \end{aligned}$$

Since LHS does not depend on the orthogonal transformation $U \leftarrow UO$ where $O \in \mathbb{O}_{r,r}$, we obtain

$$\begin{aligned}
 &\|\nu^2 (U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star\|_2 \\
 &= O\left(\frac{\nu^2}{\nu^2 (1 - \|\sin \Theta(U, U^\star)\|_2^2) + \sigma_{(d)}^2} \frac{\nu^2 + \sigma_{(1)}^2}{\nu^2 + \sigma_{(d)}^2} \inf_{O \in \mathbb{O}_{r,r}} \|UO - U^\star\|_2\right).
 \end{aligned}$$

Combined again with Lemma A.3, we obtain the desired result. \blacksquare

Lemma B.15 For any $U \in \mathbb{O}_{d,r}$,

$$\begin{aligned}
 &\lambda_{\min}(\nu^2 (U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star) \\
 &\geq \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} - \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^\star)\|_F^2).
 \end{aligned}$$

Proof Observe

$$\begin{aligned}
 &\lambda_{\min}(\nu^2 (U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star) \\
 &\geq \lambda_{\min}((I + (1/\nu^2) U^{\star\top} \Sigma U^\star)^{-1}) - \|U^{\star\top} U (U^\top U^\star U^{\star\top} U + (1/\nu^2) U^\top \Sigma U)^{-1} U^\top U^\star\|_2.
 \end{aligned}$$

Since $U^\top U^\star U^{\star\top} U \succeq 0$, it follows that $(U^\top U^\star U^{\star\top} U + (1/\nu^2) U^\top \Sigma U)^{-1} \preceq \nu^2 (U^\top \Sigma U)^{-1}$. Thus

$$\begin{aligned}
 &\|U^{\star\top} U (U^\top U^\star U^{\star\top} U + (1/\nu^2) U^\top \Sigma U)^{-1} U^\top U^\star\|_2 \\
 &\leq \nu^2 \lambda_{\max}((U^\top \Sigma U)^{-1}) \|U^{\star\top} U\|_2^2 \\
 &\leq \frac{\nu^2}{\sigma_{(d)}^2} \|U^{\star\top} U\|_F^2 \\
 &= \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^\star)\|_F^2),
 \end{aligned}$$

where we used $\lambda_{\max}((U^\top \Sigma U)^{-1}) \leq 1/\lambda_{\min}(U^\top \Sigma U) \leq 1/\sigma_{(d)}^2$ and $\|\sin \Theta(U_1, U_2)\|_F^2 = r - \|U_1^\top U_2\|_F^2$ from Proposition A.1. Combined with Lemma B.13, we obtain

$$\begin{aligned} & \lambda_{\min}(\nu^2(U^{\star\top} \Sigma_x U^\star)^{-1} - \nu^2 U^{\star\top} U (U^\top \Sigma_x U)^{-1} U^\top U^\star) \\ & \geq \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} - \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^\star)\|_F^2). \end{aligned}$$

■

Lemma B.16 *Suppose the conditions in Theorem 13 hold. Fix $c_1 \in (0, 1)$. There exists a constant $c_2 > 0$ such that if $\sqrt{r \log d/d} \vee \rho^2 \vee d/n < c_2$, then,*

$$\mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F^2 \geq c_1 r,$$

where $c_1 \in (0, 1)$ is a universal constant.

Proof By Cauchy-Schwartz inequality,

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F^2 - r \\ & \geq (\mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F)^2 - r \\ & = (\mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F - \sqrt{r})(\mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F + \sqrt{r}). \end{aligned}$$

From Theorem 13, there exists a constant $c_3 > 0$ such that we have

$$\mathbb{E}_{\mathcal{D}} \|\sin \Theta(U^\star, U_{AE})\|_F \geq \sqrt{r} - c_3 \frac{r}{\sqrt{d}} \sqrt{\log d} - c_3 \sqrt{r} \left(\rho^2 + \sqrt{\frac{d}{n}} + \rho \sqrt{\frac{d}{n}} \right).$$

Therefore combined with a trivial bound $\|\sin \Theta(U_{AE}, U^\star)\|_F \leq \sqrt{r}$,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \|\sin \Theta(U_{AE}, U^\star)\|_F^2 - r & \geq -rc_3 \frac{r^{1/2}}{\sqrt{d}} \sqrt{\log d} + \rho^2 + \sqrt{\frac{d}{n}} + \rho \sqrt{\frac{d}{n}} \\ & \geq -rc_3 \left(2 \frac{r^{1/2}}{\sqrt{d}} \sqrt{\log d} \vee 6\rho^2 \vee 6\sqrt{\frac{d}{n}} \right), \end{aligned}$$

where we used $\rho \sqrt{d/n} \leq \rho^2 \vee d/n \leq \rho^2 \vee \sqrt{d/n}$ since $d < n$. Thus we can take $c_2 = 6(1 - c_1)/c_3$. This concludes the proof. ■

Lemma B.17 *For any $U \in \mathbb{O}_{d,r}$,*

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^\star,w}) \right] \\ & \leq ((\kappa(1 + \rho^2))^3 + \kappa\rho^2(1 + \rho^{-2})^2 + (\kappa\rho^2 \vee 1)^{-1}) \mathbb{E}_{\mathcal{D}} [\|\sin \Theta(U, U^\star)\|_2]. \end{aligned}$$

Proof Recall that we are considering the class of linear classifiers $\{\delta_{U,w} : w \in \mathbb{R}^r\}$, where $\delta_{U,w}(\tilde{x}) = \mathbb{I}\{F(\tilde{x}^\top U w) > 1/2\}$. For notational simplicity, write $\beta := U w$ and $\beta^* := U^* w^*$.

$$\mathcal{R}_c(\delta_{U,w}) = \mathbb{P}_\mathcal{E}(\delta_{U,w}(\tilde{x}) \neq \tilde{y}) = \mathbb{P}_\mathcal{E}(\tilde{y} = 0, F(\tilde{x}^\top \beta) > 1/2) + \mathbb{P}_\mathcal{E}(\tilde{y} = 1, F(\tilde{x}^\top \beta) \leq 1/2).$$

Since $F(0) = 1/2$ and F is monotone increasing, the false positive probability becomes

$$\begin{aligned} \mathbb{P}_\mathcal{E}(\tilde{y} = 0, F(\tilde{x}^\top \beta) > 1/2) &= \mathbb{P}_\mathcal{E}(\tilde{y} = 0, \tilde{x}^\top \beta > 0) \\ &= \mathbb{E}_\mathcal{E}[\mathbb{E}_\mathcal{E}[\mathbb{I}\{\tilde{y} = 0\} | \tilde{x}, \tilde{z}] \mathbb{I}\{\tilde{x}^\top \beta > 0\}] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\nu^{-1} \tilde{z}^\top U^{*\top} \beta^*)) \mathbb{I}\{\tilde{x}^\top \beta > 0\}]. \end{aligned}$$

Write $\omega := \tilde{x}^\top \beta$ and $\omega^* := \nu^{-1} \tilde{z}^\top U^{*\top} \beta^*$. From assumption, (ω^*, ω) jointly follows a normal distribution with mean 0. Write $v^{*2} := \text{Var}(\omega^*) = w^{*\top} w^*$, $v^2 := \text{Var}(\omega) = \beta^\top \Sigma_x \beta$, where $\Sigma_x := \nu^2 U^* U^{*\top} + \Sigma$. Let $\tau := \text{Cor}(\omega^*, \omega) = \nu w^{*\top} U^{*\top} \beta / (v^* v)$. By a formula for conditional normal distribution, we have $\omega | \omega^* \sim N(\tau v \omega^* / v^*, v^2(1 - \tau^2))$. This gives

$$\begin{aligned} \mathbb{P}_\mathcal{E}(\tilde{y} = 0, F(\tilde{x}^\top \beta) > 1/2) &= \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \mathbb{I}\{\omega > 0\}] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \mathbb{E}_\mathcal{E}[\mathbb{I}\{\omega > 0\} | \omega^*]] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \mathbb{P}_\mathcal{E}(\omega > 0 | \omega^*)] \\ &= \mathbb{E}_\mathcal{E} \left[(1 - F(\omega^*)) \mathbb{P}_\mathcal{E} \left(\frac{\omega - \tau v \omega^* / v^*}{v(1 - \tau^2)^{1/2}} > -\frac{\tau v \omega^* / v^*}{v(1 - \tau^2)^{1/2}} \middle| \omega^* \right) \right] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \Phi(\alpha \omega^* / v^*)] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \Phi(\alpha \omega^* / v^*) \mathbb{I}\{\omega^* > 0\}] + \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \Phi(\alpha \omega^* / v^*) \mathbb{I}\{\omega^* < 0\}], \end{aligned}$$

where Φ is cumulative distribution function of $N(0, 1)$ and $\alpha := \tau / (1 - \tau^2)^{1/2}$. We define Ψ_F as $\Psi_F(s^2) := 2E_{u \sim N(0, s^2)}[F(u) \mathbb{I}\{u > 0\}]$. When $F(u) = 1/(1 + e^{-u})$, $\Psi_F(s^2)$ is called the logistic-normal integral, whose analytical form is not known (Pirjol, 2013). Since a random variable ω^* is symmetric about mean 0 and $F(u) = 1 - F(-u)$,

$$\begin{aligned} \mathbb{E}_\mathcal{E}[(1 - F(\omega^*)) \Phi(\alpha \omega^* / v^*) \mathbb{I}\{\omega^* < 0\}] &= \mathbb{E}_\mathcal{E}[(1 - F(-\omega^*)) (1 - \Phi(\alpha \omega^* / v^*)) \mathbb{I}\{\omega^* > 0\}] \\ &= \mathbb{E}_\mathcal{E}[F(\omega^*) (1 - \Phi(\alpha \omega^* / v^*)) \mathbb{I}\{\omega^* > 0\}]. \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{P}_\mathcal{E}(\tilde{y} = 0, F(\tilde{x}^\top \beta) > 1/2) &= \mathbb{E}_\mathcal{E}[(\Phi(\alpha \omega^* / v^*) + F(\omega^*) - 2F(\omega^*) \Phi(\alpha \omega^* / v^*)) \mathbb{I}\{\omega^* > 0\}] \\ &= \frac{1}{2} \Psi_F(v^{*2}) - \mathbb{E}_\mathcal{E}[(2F(\omega^*) - 1) \Phi(\alpha \omega^* / v^*) \mathbb{I}\{\omega^* > 0\}]. \end{aligned}$$

Note that the true negative probability is exactly the same as the false positive probability under our settings:

$$\begin{aligned} \mathbb{P}_\mathcal{E}(\tilde{y} = 1, F(\tilde{x}^\top \beta) \leq 1/2) &= \mathbb{E}_\mathcal{E}[F(\tilde{x}^\top \beta^*) \mathbb{I}\{\tilde{x}^\top \beta \leq 0\}] \\ &= \mathbb{E}_\mathcal{E}[F(-\tilde{x}^\top \beta^*) \mathbb{I}\{\tilde{x}^\top \beta \geq 0\}] \\ &= \mathbb{E}_\mathcal{E}[(1 - F(\tilde{x}^\top \beta^*)) \mathbb{I}\{\tilde{x}^\top \beta \geq 0\}] \\ &= \mathbb{P}_\mathcal{E}(\tilde{y} = 0, F(\tilde{x}^\top \beta) > 1/2). \end{aligned}$$

Therefore

$$\mathcal{R}_c(\delta_{U,w}) = \Psi_F(v^{*2}) - 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)\Phi(\alpha\omega^*/v^*)\mathbb{I}\{\omega^* > 0\}].$$

Let

$$\begin{aligned}\tau_{\max,U} &:= \sup_{w \in \mathbb{R}^r} \nu w^{*\top} U^{*\top} U w / (w^{*\top} w^* w^\top U^\top \Sigma_x U w)^{1/2}, \\ \tau_{\max,U^*} &:= \sup_{w \in \mathbb{R}^r} \nu w^{*\top} w / (w^{*\top} w^* w^\top U^{*\top} \Sigma_x U^* w)^{1/2}.\end{aligned}$$

From Cauchy-Schwartz inequality,

$$\begin{aligned}\tau_{\max,U}^2 &= \frac{\nu^2 w^{*\top} U^{*\top} U (U^\top \Sigma_x U)^{-1} U^\top U^* w^*}{w^{*\top} w^*}, \\ \tau_{\max,U^*}^2 &= \frac{\nu^2 w^{*\top} (U^{*\top} \Sigma_x U^*)^{-1} w^*}{w^{*\top} w^*}.\end{aligned}$$

Define $\alpha_{\max,U} := \tau_{\max,U} / (1 - \tau_{\max,U}^2)^{1/2}$ and $\alpha_{\max,U^*} := \tau_{\max,U^*} / (1 - \tau_{\max,U^*}^2)^{1/2}$. Then, since on the event where $\omega^* > 0$, $\alpha \mapsto \Phi(\alpha\omega^*/v^*)$ is monotone increasing and $2F(\omega^*) - 1$ is non-negative, we have

$$\begin{aligned}\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) &= \Psi_F(v^{*2}) - 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)\Phi(\alpha_{\max,U}\omega^*/v^*)\mathbb{I}\{\omega^* > 0\}] \\ \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) &= \Psi_F(v^{*2}) - 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)\Phi(\alpha_{\max,U^*}\omega^*/v^*)\mathbb{I}\{\omega^* > 0\}].\end{aligned}$$

This yields

$$\begin{aligned}\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\ = 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)(\Phi(\alpha_{\max,U^*}\omega^*/v^*) - \Phi(\alpha_{\max,U}\omega^*/v^*))\mathbb{I}\{\omega^* > 0\}].\end{aligned}$$

Note that for any $a, b \geq 0$,

$$|\Phi(b) - \Phi(a)| \leq \phi(a \wedge b)|b - a|,$$

where ϕ is a density function of standard normal distribution. Observe

$$\begin{aligned}& \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\ & \leq 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)|\Phi(\alpha_{\max,U^*}\omega^*/v^*) - \Phi(\alpha_{\max,U}\omega^*/v^*)|\mathbb{I}\{\omega^* > 0\}] \\ & \lesssim \frac{2}{v^*} \int_0^\infty (2F(\omega^*) - 1)|\alpha_{\max,U^*} - \alpha_{\max,U}| \omega^* \phi((\alpha_{\max,U^*} \wedge \alpha_{\max,U})\omega^*/v^*) \frac{\phi(\omega^*/v^*)}{v^*} d\omega^* \\ & \lesssim \frac{|\alpha_{\max,U^*} - \alpha_{\max,U}|}{v^*} \int_0^\infty (2F(\omega^*) - 1)\phi((\alpha_{\max,U^*} \wedge \alpha_{\max,U})\omega^*/v^*) d\omega^* \\ & = \frac{|\alpha_{\max,U^*} - \alpha_{\max,U}|}{\alpha_{\max,U^*} \wedge \alpha_{\max,U}} \int_0^\infty (2F(\omega^*) - 1) \frac{\exp(-1/(2((\alpha_{\max,U^*} \wedge \alpha_{\max,U})^{-2}v^{*2}))\omega^{*2})}{\sqrt{2\pi((\alpha_{\max,U^*} \wedge \alpha_{\max,U})^{-2}v^{*2})}} d\omega^* \\ & = \frac{|\alpha_{\max,U^*} - \alpha_{\max,U}|}{\alpha_{\max,U^*} \wedge \alpha_{\max,U}} (\Psi_F(((\alpha_{\max,U^*} \wedge \alpha_{\max,U})^{-2}v^{*2})) - 1/2),\end{aligned}$$

where we used $\sup_{u>0} u\phi(u) < \infty$. Since $(a - b) = (a^2 - b^2)/(a + b) \leq (a^2 - b^2)/(a \wedge b)$ for $a, b > 0$, and $\Psi_F \leq 1$, we obtain

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \lesssim \frac{|\alpha_{\max,U^*}^2 - \alpha_{\max,U}^2|}{\alpha_{\max,U^*}^2 \wedge \alpha_{\max,U}^2}.$$

When $\tau_{\max,U^*} \geq \tau_{\max,U}$, since $\tau \mapsto \tau^2/(1 - \tau^2)$ is increasing in $\tau > 0$,

$$\begin{aligned} \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) &\lesssim \frac{\alpha_{\max,U^*}^2 - \alpha_{\max,U}^2}{\alpha_{\max,U}^2} \\ &= \frac{\tau_{\max,U^*}^2 - \tau_{\max,U}^2}{(1 - \tau_{\max,U^*}^2)\tau_{\max,U}^2}. \end{aligned} \quad (63)$$

From Lemma B.12 and B.13, we have

$$\begin{aligned} \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} (1 - \|\sin \Theta(U, U^*)\|_2^2) &\leq \tau_{\max,U}^2 \leq \frac{\nu^2}{\nu^2 (1 - \|\sin \Theta(U, U^*)\|_2^2) + \sigma_{(d)}^2}, \\ \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} &\leq \tau_{\max,U^*}^2 \leq \frac{\nu^2}{\nu^2 + \sigma_{(d)}^2}. \end{aligned} \quad (64)$$

Then, Equation (63) becomes

$$\begin{aligned} &\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\ &\lesssim \frac{\nu^2 + \sigma_{(d)}^2}{\sigma_{(d)}^2} \frac{\nu^2 + \sigma_{(1)}^2}{\nu^2 (1 - \|\sin \Theta(U, U^*)\|_2^2)} (\tau_{\max,U^*}^2 - \tau_{\max,U}^2) \\ &\leq \frac{\nu^2 + \sigma_{(d)}^2}{\sigma_{(d)}^2} \frac{\nu^2 + \sigma_{(1)}^2}{\nu^2 (1 - \|\sin \Theta(U, U^*)\|_2^2)} \|\nu^2 (U^{*\top} \Sigma_x U^*)^{-1} - \nu^2 U^{*\top} U (U^\top \Sigma_x U)^{-1} U^\top U^*\|_2 \\ &\leq \frac{(\kappa \rho^2 + 1)(\rho^{-2} + 1)^2}{(1 + \kappa^{-1} \rho^{-2})(1 - \|\sin \Theta(U, U^*)\|_2^2)^2} \|\sin \Theta(U, U^*)\|_2 \\ &= \frac{\kappa \rho^2 (\rho^{-2} + 1)^2}{(1 - \|\sin \Theta(U, U^*)\|_2^2)^2} \|\sin \Theta(U, U^*)\|_2. \end{aligned}$$

where the last inequality follows from Lemma B.14.

On the event where $\|\sin \Theta(U, U^*)\|_2^2 \leq 1/2$,

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \lesssim \kappa \rho^2 (1 + \rho^{-2})^2 \|\sin \Theta(U, U^*)\|_2.$$

When $\tau_{\max, U^*} < \tau_{\max, U}$, on the event where $\|\sin \Theta(U, U^*)\|_2 \leq \kappa^{-1} \rho^{-2}/2$,

$$\begin{aligned}
 & \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\
 & \lesssim \frac{\nu^2 + \sigma_{(1)}^2}{\nu^2} \frac{\nu^2(1 - \|\sin \Theta(U, U^*)\|_2^2) + \sigma_{(d)}^2}{-\nu^2 \|\sin \Theta(U, U^*)\|_2^2 + \sigma_{(d)}^2} (\tau_{\max, U}^2 - \tau_{\max, U^*}^2) \\
 & \leq \frac{(\nu^2 + \sigma_{(1)}^2)^2}{\nu^2} \frac{1}{-\nu^2 \|\sin \Theta(U, U^*)\|_2^2 + \sigma_{(d)}^2} \\
 & \quad \times \|\nu^2(U^{*\top} \Sigma_x U^*)^{-1} - \nu^2 U^{*\top} U (U^\top \Sigma_x U)^{-1} U^\top U^*\|_2 \\
 & \leq \frac{(1 + \rho^{-2})^3}{(-\|\sin \Theta(U, U^*)\|_2^2 + \kappa^{-1} \rho^{-2})^3} \|\sin \Theta(U, U^*)\|_2 \\
 & \lesssim (\kappa(1 + \rho^2))^3 \|\sin \Theta(U, U^*)\|_2,
 \end{aligned}$$

where we used Lemma B.14 again.

In summary, on the event where $\|\sin \Theta(U, U^*)\|_2 \leq \kappa^{-1} \rho^{-2}/2 \wedge 1/2$,

$$\begin{aligned}
 & \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\
 & \lesssim ((\kappa(1 + \rho^2))^3 + \kappa \rho^2 (1 + \rho^{-2})^2) \|\sin \Theta(U, U^*)\|_2.
 \end{aligned}$$

On the other hand, on the event where $\|\sin \Theta(U, U^*)\|_2 > \kappa^{-1} \rho^{-2}/2 \wedge 1/2$, we have a trivial inequality $\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \leq 1$. This gives

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w})] \\
 & \lesssim ((\kappa(1 + \rho^2))^3 + \kappa \rho^2 (1 + \rho^{-2})^2) \mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2] \\
 & \quad + \mathbb{P}_{\mathcal{D}}(\|\sin \Theta(U, U^*)\|_2 > \kappa^{-1} \rho^{-2}/2 \wedge 1/2) \\
 & \lesssim ((\kappa(1 + \rho^2))^3 + \kappa \rho^2 (1 + \rho^{-2})^2 + (\kappa \rho^2 \vee 1)) \mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2],
 \end{aligned}$$

where the last inequality follows from Markov's inequality. \blacksquare

Lemma B.18 *Suppose $U \in \mathbb{O}_{d,r}$ satisfies $1/(1 + \rho^2) - \kappa(r - \|\sin \Theta(U, U^*)\|_F^2) \geq 0$. Then,*

$$\begin{aligned}
 & \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\
 & \gtrsim \frac{(1 + \rho^2)^{3/2}}{(1 + \kappa \rho^2)^{3/2}} \rho^2 \left(\frac{1}{1 + \rho^2} - \kappa(r - \|\sin \Theta(U, U^*)\|_F^2) \right).
 \end{aligned}$$

Proof We firstly bound the term $\tau_{\max, U^*}^2 - \tau_{\max, U}^2$. From Lemma B.15,

$$\begin{aligned}
 \tau_{\max, U^*}^2 - \tau_{\max, U}^2 & \geq \lambda_{\min}(\nu^2(U^{*\top} \Sigma_x U^*)^{-1} - \nu^2 U^{*\top} U (U^\top \Sigma_x U)^{-1} U^\top U^*) \\
 & \geq \frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} - \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^*)\|_F^2). \tag{65}
 \end{aligned}$$

From assumption, RHS of Equation (65) is non-negative. Then using the inequality $a - b = (a^2 - b^2)/(a + b) \geq (a^2 - b^2)/(2a)$ for $a \geq b \geq 0$,

$$\begin{aligned} \alpha_{\max, U^*} - \alpha_{\max, U} &\gtrsim \frac{1}{\alpha_{\max, U^*}} (\alpha_{\max, U^*}^2 - \alpha_{\max, U}^2) \\ &\geq \frac{(1 - \tau_{\max, U^*}^2)^{1/2}}{\tau_{\max, U^*}} \frac{\tau_{\max, U^*}^2 - \tau_{\max, U}^2}{(1 - \tau_{\max, U^*}^2)(1 - \tau_{\max, U}^2)}. \end{aligned}$$

From Equation (64) and Equation (65),

$$\begin{aligned} &\alpha_{\max, U^*} - \alpha_{\max, U} \\ &\gtrsim \left(\frac{\nu^2 + \sigma_{(d)}^2}{\nu^2} \right)^{1/2} \left(\frac{\nu^2 + \sigma_{(1)}^2}{\sigma_{(1)}^2} \right)^{3/2} \left(\frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} - \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^*)\|_F^2) \right) \\ &= (1 + \kappa^{-1} \rho^{-2})^{1/2} (1 + \rho^2)^{3/2} \left(\frac{\nu^2}{\nu^2 + \sigma_{(1)}^2} - \frac{\nu^2}{\sigma_{(d)}^2} (r - \|\sin \Theta(U, U^*)\|_F^2) \right). \end{aligned} \quad (66)$$

From the proof of Lemma B.17,

$$\begin{aligned} &\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*, w}) \\ &= 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)(\Phi(\alpha_{\max, U^*} \omega^*/v^*) - \Phi(\alpha_{\max, U} \omega^*/v^*))\mathbb{I}\{\omega^* > 0\}]. \end{aligned}$$

Note that for any $b \geq a \geq 0$, $\Phi(b) - \Phi(a) \geq \phi(b)(b - a)$. Since we assume RHS of Equation (65) is positive, $\alpha_{\max, U^*} \geq \alpha_{\max, U}$. Thus on the event where $\omega^* > 0$, $\alpha_{\max, U^*} \omega^*/v^* \geq \alpha_{\max, U} \omega^*/v^*$. Observe

$$\begin{aligned} &\inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U, w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*, w}) \\ &\geq 2\mathbb{E}_{\mathcal{E}}[(2F(\omega^*) - 1)\phi(\alpha_{\max, U^*} \omega^*/v^*)(\alpha_{\max, U^*} \omega^*/v^* - \alpha_{\max, U} \omega^*/v^*)\mathbb{I}\{\omega^* > 0\}] \\ &= \frac{2}{v^*} (\alpha_{\max, U^*} - \alpha_{\max, U}) \int_0^\infty (2F(\omega^*) - 1) \omega^* \frac{\phi(\omega^*/v^*)}{v^*} \phi(\alpha_{\max, U^*} \omega^*/v^*) d\omega^* \\ &\simeq \frac{\alpha_{\max, U^*} - \alpha_{\max, U}}{v^*} \int_0^\infty (2F(\omega^*) - 1) \omega^* \exp\left(-\frac{1}{2}(1 + \alpha_{\max, U^*}^2) \omega^{*2}/v^{*2}\right) d\omega^* \\ &\simeq \frac{\alpha_{\max, U^*} - \alpha_{\max, U}}{1 + \alpha_{\max, U^*}^2} \int_0^\infty (2F((1 + \alpha_{\max, U^*}^2)^{-1/2} v^* \omega^*) - 1) \omega^* \exp\left(-\frac{1}{2} \omega^{*2}\right) d\omega^*, \end{aligned}$$

where in the last equality we transformed $w^* \rightarrow (1 + \alpha_{\max, U^*}^2)^{1/2} w^*/v^*$. Since $F(u)$ is differentiable at 0 and $F(0) = 1/2$,

$$F(u) - 1/2 = F'(0)u + o(u).$$

Thus there exists a constant $\epsilon > 0$ only depending on F such that $2(F(u) - 1/2) \geq F'(0)u$ for all $u \in [0, \epsilon]$ since $F'(0) > 0$. This gives

$$\begin{aligned}
 & \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\
 & \gtrsim \frac{\alpha_{\max,U^*} - \alpha_{\max,U}}{1 + \alpha_{\max,U^*}^2} F'(0) (1 + \alpha_{\max,U^*}^2)^{-1/2} v^* \\
 & \quad \times \int_0^{\epsilon(1 + \alpha_{\max,U^*}^2)^{1/2} v^*} \omega^{*2} \exp\left(-\frac{1}{2}\omega^{*2}\right) d\omega^* \\
 & \gtrsim \frac{\alpha_{\max,U^*} - \alpha_{\max,U}}{1 + \alpha_{\max,U^*}^2} (1 + \alpha_{\max,U^*}^2)^{-1/2} v^* \int_0^{\epsilon v^*} \omega^{*2} \exp\left(-\frac{1}{2}\omega^{*2}\right) d\omega^* \\
 & \gtrsim \frac{\alpha_{\max,U^*} - \alpha_{\max,U}}{1 + \alpha_{\max,U^*}^2} (1 + \alpha_{\max,U^*}^2)^{-1/2}.
 \end{aligned}$$

The last inequality follows since $v^* = \|w^*\| = 1$ by assumption. It is noted that $\alpha_{\max,U^*}^2 \leq \nu^2/\sigma_{(d)}^2$ from Equation (64). Therefore with Equation (66),

$$\begin{aligned}
 & \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_c(\delta_{U^*,w}) \\
 & \gtrsim \frac{1}{(1 + \kappa\rho^2)^{3/2}} (1 + \kappa^{-1}\rho^{-2})^{1/2} (1 + \rho^2)^{3/2} \left(\frac{1}{1 + \rho^{-2}} - \kappa\rho^2(r - \|\sin \Theta(U, U^*)\|_F^2) \right) \\
 & \gtrsim \frac{(1 + \rho^2)^{3/2}}{(1 + \kappa\rho^2)^{3/2}} \rho^2 \left(\frac{1}{1 + \rho^2} - \kappa(r - \|\sin \Theta(U, U^*)\|_F^2) \right).
 \end{aligned}$$

■

Proposition B.19 For any $U \in \mathbb{O}_{d,r}$,

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) = \nu^2 w^{*\top} (I - \nu^2 U^{*\top} U (\nu^2 U^\top U^* U^{*\top} U + U^\top \Sigma U)^{-1} U^\top U^*) w^* + \sigma_\epsilon^2.$$

Proof [Proof of Proposition B.19] Generate random variables $(\check{x}, \check{z}, \check{\xi}, \check{\epsilon})$ following the model (19). We calculate the prediction risk of $\delta_{U,w}$ as:

$$\begin{aligned}
 \mathcal{R}_r(\delta_{U,w}) & := \mathbb{E}_{\mathcal{E}} (\check{y} - \check{x}^\top U w)^2 \\
 & = \text{Var}_{\mathcal{E}} (\nu^{-1} \check{z}^\top w^* + \check{\epsilon})^2 - 2 \text{Cov}_{\mathcal{E}} (\nu^{-1} \check{z}^\top w^* + \check{\epsilon}, U^* \check{z} + \check{\xi}) U w \\
 & \quad + w^\top U^\top \text{Var}_{\mathcal{E}} (U^* \check{z} + \check{\xi}) U w \\
 & = \|w^*\|^2 + \sigma_\epsilon^2 - 2\nu w^{*\top} U^{*\top} U w + w^\top (\nu^2 U^\top U^* U^{*\top} U + U^\top \Sigma U) w \\
 & = (w - A^{-1}b)^\top A (w - A^{-1}b) - b^\top A^{-1}b + \|w^*\|^2 + \sigma_\epsilon^2,
 \end{aligned}$$

where $A := \nu^2 U^\top U^* U^{*\top} U + U^\top \Sigma U$ and $b := \nu U^\top U^* w^*$. From this, we obtain

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) = w^{*\top} \left(I - U^{*\top} U (U^\top U^* U^{*\top} U + (1/\nu^2) U^\top \Sigma U)^{-1} U^\top U^* \right) w^* + \sigma_\epsilon^2.$$

■

Lemma B.20 For any $U \in \mathbb{O}_{d,r}$,

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) = O((1 + \rho^{-2}) \mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2] \|w^*\|^2).$$

Proof [Proof of Lemma B.20] From proposition B.19, we have

$$\begin{aligned} & \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) \\ &= w^{*\top} \left((I + (1/\nu^2)U^{*\top} \Sigma U^*)^{-1} - U^{*\top} U (U^\top U^* U^{*\top} U + (1/\nu^2)U^\top \Sigma U)^{-1} U^\top U^* \right) w^*. \end{aligned}$$

Note that $\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) \equiv \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{UO,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w})$ for any orthogonal matrix $O \in \mathbb{O}_{r,r}$. Take $\tilde{O} \in \mathbb{O}_{r,r}$ such that $\|U\tilde{O} - U^*\|_2 \leq \sqrt{2} \|\sin \Theta(U, O)\|_2$ without loss of generality, since we can always take a sequence $(\tilde{O}_m)_{m \geq 1}$ such that $\|UO_m - U^*\|_2 \leq \sqrt{2} \|\sin \Theta(U, O)\|_2 + 1/m$ from Lemma A.3.

Lemma B.14 gives

$$\begin{aligned} & \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) \\ &= O \left(\frac{1}{1 - \|\sin \Theta(U, U^*)\|_2^2 + \kappa^{-1} \rho^{-2}} \frac{1 + \rho^{-2}}{1 + \kappa^{-1} \rho^{-2}} \|\sin \Theta(U, U^*)\|_2 \|w^*\|^2 \right). \end{aligned}$$

On the event where $\|\sin \Theta(U, U^*)\|_2^2 < 1/2$,

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) = O \left(\frac{1 + \rho^{-2}}{(1 + \kappa^{-1} \rho^{-2})^2} \|\sin \Theta(U, U^*)\|_2 \|w^*\|^2 \right).$$

On the event where $\|\sin \Theta(U, U^*)\|_2^2 \geq 1/2$, we utilize the trivial upper bound

$$\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w}) \leq \|(I + \nu^{-2}U^{*\top} \Sigma U^*)^{-1}\|_2 \|w^*\|^2 \leq \frac{\nu^2}{\nu^2 + \sigma_{(d)}^2} \|w^*\|^2.$$

Combining these results, we have

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}}[\inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U,w}) - \inf_{w \in \mathbb{R}^r} \mathcal{R}_r(\delta_{U^*,w})] \\ & \lesssim \frac{1 + \rho^{-2}}{(1 + \kappa^{-1} \rho^{-2})^2} \mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2] \|w^*\|^2 \\ & \quad + \frac{1}{1 + \kappa^{-1} \rho^{-2}} \|w^*\|^2 \mathbb{P}_{\mathcal{D}}(\|\sin \Theta(U, U^*)\|_2 \geq 1/\sqrt{2}) \\ & \lesssim \frac{1 + \rho^{-2}}{(1 + \kappa^{-1} \rho^{-2})^2} \mathbb{E}_{\mathcal{D}}[\|\sin \Theta(U, U^*)\|_2] \|w^*\|^2, \end{aligned}$$

where the last inequality follows by Markov's inequality. ■

Appendix C. Discussion about Autoencoders and random masking augmentation

In the following, we show that our results do not change if we applied the same augmentation (2) for autoencoders. As discussed in Section 3.1, we can ignore the bias term in autoencoders for simplicity, which only serves as centralization of the data matrix. In that case, we applied random augmentation $g_1(x) = Ax$ and $g_2(x) = (I - A)x$ to the original data $\{x_i\}_{i=1}^n$, and the optimization problem can be formulated as follows:

$$\min_{W_{AE}, W_{DE}} \frac{1}{2n} \mathbb{E}_A [\|AX - W_{DE}W_{AE}AX\|_F^2 + \|(I - A)X - W_{DE}W_{AE}(I - A)X\|_F^2]. \quad (67)$$

Then, similar to Theorem 5 for contrastive learning, we can also obtain an explicit solution for this optimization problem.

Theorem C.1 *The optimal solution of autoencoders with random masking augmentation (67) is given by:*

$$W_{AE} = W_{DE}^\top = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$\frac{1}{2} \Delta(XX^\top) + D(XX^\top), \quad (68)$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

Proof We first derive the equivalent form for this objective function:

$$\begin{aligned} & \frac{1}{2n} \mathbb{E}_A [\|AX - W_{DE}W_{AE}AX\|_F^2 + \|(I - A)X - W_{DE}W_{AE}(I - A)X\|_F^2] \\ &= \frac{1}{2n} \mathbb{E}_A [\text{tr}(X^\top A^\top AX) + \text{tr}(X^\top A^\top W_{DE}W_{AE}AX) + \text{tr}(X^\top A^\top W_{AE}^\top W_{DE}^\top W_{DE}W_{AE}AX) \\ & \quad + \text{tr}(X^\top (I - A)^\top (I - A)X) + \text{tr}(X^\top (I - A)^\top W_{DE}W_{AE}(I - A)X) \\ & \quad + \text{tr}(X^\top (I - A)^\top W_{AE}^\top W_{DE}^\top W_{DE}W_{AE}(I - A)X)] \\ &= \frac{1}{2n} \mathbb{E}_A [\text{tr}(X^\top AX) + \text{tr}(AXX^\top A^\top W_{DE}W_{AE}) + \text{tr}(AXX^\top A^\top W_{AE}^\top W_{DE}^\top W_{DE}W_{AE}) \\ & \quad + \text{tr}(X^\top (I - A)X) + \text{tr}((I - A)XX^\top (I - A)^\top W_{DE}W_{AE}) \\ & \quad + \text{tr}((I - A)XX^\top (I - A)^\top W_{AE}^\top W_{DE}^\top W_{DE}W_{AE})] \\ &= \frac{1}{2n} \mathbb{E}_A [\text{tr}(X^\top X) + \text{tr}(\hat{M}W_{DE}W_{AE}) + \text{tr}(\hat{M}W_{AE}^\top W_{DE}^\top W_{DE}W_{AE})], \end{aligned} \quad (69)$$

where $\hat{M} := AXX^\top A^\top + (I - A)XX^\top (I - A)^\top$. Note that by Definition 2 we have $A = \text{diag}(a_1, \dots, a_d)$ and a_i follows the Bernoulli distribution, so we have:

$$\mathbb{E}_A \hat{M} = \frac{1}{2} \Delta(XX^\top) + D(XX^\top) \quad (70)$$

Again, by Theorem 2.4.8 in Golub and Loan (1996), the optimal solution of Eq.(67) is given by the eigenvalue decomposition of $\mathbb{E}_A \hat{M} = \frac{1}{2}\Delta(XX^\top) + D(XX^\top)$, up to an orthogonal transformation, which finishes the proof. \blacksquare

With Theorem C.1 established, we can now derive the space distance for autoencoders with random masking augmentation.

Theorem C.2 *Consider the spiked covariance model Eq.(5), under Assumptions 8-10 and $n > d \gg r$, let W_{AE} be the learned representation of augmented autoencoder with singular value decomposition $W_{AE} = (U_{AE}\Sigma_{AE}V_{AE}^\top)^\top$ (i.e., the optimal solution of optimization problem 67). If we further assume $\{\sigma_i^2\}_{i=1}^d$ are different from each other and $\sigma_{(1)}^2/(\sigma_{(r)}^2 - \sigma_{(r+1)}^2) < C_\sigma$ for some universal constant C_σ . Then there exist two universal constants $C_\rho > 0, c \in (0, 1)$, such that when $\rho < C_\rho$, we have*

$$\mathbb{E} \|\sin \Theta(U^*, U_{AE})\|_F \geq c\sqrt{r}. \quad (71)$$

Proof Step1, similar to the proof of Theorem B.7, we first bound the difference between $\hat{M} := \Delta(XX^\top) + 2D(XX^\top)$ and $\Sigma := \text{Cov}(\xi\xi^\top)$. Note that:

$$\|\hat{M} - \Sigma\|_2 = \|XX^\top - \Sigma - \frac{1}{2}\Delta(XX^\top)\|_2 \leq \|XX^\top - \Sigma\|_2 + \frac{1}{2}\|\Delta(XX^\top - \Sigma)\|_2 + \frac{1}{2}\|\Delta(\Sigma)\|_2 \quad (72)$$

Since Σ is a diagonal matrix, then by Lemma B.8 we have:

$$\|\hat{M} - \Sigma\|_2 \leq 2\|XX^\top - \Sigma\|_2 \quad (73)$$

Now, directly apply equation (41)(42)(43) we can obtain that:

$$\mathbb{E}\|\hat{M} - \Sigma\|_2 \lesssim \nu^2 \left(1 + \sqrt{\frac{r}{n}} + \frac{r}{n}\right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right) + \sqrt{\frac{d}{n}}\sigma_{(1)}\nu. \quad (74)$$

Step 2, bound the $\sin \Theta$ distance between eigenspaces. As we have shown in step 1, the target matrix of the autoencoder is close to the covariance matrix of random noise, i.e., Σ . Note that Σ is assumed to be a diagonal matrix with different elements, hence its eigenspace only consists of canonical basis e_i . Denote U_Σ to be the top- r eigenspace of Σ and $\{e_i\}_{i \in C}$ to be its corresponding basis vectors, apply the Davis-Kahan Theorem E.1 we can conclude that:

$$\begin{aligned} \mathbb{E} \|\sin \Theta(U_{AE}, U_\Sigma)\|_F &\leq \frac{2\sqrt{r}\mathbb{E}\|\hat{M} - \Sigma\|_2}{\sigma_{(r)}^2 - \sigma_{(r+1)}^2} \\ &\lesssim \sqrt{r} \frac{1}{\sigma_{(1)}^2} \left(\nu^2 \left(1 + \sqrt{\frac{r}{n}} + \frac{r}{n}\right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right) + \sqrt{\frac{d}{n}}\sigma_{(1)}\nu \right) \\ &\lesssim \sqrt{r} \left(\rho^2 + \sqrt{\frac{d}{n}} + \rho\sqrt{\frac{d}{n}} \right). \end{aligned}$$

Step 3, obtain the final result by triangular inequality. By Assumption 10 we know that the distance between canonical basis and the eigenspace of core features can be large:

$$\begin{aligned} \|\sin \Theta(U^*, U_\Sigma)\|_F^2 &= \|U_{\Sigma^\perp}^\top U^*\|_F^2 = \sum_{i \in [d]/C} \|e_i^\top U^*\|^2 = \|U^*\|_F^2 - \sum_{i \in C} \|e_i^\top U^*\|^2 \\ &\geq r - rI(U^*) = r - O\left(\frac{r^2}{d} \log d\right). \end{aligned} \quad (75)$$

Then apply the triangular inequality of $\sin \Theta$ distance (Proposition A.5) we can obtain the lower bound of the autoencoder.

$$\begin{aligned} \mathbb{E} \|\sin \Theta(U_{AE}, U^*)\|_F &\geq \mathbb{E} \|\sin \Theta(U^*, U_\Sigma)\|_F - \mathbb{E} \|\sin \Theta(U_{AE}, U_\Sigma)\|_F \\ &\geq \sqrt{r} - O\left(\frac{r}{\sqrt{d}} \sqrt{\log d}\right) - O\left(\sqrt{r} \left(\rho^2 + \sqrt{\frac{d}{n}} + \rho \sqrt{\frac{d}{n}}\right)\right). \end{aligned}$$

By Assumption 9, it implies that when n and d are sufficiently large and ρ is sufficiently small (smaller than a given constant $C_\rho > 0$), there exists a universal constant $c \in (0, 1)$ such that:

$$\mathbb{E} \|\sin \Theta(U_{AE}, U^*)\|_F \geq c\sqrt{r}. \quad \blacksquare$$

Compared with Theorem 13, we can find that random masking augmentation makes no difference to autoencoders, which justifies the fairness of our comparison between contrastive learning and autoencoders.

However, contrary to the autoencoders with random-masking augmentation, we show that the representations obtained by DAEs behave as the representations obtained by contrastive learning.

Proof [Proof of Remark 16] Let $\mathcal{L} := (1/n) \|X - W^\top W A X\|_F^2$ be the loss function of DAEs. Then,

$$E_A \mathcal{L} = \frac{1}{n} \operatorname{tr} \left(W^\top W \left(\frac{1}{2} D(XX^\top) + \frac{1}{4} \Delta(XX^\top) \right) W^\top W - W^\top W X X^\top \right) + (\text{const.}). \quad (76)$$

We minimize the loss over W such that $W W^\top = 2I_r$. Then, the loss becomes

$$\begin{aligned} \arg \min_{W W^\top = 2I_r} E_A \mathcal{L} &= \arg \min_{W W^\top = 2I_r} \frac{1}{n} \operatorname{tr} \left(W \left(D(XX^\top) + \frac{1}{2} \Delta(XX^\top) \right) W^\top - W X X^\top W^\top \right) \\ &= \arg \max_{W W^\top = 2I_r} \operatorname{tr} \left(W \frac{1}{n} \Delta(XX^\top) W^\top \right). \end{aligned}$$

Thus, the solution to the (expected) loss minimization problem is the top- r eigenvectors of $\Delta(n^{-1} X X^\top)$, i.e., $W^\top = \sqrt{2} O P_r(\Delta(n^{-1} X X^\top))$, where O is any orthogonal matrix from $\mathcal{O}_{r,r}$. We use the same argument as in the proof of Theorem B.9. First note that

$$\frac{1}{n} \Delta(X X^\top) = \frac{1}{n} \Delta(U^* Z Z^\top U^{*\top}) + \frac{1}{n} \Delta(U^* Z E + E Z^\top U^{*\top}) + \frac{1}{n} \Delta(E E^\top).$$

By Lemmas B.8, E.3 and the incoherent condition $I(U) = O(\frac{r}{d} \log d)$, we have:

$$\begin{aligned} &\mathbb{E} \left\| \Delta \left(\frac{1}{n} U^* Z Z^\top U^{*\top} \right) - \nu^2 U^* U^{*\top} \right\|_2 \\ &\leq 2 \mathbb{E} \left\| \frac{1}{n} U^* Z Z^\top U^{*\top} - \nu^2 U^* U^{*\top} \right\|_2 + \mathbb{E} \left\| \Delta \left(\nu^2 U^* U^{*\top} \right) - \nu^2 U^* U^{*\top} \right\|_2 \\ &\lesssim 2 \left(\sqrt{\frac{r}{n}} + \frac{r}{n} \right) \nu^2 + \frac{r}{d} \log d \nu^2. \end{aligned} \quad (77)$$

For the second term, applying equation (42) yields:

$$\frac{1}{n} \mathbb{E} \|\Delta(U^* Z E^\top + E Z^\top U^{*\top})\|_2 \leq \frac{4}{n} \mathbb{E} \|E Z^\top U^{*\top}\|_2 \lesssim \frac{\sqrt{d}}{\sqrt{n}} \sigma_{(1)} \nu. \quad (78)$$

For the third term, applying equation (43) yields:

$$\mathbb{E} \left\| \frac{1}{n} \Delta(E E^\top) \right\|_2 = \mathbb{E} \left\| \Delta\left(\frac{1}{n} E E^\top - \Sigma\right) \right\|_2 \leq 2 \left\| \frac{1}{n} E E^\top - \Sigma \right\|_2 \lesssim \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right) \sigma_{(1)}^2. \quad (79)$$

Combining equations (77)(78)(79) gives

$$\mathbb{E} \left\| \Delta\left(\frac{1}{n} X X^\top - \nu^2 U^* U^{*\top}\right) \right\|_2 \lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n}} + \frac{r}{n}\right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n}} + \frac{d}{n}\right) + \sigma_{(1)} \nu \sqrt{\frac{d}{n}}.$$

From Lemma E.1, we obtain the desired bound:

$$\mathbb{E} \|\sin \Theta(U_{\text{DAE}}, U^*)\|_F \leq \frac{2\sqrt{r}}{\nu^2} \mathbb{E} \left\| \Delta\left(\frac{1}{n} X X^\top - \nu^2 U^* U^{*\top}\right) \right\|_2 \lesssim \frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}}. \quad \blacksquare$$

Here we provide some experimental results about DAEs on synthetic datasets as analog to Figure 1 and 2, the settings are the same as described in Section 5.1. The results are summarized in Figure 3, as we can observe, the performance of DAEs is comparable with contrastive learning, which aligns with our theoretical results above.

Appendix D. Omitted proofs for Section 4

D.1 Proofs for Section 4.1

In this section, we will provide the proof of a generalized version of Theorem 20 to cover the imbalanced setting, the statement and the detailed proof can be found in Theorem D.2.

In the main body, we assume the unlabeled data and labeled data are both balanced for the sake of clarity and simplicity. Now we allow them to be imbalanced and provide a more general analysis. Suppose we have n unlabeled data $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$ and n_k labeled data $X_k = [x_k^1, \dots, x_k^{n_k}] \in \mathbb{R}^{d \times n_k}$ for class k , the contrastive learning task can be formulated as:

$$\min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}(W) := \min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}_{\text{SelfCon}}(W) + \mathcal{L}_{\text{SupCon}}(W; \alpha). \quad (80)$$

In addition, we write a generalized version of the supervised contrastive loss function to cover the imbalanced cases:

$$\mathcal{L}_{\text{SupCon}}(W; \alpha) = -\frac{1}{r+1} \sum_{k=1}^{r+1} \frac{\alpha_k}{n_k} \sum_{i=1}^{n_k} \left[\sum_{j \neq i} \frac{\langle W x_i^k, W x_j^k \rangle}{n_k - 1} - \frac{\sum_{j=1}^n \sum_{s \neq k} \langle W x_i^k, W x_j^s \rangle}{\sum_{s \neq k} n_s} \right] + \frac{\lambda}{2} \|W W^\top\|_F^2, \quad (81)$$

where $\alpha_k > 0$ is the weight for supervised loss of class k . Again we first provide a theorem to give the optimal solution to the contrastive learning problem.

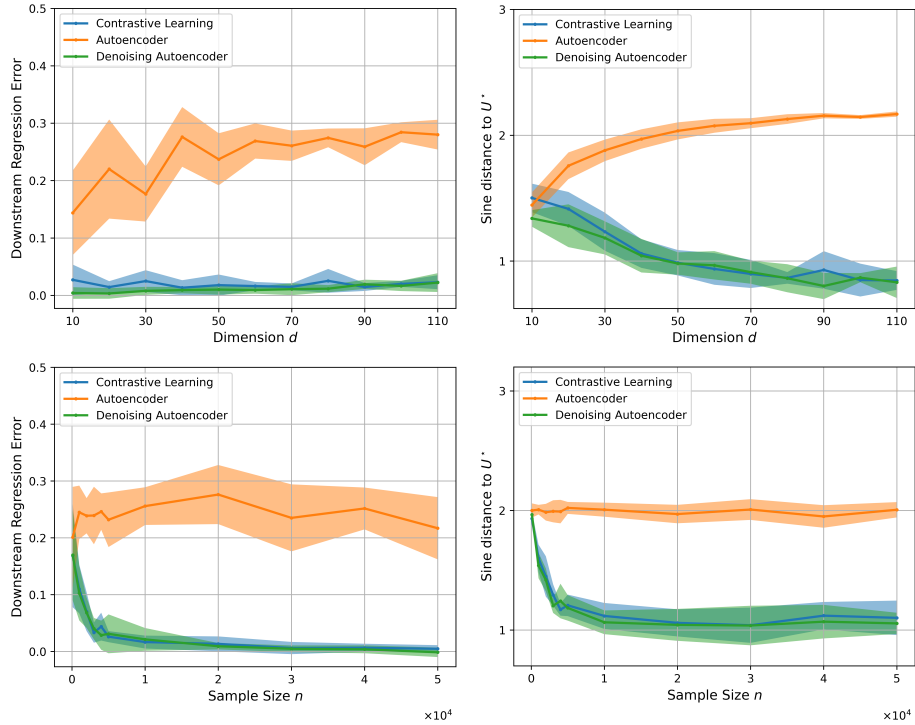


Figure 3: Comparison of denoising autoencoders, autoencoders, and contrastive learning on synthetic datasets. **Left Column:** The vertical axes indicate the downstream regression error. We subtract the regression error of the ground truth features to measure the excess error. **Top Row:** Comparison of in-domain downstream task performance of autoencoders, contrastive learning, and denoising autoencoders the dimension d . The sample size n is set as $n = 20000$. **Bottom Row:** Comparison of in-domain downstream task performance of autoencoders, contrastive learning, and denoising autoencoders the dimension n . The dimension d is set as $d = 40$.

Theorem D.1 *The optimal solution of the supervised contrastive learning problem (80) is given by :*

$$W_{CL} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a positive constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$\begin{aligned} & \frac{1}{4n} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \\ & + \frac{1}{r+1} \sum_{k=1}^{r+1} \frac{\alpha_k}{n_k} \left[\frac{1}{n_k-1} X_k(1_{n_k} 1_{n_k}^\top - I_{n_k}) X_k^\top - \frac{1}{\sum_{t \neq k} n_t} X_k 1_k 1_s^\top X_s^\top \right], \end{aligned}$$

u_i is the corresponding eigenvector and $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthonormal matrix.

Proof Under this setting, combined with the result obtained in Corollary 6, the contrastive loss can be rewritten as:

$$\begin{aligned} \mathcal{L}(W) &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{2n} \operatorname{tr} \left(\left(\frac{1}{2} \Delta(XX^\top) - \frac{1}{2(n-1)} X(1_n 1_n^\top - I_n) X^\top \right) W^\top W \right) \\ & - \frac{1}{r+1} \sum_{k=1}^{r+1} \frac{\alpha_k}{n_k} \sum_{i=1}^{n_k} \left[\frac{1}{n_k-1} \sum_{j \neq i} \langle W x_i^k, W x_j^k \rangle - \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \sum_{j=1}^{n_s} \langle W x_i^k, W x_j^s \rangle \right]. \end{aligned}$$

Then we deal with the last term independently, note that:

$$\begin{aligned} & \sum_{i=1}^{n_k} \left[\frac{1}{n_k-1} \sum_{j \neq i} \langle W x_i^k, W x_j^k \rangle - \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \sum_{j=1}^{n_s} \langle W x_i^k, W x_j^s \rangle \right] \\ &= \frac{1}{n_k-1} \sum_{i=1}^{n_k} \sum_{j \neq i} \langle W x_i^k, W x_j^k \rangle - \frac{1}{\sum_{t \neq k} n_t} \sum_{i=1}^{n_k} \sum_{s \neq k} \sum_{j=1}^{n_s} \langle W x_i^k, W x_j^s \rangle \\ &= \frac{1}{n_k-1} \operatorname{tr} \left(X_k(1_{n_k} 1_{n_k}^\top - I_{n_k}) X_k^\top W^\top W \right) - \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \operatorname{tr} \left(X_k 1_k 1_s^\top X_s^\top W^\top W \right). \end{aligned}$$

Thus we have:

$$\begin{aligned} \mathcal{L}(W) &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{4n} \operatorname{tr} \left((\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top) W^\top W \right) \\ & - \frac{1}{r+1} \sum_{k=1}^{r+1} \frac{\alpha_k}{n_k} \left[\frac{1}{n_k-1} \operatorname{tr} \left(X_k(1_{n_k} 1_{n_k}^\top - I_{n_k}) X_k^\top W^\top W \right) \right. \\ & \left. - \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \operatorname{tr} \left(X_k 1_k 1_s^\top X_s^\top W^\top W \right) \right]. \end{aligned}$$

Then by a similar argument as in the proof of Proposition 5, we can conclude that the optimal solution W_{CL} must satisfy the desired conditions. \blacksquare

With the optimal solution obtained in Theorem D.1, we can provide a generalized version of Theorem 20 to cover the imbalance cases.

Theorem D.2 (Generalized version of Theorem 20) *If Assumptions 8-10 hold, $n > d \gg r$ and let W_{CL} be any solution that minimizes the supervised contrastive learning problem in Equation (80), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have*

$$\begin{aligned} \mathbb{E}\|\sin \Theta(U_{CL}, U)\|_F &\lesssim \frac{\nu^2}{\lambda_r(T)} \left(\frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}} \right. \\ &\quad \left. + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k \left[\sum_{s \neq k} \frac{\sqrt{n_s d}}{\sum_{t \neq k} n_t} \left(\sqrt{\frac{d}{n_k}} + \sqrt{r} \right) + \sqrt{\frac{dr}{n_k}} \right] \right), \end{aligned}$$

where $T \triangleq \frac{1}{4} \sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top} + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k (\mu^k \mu^{k\top} - \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \frac{1}{2} (\mu^k \mu^{s\top} + \mu^s \mu^{k\top}))$.

Proof [Proof of Theorem D.2] For labeled data $X = [x_1, \dots, x_n]$, we write it to be $X = M + E$, where $M = [\mu_1, \dots, \mu_n]$ and $E = [\xi_1, \dots, \xi_n]$ are two matrices consisting of class mean and random noise. To be more specific, if x_i subject to the k -th cluster, then $\mu_i = \mu^k$ and $\xi_i \sim \mathcal{N}(0, \Sigma^k)$. Since the data is randomly drawn from each class, μ_i follows the multinomial distribution over μ^1, \dots, μ^r with probability p_1, \dots, p_{r+1} . Thus μ_i follows a subgaussian distribution with covariance matrix $N = \sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top}$.

As shown in Theorem D.1, the optimal solution of contrastive learning is equivalent to PCA of the following matrix:

$$\begin{aligned} \hat{T} &\triangleq \frac{1}{4n} (\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n)X^\top) \\ &\quad + \frac{1}{r+1} \sum_{k=1}^{r+1} \frac{\alpha_k}{n_k} \left[\frac{1}{n_k-1} X_k(1_{n_k} 1_{n_k}^\top - I_{n_k})X_k^\top \right. \\ &\quad \left. - \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \frac{1}{2} (X_k 1_k 1_s^\top X_s^\top + X_s 1_s 1_k^\top X_k^\top) \right]. \end{aligned}$$

Again we will deal with these terms separately,

1. For the first term, as we have discussed, X can be divided into two matrices M and E , each of them consisting of sub-gaussian columns. Again we can obtain the result as in (56) (the proof is totally the same):

$$\mathbb{E}\left\| \frac{1}{n} (\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n)X^\top) - N \right\|_2 \lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n}} \right) + \sigma^2(1) \sqrt{\frac{d}{n}}. \quad (82)$$

2. For the second term, notice that:

$$\begin{aligned} X_k(1_{n_k} 1_{n_k}^\top - I_{n_k})X_k^\top &= \sum_{i=1}^{n_k} \sum_{j \neq i}^{n_k} (\mu^k + \xi_i^k)(\mu^k + \xi_j^k)^\top \\ &= n_k(n_k - 1)\mu^k \mu^{k\top} + (n_k - 1)\mu^k \left(\sum_{i=1}^{n_k} \xi_i^k \right)^\top + (n_k - 1) \left(\sum_{i=1}^{n_k} \xi_i^k \right) \mu^{k\top} + \sum_{i=1}^{n_k} \sum_{j \neq i}^{n_k} \xi_i^k \xi_j^{k\top}, \end{aligned} \quad (83)$$

and that:

$$\begin{aligned} \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} X_k 1_k 1_s^\top X_s^\top &= \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \sum_{i=1}^{n_k} (\mu^k + \xi_i^k) \sum_{j=1}^{n_s} (\mu^s + \xi_j^s)^\top \\ &= \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} [n_k n_s \mu^k \mu^{s\top} + n_k \mu^k (\sum_{j=1}^{n_s} \xi_j^s)^\top + n_s \sum_{i=1}^{n_k} \xi_i^k \mu^{s\top} + \sum_{i=1}^{n_k} \xi_i^k \sum_{j=1}^{n_s} \xi_j^{s\top}]. \end{aligned} \quad (84)$$

Since $\xi_i^k \sim \mathcal{N}(0, \Sigma^k)$, we can conclude that:

$$\mathbb{E} \left\| \frac{1}{n_k} \sum_{i=1}^{n_k} \xi_i^k \right\|_2 \leq \sqrt{\mathbb{E} \left\| \frac{1}{n_k} \sum_{i=1}^{n_k} \xi_i^k \right\|_2^2} = \sqrt{\frac{d}{n_k}} \sigma_{(1)}. \quad (85)$$

Moreover, we have

$$\begin{aligned} \frac{1}{n_k(n_k-1)} \mathbb{E} \left\| \sum_{i=1}^{n_k} \sum_{j \neq i} \xi_i^k \xi_j^{k\top} \right\|_2 &\leq \frac{1}{n_k(n_k-1)} \mathbb{E} \|E_k E_k^\top\|_2 + \frac{n_k}{n_k-1} \mathbb{E} \|\bar{\xi}^k \bar{\xi}^{k\top}\|_2 \\ &\lesssim \frac{d}{n_k} \sigma_{(1)}^2. \end{aligned} \quad (86)$$

Take equation (85) and (86) back into (83) we can conclude:

$$\mathbb{E} \left\| \frac{1}{n_k(n_k-1)} X_k (1_{n_k} 1_{n_k}^\top - I_{n_k}) X_k^\top - \mu^k \mu^{k\top} \right\|_2 \lesssim \sqrt{\frac{d}{n_k}} \sigma_{(1)} \sqrt{r\nu} + \frac{d}{n_k} \sigma_{(1)}^2. \quad (87)$$

On the other hand, by equation (85) we know:

$$\mathbb{E} \left\| \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} \sum_{j=1}^{n_s} \xi_j^s \right\|_2 \leq \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \mathbb{E} \left\| \frac{1}{n_s} \sum_{i=1}^{n_s} \xi_i^s \right\|_2 \lesssim \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \sqrt{\frac{d}{n_s}} \sigma_{(1)}. \quad (88)$$

Notice that:

$$\begin{aligned} \mathbb{E} \left\| \frac{1}{\sum_{t \neq k} n_t} \frac{1}{n_k} \sum_{s \neq k} \sum_{i=1}^{n_k} \xi_i^k \sum_{j=1}^{n_s} \xi_j^{s\top} \right\|_2 &\leq \mathbb{E} \left\| \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \bar{\xi}^k \bar{\xi}^{s\top} \right\|_2 \\ &\leq \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \mathbb{E} \|\bar{\xi}^k \bar{\xi}^{s\top}\|_2 \lesssim \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \frac{d}{\sqrt{n_k n_s}} \sigma_{(1)}^2. \end{aligned} \quad (89)$$

Thus take equations (88) and (89) back into equation (84) we have:

$$\mathbb{E} \left\| \frac{1}{n_k} \frac{1}{\sum_{t \neq k} n_t} \sum_{s \neq k} X_k 1_k 1_s^\top X_s^\top - \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \mu^k \mu^{s\top} \right\|_2 \quad (90)$$

$$\lesssim \sum_{s \neq k} \frac{\sqrt{n_s d}}{\sum_{t \neq k} n_t} \left(\sqrt{\frac{d}{n_k}} \sigma_{(1)}^2 + \sigma_{(1)} \sqrt{r\nu} \right). \quad (91)$$

Then combine equations (82)(87)(90) together, we can obtain the following result:

$$\begin{aligned} & \mathbb{E} \left\| \hat{T} - \frac{1}{4}N - \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k (\mu^k \mu^{k\top} - \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \frac{1}{2} (\mu^k \mu^{s\top} + \mu^s \mu^{k\top})) \right\|_2 \\ & \lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n}} \right) + \sigma_{(1)}^2 \sqrt{\frac{d}{n}} \\ & \quad + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k \left[\sum_{s \neq k} \frac{\sqrt{n_s d}}{\sum_{t \neq k} n_t} \left(\sqrt{\frac{d}{n_k}} \sigma_{(1)}^2 + \sqrt{r} \sigma_{(1)} \nu \right) + \sqrt{\frac{d}{n_k}} \sigma_{(1)} \sqrt{r} \nu + \frac{d}{n_k} \sigma_{(1)}^2 \right]. \end{aligned}$$

Since we have assumed that $\text{rank}(\sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top}) = r$ we can find that the top- r eigenspace of matrix:

$$T = \frac{1}{4} \sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top} + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k \left(\mu^k \mu^{k\top} - \sum_{s \neq k} \frac{n_s}{\sum_{t \neq k} n_t} \frac{1}{2} (\mu^k \mu^{s\top} + \mu^s \mu^{k\top}) \right)$$

is spanned by U^* , then apply Lemma E.1 again we have:

$$\begin{aligned} & \mathbb{E} \left\| \sin \Theta(U_{\text{SCL}}, U) \right\|_F \leq \frac{2\sqrt{r} \mathbb{E} \|\hat{N} - N\|_2}{\lambda_r(N)} \\ & \lesssim \frac{\sqrt{r}}{\lambda_r(T)} \left[\nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n}} \right) + \sigma_{(1)}^2 \sqrt{\frac{d}{n}} \right. \\ & \quad \left. + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k \left[\sum_{s \neq k} \frac{\sqrt{n_s d}}{\sum_{t \neq k} n_t} \left(\sqrt{\frac{d}{n_k}} \sigma_{(1)}^2 + \sqrt{r} \sigma_{(1)} \nu \right) + \sqrt{\frac{d}{n_k}} \sqrt{r} \sigma_{(1)} \nu + \frac{d}{n_k} \sigma_{(1)}^2 \right] \right] \\ & \lesssim \frac{\nu^2}{\lambda_r(T)} \left(\frac{r^{3/2}}{d} \log d + \sqrt{\frac{dr}{n}} + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha_k \left[\sum_{s \neq k} \frac{\sqrt{n_s d}}{\sum_{t \neq k} n_t} \left(\sqrt{\frac{d}{n_k}} + \sqrt{r} \right) + \sqrt{\frac{dr}{n_k}} \right] \right). \end{aligned}$$

■

Now we use this result to derive Theorem 20. Since $\|\mu^k\| = O(\sqrt{r}\nu)$ and $\sum_{k=1}^{r+1} p_k \mu^k = 0$, approximately we have $\frac{\nu^2}{\lambda_r(N)} \approx \frac{1}{\min_{k \in [r]} [1 + \alpha_k]}$. Although we can not obtain the closed-form eigenvalue in general, in a special case, where $\alpha = \alpha_1 = \dots = \alpha_{r+1}$, $m = n_1 = n_2 = \dots = n_{r+1}$ and $\frac{1}{r+1} = p_1 = p_2 = \dots = p_{r+1}$, it is easy to find that:

$$\sum_{s \neq k} \frac{1}{2} (\mu^k \mu^{s\top} + \mu^s \mu^{k\top}) = -\mu^k \mu^{k\top},$$

which further implies that:

$$T = \frac{1}{4} \sum_{k=1}^{r+1} p_k \mu^k \mu^{k\top} + \frac{1}{r+1} \sum_{k=1}^{r+1} \alpha \left(1 + \frac{1}{r} \right) \mu^k \mu^{k\top}, \quad \lambda_r(T) = \left[\frac{1}{4} + \alpha \left(1 + \frac{1}{r} \right) \right] \lambda(N).$$

and we can obtain the result in Theorem 20.

D.2 Proofs for Section 4.2

In this section, we will provide the proof of the generalized version of Theorems 23 and 26 to cover the imbalanced setting, the statement and detailed proof can be found in Theorems D.6 and D.8. With the two generalized theorems proven, Theorems 23, 24, 26, 27 holds immediately.

First, we prove a useful lemma to illustrate that the supervised loss function only yields estimation along a 1-dimensional space. Consider a single source task, where the data $x = U^*z + \xi$ is generated by the spiked covariance model and the label is generated by

$$y = \langle w^*, z \rangle / \nu$$

suppose we have collect n labeled data from this task, denote the data as $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{d \times n}$ and the label $y = [y_1, y_2, \dots, y_n] \in \mathbb{R}^n$, then we have the following result.

Lemma D.3 *Under the conditions similar to Theorem 14, we can find an event A such that $\mathbb{P}(A^C) = O(\sqrt{d/n})$ and:*

$$\mathbb{E} \left[\left\| \frac{1}{(n-1)^2} XHyy^\top HX^\top - \nu^2 U^* w^* w^{*\top} U^{*\top} \right\|_F \mathbb{I}\{A\} \right] \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)} \nu. \quad (92)$$

The proof strategy is to estimate the difference between the two rank-1 matrices via bounding the difference of the corresponding vector component. We first provide a simple lemma to illustrate the technique:

Lemma D.4 *Suppose $\alpha, \beta \in \mathbb{R}^d$ are two vectors, then we have:*

$$\|\alpha\alpha^\top - \beta\beta^\top\|_F \leq \sqrt{2}(\|\alpha\|_2 + \|\beta\|_2)\|\alpha - \beta\|_2.$$

Proof Denote $\alpha = (\alpha_1, \dots, \alpha_d), \beta = (\beta_1, \dots, \beta_d)$, then we have:

$$\begin{aligned} \|\alpha\alpha^\top - \beta\beta^\top\|_F^2 &\leq \sum_{i=1}^d \sum_{j=1}^d |\alpha_i\alpha_j - \beta_i\beta_j|^2 \leq 2 \sum_{i=1}^d \sum_{j=1}^d |\alpha_i\alpha_j - \alpha_i\beta_j|^2 + |\alpha_i\beta_j - \beta_i\beta_j|^2 \\ &\leq 2 \sum_{i=1}^d \sum_{j=1}^d |\alpha_i|^2 |\alpha_j - \beta_j|^2 + |\beta_j|^2 |\alpha_i - \beta_i|^2 \leq 2(\|\alpha\|_2^2 + \|\beta\|_2^2) \|\alpha - \beta\|_2^2 \\ &\leq 2(\|\alpha\|_2 + \|\beta\|_2)^2 \|\alpha - \beta\|_2^2. \end{aligned}$$

Take square root on both sides we can finish the proof. ■

Now we can prove the Lemma D.3.

Proof [Proof of Lemma D.3] Clearly, we have:

$$\begin{aligned} &\left\| \frac{1}{(n-1)^2} XHyy^\top HX^\top - \nu^2 U^* w^* w^{*\top} U^{*\top} \right\|_F \\ &\leq \frac{n^2}{(n-1)^2} \left\| \frac{1}{n^2} XHyy^\top HX^\top - \nu^2 U^* w^* w^{*\top} U^{*\top} \right\|_F + \frac{2n+1}{(n-1)^2} \|\nu^2 U^* w^* w^{*\top} U^{*\top}\|_F \quad (93) \\ &\lesssim \left\| \frac{1}{n^2} XHyy^\top HX^\top - \nu^2 U^* w^* w^{*\top} U^{*\top} \right\|_F + \frac{r}{n} \nu^2, \end{aligned}$$

thus we can replace the $\frac{1}{(n-1)^2}$ with $\frac{1}{n}$ in equation (92) and conclude the proof. Denote $\hat{N} \triangleq \frac{1}{n^2} XHy y^\top HX^\top$, note that both of \hat{N} and $Uw^*w^{*\top}U^\top$ are rank-1 matrices. We first bound the difference between $\frac{1}{n}XHy$ and Uw^* :

$$\begin{aligned} \left\| \frac{1}{n}XHy - \nu U^*w^* \right\| &= \left\| \frac{1}{n\nu} (U^*Z + E)HZ^\top w^* - \nu U^*w^* \right\| \\ &\leq \left\| \frac{1}{n\nu} (U^*Z + E)HZ^\top - \nu U^* \right\|_2 \\ &\leq \frac{1}{\nu} \left(\left\| \frac{1}{n}U^*ZZ^\top - \nu^2 U^* \right\|_2 + \frac{1}{n} \|EZ^\top\|_2 + \frac{1}{n} \|U^*Z\bar{Z}^\top\|_2 + \frac{1}{n} \|E\bar{Z}^\top\|_2 \right). \end{aligned} \quad (94)$$

We deal with the four terms in (94) separately:

1. For the first term, apply Lemma E.3 we have:

$$\mathbb{E} \left\| \frac{1}{n}U^*ZZ^\top - \nu^2 U^* \right\|_2 \leq \mathbb{E} \left\| \frac{1}{n}ZZ^\top - \nu^2 I_r \right\|_2 \leq \left(\frac{r}{n} + \sqrt{\frac{r}{n}} \right) \nu^2. \quad (95)$$

2. For the second term, apply Lemma E.2 twice we have:

$$\begin{aligned} \frac{1}{n} \mathbb{E} \|EZ^\top\|_2 &= \frac{1}{n} \mathbb{E}_Z [\mathbb{E}_E [\|EZ^\top\|_2 | Z]] \\ &\lesssim \frac{1}{n} \mathbb{E}_Z [\|Z\|_2 (\sigma_{\text{sum}} + r^{1/4} \sqrt{\sigma_{\text{sum}} \sigma_{(1)}} + \sqrt{r} \sigma_{(1)})] \\ &\lesssim \frac{1}{n} \mathbb{E}_Z [\|Z\|_2] \sqrt{d} \sigma_{(1)} \\ &\lesssim \frac{1}{n} \sqrt{d} \sigma_{(1)} (r^{1/2} \nu + (nr)^{1/4} \nu + n^{1/2} \nu) \\ &\lesssim \frac{\sqrt{d}}{\sqrt{n}} \sigma_{(1)} \nu. \end{aligned} \quad (96)$$

3. For the third term and fourth term, from equation (44) we know:

$$\mathbb{E} \frac{1}{n} \|U^*Z\bar{Z}^\top\|_2 + \mathbb{E} \frac{1}{n} \|E\bar{Z}^\top\|_2 \leq \mathbb{E} \|\bar{z}\bar{z}^\top\|_2 + \mathbb{E} \|\bar{\xi}\bar{z}^\top\|_2 \leq \frac{r}{n} \nu^2 + \sqrt{\frac{d}{n}} \nu \sigma_{(1)}. \quad (97)$$

Combine these three equations (95)(96)(97) together we have:

$$\mathbb{E} \left\| \frac{1}{n}XHy - \nu U^*w^* \right\| \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)}. \quad (98)$$

With equation (98), we can now turn to the difference between \hat{N} and $Uw^*w^{*\top}U^\top$. By Lemma D.4 we know that:

$$\|\hat{N} - \nu^2 U^*w^*w^{*\top}U^{*\top}\|_F \lesssim \left(\left\| \frac{1}{n}XHy \right\| + \|\nu U^*w^*\| \right) \left\| \frac{1}{n}XHy - \nu U^*w^* \right\|.$$

Using Markov's inequality, we can conclude from (98) that:

$$\mathbb{P} \left(\left\| \frac{1}{n}XHy - \nu U^*w^* \right\| \geq \nu \right) \leq \frac{\mathbb{E} \left\| \frac{1}{n}XHy - \nu U^*w^* \right\|}{\nu} \lesssim \sqrt{\frac{d}{n}}.$$

Then denote $A = \{\omega : \|\frac{1}{n}XHy - \nu^2 U^* w^*\|_2 < \nu\}$ we have:

$$\begin{aligned} \mathbb{E}\|\hat{N} - \nu^2 U^* w^* w^{*\top} U^{*\top}\|_F \mathbb{I}\{A\} &\lesssim \mathbb{E}(\|\frac{1}{n}XHy\| + \|\nu U^* w^*\|) \|\frac{1}{n}XHy - \nu U^* w^*\| \mathbb{I}\{A\} \\ &\lesssim \nu (\mathbb{E}\|\frac{1}{n}XHy - \nu U^* w^*\|) \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)} \nu. \end{aligned}$$

which finished the proof. \blacksquare

In the main body, we assume the number of labeled data and the ratio of the loss function is both balanced. Now we will provide a more general result to cover the imbalance occasions. Formally, suppose we have n unlabeled data $X = [x_1, \dots, x_n] \in \mathbb{R}^{d \times n}$ and n_i labeled data \mathcal{S}_i $X_i = [x_i^1, \dots, x_i^{n_i}]$, $y_i = [y_i^1, \dots, y_i^{n_i}]$, $\forall i = 1, \dots, T$ for source task, we learn the linear representation via joint optimization:

$$\min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}(W) := \min_{W \in \mathbb{R}^{r \times d}} \mathcal{L}_{\text{SelfCon}}(W) - \sum_{t=1}^T \alpha_t \text{HSIC}(\hat{X}^t, y^t; W), \quad (99)$$

To investigate its feature recovery ability, we first give the following result.

Theorem D.5 *For the optimization problem (99), if we apply augmented pairs generation in Definition 1 with random masking augmentation 2 for unlabeled data, then the optimal solution is given by:*

$$W_{CL} = C \left(\sum_{i=1}^r u_i \sigma_i v_i^\top \right)^\top,$$

where $C > 0$ is a constant, σ_i is the i -th largest eigenvalue of the following matrix:

$$\frac{1}{4n} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) + \sum_{i=1}^T \frac{\alpha_i}{(n_i - 1)^2} X_i H_{n_i} y_i y_i^\top H_{n_i} X_i^\top,$$

u_i is the corresponding eigenvector, $V = [v_1, \dots, v_r] \in \mathbb{R}^{r \times r}$ can be any orthogonal matrix and $H_{n_i} = I_{n_i} - \frac{1}{n_i} 1_{n_i} 1_{n_i}^\top$ is the centering matrix.

Proof Under this setting, combined with the result obtained in Corollary 6, the loss function can be rewritten as:

$$\begin{aligned}
 \mathcal{L}(W) &= \frac{\lambda}{2} \|WW^\top\|_F^2 - \frac{1}{2n} \operatorname{tr} \left(\left(\frac{1}{2} \Delta(XX^\top) - \frac{1}{2(n-1)} X(1_n 1_n^\top - I_n) X^\top \right) W^\top W \right) \\
 &\quad - \sum_{i=1}^T \alpha_i \frac{1}{(n_i - 1)^2} \operatorname{tr} \left(X_i^\top W^\top W X_i H y_i y_i^\top H \right) \\
 &= \frac{\lambda}{2} \left\| WW^\top - \frac{1}{4n\lambda} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \right. \\
 &\quad \left. - \sum_{i=1}^T \frac{\alpha_i}{\lambda(n_i - 1)^2} X_i H_{n_i} y_i y_i^\top H_{n_i} X_i^\top \right\|_F^2 \\
 &\quad - \frac{\lambda}{2} \left\| \frac{1}{4n\lambda} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) \right. \\
 &\quad \left. + \sum_{i=1}^T \frac{\alpha_i}{\lambda(n_i - 1)^2} X_i H_{n_i} y_i y_i^\top H_{n_i} X_i^\top \right\|_F^2.
 \end{aligned}$$

Then by a similar argument as in the proof of Proposition 5, we can conclude that the optimal solution W_{CL} must satisfy the desired conditions. \blacksquare

Then we can give the proofs of Theorem 23 and Theorem 26 under our generalized setting, one can easily obtain those under balanced setting by simply setting $\alpha = \alpha_1 = \dots = \alpha_T$ and $m = n_1 = \dots = n_T$, which is consistent with Theorem 23 and Theorem 26 in the main body.

Theorem D.6 (Generalized version of Theorem 23) *In the regression setting where $y^t = \langle w_t, z \rangle / \nu$, suppose Assumptions 8-10 hold for spiked covariance model (5) and $n > d \gg r$, if we further assume that $T < r$ and w_t 's are orthogonal to each other, and let W^{CL} be any solution that optimizes the problem in Equation (99), and denote its singular value decomposition as $W^{CL} = (U_{CL} \Sigma_{CL} V_{CL}^\top)^\top$, then we have:*

$$\begin{aligned}
 \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F &\lesssim \left(\frac{\sqrt{r-T}}{\min_{i \in [T]} \{\alpha_i, 1\}} + \frac{\sqrt{T}}{\min_{i \in [T]} \alpha_i} \right) \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) \\
 &\quad + \sum_{i=1}^T \left(\sqrt{r-T} \frac{\alpha_i + \min_{i \in [T]} \{\alpha_i, 1\}}{\min_{i \in [T]} \{\alpha_i, 1\}} + \sqrt{T} \frac{\alpha_i + \min_{i \in [T]} \alpha_i}{\min_{i \in [T]} \alpha_i} \right) \sqrt{\frac{d}{n_i}}.
 \end{aligned}$$

Proof [Proof of Theorem D.6] As shown in Theorem D.5, optimizing loss function (99) is equivalent to find the top- r eigenspace of matrix

$$\frac{1}{4n} \left(\Delta(XX^\top) - \frac{1}{n-1} X(1_n 1_n^\top - I_n) X^\top \right) + \sum_{i=1}^T \frac{\alpha_i}{(n_i - 1)^2} X_i H_{n_i} y_i y_i^\top H_{n_i} X_i^\top.$$

Again denote $\hat{M}_2 \triangleq \frac{1}{n}(\Delta(XX^\top) - \frac{1}{n-1}X(1_n 1_n^\top - I_n)X^\top)$ and $\hat{N}_i \triangleq \frac{1}{(n_i-1)^2}X_i H y_i y_i^\top H X_i^\top$. By equation (56) we know that:

$$\mathbb{E}\|\hat{M}_2 - M\|_2 \lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sigma_{(1)} \nu \sqrt{\frac{d}{n}}.$$

By Theorem D.3 we know that for each task \mathcal{S}_i , we can find an event A_i such that $\mathbb{P}(A_i) = O(\sqrt{\frac{d}{n}})$:

$$\mathbb{E}\|\hat{N}_i - \nu^2 U^* w_i w_i^\top U^{*\top}\|_F \mathbb{I}\{A_i\} \lesssim \sqrt{\frac{d}{n_i}} \sigma_{(1)} \nu.$$

The target matrix is $N = \nu^2 U^* U^{*\top} + \sum_{i=1}^T \alpha_i \nu^2 U^* w_i w_i^\top U^{*\top}$, and we can obtain the upper bound for the difference between N and \hat{N} :

$$\begin{aligned} \mathbb{E}\|\hat{N} - N\|_2 \mathbb{I}\{\cap_{i=1}^T A_i\} &\leq \frac{1}{4} \mathbb{E}\|\hat{M}_2 - M\|_2 + \sum_{i=1}^T \alpha_i \mathbb{E}\|\hat{N}_i - \nu^2 U^* w_i w_i^\top U^{*\top}\|_F \mathbb{I}\{A_i\} \\ &\lesssim \nu^2 \left(\frac{r}{d} \log d + \sqrt{\frac{r}{n} + \frac{r}{n}} \right) + \sigma_{(1)}^2 \left(\sqrt{\frac{d}{n} + \frac{d}{n}} \right) + \sigma_{(1)} \nu \sqrt{\frac{d}{n}} + \sum_{i=1}^T \left[\alpha_i \sqrt{\frac{d}{n_i}} \sigma_{(1)} \nu \right]. \end{aligned} \quad (100)$$

We divide the top- r eigenspace U_{CL} of $W_{\text{CL}} W_{\text{CL}}^\top$ into two parts: the top- T eigenspace $U_{\text{CL}}^{(1)}$ and top- $(T+1)$ to top- r eigenspace $U_{\text{CL}}^{(2)}$. Similarly, we also divide the top- r eigenspace U^* of N into two parts: $U^{*(1)}$ and $U^{*(2)}$. Then applying Lemma E.1 we can bound the sine distance for each part: on the one hand,

$$\begin{aligned} &\mathbb{E}\|\sin\left(\Theta(U_{\text{CL}}^{(1)}, U^{*(1)})\right)\|_F \\ &= \mathbb{E}\|\sin\left(\Theta(U_{\text{CL}}^{(1)}, U^{*(1)})\right)\|_F \mathbb{I}\{\cap_{i=1}^T A_i\} + \mathbb{E}\|\sin\left(\Theta(U_{\text{CL}}^{(1)}, U^{*(1)})\right)\|_F \mathbb{I}\{\cup_{i=1}^T A_i^C\} \\ &\leq \frac{\sqrt{T} \mathbb{E}\|\hat{N} - N\|_2 \mathbb{I}\{\cap_{i=1}^T A_i\}}{\lambda_{(T)}(N) - \lambda_{(T+1)}(N)} + \sqrt{T} \mathbb{P}(\cup_{i=1}^T A_i^C) \\ &\lesssim \frac{\sqrt{T}}{\min_{i \in [T]} \alpha_i \nu^2} \left(\nu^2 \frac{r}{d} \log d + \sigma_{(1)}^2 \sqrt{\frac{d}{n}} + \sum_{i=1}^T \alpha_i \sqrt{\frac{d}{n_i}} \sigma_{(1)} \nu \right) + \sqrt{T} \sum_{i=1}^T \sqrt{\frac{d}{n_i}} \\ &\lesssim \frac{\sqrt{T}}{\min_{i \in [T]} \alpha_i} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + \sqrt{T} \sum_{i=1}^T \frac{\alpha_i + \min_{i \in [T]} \alpha_i}{\min_{i \in [T]} \alpha_i} \sqrt{\frac{d}{n_i}}. \end{aligned}$$

On the other hand,

$$\begin{aligned}
 & \mathbb{E} \|\sin(\Theta(U_{\text{CL}}^{(2)}, U^{*(2)}))\|_F \\
 &= \mathbb{E} \|\sin(\Theta(U_{\text{CL}}^{(2)}, U^{*(2)}))\|_F \mathbb{I}\{\cap_{i=1}^T A_i\} + \mathbb{E} \|\sin(\Theta(U_{\text{CL}}^{(2)}, U^{*(2)}))\|_F \mathbb{I}\{\cup_{i=1}^T A_i^C\} \\
 &\leq \frac{\sqrt{r-T} \mathbb{E} \|\hat{N} - N\|_2 \mathbb{I}\{\cap_{i=1}^T A_i\}}{\min\{\lambda_{(T)}(N) - \lambda_{(T+1)}(N), \lambda_{(r)}(N)\}} + \sqrt{r-T} \mathbb{P}(\cup_{i=1}^T A_i^C) \\
 &\lesssim \frac{\sqrt{r-T}}{\min_{i \in [T]} \{\alpha_i, 1\} \nu^2} \left(\nu^2 \frac{r}{d} \log d + \sigma_{(1)}^2 \sqrt{\frac{d}{n}} + \sum_{i=1}^T \alpha_i \sqrt{\frac{d}{n_i}} \sigma_{(1)} \nu \right) + \sqrt{r-T} \sum_{i=1}^T \sqrt{\frac{d}{n_i}} \\
 &\lesssim \frac{\sqrt{r-T}}{\min_{i \in [T]} \{\alpha_i, 1\}} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + \sqrt{r-T} \sum_{i=1}^T \left(\frac{\alpha_i}{\min_{i \in [T]} \{\alpha_i, 1\}} + 1 \right) \sqrt{\frac{d}{n_i}}.
 \end{aligned}$$

Note that:

$$\begin{aligned}
 & \|\sin(\Theta(U_{\text{CL}}, U^*))\|_F^2 \\
 &= r - \|U_{\text{CL}}^\top U^*\|_F^2 \\
 &\leq r - \|U_{\text{CL}}^{(1)\top} U^{*(1)}\|_F^2 - \|U_{\text{CL}}^{(2)\top} U^{*(2)}\|_F^2 \\
 &\leq T - \|U_{\text{CL}}^{(1)\top} U^{*(1)}\|_F^2 + (r - T) - \|U_{\text{CL}}^{(2)\top} U^{*(2)}\|_F^2 \\
 &\leq \|\sin \Theta(U_{\text{CL}}^{(1)}, U^{*(1)})\|_F^2 + \|\sin \Theta(U_{\text{CL}}^{(2)}, U^{*(2)})\|_F^2,
 \end{aligned}$$

and the sine distance has trivial upper bounds:

$$\|\sin \Theta(U_{\text{CL}}^{(1)}, U^{*(1)})\|_F^2 \leq T, \quad \|\sin \Theta(U_{\text{CL}}^{(2)}, U^{*(2)})\|_F^2 \leq r - T$$

Thus we can conclude:

$$\begin{aligned}
 & \mathbb{E} \|\sin(\Theta(U_{\text{CL}}, U^*))\|_F \\
 &\leq \mathbb{E} \|\sin(\Theta(U_{\text{CL}}^{(1)}, U^{*(1)}))\|_F + \mathbb{E} \|\sin(\Theta(U_{\text{CL}}^{(2)}, U^{*(2)}))\|_F \\
 &\lesssim \left(\frac{\sqrt{r-T}}{\min_{i \in [T]} \{\alpha_i, 1\}} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + \sum_{i=1}^T \sqrt{r-T} \frac{\alpha_i + \min_{i \in [T]} \{\alpha_i, 1\}}{\min_{i \in [T]} \{\alpha_i, 1\}} \sqrt{\frac{d}{n_i}} \right) \wedge \sqrt{r-T} \\
 &+ \left(\frac{\sqrt{T}}{\min_{i \in [T]} \alpha_i} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + \sum_{i=1}^T \sqrt{T} \frac{\alpha_i + \min_{i \in [T]} \alpha_i}{\min_{i \in [T]} \alpha_i} \sqrt{\frac{d}{n_i}} \right) \wedge \sqrt{T}.
 \end{aligned}$$

■

Theorem D.7 (Restatement of Theorem 24) *Suppose the conditions in Theorem 23 hold. Then,*

$$\mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{W_{\text{CL}}, w})] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{U^{*\top}, w})] \right] \quad (101)$$

$$\lesssim \sqrt{r-T} \left(\frac{r \log d}{d} + \sqrt{\frac{d}{n}} + \alpha T \sqrt{\frac{d}{m}} \wedge 1 \right) + \sqrt{T} \left(\frac{r \log d}{\alpha d} + \frac{1}{\alpha} \sqrt{\frac{d}{n}} + T \sqrt{\frac{d}{m}} \right). \quad (102)$$

Proof [Proof of Theorem D.7] Theorem 24 follows directly from Lemma B.20 and Theorem 23. \blacksquare

Theorem D.8 (Generalized version of Theorem 26) *In the regression setting where $y^t = \langle w_t, z \rangle / \nu$, suppose Assumptions 8-10 hold for spiked covariance model (5) and $n > d \gg r$, if we further assume that $T \geq r$ and $\sum_{i=1}^T \alpha_i w_i w_i^\top$ is full rank, suppose W^{CL} is the optimal solution of optimization problem Equation (99), and denote its singular value decomposition as $W_{CL} = (U_{CL} \Sigma_{CL} V_{CL}^\top)^\top$, then we have:*

$$\begin{aligned} \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F &\lesssim \frac{\sqrt{r}}{1 + \nu^2 \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) \\ &\quad + \sqrt{r} \sum_{i=1}^T \left(\frac{\alpha_i}{1 + \nu^2 \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} + 1 \right) \sqrt{\frac{d}{n_i}}. \end{aligned}$$

Proof [Proof of Theorem D.8] The proof strategy is similar to that of Theorem 23, here the difference is that each direction can be accurately estimated by the labeled data and we do not need to separate the eigenspace. Directly applying Lemma E.1 and equation (100) we have:

$$\begin{aligned} &\mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F \\ &= \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F \mathbb{I}\{\cap_{i=1}^T A_i\} + \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F \mathbb{I}\{\cup_{i=1}^T A_i^C\} \\ &\lesssim \frac{\sqrt{r} \mathbb{E} \|\hat{N} - N\|_2 \mathbb{I}\{\cap_{i=1}^T A_i\}}{\lambda_{(r)}(N)} + \sqrt{r} \mathbb{P}(\cup_{i=1}^T A_i^C) \\ &\lesssim \frac{\sqrt{r}}{\nu^2 + \nu^2 \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} \left(\nu^2 \frac{r}{d} \log d + \sigma_{(1)}^2 \sqrt{\frac{d}{n}} + \sum_{i=1}^T \alpha_i \sqrt{\frac{d}{n_i}} \sigma_{(1)} \nu \right) + \sqrt{r} \sum_{i=1}^T \sqrt{\frac{d}{n_i}} \\ &\lesssim \frac{\sqrt{r}}{1 + \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + \sqrt{r} \sum_{i=1}^T \left(\frac{\alpha_i}{1 + \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} + 1 \right) \sqrt{\frac{d}{n_i}}. \end{aligned}$$

Theorem D.9 (Restatement of Theorem 27) *Suppose the conditions in Theorem 26 hold. Then,*

$$\mathbb{E}_{\mathcal{D}} \left[\inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{W_{CL}, w})] \right] - \inf_{w \in \mathbb{R}^r} \mathbb{E}_{\mathcal{E}} [\ell_r(\delta_{U^{\star\top}, w})] \lesssim \frac{\sqrt{r}}{\alpha + 1} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) + T \sqrt{\frac{dr}{m}}. \quad (103)$$

Proof [Proof of Theorem 27] Theorem 27 follows directly from Lemma B.20 and Theorem 26. \blacksquare

Now we move to a binary classification setting, where labels y are generated by $y = \text{sign}(\langle w^*, z \rangle)$ instead of $y = \langle w^*, z \rangle / \nu$ in previous regression setting. We first give the corresponding generalized version of Theorem 28 and Theorem 29 to cover the general imbalanced settings.

Theorem D.10 (Generalized version of Theorem 28) *In the classification setting where $y^t = \text{sign}(\langle w_t, z \rangle)$, suppose Assumptions 8-10 hold for spiked covariance model (5), z follows a Gaussian distribution, and $n > d \gg r$, if we further assume that $T < r$ and w_t 's are orthogonal to each other, and let W^{CL} be any solution that optimizes the problem in Equation (99), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have:*

$$\begin{aligned} \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F &\lesssim \left(\frac{\sqrt{r-T}}{\min_{i \in [T]} \{\alpha_i, 1\}} + \frac{\sqrt{T}}{\min_{i \in [T]} \alpha_i} \right) \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) \\ &\quad + \sum_{i=1}^T \left(\sqrt{r-T} \frac{\alpha_i + \min_{i \in [T]} \{\alpha_i, 1\}}{\min_{i \in [T]} \{\alpha_i, 1\}} + \sqrt{T} \frac{\alpha_i + \min_{i \in [T]} \alpha_i}{\min_{i \in [T]} \alpha_i} \right) \sqrt{\frac{d}{n_i}}. \end{aligned}$$

Theorem D.11 (Generalized version of Theorem 29) *In the classification setting where $y^t = \text{sign}(\langle w_t, z \rangle)$, suppose Assumptions 8-10 hold for spiked covariance model (5), z follows a Gaussian distribution, and $n > d \gg r$, if we further assume that $T \geq r$ and $\sum_{i=1}^T \alpha_i w_i w_i^\top$ is full rank, suppose W^{CL} is the optimal solution of optimization problem Equation (99), and denote its singular value decomposition as $W_{CL} = (U_{CL}\Sigma_{CL}V_{CL}^\top)^\top$, then we have:*

$$\begin{aligned} \mathbb{E} \|\sin(\Theta(U_{CL}, U^*))\|_F &\lesssim \frac{\sqrt{r}}{1 + \nu^2 \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} \left(\frac{r}{d} \log d + \sqrt{\frac{d}{n}} \right) \\ &\quad + \sqrt{r} \sum_{i=1}^T \left(\frac{\alpha_i}{1 + \nu^2 \lambda_{(r)}(\sum_{i=1}^T \alpha_i w_i w_i^\top)} + 1 \right) \sqrt{\frac{d}{n_i}}. \end{aligned}$$

The only difference between these two settings is the distribution of labels y . Thus to prove Theorem D.10 and Theorem D.11, we only need to recover Lemma D.3 in this binary classification setting. Since in the classification setting the labels are discrete and could be harder to analyze, we make the Gaussian assumption on z to make problems mathematically tractable in these two Theorems.

Lemma D.12 (Classification version of Lemma D.3) *In the binary classification setting, under the conditions similar to Theorem 14 and assume z in the spiked covariance model (5) follows a Gaussian distribution, we can find an event A such that $\mathbb{P}(A^C) = O(\sqrt{d/n})$ and:*

$$\mathbb{E} \left[\left\| \frac{1}{(n-1)^2} X H y y^\top H X^\top - \frac{2\nu^2}{\pi} U^* w^* w^{*\top} U^{*\top} \right\|_F \mathbb{I}\{A\} \right] \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)} \nu. \quad (104)$$

Proof Again, by (93) we have:

$$\begin{aligned} & \left\| \frac{1}{(n-1)^2} XHy y^\top HX^\top - \frac{2\nu^2}{\pi} U^* w^* w^{*\top} U^{*\top} \right\|_F \\ & \lesssim \left\| \frac{1}{n^2} XHy y^\top HX^\top - \frac{2\nu^2}{\pi} U^* w^* w^{*\top} U^{*\top} \right\|_F + \frac{r}{n} \nu^2, \end{aligned}$$

thus we can replace the $\frac{1}{(n-1)^2}$ with $\frac{1}{n}$ in equation (104) and conclude the proof. Denote $\hat{N} \triangleq \frac{1}{n^2} XHy y^\top HX^\top$, note that both of \hat{N} and $Uw^*w^{*\top}U^\top$ are rank-1 matrices. We first bound the difference between $\frac{1}{n}XHy$ and $\sqrt{\frac{2\nu^2}{\pi}}Uw^*$:

$$\begin{aligned} \left\| \frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| &= \left\| \frac{1}{n}(U^*Z + E)Hy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| \\ &\leq \left\| \frac{1}{n}U^*Zy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| + \frac{1}{n}\|Ey\| + \frac{1}{n}\|U^*Z\bar{y}\| + \frac{1}{n}\|E\bar{y}\|. \end{aligned} \quad (105)$$

We deal with the four terms in (105) separately:

1. For the first term, note that: $\frac{1}{n}Zy = \frac{1}{n} \sum_{i=1}^n z_i \text{sign}(z_i^\top w^*)$ and $z_i \sim \mathcal{N}(0, \nu^2 I_r)$, thus $z_i \text{sign}(z_i^\top w^*)$ follows a folded Gaussian distribution, which is a reflection of standard Gaussian distribution along the normal plane of w^* , thus

$$\begin{aligned} \mathbb{E} \left\| \frac{1}{n}U^*Zy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| &\leq \mathbb{E} \left\| \frac{1}{n}Zy - \sqrt{\frac{2\nu^2}{\pi}}w^* \right\| \leq \sqrt{\mathbb{E} \left\| \frac{1}{n}Zy - \sqrt{\frac{2\nu^2}{\pi}}w^* \right\|^2} \\ &\leq \sqrt{\frac{r}{n}} \nu \end{aligned} \quad (106)$$

2. For the second term, note that y and E are independent and $|y| = 1$ almost surely

$$\frac{1}{n} \mathbb{E} \|Ey\| = \frac{1}{n} \mathbb{E} \left\| \sum_{i=1}^n \xi_i \right\| \leq \frac{1}{n} \sqrt{\mathbb{E} \left\| \sum_{i=1}^n \xi_i \right\|^2} \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)} \quad (107)$$

3. For the third term and fourth terms, we have:

$$\mathbb{E} \frac{1}{n} \|U^*Z\bar{y}\| + \mathbb{E} \frac{1}{n} \|E\bar{y}\| \leq \mathbb{E} \frac{1}{n} \left\| \sum_{i=1}^n z_i \right\| + \mathbb{E} \frac{1}{n} \left\| \sum_{i=1}^n \xi_i \right\| \lesssim \sqrt{\frac{r}{n}} \nu + \sqrt{\frac{d}{n}} \sigma_{(1)}. \quad (108)$$

Combine these three equations (106)(107)(108) together we have:

$$\mathbb{E} \left\| \frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| \lesssim \sqrt{\frac{d}{n}} \sigma_{(1)}. \quad (109)$$

With equation (109), we can now turn to the difference between \hat{N} and $\frac{2\nu^2}{\pi}Uw^*w^{*\top}U^\top$. By Lemma D.4 we know that:

$$\left\| \hat{N} - \frac{2\nu^2}{\pi}Uw^*w^{*\top}U^\top \right\|_F \lesssim \left(\left\| \frac{1}{n}XHy \right\| + \left\| \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\| \right) \left\| \frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}Uw^* \right\|.$$

Using Markov's inequality, we can conclude from (109) that:

$$\mathbb{P}(\|\frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}U^*w^*\| \geq \nu) \leq \frac{\mathbb{E}\|\frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}U^*w^*\|}{\nu} \lesssim \sqrt{\frac{d}{n}}.$$

Then denote $A = \{\omega : \|\frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}U^*w^*\|_2 < \nu\}$ we have:

$$\begin{aligned} \mathbb{E}\|\hat{N} - \frac{2\nu^2}{\pi}U^*w^*w^{*\top}U^{*\top}\|_{F|\mathbb{I}\{A\}} &\lesssim \mathbb{E}(\|\frac{1}{n}XHy\| + \|\sqrt{\frac{2\nu^2}{\pi}}U^*w^*\|)\|\frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}U^*w^*\|_{\mathbb{I}\{A\}} \\ &\lesssim \nu \mathbb{E}\|\frac{1}{n}XHy - \sqrt{\frac{2\nu^2}{\pi}}U^*w^*\| \lesssim \sqrt{\frac{d}{n}}\sigma_{(1)}\nu. \end{aligned}$$

which finished the proof. \blacksquare

With Lemma D.12 established, it is straightforward to obtain the same results as in Theorem 23, Theorem 24, Theorem 26 and Theorem 27 for this binary classification setting.

Appendix E. Useful lemmas

In this section, we list some of the main techniques that have been used in the proof of the main results.

Lemma E.1 (Theorem 2 in Yu et al. (2015)) *Let $\Sigma, \hat{\Sigma} \in \mathbb{R}^{p \times p}$ be symmetric, with eigenvalues $\lambda_1 \geq \dots \geq \lambda_p$ and $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_p$ respectively. Fix $1 \leq r \leq s \leq p$ and assume that $\min(\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1}) > 0$ where $\lambda_0 := \infty$ and $\lambda_{p+1} := -\infty$. Let $d := s - r + 1$, and let $V = (v_r, v_{r+1}, \dots, v_s) \in \mathbb{R}^{p \times d}$ and $\hat{V} = (\hat{v}_r, \hat{v}_{r+1}, \dots, \hat{v}_s) \in \mathbb{R}^{p \times d}$ have orthonormal columns satisfying $\Sigma v_j = \lambda_j v_j$ and $\hat{\Sigma} \hat{v}_j = \hat{\lambda}_j \hat{v}_j$ for $j = r, r+1, \dots, s$. Then*

$$\|\sin \Theta(\hat{V}, V)\|_F \leq \frac{2 \min(d^{1/2}\|\hat{\Sigma} - \Sigma\|_2, \|\hat{\Sigma} - \Sigma\|_F)}{\min(\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1})}.$$

Moreover, there exists an orthogonal matrix $\hat{O} \in \mathbb{R}^{d \times d}$ such that

$$\|\hat{V}\hat{O} - V\|_F \leq \frac{2^{3/2} \min(d^{1/2}\|\hat{\Sigma} - \Sigma\|_2, \|\hat{\Sigma} - \Sigma\|_F)}{\min(\lambda_{r-1} - \lambda_r, \lambda_s - \lambda_{s+1})}.$$

Lemma E.2 (Lemma 2 in Zhang et al. (2018)) *Assume that $E \in \mathbb{R}^{p_1 \times p_2}$ has independent sub-Gaussian entries, $\text{Var}(E_{ij}) = \sigma_{ij}^2$, $\sigma_C^2 = \max_j \sum_i \sigma_{ij}^2$, $\sigma_R^2 = \max_i \sum_j \sigma_{ij}^2$, $\sigma_{(1)}^2 = \max_{i,j} \sigma_{ij}^2$. Assume that*

$$\|E_{ij}/\sigma_{ij}\|_{\psi_2} := \max_{q \geq 1} q^{-1/2} \{\mathbb{E}(|E_{ij}|/\sigma_{ij})^q\}^{1/q} \leq \kappa.$$

Let $V \in \mathbb{O}_{p_2, r}$ be a fixed orthogonal matrix. Then

$$\begin{aligned} \mathbb{P}(\|EV\|_2 \geq 2(\sigma_C + x)) &\leq 2 \exp\left(5r - \min\left\{\frac{x^4}{\kappa^4 \sigma_{(1)}^2 \sigma_C^2}, \frac{x^2}{\kappa^2 \sigma_{(1)}^2}\right\}\right), \\ \mathbb{E}\|EV\|_2 &\lesssim \sigma_C + \kappa r^{1/4} (\sigma_{(1)} \sigma_C)^{1/2} + \kappa r^{1/2} \sigma_{(1)}. \end{aligned}$$

Lemma E.3 (Theorem 6 in Cai et al. (2020)) *Suppose Z is a p_1 -by- p_2 random matrix with independent mean-zero sub-Gaussian entries. If there exist $\sigma_1, \dots, \sigma_p \geq 0$ such that $\|Z_{ij}/\sigma_i\|_{\psi_2} \leq C_K$ for constant $C_K > 0$, then*

$$\mathbb{E} \left\| ZZ^\top - \mathbb{E}ZZ^\top \right\|_2 \lesssim \sum_i \sigma_i^2 + \sqrt{p_2 \sum_i \sigma_i^2} \cdot \max_i \sigma_i.$$

Lemma E.4 (The Eckart-Young-Mirsky Theorem (Eckart and Young, 1936)) *Suppose that $A = U\Sigma V^\top$ is the singular value decomposition of A . Then the best rank- k approximation of the matrix A w.r.t the Frobenius norm, $\|\cdot\|_F$, is given by*

$$A_k = \sum_{i=1}^k \sigma_i u_i v_i^\top.$$

that is, for any matrix B of rank at most k

$$\|A - A_k\|_F \leq \|A - B\|_F.$$

References

- Anish Agarwal, Devavrat Shah, and Dennis Shen. On principal component regression in a high-dimensional error-in-variables setting. *arXiv preprint arXiv:2010.14449*, 2020.
- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- Zhidong Bai and Jianfeng Yao. On sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*, 106:167–177, 2012.
- Dana H Ballard. Modular learning in neural networks. In *AAAI*, volume 647, pages 279–284, 1987.
- Elnaz Barshan, Ali Ghodsi, Zohreh Azimifar, and Mansoor Zolghadri Jahromi. Supervised principal component analysis: Visualization, classification and regression on subspaces and submanifolds. *Pattern Recognition*, 44(7):1357–1371, 2011.
- Xin Bing, Florentina Bunea, Seth Strimas-Mackey, and Marten Wegkamp. Prediction under latent factor regression: Adaptive pcr, interpolating predictors and beyond. *Journal of Machine Learning Research*, 22(177):1–50, 2021.
- Hervé Boursard and Yves Kamp. Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, 59(4):291–294, 1988.
- T Tony Cai and Anru Zhang. Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *The Annals of Statistics*, 46(1):60–89, 2018.

- T Tony Cai, Rungang Han, and Anru R Zhang. On the non-asymptotic concentration of heteroskedastic wishart-type matrix. *arXiv preprint arXiv:2008.12434*, 2020.
- Emmanuel J Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- Domagoj Čevič, Peter Bühlmann, and Nicolai Meinshausen. Spectral deconfounding via perturbed sparse linear models. *Journal of Machine Learning Research*, 21:232, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020a.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020b.
- Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020c.
- Adam Coates, A. Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *AISTATS*, 2011.
- Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse pca. In *2014 IEEE International Symposium on Information Theory*, pages 2197–2201. IEEE, 2014.
- Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. *Advances in neural information processing systems*, 27:766–774, 2014.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- Carl Eckart and G. Marion Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936.
- Jianqing Fan, Cong Ma, and Yiqiao Zhong. A selective overview of deep learning. *arXiv preprint arXiv:1904.05526*, 2019.
- Soheil Feizi, Farzan Farnia, Tony Ginart, and David Tse. Understanding gans in the lqg setting: Formulation, generalization and stability. *IEEE Journal on Selected Areas in Information Theory*, 1(1):304–311, 2020.

- Daniel Y Fu, Mayee F Chen, Michael Zhang, Kayvon Fatahalian, and Christopher Ré. The details matter: Preventing class collapse in supervised contrastive learning. In *Computer Sciences & Mathematics Forum*, volume 3, page 4. MDPI, 2022.
- Siddhant Garg and Yingyu Liang. Functional regularization for representation learning: A unified theoretical perspective. *Advances in Neural Information Processing Systems*, 33: 17187–17199, 2020.
- Gene H. Golub and Charles Van Loan. *Matrix computations* (3rd ed.). 1996.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Florian Graf, Christoph Hofer, Marc Niethammer, and Roland Kwitt. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, pages 3821–3830. PMLR, 2021.
- Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *International conference on algorithmic learning theory*, pages 63–77. Springer, 2005.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006.
- Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34): 8505–8510, 2018.
- Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *arXiv preprint arXiv:2106.04156*, 2021.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.

- Xinwei He, Yang Zhou, Zhichao Zhou, Song Bai, and Xiang Bai. Triplet-center loss for multi-view 3d object retrieval. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1945–1954, 2018.
- Aapo Hyvärinen, Jarmo Hurri, Patrik O Hoyer, Aapo Hyvärinen, Jarmo Hurri, and Patrik O Hoyer. *Independent component analysis*. Springer, 2009.
- Ashraf Islam, Chun-Fu Chen, Rameswar Panda, Leonid Karlinsky, Richard J. Radke, and Rogério Schmidt Feris. A broad study on the transferability of visual representations with contrastive learning. *ArXiv*, abs/2103.13517, 2021.
- Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2021.
- Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348*, 2021.
- Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.
- Ian T Jolliffe. A note on the use of principal components in regression. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 31(3):300–303, 1982.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323, 2021.
- Yazhe Li, Roman Pogodin, Danica J Sutherland, and Arthur Gretton. Self-supervised learning with kernel dependence maximization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- George Marsaglia. Choosing a point from the surface of a sphere. *Annals of Mathematical Statistics*, 43:645–646, 1972.

- Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Dan Pirjol. The logistic-normal integral and its generalizations. *Journal of Computational and Applied Mathematics*, 237(1):460–469, 2013.
- Elad Plaut. From principal subspaces to principal components with linear autoencoders. *arXiv preprint arXiv:1804.10253*, 2018.
- Maria Refinetti and Sebastian Goldt. The dynamics of representation learning in shallow, non-linear autoencoders. *arXiv preprint arXiv:2201.02115*, 2022.
- Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. In *International Conference on Machine Learning*, pages 19250–19286. PMLR, 2022.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015.
- Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NIPS*, 2016.
- Le Song, Arthur Gretton, Karsten Borgwardt, and Alex Smola. Colored maximum variance unfolding. *Advances in Neural Information Processing Systems*, 20:1385–1392, 2007a.
- Le Song, Alex Smola, Arthur Gretton, and Karsten M Borgwardt. A dependence maximization view of clustering. In *Proceedings of the 24th international conference on Machine learning*, pages 815–822, 2007b.
- Le Song, Alex Smola, Arthur Gretton, Karsten M Borgwardt, and Justin Bedo. Supervised feature selection via dependence estimation. In *Proceedings of the 24th international conference on Machine learning*, pages 823–830, 2007c.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020.
- Yuandong Tian. Understanding deep contrastive learning via coordinate-wise optimization. In *Advances in Neural Information Processing Systems*, 2022.
- Yuandong Tian, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning dynamics without contrastive pairs. *arXiv preprint arXiv:2102.06810*, 2021.

- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pages 1179–1206. PMLR, 2021.
- Nilesh Tripuraneni, Chi Jin, and Michael Jordan. Provable meta-learning of linear representations. In *International Conference on Machine Learning*, pages 10434–10443. PMLR, 2021.
- Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Self-supervised learning from a multi-view perspective. *arXiv preprint arXiv:2006.05576*, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008.
- Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020.
- Xiang Wang, Xinlei Chen, Simon S Du, and Yuandong Tian. Towards demystifying representation learning with non-contrastive self-supervision. *arXiv preprint arXiv:2110.04947*, 2021.
- Zixin Wen and Yuanzhi Li. Toward understanding the feature learning process of self-supervised contrastive learning. *arXiv preprint arXiv:2105.15134*, 2021.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.
- Jianfeng Yao, Shurong Zheng, and ZD Bai. *Sample covariance matrices and high-dimensional data analysis*. Cambridge University Press Cambridge, 2015.
- Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6210–6219, 2019.
- Yi Yu, Tengyao Wang, and Richard J Samworth. A useful variant of the davis–kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.
- Anru R Zhang, T Tony Cai, and Yihong Wu. Heteroskedastic pca: Algorithm, optimality, and applications. *arXiv preprint arXiv:1810.08316*, 2018.
- Roland S Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021.