
Which Features are Learnt by Contrastive Learning? On the Role of Simplicity Bias in Class Collapse and Feature Suppression

Yihao Xue¹ Siddharth Joshi¹ Eric Gan¹ Pin-Yu Chen² Baharan Mirzasoleiman¹

Abstract

Contrastive learning (CL) has emerged as a powerful technique for representation learning, with or without label supervision. However, supervised CL is prone to collapsing representations of subclasses within a class by not capturing all their features, and unsupervised CL may suppress harder class-relevant features by focusing on learning easy class-irrelevant features; both significantly compromise representation quality. Yet, there is no theoretical understanding of *class collapse* or *feature suppression* at *test* time. We provide the first unified theoretically rigorous framework to determine *which* features are learnt by CL. Our analysis indicate that, perhaps surprisingly, bias of (stochastic) gradient descent towards finding simpler solutions is a key factor in collapsing subclass representations and suppressing harder class-relevant features. Moreover, we present increasing embedding dimensionality and improving the quality of data augmentations as two theoretically motivated solutions to feature suppression. We also provide the first theoretical explanation for why employing supervised and unsupervised CL together yields higher-quality representations, even when using commonly-used stochastic gradient methods.

1. Introduction

Learning high-quality representations that generalize well to a variety of downstream prediction tasks has been a long-standing goal of machine learning (Hinton et al., 2006; Ranzato et al., 2006). Contrastive learning (CL) has emerged as an effective approach for solving this problem, both with and without supervision (Chen et al., 2020;

Chuang et al., 2020; Grill et al., 2020; Khosla et al., 2020). Unsupervised CL learns representations of training examples by maximizing agreement between augmented views of the same example. Similarly, supervised CL maximizes agreement between augmented views of examples in the same class. Despite their empirical success, both supervised and unsupervised contrastive learning fail to capture *all* semantically relevant features in the data. In particular, *supervised* CL can fall prey to *class collapse* (Graf et al., 2021; Chen et al., 2022), where representations of *subclasses* within a class may no longer be distinguishable from each other; thus, yielding a poor classification performance at the subclass level. Similarly, *unsupervised* CL can be afflicted with *feature suppression* (Chen et al., 2021; Robinson et al., 2021) where easy but class-irrelevant features suppress the learning of harder class-relevant ones; deteriorating the generalizability of the obtained representations.

In spite of the significance of these failure modes, there is no clear theoretical understanding of them and consequently, no rigorous solution. Feature suppression has not been studied theoretically by prior work and the only theoretical work on class collapse (Graf et al., 2021) cannot explain why we observe class collapse at *test time*.

Addressing class collapse and feature suppression requires a theoretical understanding of *which* features CL learns. However, existing CL theory (Wang & Isola, 2020; Graf et al., 2021; Lee et al., 2021; Tosh et al., 2021a;b; Arora et al., 2019b; Tsai et al., 2020; HaoChen et al., 2021; Wen & Li, 2021; Ji et al., 2021) only explains how semantically relevant features are learned. The implicit assumption is that *all* semantically relevant features are learned, but the occurrence of class collapse and feature suppression proves otherwise. We propose the first unified (i.e. for both supervised and unsupervised CL) framework to answer *which* semantically relevant features are learned. We then leverage this framework to characterize class collapse and feature suppression. Table 1 summarizes the main findings in this paper, which are detailed below.

Class Collapse in Supervised CL. We prove that, perhaps surprisingly and in contrast to the current understanding (Graf et al., 2021), global minimizers of the supervised contrastive loss do not necessarily collapse the representations

*Equal contribution ¹Department of Computer Science, University of California, Los Angeles, USA ²IBM Research. Correspondence to: Yihao Xue <yihaoxue@g.ucla.edu>.

Table 1. A concise overview of the key findings in this research. In the table, ‘CC’ and ‘FS’ refers to class collapse and feature suppression, respectively. ‘Thm’ and ‘Exp’ refers to theorem and experiment, respectively.

Loss	Finding	Thm/Exp	Implication
SCL	min loss $\not\Rightarrow$ CC	Thm 4.3	Simplicity bias of (S)GD contributes to CC
	(min loss & min norm) \Rightarrow CC	Thm 4.4 & 4.7	
	(S)GD learns subclasses early in training	Thm 4.5 & Exp	
	(S)GD eventually unlearns subclasses, leading to CC	Exp	
UCL	With insufficient embedding size, (min loss & min norm) \Rightarrow FS	Thm 5.1 & Exp	Simplicity bias of (S)GD contributes to FS; Larger embedding size/better augmentation alleviates FS
	With imperfect data augmentation, (min loss & min norm) \Rightarrow FS, even with sufficient embedding size	Thm 5.4	
Joint	Joint loss can avoid both CC and FS	Thm 6.1 & Exp	Justification of joint loss

of the subclasses at *test* time. We find, however, that the *minimum norm* global minimizer does suffer from class collapse on test data.

We then study minimizing the supervised contrastive loss using (S)GD and show that, interestingly, subclass features are learned early in training. However, we verify empirically, that as training proceeds, (S)GD forgets the learned subclass features and collapses class representations.

Altogether, our findings indicate that the bias of SGD towards finding simpler solutions (Lyu et al., 2021) is the main deriving factor in collapsing class representations.

Feature Suppression in Unsupervised CL. We provide the first theoretical characterization of feature suppression in unsupervised CL. In particular, we show that the *minimum norm* global minimizer of the unsupervised contrastive loss results in feature suppression, when the embedding dimensionality is small *or* when data augmentations preserve class-irrelevant features better than class-relevant features. Again, our results identify the simplicity bias of (S)GD as a key factor in suppressing features of the input data. In addition, our findings suggest practical solutions to the problem of feature suppression: increasing embedding dimensionality and/or improving the quality of data augmentations.

Theoretical Justification for Combining Supervised and Unsupervised CL to Obtain Superior Representations. Finally, we prove that the *minimum norm* global minimizer of the joint loss (weighted sum of the supervised and unsupervised contrastive loss) does not suffer from class collapse or feature suppression, explaining why Chen et al. (2022); Islam et al. (2021) observes this empirically (i.e. even when using SGD).

2. Related Work

Theory of CL. While there has been much progress in theoretically understanding CL, most prior work (Wang & Isola, 2020; Graf et al., 2021; Lee et al., 2021; Tosh et al., 2021a;b; Arora et al., 2019b; Tsai et al., 2020; HaoChen

et al., 2021) are focused on understanding how CL clusters examples using semantically meaningful information or providing generalization guarantees on downstream tasks. Feature learning has only been studied by (Wen & Li, 2021; Ji et al., 2021) which show that CL learns semantically meaningful features from the data. In contrast, we show that CL may not learn *all* semantically relevant features.

Other important recent work (Saunshi et al., 2022; HaoChen & Ma, 2022) studied the role of inductive bias of the function class in the success of CL. Our analysis, however, is focused on understanding failure modes of CL i.e. class collapse and feature suppression.

Class Collapse in Supervised CL. Chen et al. (2022) empirically demonstrates *class collapse* on test data, but does not offer any rigorous theoretical explanation. Graf et al. (2021) proves that optimizing the supervised contrastive loss leads to class-collapsed training set representations. However, we show that there exist many minimizers with such class-collapsed training set representations and not all of them suffer from class collapse at *test time*. We also present the first theoretical characterization of class collapse at test time.

Feature Suppression in Unsupervised CL. Feature suppression has been empirically observed by Tian et al. (2020); Chen et al. (2021); Robinson et al. (2021) but we lack a theoretical formulation of this phenomenon. Li et al. (2020) shows that InfoNCE has local minimums that exhibit feature suppression, thus attributing this phenomenon to failure of optimizing the loss. However, Robinson et al. (2021) shows that the InfoNCE loss can be minimized by many models, some of which learn all task-relevant features, while others do not. We put forth the only theoretical characterization of feature suppression and consequently, use this understanding to suggest practical solutions to remedy this problem.

Joint Supervised and Unsupervised Contrastive Loss. Recently, several versions of loss functions that combine supervised and unsupervised contrastive losses (Islam et al., 2021; Chen et al., 2022) have been empirically observed to have superior transfer learning performance, by avoiding

class collapse. We provide the first theoretically rigorous analysis of which features the *minimum norm* global minimizer of the joint loss learns, provably demonstrating that it can avoid class collapse and feature suppression. To the best of our knowledge, this is the only theoretical result that can be used to understand the empirical success of joint losses.

3. Problem Formulation

3.1. Data distribution

We define data distribution $\mathcal{D}_{\text{orig}}$ below. Each example $(\mathbf{x}, y, y_{\text{sub}}) \in \mathcal{D}_{\text{orig}}$ is generated as follows:

$$\mathbf{x} = \mathbf{u} + \boldsymbol{\xi}, \quad \text{where}$$

$$\mathbf{u} = (y\phi_1 + \mu_1)\mathbf{v}_1 + (y_{\text{sub}}\phi_2 + \mu_2)\mathbf{v}_2 + (\rho_k\phi_k + \mu_k)\mathbf{v}_k,$$

and k is uniformly selected from $3, \dots, K$; and $y, y_{\text{sub}}, \rho_k$ are uniformly sampled from $\{-1, 1\}$.

Features and Noise. We assume features and noise form an orthonormal basis of \mathbb{R}^d , i.e., a set of unit orthogonal vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_d\}$ in \mathbb{R}^d . W.l.o.g., one can let \mathbf{v} 's be the standard basis, where the first K basis are feature vectors. $\{\phi_1, \dots, \phi_K\}$ are constants indicating the strength of each feature, and $\{\mu_1, \dots, \mu_K\}$ are the means of the corresponding entries in the feature vectors. In particular:

- Class Feature: \mathbf{v}_1 .
- Subclass Feature: \mathbf{v}_2 .
- (Class and subclass) irrelevant features:¹ $\mathbf{v}_3, \dots, \mathbf{v}_K$.
- Noise $\boldsymbol{\xi} \sim \mathcal{D}_{\boldsymbol{\xi}}$: $\mathcal{D}_{\boldsymbol{\xi}}$ is a uniform distribution over features $\sigma_{\boldsymbol{\xi}}\mathbf{v}_1, \dots, \sigma_{\boldsymbol{\xi}}\mathbf{v}_d$, where $\sigma_{\boldsymbol{\xi}}$ indicates the variance of the noise.²

We sample n examples from $\mathcal{D}_{\text{orig}}$ to form the original dataset $\hat{\mathcal{D}}_{\text{orig}}$.

Assumption 3.1 (Balanced Dataset). All combinations of $(y_i, y_{\text{sub},i}, k_i, \rho_i)$ are equally represented in $\hat{\mathcal{D}}_{\text{orig}}$.³

A Concrete Example of the Above Data Distribution. Let $y = 1$ be dogs and $y = -1$ be cats, $y_{\text{sub}} = 1$ if they are fluffy and $y_{\text{sub}} = -1$ if they are not-fluffy. Then $(\phi_1 + \mu_1)\mathbf{v}_1 + (\phi_2 + \mu_2)\mathbf{v}_2$ denotes a fluffy dog. Here, the background can be interpreted as an irrelevant feature: let $\rho_3 = 1$ for grass and $\rho_3 = -1$ for forest. Then $(\phi_1 + \mu_1)\mathbf{v}_1 + (\phi_2 + \mu_2)\mathbf{v}_2 + (\phi_3 + \mu_3)\mathbf{v}_3$ represents a fluffy dog on grass. Note that each

¹In the rest of the paper, we use irrelevant features to refer to features that may have semantic meaning but are irrelevant to class and subclass.

²This definition of noise is nearly identical to Gaussian noise $\mathcal{N}(0, \frac{\sigma_{\boldsymbol{\xi}}^2}{d}\mathbf{I}_d)$ in the high-dimensional regime but keeps the analysis clear. Our results can be extended to the Gaussian noise setting.

³This can be approximately achieved when n is sufficiently larger than K . While our analysis can be generalized to consider imbalanced data, this is outside the scope of this work.

example only selects one irrelevant feature, which mimics the real world, where examples do not necessarily have all types of objects in the background i.e. many examples have neither grass or forests as their background.

Rationale for Including Feature Means μ_i . In general, it is unreasonable to expect all features to have 0 expectation over entire data, thus we introduce μ to further generalize our analysis. We find that considering a non-zero mean for the subclass feature is sufficient to provide novel insights into class collapse (Theorem 4.5). Therefore, for clarity, we set all the μ 's except μ_2 to zero.

Relation to Sparse Coding Model. This data distribution is a variant of the sparse coding model which is usually considered as a provision model for studying the feature learning process in machine learning (e.g., (Zou et al., 2021; Wen & Li, 2021; Liu et al., 2021)). It naturally fits into many settings in machine learning, and in general mimics the outputs of intermediate layers of neural networks which have been shown to be sparse (Papayan et al., 2017). It is also used to model the sparse occurrences of objects in image tasks (Olshausen & Field, 1997; Vinje & Gallant, 2000; Foldiak, 2003; Protter & Elad, 2008; Yang et al., 2009; Mairal et al., 2014) and polysemy of words in language tasks (Arora et al., 2018).

3.2. Data Augmentation $\mathcal{A}(\cdot)$

For each example in $\hat{\mathcal{D}}_{\text{orig}}$, we generate m augmentations to form $\hat{\mathcal{D}}_{\text{aug}}$. We consider the following augmentation strategy: given an example $\mathbf{x} = \mathbf{u} + \boldsymbol{\xi}$, its augmentation is given by $\mathcal{A}(\mathbf{x}) = \mathbf{u} + \boldsymbol{\xi}'$, where $\boldsymbol{\xi}'$ is a new random variable from $\mathcal{D}_{\boldsymbol{\xi}}$ independent of $\boldsymbol{\xi}$. This is an abstract of augmentations used in practice where two augmentations from the same example share certain parts of the features and have the correlation between their noise parts removed or weakened.

Assumption 3.2 (High dimensional regime). d is at least $\omega(n^2m^2)$.

Assumption 3.3 (Sufficient sample size). The noise-to-sample-size ratio is not too large $\frac{\sigma_{\boldsymbol{\xi}}^2}{mn} = o(1)$.

3.3. Linear Model

We consider a linear model with p outputs. The model has weights $\mathbf{W} \in \mathbb{R}^{p \times d}$ and bias $\mathbf{b} \in \mathbb{R}^p$ where $p \geq 3$. The function represented by the model is $f_{\Theta}(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b}$, where we define $\Theta \in \mathbb{R}^{p \times (d+1)}$ as the concatenated parameter $[\mathbf{W} \ \mathbf{b}]$. We establish theoretical proofs of class collapse and feature suppression for linear model, and also empirically verified them for (non-linear) deep neural networks.

3.4. Loss function

For unsupervised contrastive learning, we use the unsupervised spectral contrastive loss popular in prior theoretical and empirical work (HaoChen et al., 2021; Saunshi et al., 2022; HaoChen & Ma, 2022) and for supervised contrastive learning, we consider the natural generalization of this loss to incorporate supervision. Let \mathcal{A}_i denote the set of augmentations in $\hat{\mathcal{D}}_{\text{aug}}$ generated from the i -th original example with $\mathcal{A}(\cdot)$. Let \mathcal{S}_{+1} and \mathcal{S}_{-1} denote the set of augmentations in $\hat{\mathcal{D}}_{\text{aug}}$ with class labels $+1$ and -1 , respectively. Let $\hat{\mathbb{E}}$ denote the empirical expectation. Then we have the following loss functions:

$$\mathcal{L}_{\text{UCL}}(\Theta) = -2\hat{\mathbb{E}}_{i \in [n], \mathbf{x} \in \mathcal{A}_i, \mathbf{x}^+ \in \mathcal{A}_i} [f_{\Theta}(\mathbf{x})^{\top} f_{\Theta}(\mathbf{x}^+)] \\ + \hat{\mathbb{E}}_{\mathbf{x} \in \hat{\mathcal{D}}_{\text{aug}}, \mathbf{x}^- \in \hat{\mathcal{D}}_{\text{aug}}} [(f_{\Theta}(\mathbf{x})^{\top} f_{\Theta}(\mathbf{x}^-))^2] \quad (1)$$

$$\mathcal{L}_{\text{SCL}}(\Theta) = -2\hat{\mathbb{E}}_{c \in \{-1, 1\}, \mathbf{x} \in \mathcal{S}_c, \mathbf{x}^+ \in \mathcal{S}_c} [f_{\Theta}(\mathbf{x})^{\top} f_{\Theta}(\mathbf{x}^+)] \\ + \hat{\mathbb{E}}_{\mathbf{x} \in \hat{\mathcal{D}}_{\text{aug}}, \mathbf{x}^- \in \hat{\mathcal{D}}_{\text{aug}}} [(f_{\Theta}(\mathbf{x})^{\top} f_{\Theta}(\mathbf{x}^-))^2]. \quad (2)$$

4. Simplicity Bias Contributes to Class Collapse in Supervised CL

We make two key observations through our theoretical analysis and experiments (henceforth we refer to class collapse at *test time* simply as ‘class collapse’):

1. Theoretically, not all global minimizers exhibit class collapse, but the *minimum norm* minimizer does.
2. Theoretically and empirically, when the model is trained using (S)GD, some subclasses are *provably* learned early in training. Empirically, however, those subclasses will eventually be unlearned i.e. S(GD) converges to minimizers that exhibit class collapse.

Altogether, these observations suggest that class collapse, which has been observed in practice when certain gradient-based algorithms are used to minimize the loss, cannot be explained by simply analyzing the loss function. This highlights the importance of studying the dynamics and inductive bias of training algorithms in contrastive learning.

4.1. What Minimizers Have Class Collapse?

We first define class collapse in terms of the alignment between the model weights and the subclass feature.

Definition 4.1 (Exact class collapse). We say exact class collapse happens at test time when:

$$\forall \beta \in \mathbb{R}^p, \Pr_{(\mathbf{x}, y, y_{\text{sub}}) \sim \mathcal{D}_{\text{orig}}} (y_{\text{sub}} \beta^{\top} f_{\Theta}(\mathbf{x}) > 0) = 1/2.$$

The definition means that no linear classifier on the embeddings of examples drawn from $\mathcal{D}_{\text{orig}}$ can predict the subclass

label with accuracy beyond random guess.⁴

This is different from class collapse on the training set which is not defined on the population set $\mathcal{D}_{\text{orig}}$ but on the training samples $\hat{\mathcal{D}}_{\text{orig}}$.

Proposition 4.2. For any $\Theta^* \in \min_{\Theta} \mathcal{L}_{\text{SCL}}(\Theta)$, we have $f_{\Theta^*}(\mathbf{x}_i) = f_{\Theta^*}(\mathbf{x}_j)$ for all $\mathbf{x}_i, \mathbf{x}_j$ in the training set $\hat{\mathcal{D}}_{\text{aug}}$ such that $y_i = y_j$.

This directly implies that minimizing the loss results in class collapse on the training set. However, the following theorem 4.3 shows that minimizing the loss does not necessarily lead to class collapse on the test set. To determine whether class collapse occurs, we need to determine whether the model learns the subclass feature. With a linear model, this exactly corresponds to constant alignment between weights and the subclass feature.

Theorem 4.3 (Minimizing $\mathcal{L}_{\text{SCL}} \not\Rightarrow$ Class Collapse). With high probability i.e. at least $1 - O(\frac{m^2 n^2}{d}) = 1 - o(1)$, there exists $\Theta^* = [\mathbf{W}^* \ \mathbf{b}^*]$ such that $\Theta^* \in \min_{\Theta} \mathcal{L}_{\text{SCL}}(\Theta)$ \mathbf{W}^* has constant alignment with subclass feature \mathbf{v}_2 i.e.

$$\|\mathbf{W}^* \mathbf{v}_2\| = \Omega(1).$$

Hence, there exists a linear classifier in the embedding space that can predict subclass labels almost perfectly. I.e.,

$$\exists \beta, \text{ s.t. } \Pr_{(\mathbf{x}, y, y_{\text{sub}}) \sim \mathcal{D}_{\text{orig}}} (y_{\text{sub}} \beta^{\top} \mathbf{W}^* \mathbf{x} > 0 | y) = 1 - o(1).$$

We prove the theorem in Appendix D. The proof utilizes Lemma C.1 which implies that, due to the high-dimensionality, the noise vectors have non-trivial effects on the empirical covariance matrix by rotating its kernel space. This results in the kernel space to have a $\Theta(\frac{\sigma_{\xi}}{\sqrt{mn}})$ alignment with the subclass feature. Since minimizers of the loss can behave arbitrarily on this kernel space, without any additional restriction, they can have any alignment with the subclass feature.

Next, we show that, the *minimum norm* minimizer exhibits class collapse.

Theorem 4.4 (Minimizing $\mathcal{L}_{\text{SCL}} + \text{Minimum Norm} \Rightarrow$ Class Collapse). Assume $\mu_2 = 0$. Let $\Theta^{**} = [\mathbf{W}^{**} \ \mathbf{b}^{**}]$ be the minimum norm minimizer of \mathcal{L}_{SCL} , i.e.,

$$\Theta^{**} = \arg \min_{\Theta^*} \|\Theta^*\|_F \text{ s.t. } \Theta^* \in \arg \min_{\Theta} \mathcal{L}_{\text{SCL}}(\Theta).$$

Then with high probability i.e. at least $1 - O(\frac{m^2 n^2}{d}) = 1 - o(1)$, \mathbf{W}^{**} has no alignment with subclass feature \mathbf{v}_2 i.e.

$$\|\mathbf{W}^{**} \mathbf{v}_2\| = 0.$$

⁴Actually we are able to analyze a stronger version of class collapse: $\Pr_{(\mathbf{x}, y, y_{\text{sub}}) \sim \mathcal{D}_{\text{orig}}} (f_{\Theta}(\mathbf{x}) | y_{\text{sub}}) = \Pr_{(\mathbf{x}, y, y_{\text{sub}}) \sim \mathcal{D}_{\text{orig}}} (f_{\Theta}(\mathbf{x}))$, which means the distributions of embeddings given and not given the subclass label are exactly the same. Nonetheless, we present this simpler formulation for clarity.

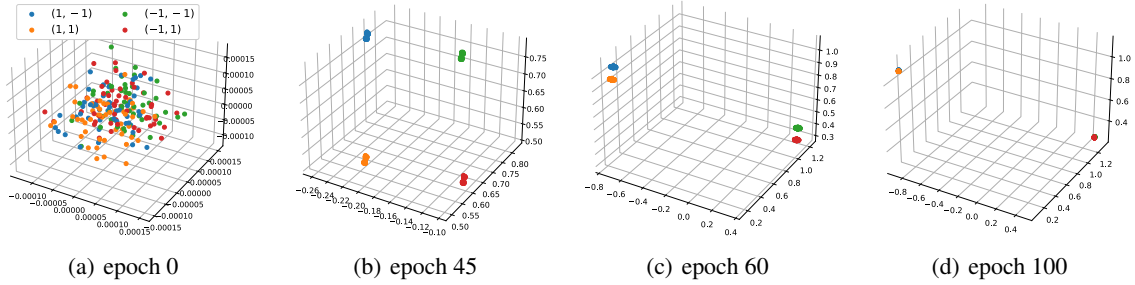


Figure 1. Visualization of the embedding space at different epochs. We let $p = 3$ so that we can see the whole embedding space from a 3D plot. Other parameters: $n = 1000, m = 5, d = 2000, K = 4, \phi_1 = \phi_2 = \phi_3 = \phi_4 = 1, \mu = 1, \sigma = 2, \sigma_0 = 0.001, \eta = 0.05$. Colors represent combinations of class and subclass labels (y, y_{sub}) . We use test examples for the plots. At epoch 45, the four groups of examples are well separated in the embedding space. However groups in the same classes are merged afterwards.




Figure 2. $\|\mathbf{W}_t \mathbf{v}_1\|$ and $\|\mathbf{W}_t \mathbf{v}_2\|$ at different epochs. Both features are learned early in training, but \mathbf{v}_2 is unlearned later.

This means class collapse occurs at test time (Definition 4.1), and no linear classifier does better than random guess for predicting subclass labels.

Theorems 4.3 and 4.4 show that minimizing the training loss does not necessarily lead to class collapse on test data, but does with additional constraint on the weights of the model. This is not due to a degenerate solution, as we show that both minimizers learn the class feature \mathbf{v}_1 (see corollary C.5).

4.2. Intriguing Properties of GD

We now further our theoretical characterization of class collapse by investigating the setting where \mathcal{L}_{SCL} is minimized by GD. This is an important step toward understanding class collapse in practice, where similar optimization algorithms are used to minimize the loss. Our findings indicate that it is likely the simplicity bias of commonly used optimization algorithms that eventually leads to class collapse.

We consider GD with a constant learning rate η . The weights are initialized from a Gaussian distribution, i.e., the initial weight Θ_0 has each of its element drawn from $\mathcal{N}(0, \frac{\sigma_0^2}{d})$. And the weights at training epoch t are given by:

$$\Theta_t = \Theta_{t-1} - \eta \nabla_{\Theta} \mathcal{L}_{\text{SCL}}(\Theta_{t-1}).$$

Early in Training Some Subclasses are Provably Learned. By analyzing the training dynamics of GD, we find that subclasses are learned early in training.


Theorem 4.5 (Early in training subclass features are learned). *Assume $\sigma_0 \sqrt{\frac{p}{d}} = o(1)$ and $\sigma_{\xi} = o(1)$. If the subclass feature has a constant non-zero mean such that $1 + \mu^2 > \phi_1^2$, then with probability at least $1 - O(\frac{m^2 n^2}{\text{poly}(p)} + \frac{1}{\text{poly}(p)}) = 1 - o(1)$ the following holds:*

- $\|\mathbf{W}_0 \mathbf{v}_2\| = o(1)$.
- $\exists t = O(\ln(\frac{1}{\sigma_0} \sqrt{\frac{d}{p}}))$, s.t. $\|\mathbf{W}_t \mathbf{v}_2\| = \Omega(1)$, and
- $\exists \beta$, s.t. $\Pr_{(\mathbf{x}, y, y_{\text{sub}}) \sim \mathcal{D}_{\text{orig}}}(y_{\text{sub}} \beta^{\top} \mathbf{W}_t \mathbf{x} > 0 | y) = 1 - o(1)$.

The above theorem shows that there exists an epoch where the weights have constant alignment with the subclass feature and produce distinguishable subclass embeddings (proof in Appendix G).

The key step of our analysis is showing that early in training, GD aligns the weights with the first eigenvector of the covariance matrix of class centers. This alignment grows exponentially faster than alignments with any other directions. When $1 + \mu^2 > \phi_1^2$, the subclass feature has a constant projection onto the first eigenvector and is therefore learned by the model.

More importantly, the same phenomenon can be observed in *neural networks*. We use SGD to train a ResNet18 (He et al., 2016) on CIFAR-100 (Krizhevsky et al., 2009) with supervised CL loss (Khosla et al., 2020) with 20 class (superclass) labels, and perform linear evaluation on embeddings of test data with 100 subclass (class) labels (see details in Appendix H). We observe that the subclass accuracy increases during the first 200 epochs before it starts to drop (Figure 3(a)). Some subclasses can even achieve a high accuracy around 80% (Figure 3(b)). This is surprising as it confirms that models trained with commonly used loss functions *do* learn subclass features early in training.



(a) Average subclass accuracy and class accuracy (b) Subclasses are initially learned well but later unlearned

Figure 3. (a) Average subclass accuracy and class accuracy. (b) Accuracy in subclasses ‘road’, ‘rocket’ and ‘sea’. In both plots, the subclass accuracy increases and then decreases, which confirms that subclasses are learned early in training before class collapse happens. The class accuracy only increases during training.

Empirical Evidence Showing that Class Collapse Eventually Happens in (S)GD. We simulate our theoretical analysis using numerical experiments to show that gradient descent converges to a minimizer that exhibits class collapse, despite learning subclasses early in training. We visualize the embeddings of test data at different epochs in Figure 1, and plot the alignment between weights and class/subclass features in Figure 2. Subclasses are perfectly separated and the weights align with both v_1 and v_2 after around 100 epochs of training. The model then starts unlearning v_2 which causes the alignment to drop, thus subclasses are merged in the embedding space. We also confirm that same conclusion holds for neural networks in realistic settings. In Figure 3, we see that the subclass accuracy drops after around 200 epochs of training and eventually reaches a low value. In contrast, the class accuracy does not drop during training.

Minimum Norm Minimizer Exhibits Class Collapse. Note that in Theorem 4.5, assuming $\mu \neq 0$ leads us to discovering that subclasses are learned early in training. Here, we extend Theorem 4.4 to this setting under asymptotic class collapse.

Definition 4.6 (Asymptotic Class Collapse). We say asymptotic class collapse happens when $\|Wv_2\| = O(\frac{\sigma_\xi}{\sqrt{mn}}) = o(1)$.

This definition implies that: (1) representations of subclasses are not well separated, hence it is nearly impossible to distinguish between them, and (2) the distinguishability of subclasses is at odds with generalization, which improves as number of augmented views per example m and size of training data n increase. Thus, while this definition is a relaxation of Definition 4.1, practically, this results in equally severe class collapse.

Theorem 4.7 (Extension of Theorem 4.4 for $\mu_2 \neq 0$). Let $\Theta^{**} = [W^{**} b^{**}]$ be the minimum norm minimizer of \mathcal{L}_{SCL} :

$$\Theta^{**} = \arg \min_{\Theta^*} \|\Theta^*\|_F \text{ s.t. } \Theta^* \in \arg \min_{\Theta} \mathcal{L}_{SCL}(\Theta).$$

Then with probability at least $1 - O(\frac{m^2 n^2}{d}) = 1 - o(1)$, asymptotic class collapse happens, i.e.,

$$\|W^{**}v_2\| = O(\frac{\sigma_\xi}{\sqrt{mn}}) = o(1).$$

4.3. Simplicity Bias of (S)GD

We reiterate our main findings:

1. Minimizing the supervised contrastive loss *does not* necessarily lead to class collapse.
2. However, *simpler* minimizers of the supervised contrastive loss (e.g. *minimum norm*) do suffer from class collapse.
3. Optimizing with (S)GD does learn the subclass features early in training, but eventually unlearns them, resulting in class collapse.

These coupled with the fact that (S)GD is known to have a bias towards simpler solutions (Kalimeris et al., 2019) prompt us to conjecture:

The simplicity bias of (S)GD leads it to unlearn subclass features, thus causing class collapse.

The simplicity bias of (S)GD has not been rigorously studied for CL, and our results indicate the surprising role it may play in class collapse. Note that, the supervised contrastive loss is different than common supervised objectives, where the role of such bias of (S)GD is understood better (Gunasekar et al., 2018; Soudry et al., 2018; Ji & Telgarsky, 2019; Wu et al., 2019; Lyu et al., 2021). Rather, the supervised CL objective can be reformulated as a matrix factorization objective (Eq. 39), where the debate on the bias of (S)GD (e.g., minimum norm (Gunasekar et al., 2017) or rank (Arora et al., 2019a; Razin & Cohen, 2020)) is still ongoing.

5. Understanding Feature Suppression in Unsupervised CL

Empirically, feature suppression can be observed due to a variety of reasons (Li et al., 2020; Chen et al., 2021; Robinson et al., 2021). Easy features for unsupervised CL are those that allow the model to discriminate between examples (highly discriminative). Here, we consider different ways irrelevant features can be easy (highly discriminative) and characterize how this can lead to feature suppression. We show that the types of feature suppression we consider can be largely attributed to insufficient embedding dimensionality and/or poor data augmentations. Surprisingly, we find again that the minimum norm simplicity bias is critical in explaining this phenomenon.

5.1. Feature Suppression due to Easy Irrelevant Features and Limited Embedding Space

In Theorem 5.1, we show that easy (discriminative) irrelevant features can suppress the class feature when the embedding dimensionality is limited. For clarity, we let $\mu_2 = 0$.

Theorem 5.1 (Feature Suppression 1). *Assume $p \leq K$. Let L be the $(K+1)$ -element tuple $[1, \phi_1^2, \phi_2^2, \frac{\phi_3^2}{K-2}, \dots, \frac{\phi_K^2}{K-2}]$ whose last K elements are the variances of features. If ϕ_1^2 is not among the p largest elements in L , then with probability at least $1 - O(\frac{m^2 n^2}{d}) = 1 - o(1)$: (1) there exists a global minimizer Θ^* of \mathcal{L}_{UCL} such that $\|\mathbf{W}^* \mathbf{v}_1\| = \Omega(1)$, (2) However, the minimum norm minimizer Θ^{**} satisfies $\|\mathbf{W}^{**} \mathbf{v}_1\| = 0$.*

We prove the theorem in Appendix E. The elements except the first one in tuple L can be interpreted as the variance of examples at each coordinate $\mathbf{v}_k, k = 1, 2, \dots, K$, which indicates how much the examples are discriminated by each feature. The theorem shows that when the embedding space is not large enough to represent all the K features (which requires $K+1$ dimensions), the minimum norm minimizer only picks the most discriminative ones. In practice, the embedding space in unsupervised CL is relatively low-dimensional (compared to input dimensionality) and thus the model cannot fit all the information about inputs into the embedding space. As is suggested by Theorem 5.1, if the training algorithm prefers functions with certain simple structures, only the easiest (most discriminative) features that can be mapped into the embedding space by less complex functions (e.g., smaller norm) are learned. The class features are suppressed if they are not amongst the easiest ones.

Remark 5.2. Following the same analysis we can also show that when ϕ_1 is among the p largest elements in L , i.e., the class feature is among the easiest (most discriminative) ones, the class feature \mathbf{v}_1 is learned by the minimum norm minimizer; when ϕ_1 is exactly on par with some other element as the p -th largest, there exist both minimum norm minimizers that learn and do not learn the class feature \mathbf{v}_1 .

Numerical Experiments with GD. Our theory for the minimum norm minimizer matches the experimental results for models trained with GD. We let $p = K$ and let $1 \geq \phi_2^2 \geq \frac{\phi_3^2}{K-2} \geq \dots \geq \frac{\phi_{K-1}^2}{K-2} > \phi_1^2$ so that ϕ_1^2 must be among the smallest two variances i.e. \mathbf{v}_1 is among the two most difficult features. Then we vary ϕ_K and see how the trained weights align with \mathbf{v}_1 . Consistent with Theorem 1, Figure 4 shows that \mathbf{v}_1 is suppressed when $\frac{\phi_K^2}{K-2} > \phi_1^2$. Interestingly, we also see that the result at $\frac{\phi_K^2}{K-2} = \phi_1^2$ diverges, indicating that GD can find both minimizers that learn and do not learn \mathbf{v}_1 when the variances at \mathbf{v}_1 and \mathbf{v}_K are the same.

Empirically Verifying Benefits of Larger Embedding Size. Theorem 5.1 also provides one practical solution for




Figure 4. **The irrelevant feature suppresses the class feature when its variance is beyond the variance of the class feature (the red vertical line).** We let $d = 2000, p = K, \phi_1 = 0.8, \phi_2 = 1, \mu = 0, \frac{\phi_k^2}{K-2} > \phi_1, \forall k \in [K-1]$ and vary ϕ_K . Thus whether ϕ_1^2 is among the p largest variances only depends on ϕ_K . We train the linear model to convergence. Plots show that the alignment between the trained weights and \mathbf{v}_1 drops when ϕ_K increases. We report the average of 10 runs. The result diverges at $\frac{\phi_K^2}{K-2} = \phi_1^2$ indicating that the model can learn either \mathbf{v}_1 or \mathbf{v}_K in this case.




Figure 5. Effect of embedding size on feature suppression in MNIST RandBit (Chen et al., 2021). Legends show the number of bits in the extra channel which indicates how easy (discriminative) the irrelevant features are. We observe that (1) increasing the easiness of irrelevant features exacerbates feature suppression; (2) increasing the embedding size alleviates feature suppression.

Table 2. Effect of embedding size on feature suppression in CIFAR-10/100 RandBit. ‘Acc’ refers to class accuracy and ‘Sub Acc’ refers to subclass accuracy. We see that increasing embedding size alleviates feature suppression, improving class/subclass accuracy.

w	CIFAR-10 RandBit		CIFAR-100 RandBit	
	Sub Acc	Acc	Sub Acc	Acc
4	34.38	86.73	11.67	23.53
64	71.96	96.82	34.11	52.32
128	76.69	97.65	38.51	57.40

feature suppression due to limited embedding size: increasing the embedding size so that every feature can be learned by the model. To provide empirical evidence for this, we conduct two sets of experiments:

First, we train 5-layer convolutional networks on the RandomBit dataset with the same setup as in (Chen et al., 2021), but we vary the embedding size (see details in Appendix H). Varying the # bits in the extra channel intuitively controls how discriminative the irrelevant feature are, i.e., how easy-to-learn it is for CL. In this setting, the

random bit can suppress the MNIST digits. We make two observations in Figure 5.1: (1) with a fixed embedding size, increasing easiness (number of random bits) of the irrelevant features exacerbates feature suppression; (2) with a fixed easiness of irrelevant features, increasing the embedding size alleviates feature suppression.

Second, we train ResNet18 (He et al., 2016) on the CIFAR-10/100 RandBit Dataset, constructed similarly to the MNIST RandBit dataset but with images from CIFAR-10/100 (Krizhevsky et al., 2009) (see Appendix H.1). For CIFAR-10, we use 2 random bits, and for CIFAR-100, we use one random bit as the class irrelevant features. Table 2 presents the test performance for different values of the model width w , where a larger w indicates a larger embedding size (see Appendix H.3 for details). On both datasets, increasing the embedding size alleviates feature suppression, leading to improvements in both class and subclass accuracies. We also provide additional experiments and discussion in Appendix H.3. Both experimental results confirm the conclusion drawn from the theoretical analysis.

5.2. Feature Suppression due to High-dimensional Irrelevant Features and Imperfect Augmentation

Empirically, another form of feature suppression has been observed that cannot be remedied by larger embedding dimensionality (Li et al., 2020). We characterize this form of feature suppression by defining easy irrelevant features as being: (1) drawn from a high dimensional space so that the collection of irrelevant features is large and discriminating based on irrelevant features is easier, (2) less altered by data augmentation compared to the class feature.

For (1), formally we assume $K = \omega(n^2)$, as opposed to assumption 3.1 which implies that K is smaller than n . A consequence of this assumption is that with high probability the n original examples each have a unique irrelevant feature. For (2) we consider the following imperfect data augmentation:

Definition 5.3 (Imperfect data augmentation $\mathcal{A}'(\cdot)$). For a given example $\mathbf{x} = \boldsymbol{\mu} + \boldsymbol{\xi} \in \hat{\mathcal{D}}_{\text{orig}}$,

$$\mathcal{A}'(\mathbf{x}) = \mathbf{u} + \zeta' \mathbf{v}_1 + \zeta'' \mathbf{v}_2 + \boldsymbol{\xi}',$$

where $\zeta' \sim \mathcal{N}(0, \sigma_{\zeta'}^2)$, $\zeta'' \sim \mathcal{N}(0, \sigma_{\zeta''}^2)$, $\sigma_{\zeta'}^2, \sigma_{\zeta''}^2 \neq 0$ and $\boldsymbol{\xi}'$ is a new random variable drawn from $\mathcal{N}(\boldsymbol{\xi}, \boldsymbol{\Sigma}_{\boldsymbol{\xi}})$ with $\text{rank}(\boldsymbol{\Sigma}_{\boldsymbol{\xi}}) \leq \frac{m}{2}$.

In the definition, the data augmentation adds small perturbations (ζ' and ζ'') to class and subclass features, weakly alters the noise, but preserves the irrelevant features. For example, on Colorful-Moving-MNIST (Tian et al., 2020) constructed by assigning each MNIST digit a background object image selected randomly from STL-10, the colorful background objects are high-dimensional and the colors are invariant to data augmentations without color distortion.

Theorem 5.4 (Feature Suppression 2). *If $K = \omega(n^2)$ and augmentation is $\mathcal{A}'(\cdot)$, with probability $\geq 1 - o(\frac{n^2 m^2}{d} + \frac{1}{n}) = 1 - o(1)$, the minimum norm minimizer $\Theta^* = [\mathbf{W}^*, \mathbf{b}^*]$ satisfies $\|\mathbf{W}^* \mathbf{v}_1\| = 0$.*

This theorem shows that feature suppression can happen even when embedding dimensionality p is arbitrarily large and helps understand empirical observations made both in our work (Figure 5.1, the line with 15 bits) and previous work. For example Li et al. (2020) showed that on Colorful-Moving-MNIST, the colorful background can suppress learning the digits especially when color distortion is not used in augmentation, and increasing embedding size does not address the issue.

In conclusion, Theorem 5.4 highlights that designing data augmentations that disrupt the highly-discriminative irrelevant features is a key to addressing feature suppression.

6. Combining Supervised and Unsupervised CL Losses Can Avoid Both Class Collapse and Feature Suppression

We now consider the following loss which is a weighted sum of the supervised and unsupervised CL loss functions:

$$\mathcal{L}_{\text{joint}, \beta}(\Theta) = \beta \mathcal{L}_{\text{SCL}}(\Theta) + (1 - \beta) \mathcal{L}_{\text{UCL}}(\Theta).$$

Similar loss functions have been proposed recently with notable empirical success. For example, Chen et al. (2022) put forth a weighted sum of supervised CL loss and class-conditional InfoNCE (which has similar effect as \mathcal{L}_{UCL} in our setting) to avoid class collapse. Islam et al. (2021) empirically observed that the joint objective of supervised and unsupervised contrastive loss leads to better transferability of the learned models than their supervised counterparts. However, we still lack a theoretical understanding of why this weighted sum of losses can outperform both losses.

From our investigation of class collapse and feature suppression, the benefit of the joint objective $\mathcal{L}_{\text{joint}}$ becomes evident: the unsupervised term in $\mathcal{L}_{\text{joint}}$ increases the chance of learning features that do not appear relevant to the labels but might be useful for downstream tasks, while the supervised term in $\mathcal{L}_{\text{joint}}$ ensures that even hard-to-learn class features are learnt. Thus, $\mathcal{L}_{\text{joint}}$ can learn rich representations capturing more task relevant information than either $\mathcal{L}_{\text{UCL}}(\Theta)$ or $\mathcal{L}_{\text{SCL}}(\Theta)$. We show below that with an appropriate choice of β , $\mathcal{L}_{\text{joint}}$ can provably succeed where \mathcal{L}_{SCL} fails due to collapse and \mathcal{L}_{UCL} fails due to feature suppression (for clarity, we let $\mu = 0$).

Theorem 6.1. *W.L.O.G., assume $\phi_3 \geq \phi_4 \geq \dots \geq \phi_K$. If $p \leq K$, $\phi_2^2 > \frac{\phi_{p-2}^2}{K-2}$ and $\phi_1^2 < \frac{\phi_{p-1}^2}{K-2}$, then by Theorem 4.4 the minimum norm minimizer of \mathcal{L}_{SCL} suffers from class collapse and by Theorem 5.1 the minimum norm minimizer*

Table 3. Joint loss alleviates class collapse on CIFAR-100.

Loss	Subclass Acc
SCL	26.11
Joint loss ($\beta = 0.8$)	41.37

Table 4. Joint loss alleviates feature suppression on MNIST RandBit.

Loss	Class Acc
UCL	61.21
Joint loss ($\beta = 0.5$)	79.37

Table 5. Joint loss alleviates both class collapse and feature suppression on CIFAR-100 RandBit.

Loss	Subclass Acc	Class Acc
SCL	28.13	61.10
UCL	34.11	52.32
Joint loss ($\beta = 0.8$)	35.72	63.94

of \mathcal{L}_{UCL} suffers from feature suppression. However, for constant $\beta \in (0, 1)$, the minimum norm minimizer of $\mathcal{L}_{joint, \beta}$, denoted by $\Theta^* = [\mathbf{W}^* \ \mathbf{b}^*]$, satisfies $\|\mathbf{W}^* \mathbf{v}_1\| = \Omega(1)$ and $\|\mathbf{W}^* \mathbf{v}_2\| = \Omega(1)$.

Empirically Verifying Benefits of the Joint Loss. We empirically examine the impact of the joint loss on MNIST RandBit, CIFAR-100, and CIFAR-100 RandBit. The training details are in Appendix H.2. The results indicate that the joint loss significantly improves performance in scenarios where SCL suffers from class collapse (Table 3) and UCL suffers from feature suppression (Table 4). Furthermore, on CIFAR-100 RandBit dataset, where both phenomena can occur simultaneously, the joint loss effectively alleviates both issues (Table 5).

7. Discussion

Negative Impact of Simplicity Bias in Deep Learning.

The simplicity bias of optimization algorithms has been studied as a key beneficial factor in achieving good generalization (Gunasekar et al., 2017; 2018; Soudry et al., 2018; Ji & Telgarsky, 2019; Wu et al., 2019; Lyu et al., 2021). However, our study reveals the negative impact of simplicity bias in CL. In fact, it has also been conjectured to lead to undesirable outcomes in other scenarios, such as learning spurious correlations (Sagawa et al., 2020) and shortcut solutions (Robinson et al., 2021). We hope our study can inspire further theoretical characterization of the negative role of simplicity bias in these scenarios, thereby deepening our understanding and fostering potential solutions.

Connection to Neural Collapse. Neural collapse (NC) (Papayan et al., 2020) refers to the collapse of representations within each class in supervised learning. Similar to the rationale in this study, overparameterized models that exhibit NC on training data can demonstrate different behaviors on test data due to their capacity to implement training set NC

in various ways, and it is worth considering whether current theoretical frameworks (Han et al., 2021; Zhu et al., 2021; Zhou et al., 2022b;a; Lu & Steinerberger, 2022; Fang et al., 2021) can effectively capture NC on test data. In fact, the empirical results in (Hui et al., 2022) emphasize the distinction between NC on training and test data, as there can be an inverse correlation between the two. Our results suggest that analyzing the learned features and considering the inductive bias of training algorithms can aid in this distinction.

Theoretical Characterization of Class Collapse in (S)GD.

The results in Section 4.2 highlight the need for theoretical characterization of class collapse in (S)GD. We provide two potential approaches for future investigation. (1) Given that the objective can be reformulated as matrix factorization (Eq. 39), and our Theorems 4.4 and 4.7 on minimum norm minimizer, it is reasonable to investigate whether the implicit bias of (S)GD is to seek the minimum norm solution. We note that understanding the implicit bias in matrix factorization is a longstanding pursuit in the machine learning community, with no consensus reached thus far (see Appendix I.1). Hence, further effort is still needed. (2) As elaborated in Appendix I.2, the gradient consists of two terms with distinct roles. One promotes alignment with the subclass feature, while the other counteracts its influence. The relative scale of these two terms undergoes a phase transition (Figure 6), and analyzing this can provide insights into class collapse.

8. Conclusion

To conclude, we present the first theoretically rigorous characterization of the failure modes of CL: class collapse and feature suppression at test time. We explicitly construct minimizers of supervised contrastive loss to show that optimizing this loss does not necessarily lead to class collapse. Then we show that the minimum norm minimizer does exhibit class collapse. Our analysis also reveals a peculiar phenomenon for supervised CL, when optimized with (S)GD: subclass features are learned early in training and then unlearned. To analyze feature suppression, we consider two formalisms of easy features that can prevent learning of class features and provably attribute feature suppression to insufficient embedding space and/or imperfect data augmentations; thus, motivating practical solutions to this problem. The unified framework we develop to determine which features are learnt by CL allows us to also offer the only theoretical justification for recent empirical proposals to combine unsupervised and supervised contrastive losses. Perhaps, most surprisingly, our findings from this theoretical study indicate that simplicity bias of (S)GD is likely the driving factor behind class collapse and feature suppression.

Acknowledgment. This research was supported by the National Science Foundation CAREER Award 2146492.

References

- Arora, S., Li, Y., Liang, Y., Ma, T., and Risteski, A. Linear algebraic structure of word senses, with applications to polysemy. *Transactions of the Association for Computational Linguistics*, 6:483–495, 2018.
- Arora, S., Cohen, N., Hu, W., and Luo, Y. Implicit regularization in deep matrix factorization. *Advances in Neural Information Processing Systems*, 32, 2019a.
- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019b.
- Chen, M., Fu, D. Y., Narayan, A., Zhang, M., Song, Z., Fatahalian, K., and Ré, C. Perfectly balanced: Improving transfer and robustness of supervised contrastive learning. In *International Conference on Machine Learning*, pp. 3090–3122. PMLR, 2022.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. February 2020. doi: 10.48550/arXiv.2002.05709. URL <https://arxiv.org/abs/2002.05709v3>.
- Chen, T., Luo, C., and Li, L. Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 34:11834–11845, 2021.
- Chuang, C.-Y., Robinson, J., Lin, Y.-C., Torralba, A., and Jegelka, S. Debaised Contrastive Learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 8765–8775. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/63c3d3ddcc7b23daale42dc41f9a44a873-Abstract.html>.
- Fang, C., He, H., Long, Q., and Su, W. J. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.
- Foldiak, P. Sparse coding in the primate cortex. *The handbook of brain theory and neural networks*, 2003.
- Graf, F., Hofer, C., Niethammer, M., and Kwitt, R. Dissecting supervised contrastive learning. In *International Conference on Machine Learning*, pp. 3821–3830. PMLR, 2021.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., and Valko, M. Bootstrap your own latent: A new approach to self-supervised Learning, September 2020. URL <http://arxiv.org/abs/2006.07733>. arXiv:2006.07733 [cs, stat].
- Gunasekar, S., Woodworth, B. E., Bhojanapalli, S., Neyshabur, B., and Srebro, N. Implicit regularization in matrix factorization. *Advances in Neural Information Processing Systems*, 30, 2017.
- Gunasekar, S., Lee, J. D., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. *Advances in Neural Information Processing Systems*, 31, 2018.
- Han, X., Pappayan, V., and Donoho, D. L. Neural collapse under mse loss: Proximity to and dynamics on the central path. *arXiv preprint arXiv:2106.02073*, 2021.
- HaoChen, J. Z. and Ma, T. A theoretical study of inductive biases in contrastive learning. *arXiv preprint arXiv:2211.14699*, 2022.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hinton, G. E., Osindero, S., and Teh, Y.-W. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006.
- Hui, L., Belkin, M., and Nakkiran, P. Limitations of neural collapse for understanding generalization in deep learning. *arXiv preprint arXiv:2202.08384*, 2022.
- Islam, A., Chen, C.-F. R., Panda, R., Karlinsky, L., Radke, R., and Feris, R. A broad study on the transferability of visual representations with contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8845–8855, 2021.
- Ji, W., Deng, Z., Nakada, R., Zou, J., and Zhang, L. The power of contrast for feature learning: A theoretical analysis. *arXiv preprint arXiv:2110.02473*, 2021.
- Ji, Z. and Telgarsky, M. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*, pp. 1772–1798. PMLR, 2019.
- Kalimeris, D., Kaplun, G., Nakkiran, P., Edelman, B., Yang, T., Barak, B., and Zhang, H. Sgd on neural networks learns functions of increasing complexity. *Advances in neural information processing systems*, 32, 2019.

- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Laurent, B. and Massart, P. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pp. 1302–1338, 2000.
- Lee, J. D., Lei, Q., Saunshi, N., and Zhuo, J. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323, 2021.
- Li, T., Fan, L., Yuan, Y., He, H., Tian, Y., Feris, R., Indyk, P., and Katabi, D. Addressing feature suppression in unsupervised visual representations, 2020. URL <https://arxiv.org/abs/2012.09962>.
- Liu, H., HaoChen, J. Z., Gaidon, A., and Ma, T. Self-supervised learning is more robust to dataset imbalance. *arXiv preprint arXiv:2110.05025*, 2021.
- Lu, J. and Steinerberger, S. Neural collapse under cross-entropy loss. *Applied and Computational Harmonic Analysis*, 59:224–241, 2022.
- Lyu, K., Li, Z., Wang, R., and Arora, S. Gradient descent on two-layer nets: Margin maximization and simplicity bias. *Advances in Neural Information Processing Systems*, 34:12978–12991, 2021.
- Mairal, J., Bach, F., Ponce, J., et al. Sparse modeling for image and vision processing. *Foundations and Trends® in Computer Graphics and Vision*, 8(2-3):85–283, 2014.
- Olshausen, B. A. and Field, D. J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.
- Papayan, V., Romano, Y., and Elad, M. Convolutional neural networks analyzed via convolutional sparse coding. *The Journal of Machine Learning Research*, 18(1):2887–2938, 2017.
- Papayan, V., Han, X., and Donoho, D. L. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- Protter, M. and Elad, M. Image sequence denoising via sparse and redundant representations. *IEEE transactions on Image Processing*, 18(1):27–35, 2008.
- Ranzato, M., Poultney, C., Chopra, S., and Cun, Y. Efficient learning of sparse representations with an energy-based model. *Advances in neural information processing systems*, 19, 2006.
- Razin, N. and Cohen, N. Implicit regularization in deep learning may not be explainable by norms. *Advances in neural information processing systems*, 33:21174–21187, 2020.
- Robinson, J., Sun, L., Yu, K., Batmanghelich, K., Jegelka, S., and Sra, S. Can contrastive learning avoid shortcut solutions?, 2021. URL <https://arxiv.org/abs/2106.11230>.
- Sagawa, S., Raghunathan, A., Koh, P. W., and Liang, P. An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, pp. 8346–8356. PMLR, 2020.
- Saunshi, N., Ash, J., Goel, S., Misra, D., Zhang, C., Arora, S., Kakade, S., and Krishnamurthy, A. Understanding contrastive learning requires incorporating inductive biases. *arXiv preprint arXiv:2202.14037*, 2022.
- Shorack, G. R. and Shorack, G. *Probability for statisticians*, volume 951. Springer, 2000.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive estimation reveals topic posterior information to linear models. *J. Mach. Learn. Res.*, 22:281–1, 2021a.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pp. 1179–1206. PMLR, 2021b.
- Tsai, Y.-H. H., Wu, Y., Salakhutdinov, R., and Morency, L.-P. Self-supervised learning from a multi-view perspective. *arXiv preprint arXiv:2006.05576*, 2020.
- Vinje, W. E. and Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*, 287(5456):1273–1276, 2000.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.

- Wen, Z. and Li, Y. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pp. 11112–11122. PMLR, 2021.
- Wu, Y., Póczos, B., and Singh, A. Towards understanding the generalization bias of two layer convolutional linear classifiers with gradient descent. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1070–1078. PMLR, 2019.
- Yang, J., Yu, K., Gong, Y., and Huang, T. Linear spatial pyramid matching using sparse coding for image classification. In *2009 IEEE Conference on computer vision and pattern recognition*, pp. 1794–1801. IEEE, 2009.
- Zhou, J., Li, X., Ding, T., You, C., Qu, Q., and Zhu, Z. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *International Conference on Machine Learning*, pp. 27179–27202. PMLR, 2022a.
- Zhou, J., You, C., Li, X., Liu, K., Liu, S., Qu, Q., and Zhu, Z. Are all losses created equal: A neural collapse perspective. *arXiv preprint arXiv:2210.02192*, 2022b.
- Zhu, Z., Ding, T., Zhou, J., Li, X., You, C., Sulam, J., and Qu, Q. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34:29820–29834, 2021.
- Zou, D., Cao, Y., Li, Y., and Gu, Q. Understanding the generalization of adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371*, 2021.