
Understanding Contrastive Learning Requires Incorporating Inductive Biases

Nikunj Saunshi¹ Jordan T. Ash² Surbhi Goel² Dipendra Misra² Cyril Zhang²
Sanjeev Arora¹ Sham Kakade³ Akshay Krishnamurthy²

Abstract

Contrastive learning is a popular form of self-supervised learning that encourages augmentations (views) of the same input to have more similar representations compared to augmentations of different inputs. Recent attempts to theoretically explain the success of contrastive learning on downstream classification tasks prove guarantees depending on properties of *augmentations* and the value of *contrastive loss* of representations. We demonstrate that such analyses, that ignore *inductive biases* of the function class and training algorithm, cannot adequately explain the success of contrastive learning, even *provably* leading to vacuous guarantees in some settings. Extensive experiments on image and text domains highlight the ubiquity of this problem – different function classes and algorithms behave very differently on downstream tasks, despite having the same augmentations and contrastive losses. Theoretical analysis is presented for the class of linear representations, where incorporating inductive biases of the function class allows contrastive learning to work with less stringent conditions compared to prior analyses.

1. Introduction

Recently, representation functions learned via contrastive learning have transformed machine learning. Using unlabeled data, a representation function is learnt by generating simple augmentations of each datapoint and by enforcing, via a suitable loss function, that (1) augmentations of a single datapoint tend to be clustered (2) augmentations of different datapoints tend to be far apart. Such representa-

tions give competitive classification performance —via even a linear classifier— on a host of downstream tasks, bringing us closer to the old dream of machine learners capable of generalization across different data distributions and tasks.

We lack a conceptual framework for understanding such phenomena — which is unsurprising, since good quantitative understanding of generalization is lacking even for single task and single data distribution. However, deriving even partial conceptual understanding could help push the field forward, and researchers have begun to grapple with this task (Arora et al., 2019; Tosh et al., 2021a; HaoChen et al., 2021). The current paper seeks to provide guidance for further development of this nascent theory¹ using simple experiments and theoretical analysis. A common thread in these existing theories is the following components: (1) Quantifying how data augmentations *implicitly* encode downstream class labels. (2) Demonstrating how representations with small contrastive loss can uncover this implicit structure and do well on downstream tasks.

Recent works formalize (1) via assumptions that end up implying that the augmentation distributions of inputs from the same class have significant overlap, but there is little overlap for inputs from different classes. For example, distributions of augmentations of different dog images tend to be similar to each other, but their union has little overlap with distributions of augmentations of cat images. Arora et al. (2019) — which predates the recent wave of methods — assume that points in the same class share the same augmentation distribution, and use this to show that the contrastive loss is a surrogate to the downstream performance. Since methods like SimCLR (Chen et al., 2020) do not appear to satisfy such assumptions, recently HaoChen et al. (2021) gave a more refined analysis under milder assumptions, that requires only some overlap in augmentation distributions, such that the resultant graph of connections due to overlaps within a class is dense. Again, it can be shown that a low-dimensional representation that is near-optimal in the contrastive loss is guaranteed to linearly separate the downstream classes.

¹Our title is a clear allusion to Zhang et al. (2017), which highlighted a gap between deep learning phenomena and classical ML theory, motivating development of better theoretical understanding.

¹Department of Computer Science, Princeton University
²Microsoft Research, New York City ³Departments of Computer Science & Statistics, Harvard University. Correspondence to: Nikunj Saunshi <nsaunshi@cs.princeton.edu>.




Figure 1. Cartoon of our theoretical example. The downstream labels (and thus classification loss L_{clf}) are determined by a few relevant attributes (e.g. cat or dog?), and the augmentations perturb irrelevant attributes (e.g. grayscale, random crop). Without restricting the function class for the contrastive pretraining task, there exist perfect (f) and spurious (g) augmentation-invariant representations which both minimize the contrastive loss L_{cont} . However, minimizing using a linear representation class is always guaranteed to succeed with these augmentations (Section 3).

Note that properties of the class of representation functions (architecture) or the training algorithm (SGD, Adam etc.) make no appearance in the above analyses; but only properties of the *augmentation distributions* and *value of contrastive loss* of representations. This is understandable since currently theory is unable to pinpoint why different real-life architectures differ in their capabilities, or to pinpoint the implicit bias of the training algorithms. Nevertheless it raises interesting questions: *Is the contrastive loss indeed a good indicator of downstream performance? Do augmentations overlap sufficiently enough in practice to explain the success of contrastive learning? Can contrastive learning succeed even when there is little to no overlap?* In a nutshell, the current paper suggests via experiments and simple theory that the answers are, respectively: *No, Maybe no, Yes*. In particular, ignoring the architecture and the training algorithm can make the current theoretical analyses of contrastive learning vacuous. We present three key phenomena:

- **Function class sensitivity.** Downstream performance of a representation depends not just on its contrastive loss, but it is also sensitive to the function class (architecture) and training procedure used to learn it.
- **Brittleness of transfer.** Minimizing the contrastive loss to optimality can sometimes have a non-monotonic, deleterious effect on downstream performance, despite the augmentations being effective for some function classes.
- **The disjoint augmentations regime.** When augmentation distributions for inputs do not overlap, it can be shown that any function-class-agnostic analysis (subsumes prior work) provably leads to vacuous guarantees. However on-overlapping augmentations can still be “informative”, and contrastive learning with appropriate function classes can

succeed, a phenomenon unexplained by existing theory.

Organization. We define the contrastive losses and downstream performance in Section 2, and summarize prior theoretical results and how they ignore inductive biases. In Section 3 we describe a simple synthetic setting that elucidates all of the aforementioned phenomena. A pictorial depiction in Figure 1 demonstrates the existence of bad contrastive solutions, despite the augmentations satisfying intuitive properties. These ideas are grounded through theoretical results in Section 4, which includes lower bounds for function class agnostic analyses and upper bounds that are sensitive to the function class of linear representations. Finally we describe various experimental setups in Section 5.

1.1. Related work

Contrastive learning has been very successful at solving downstream tasks by learning representations from similar pairs of data obtained using temporal information (Wang & Gupta, 2015; Logeswaran & Lee, 2018) or different views or augmentations of inputs (Dosovitskiy et al., 2014; Hjelm et al., 2019; Wu et al., 2018; Bachman et al., 2019; Tian et al., 2020a; He et al., 2020; Chen et al., 2020; Chen & He, 2021; Gao et al., 2021). Given its empirical success, there has been significant interest in the theory of contrastive learning, from various perspectives. Most relevant to us are learning theoretic analyses (Arora et al., 2019; Tosh et al., 2021b;a; HaoChen et al., 2021; Wang et al., 2022) and their follow ups (Nozawa & Sato, 2021; Ash et al., 2021). These study the downstream linear classification performance of learned representation, by making assumptions about the data and augmentation distributions.

Contrastive learning has also been studied (1) from a mutual information maximization view (Oord et al., 2018; Hjelm et al., 2019; Bachman et al., 2019); (Tschannen et al., 2020) points out certain issues with this view, (2) using an information theoretic framework Tsai et al. (2021); fails to explain downstream success via simple linear classifiers, (3) through properties like alignment and uniformity on the sphere Wang & Isola (2020), (4) under certain latent variable data generative processes (Zimmermann et al., 2021; Von Kügelgen et al., 2021), and (5) through a causality perspective (Mitrovic et al., 2021). On the optimization front, (Wen & Li, 2021) study the feature learning process of contrastive learning with gradient dynamics on a two layer network, under a sparse coding model. The theory of noise contrastive estimation (Gutmann & Hyvärinen, 2010) has been a useful motivation for negative sampling based objectives. On the empirical side, there are studies on identifying useful augmentation properties (Tian et al., 2020b).

Non-contrastive methods, with no negative samples, (Chen & He, 2021; Grill et al., 2020) rely on tricks like stop-grad

to avoid representation collapse. Dimension collapse of representations has also been studied (Jing et al., 2022). Unlike these works, the brittleness of transfer we study is neither due to ill-designed objectives nor due to training degeneracies. It is fundamental to data distributions and arises out of existence of spurious solutions. A related idea of feature suppression (Chen et al., 2021) and shortcut solutions found by contrastive learning was recently studied in Robinson et al. (2021) in certain stylized settings, with a proposed fix through better augmentations strategies. We instead study the role of inductive bias of function classes in avoiding such shortcut solutions. (Abnar et al., 2022) analyze upstream to downstream transfer for supervised pre-training, complementing our experiments for unsupervised pre-training, whereas Wu et al. (2020a) studies negative transfer for multi-task learning. Robinson et al. (2021) also observe negative transfer for contrastive learning in some settings. Finally, there are theoretical works for other types of self-supervised learning (Bansal et al., 2021), including methods like context reconstruction (Lee et al., 2021) and language model (Saunshi et al., 2021), studying their benefits on downstream tasks.

2. Preliminaries

Here we formalize the problem of learning useful representations via contrastive learning for downstream classification.

Notation. We use $[n]$ for the set $\{1, \dots, n\}$. $\mathcal{U}(S)$ denotes uniform distribution over a set S . For a vector $v \in \mathbb{R}^n$, we denote $v_{:i} \in \mathbb{R}^i$ and $v_{i:} \in \mathbb{R}^{n-i}$ to be the sub-vector of first $i \in [n]$ and last i coordinates respectively. For sets P, Q , we use P^Q to denote the set of functions from Q to P .

Augmentations. We use $\bar{\mathcal{X}}$ to denote the set of all (unaugmented) samples and denote their marginal distribution as $\mathcal{D}_{\bar{\mathcal{X}}}$. \mathcal{X} denotes the set of all augmented data. For an input $\bar{x} \in \bar{\mathcal{X}}$, we define the corresponding augmentation distribution over \mathcal{X} as $\mathcal{A}(\cdot | \bar{x})$. For instance, augmentations for an image \bar{x} can correspond to applying a sequence of random transformations such as random cropping, Gaussian blur, and color jitter. The distributions $\mathcal{D}_{\bar{\mathcal{X}}}$ and \mathcal{A} together induce a marginal distribution $\mathcal{D}_{\mathcal{X}}$ over augmentations.

Contrastive self-supervised learning. The goal is to learn a representation function $f : \mathcal{X} \rightarrow \mathbb{R}^d$ that maps augmentations to d -dimensional vectors by encouraging representations of ‘‘similar pairs’’ of augmentations to be closer to each other, compared to representations of random pairs. A common strategy to pick a similar pair (x, x^+) is to pick two augmentations of the same input. Formally we define this distribution of similar pairs \mathcal{D}_{sim} as follows

$$(x, x^+) \sim \mathcal{D}_{\text{sim}} \equiv \bar{x} \sim \mathcal{D}_{\bar{\mathcal{X}}}; x, x^+ \sim_{\text{i.i.d.}} \mathcal{A}(\cdot | \bar{x})$$

The negative sampling distribution, denoted by \mathcal{D}_{neg} , is picked to be the augmentation marginal distribution $\mathcal{D}_{\mathcal{X}}$. There are several variants of the contrastive loss, a popular one being the *SimCLR loss* (Chen et al., 2020).

$$L_{\text{SimCLR}}(f) = \mathbb{E}_{\substack{(x, x^+) \sim \mathcal{D}_{\text{sim}}, \\ x_{1:n}^- \sim \mathcal{D}_{\text{neg}}^n}} \left[-\log \left(\frac{e^{f(x)^\top f(x^+)}}{e^{f(x)^\top f(x^+)} + \sum_{i=1}^n e^{f(x)^\top f(x_i^-)}} \right) \right] \quad (1)$$

Intuitively the contrastive loss aims to make $f(x)^\top f(x^+)$ larger compared to $f(x)^\top f(x_i^-)$. Another variant proposed in HaoChen et al. (2021) is the spectral contrastive loss:

$$L_{\text{spec}}(f) = \mathbb{E}_{\substack{(x, x^+) \\ \sim \mathcal{D}_{\text{sim}}}} [-2f(x)^\top f(x^+)] + \mathbb{E}_{\substack{x, x^- \\ \sim \mathcal{D}_{\text{neg}}^d}} \left[(f(x)^\top f(x^-))^2 \right] \quad (2)$$

We will use L_{cont} to refer to a generic contrastive loss, either L_{SimCLR} or L_{spec} or something else.

Downstream task. We assume these involve binary classification². If the ground-truth labeling function is $\bar{y}^* : \bar{\mathcal{X}} \rightarrow \{\pm 1\}$, the quality of representation $f : \bar{\mathcal{X}} \rightarrow \mathbb{R}^d$ is captured by how well it allows *linear classification*:

$$L_{\text{clf}}(\bar{f}; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x}} \left[\mathbb{1} \left\{ \bar{y}^*(\bar{x}) \left(\bar{f}(\bar{x})^\top w \right) < 0 \right\} \right] \quad (3)$$

Since the representation function is trained to map augmentations to vectors, its behavior on unaugmented inputs can be undefined. We evaluate downstream performance on original inputs $\bar{\mathcal{X}}$ by using the average augmentation representation $f_{\mathcal{A}} : \bar{\mathcal{X}} \rightarrow \mathbb{R}^d$, defined as:

$$f_{\mathcal{A}}(\bar{x}) = \mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})} [f(x)], L_{\text{clf}}(f; \bar{y}^*) := L_{\text{clf}}(f_{\mathcal{A}}; \bar{y}^*) \quad (4)$$

Experimentally such an average gives better performance than the standard un-averaged approach.

Transfer Bounds. We introduce an abstraction of *transfer function* \mathcal{T} to capture prior analyses (Arora et al., 2019; Tosh et al., 2021a; HaoChen et al., 2021). It translates performance on contrastive loss to performance on the downstream task as follows:

$$L_{\text{clf}}(f; \bar{y}^*) \leq \mathcal{T}(\Gamma, L_{\text{cont}}(f), d), \Gamma = (\mathcal{D}_{\bar{\mathcal{X}}}, \mathcal{A}, \bar{y}^*, L_{\text{cont}}) \quad (5)$$

These guarantees only depend on (1) problem dependent quantities like input marginals $\mathcal{D}_{\bar{\mathcal{X}}}$, properties of augmentations \mathcal{A} , downstream label \bar{y}^* , form of contrastive loss

²We consider binary tasks mostly for simplicity. Extensions of our results (lower bounds for function class agnostic analyses and upper bound guarantees for linear representations) to more than two classes are not difficult.

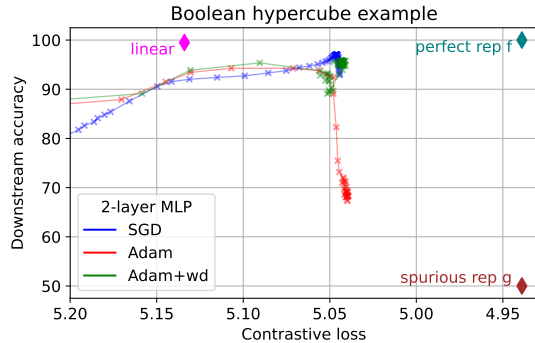


Figure 2. Contrastive loss \rightarrow accuracy transfer plots for the Boolean hypercube example. There exist global minimizers of L_{cont} with perfect (top right) and worst possible (bottom right) downstream classification error L_{clf} . The representations learned by two-layer neural networks are very sensitive to training configuration. With a smaller (linear) function class, the contrastive loss minimizer gives a nearly-perfect downstream classifier.

L_{cont} , (2) contrastive loss $L_{\text{cont}}(f)$ of the representation f and (3) its dimensionality d . Typically \mathcal{T} is monotone non-decreasing function of $L_{\text{cont}}(f)$ and so the above bound justifies minimizing the contrastive loss. A common property of augmentations and labels that transfer bounds assume is *overlap* between augmentations of images from the same class. For instance, Arora et al. (2019) effectively assume full overlap, that is, all images from the same class have identical augmentation distribution, and these distributions are different for different classes. HaoChen et al. (2021) relax this requirement to a spectral quantity that depends on the ratio of overlap between augmentation distributions of the same class and those of different classes, leading to a bound $L_{\text{clf}}(f) \leq \alpha(L_{\text{cont}}(f) - \min_{f^*} L_{\text{cont}}(f^*)) + \beta$, where f^* is a minimizer of the contrastive loss, and α, β quantify the overlap in augmentations. These bounds place a premium on the value of contrastive loss of f , but are agnostic to any other properties of f , like the representation function class \mathcal{F} it belongs to or how it was trained. We are interested in transfer bounds that also incorporate these effects. A simple abstraction that incorporates the function class is

$$L_{\text{clf}}(f; \bar{y}^*) \leq \mathcal{T}(\Gamma, L_{\text{cont}}(f), \mathcal{F}), \Gamma = (\mathcal{D}_{\bar{x}}, \mathcal{A}, \bar{y}^*, L_{\text{cont}}) \quad (6)$$

A bound like this, unlike the one in Equation (5), reflects that the downstream performance at a particular value of contrastive loss depends also on the representation function class, which we also find to be true in many experiments.

Not about generalization. The above bounds only deal with upstream and downstream *population losses*. Thus the role of function class bias is not for guaranteeing good generalization properties with finite samples, as in supervised learning. It pertains to how good performance on a pre-

Table 1. Contrastive loss and downstream accuracy of various representation classes and training procedures used to train on the hypercube example. There exist minimizers of the contrastive loss which transfer to perfect downstream classifiers, and ones which are no better than random guessing. Function class matters (MLP vs. linear vs. any representation), as does the training algorithm.

Representation	Contrastive loss	Accuracy (%)
$\exists f$ (perfect)	4.939	100
$\exists g$ (spurious)	4.939	50
MLP + Adam	5.039 ± 0.001	74.1 ± 4.3
MLP + Adam + wd	5.040 ± 0.002	89.5 ± 4.9
Linear	5.134 ± 0.002	99.5 ± 0.1

training task can guarantee good downstream performance, as will become evident in the following sections.

3. Warm-up: contrastive learning on hypercubes

In this section, we present a simple but illustrative example that succinctly highlights brittleness of transfer and the importance of incorporating inductive biases in transfer bounds. Since contrastive learning tries to make representations invariant to augmentation transformations, ideal augmentations are those that retain parts of the input that can predict the downstream label, but modify parts that are less important for the label. We now describe a simple example on the Boolean hypercube that captures these intuitions.

Example 1. The input set is $\bar{\mathcal{X}} = \{\pm 1\}^D$, the augmentation set is $\mathcal{X} = \mathbb{R}^D$. Downstream label \bar{y}^* is linear in the first $k \ll D$ coordinates.

$$\bar{y}^*(\bar{x}) = \text{sign}(w^{*\top} \bar{x}_{:k}), w^* \in \mathbb{R}^k$$

Augmentation distribution $\mathcal{A}(\cdot | \bar{x})$ for input $\bar{x} \in \bar{\mathcal{X}}$ randomly scales down the last k coordinates while keeps the first k coordinates unchanged³. Formally it is defined as

$$x \sim \mathcal{A}(\cdot | \bar{x}) \equiv \tau \sim \mathcal{U}((0, 1]), x_{:k} = \bar{x}_{:k}, x_{k:} = \tau \bar{x}_{k:}$$

where $\mathcal{U}((0, 1])$ is the uniform distribution over $(0, 1]$.

We experimentally study this example using two function classes to minimize the contrastive objective: MLP⁴, linear. Results from Table 1 and Figure 2 are summarized below.

Transfer is sensitive to function class and algorithm. Firstly, we notice that despite having much worse (higher) contrastive loss compared to MLP, linear representation has significantly better downstream performance. Secondly, Figure 2 suggests that even for the same MLP architecture, the

³Generalizable to downscaling subsets of $\bar{x}_{k:}$, analogous to downscaling different aspects of images like color, sharpness

⁴CNN behave similarly to MLP.

training algorithm (Adam v/s SGD, weight decay or not) can drastically affect the downstream accuracy.

Brittleness of transfer and disjoint augmentations. Although the augmentations in the example seem intuitively helpful, there exists a spurious representation that has much smaller contrastive loss than all architectures, but random guessing performance downstream. This, as we show in later sections, is a consequence of having disjoint augmentation distributions, since the original input can be recovered from an augmentation by simply performing $\bar{x} = \text{sign}(x)$. The existence of a bad minimizer of the contrastive loss also gives us a concrete case where contrastive learning can succeed with an appropriate function class, but the success cannot be explained by any function class agnostic analysis. With this backdrop, we present our theoretical results next.

4. Lower bounds and improved analysis

In this section we discuss the role of overlap in augmentations and function class in theoretical guarantees. We first show in the disjoint augmentation regime (augmentation distributions do not overlap), that *any* function class independent analyses will lead to vacuous bounds, which includes many previous analysis. Delving deeper into the most recent results from HaoChen et al. (2021), we discuss reasons for failure, even in approximately disjoint augmentations. Finally we present guarantees for contrastive learning with a linear representation function class that is sensitive to the function class and allows for weaker assumptions on augmentations. We instantiate this bound for the hypercube example, provably explaining the good performance of linear representations on disjoint augmentations.

4.1. Lower bound for disjoint augmentations

In this section, we prove that brittle transfer exists much more generically whenever the augmentation distributions for different inputs do not overlap, generalizing our observations from the hypercube example in the previous section.

Definition 4.1 (Disjoint augmentations). *We say the augmentation distributions are disjoint if for all distinct inputs $\bar{x}_1, \bar{x}_2 \in \bar{\mathcal{X}}$, augmentation distributions $\mathcal{A}(\cdot \mid \bar{x}_1)$ and $\mathcal{A}(\cdot \mid \bar{x}_2)$ have disjoint supports.*

Disjoint augmentations can be problematic because the contrastive loss only encourages separating individual instances, but does not encourage making classes linearly separable. We formalize this argument in the next two lemmas by showing that any representation f can be transformed — by shuffling identities of examples — to a new representation \hat{f} that has lower (or equal) contrastive loss but near-trivial downstream performance. An immediate consequence is that any function class agnostic analysis (including previ-

ous analyses) will necessarily leads to vacuous downstream guarantees. We establish this in two settings where the exact choice of contrastive objective is not critical; results hold for both L_{SimCLR} and L_{spec} and we abbreviate these by L_{cont} below. First we consider unconstrained representations.

Lemma 4.1. *Let $|\bar{\mathcal{X}}| = N$ and $d = \mathcal{O}(N/\log_2(N))$. Suppose the labeling function is balanced, i.e. $\sum_i y_i^* = 0$, and let $\mathcal{D}_{\mathcal{X}}$ be uniform over $\bar{\mathcal{X}}$. If the augmentation distribution is disjoint, then for any $f^* : \mathcal{X} \rightarrow \mathbb{R}^d$ there exists a $\hat{f} : \mathcal{X} \rightarrow \mathbb{R}^d$ such that:*

$$L_{\text{cont}}(\hat{f}) \leq L_{\text{cont}}(f^*), \text{ \& } L_{\text{clf}}(\hat{f}) \geq 1/2 - \mathcal{O}\left(\sqrt{d \log(N)/N}\right)$$

Since it is common to use normalized representations in practice (e.g., to have Euclidean norm 1), we also establish a similar result for this case.

Lemma 4.2. *In the setup of Lemma 4.1, suppose further that representations are constrained (to any given set) and that the augmentation distributions satisfy the following: There exists a fixed source of randomness W and a deterministic map $T : (\bar{x}, w) \mapsto x$ that is invertible in w for any \bar{x} such that $x \sim \mathcal{A}(\cdot \mid \bar{x}) \equiv w \sim W, x = T(\bar{x}, w)$. Then the conclusion of Lemma 4.1 holds.*

We prove both statements jointly in Appendix C. Both lemmas show that when the representation dimension is small relative to the size of the input space (as is typical) and the augmentations are disjoint, there exists a *global minimizer* of the contrastive loss with vacuous transfer to downstream. The extra assumption in Lemma 4.2 is that the augmentation generation protocol uses a common source of randomness, which is actually satisfied in many practical scenarios. For instance, the same sequence of transformations like random cropping, color jittering etc. are applied to all images to generate augmentations. The other assumptions, e.g., that the labeling function \bar{y}^* is balanced and that $\mathcal{D}_{\mathcal{X}}$ is uniform, are technical in nature and can be potentially relaxed. In Appendix C.1 we discuss some prior lower bounds (Robinson et al., 2021; Wang et al., 2022) for disjoint augmentations and normalized representations.

A corollary of these results is that any transfer learning bound that only depends on the value of the contrastive loss cannot be meaningful in the disjoint augmentations setting.

Corollary 4.1. *In the setup of Lemma 4.1 or Lemma 4.2, consider a transfer function \mathcal{T} bounding the downstream performance as $L_{\text{clf}}(f; \bar{y}^*) \leq \mathcal{T}(\Gamma, L_{\text{cont}}(f), d)$ as in Equation (5), where $\Gamma = (\mathcal{D}_{\bar{\mathcal{X}}}, \mathcal{A}, \bar{y}^*, L_{\text{cont}})$ are problem dependent but function class independent quantities. Suppose \mathcal{T} is monotonic in its second argument, then for all $f : \mathcal{X} \rightarrow \mathbb{R}^d$:*

$$\mathcal{T}(\Gamma, L_{\text{cont}}(f), d) \geq 1/2 - \tilde{\mathcal{O}}\left(\sqrt{d/|\bar{\mathcal{X}}|}\right)$$

Takeaways. The above lower bounds suggest that previous analyses for contrastive learning are vacuous in the

disjoint augmentation setting, due to existence of bad minimizers of the contrastive loss. The brittleness of transfer for disjoint augmentations is also observable in practice, as in the first row of [Table 1](#) for the hypercube example. Vision and NLP experiments in [Section 5](#) also demonstrate this phenomenon, for more expressive function classes.

4.2. Prior theoretical results and failure modes

We briefly discuss the results from [HaoChen et al. \(2021\)](#) and delve deeper into how their function class agnostic nature leads to poor guarantees even for *approximately* disjoint augmentations. Their analysis considers the spectral loss $L_{\text{spec}}(f)$ from [Equation \(2\)](#). A key component of their analysis is an *augmentation graph* constructed using \mathcal{A} , whose spectral properties characterize how much overlap there is in augmentations. This is a weighted graph on augmentations \mathcal{X} with adjacency matrix $A \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ with entries $A[x, x'] = \mathcal{D}_{\text{sim}}(x, x')$, i.e. similar augmentations have edges. The normalized adjacency matrix, a central object in spectral graph theory, is defined as $A_{\circ} \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ with entries $A_{\circ}[x, x'] = \frac{\mathcal{D}_{\text{sim}}(x, x')}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)\mathcal{D}_{\mathcal{X}}(x')}}.$

Canonical results in spectral graph theory connect the eigenvalues $\lambda_1 \leq \dots \leq \lambda_{|\mathcal{X}|}$ of the normalized Laplacian $L_{\circ} = I - A_{\circ}$ to density of edges in the graph: denser graphs have larger eigenvalues. For representation dimension d , [HaoChen et al. \(2021\)](#) roughly make two key assumptions: (1) any partition of the graph into $\mathcal{O}(d)$ partitions is dense i.e. λ_{d+1} is high, (2) the partition of downstream classes is sparse. The condition (2) is the same as saying augmentations of different classes do not overlap much. Under these assumptions, they show the following transfer bound:

Theorem 4.3 (Theorem 4.2 from [HaoChen et al. \(2021\)](#)). *If $\lambda_1 \leq \dots \leq \lambda_{|\mathcal{X}|}$ are the eigenvalues of the normalized Laplacian $L_{\circ} = I - A_{\circ}$ for the augmentation distribution \mathcal{A} , and if the augmentations can predict the original input labels with probability $1 - \alpha$, then for any $d' \in [d]$ and representation f we have*

$$L_{\text{clf}}(f; \bar{y}^*) \lesssim c_1 \frac{\alpha}{\lambda_{d'+1}} + c_2 \frac{(L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*)) d'}{(\lambda_{d+1} - \lambda_{d'})^2}$$

where $L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*)$ is the sub-optimality of f .

Firstly we note that the above bound is function class independent and fits the abstraction from [Equation \(5\)](#). If augmentations are disjoint, then augmentations of an image \bar{x} will be connected to each other in the augmentation graph, but disconnected from all other input augmentations. Thus the graph A will have $|\bar{\mathcal{X}}|$ connected components, implying that the first $|\bar{\mathcal{X}}|$ eigenvalues of the Laplacian L_{\circ} are 0, i.e. $\lambda_i = 0$ for $i \in [|\bar{\mathcal{X}}|]$.⁵ So any representation dimension

⁵Results in spectral graph theory equate number of connected

$d < |\bar{\mathcal{X}}|$ leads to vacuous bounds in [Theorem 4.3](#). This again happens because the global minimizer of L_{spec} is not unique, and some of those could be terrible on downstream, as in our proof for [Lemma 4.1](#).

Approximately disjoint augmentations. We show that the above bound does not scale well even when there is very little overlap in the augmentation distributions. To quantify approximate disjointness, we consider the problem of predicting the original input \bar{x} that could have generated an augmentation x , as a classification problem.

Definition 4.2. *Augmentation distribution \mathcal{A} is $1 - \tau$ disjoint when the minimum error achievable in the input identification task, i.e. predicting the input \bar{x} that could generate an augmentation x , is at most τ . Formally this means*

$$\inf_{g: \mathcal{X} \rightarrow \bar{\mathcal{X}}} \mathbb{E}_{\bar{x}} \left[\mathbb{E}_{x \sim \mathcal{A}(\cdot|\bar{x})} [\mathbb{1}\{g(x) \neq \bar{x}\}] \right] \leq \tau \quad (7)$$

The augmentation distributions are disjoint if and only if one can perfectly predict \bar{x} from x , i.e. under the disjoint augmentation setting from [Definition 4.1](#), it is easy to see that \mathcal{A} is 1-disjoint. The following result shows that the eigenvalues and eigen-gaps in [Theorem 4.3](#) will be small if the augmentation classification accuracy is high.

Lemma 4.4. *Suppose again that $|\bar{\mathcal{X}}| = N$ and $\mathcal{D}_{\bar{\mathcal{X}}} = \mathcal{U}(\bar{\mathcal{X}})$. If the augmentations are $1 - \tau$ disjoint (defined in [Definition 4.2](#)), i.e. average accuracy of predicting \bar{x} from x is $1 - \tau$, then for $d' \in [d]$,*

$$\lambda_{d+1} - \lambda_{d'} \leq \lambda_{d+1} \leq \frac{2\tau}{(1 - d/N)}$$

Thus for a small representation dimension $d \ll N$, the guarantees from [Theorem 4.3](#) are non-vacuous only when $L_{\text{cont}}(f) \leq \inf_{f^*} L_{\text{cont}}(f^*) + \mathcal{O}(\tau^2)$, which is a stringent condition to satisfy. The proof of this is presented in [Appendix C.2](#). Experiments, as in [Section 3](#) and later in [Section 5](#), suggest that contrastive learning can succeed even without augmentation overlap. Given that prior analysis fail, we now proceed to show function class dependent guarantees that can show tighter bounds.

4.3. Function class dependent transfer guarantees

We present guarantees for a representation that incorporates the function class in addition to the contrastive loss and augmentations. Results in this section are for the spectral contrastive loss defined in [Equation \(2\)](#). For simplicity we assume that the input and augmentation sets are finite.

We consider a representation class that is linear in fixed features $\phi: \mathcal{X} \rightarrow \mathbb{R}^D$, defined as

$$\mathcal{F}_{\phi} = \{f(\cdot) = W^{\top} \phi(\cdot) \mid W \in \mathbb{R}^{D \times d}\} \quad (8)$$

components to the multiplicity of eigenvalue 0 of the Laplacian.

A crucial property of the function class \mathcal{F}_ϕ is that it is expressive enough to solve the downstream task on *augmentations* well, even if not sample efficiently. To formalize this, we define the following metrics

Definition 4.3 (Expressivity). *For any augmentation representations $h : \mathcal{X} \rightarrow \mathbb{R}^d$ on augmentation labels $g : \mathcal{X} \rightarrow \{\pm 1\}$, the regression loss is defined as*

$$L_{reg}(h; g) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{x \sim \mathcal{X}} \left[(w^\top h(x) - g(x))^2 \right]$$

Definition 4.4 (Inconsistency). *We define inconsistency of a labeling function $g \in \{\pm 1\}^{\mathcal{X}}$ on augmentations w.r.t. ground truth labeling $\bar{y}^* \in \{\pm 1\}^{\bar{\mathcal{X}}}$ on original inputs as*

$$\Delta_{\mathcal{A}}(g, \bar{y}^*) = \mathbb{E}_{\bar{x}} \left[\mathbb{E}_{x \sim \mathcal{A}(\cdot|\bar{x})} [\mathbb{1}\{g(x) \neq \bar{y}^*(\bar{x})\}] \right] \quad (9)$$

Denote the augmentation mean features as $\phi_{\mathcal{A}} = \mathbb{E}_{x \sim \mathcal{A}(\bar{x})} [\phi(x)]$ and covariance as $\Sigma(\phi) = \mathbb{E}_x [\phi(x)\phi(x)^\top]$.

We now present the upper bound result.

Theorem 4.5. *Let $\lambda_1, \dots, \lambda_D$ be the eigenvalues of $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$ in increasing order. Then for every $d' \in [d]$, a representation $f \in \mathcal{F}_\phi$ will satisfy*

$$L_{clf}(f; \bar{y}^*) \leq \frac{\min_{g \in \{\pm 1\}^{\mathcal{X}}} 4(2\Delta_{\mathcal{A}}(g, \bar{y}^*) + \sqrt{L_{reg}(\phi; g)})}{\lambda_{d'+1}} + \frac{2d'(L_{spec}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{spec}(f^*))}{(1 - \lambda_{d'}) (\lambda_{d+1} - \lambda_{d'})^2}.$$

Firstly note that the this transfer bound is indeed of the form $L_{clf}(f; \bar{y}^*) \leq \mathcal{T}(\Gamma, L_{spec}(f), \mathcal{F}_\phi)$ as in Equation (6), connecting to the function class sensitivity discussed in Section 1. This is because the eigenvalues, and the inconsistency and regression metrics in the above bound depend on the features ϕ that defines the function class, unlike the guarantee from HaoChen et al. (2021) in Theorem 4.3, where the eigenvalues depend only on the data distributions. We discuss the result in more detail in Appendix B.4 and present its proof in Appendix B. The result can in fact recover Theorem 4.3, in the special case of ϕ being full rank, i.e. $D = |\mathcal{X}|$; we discuss this in Appendix B.4.1. However when ϕ is not full rank, we get a function class dependent bound that can potentially provide non-vacuous guarantees under weaker assumptions, as will be evident in the next part.

Revisiting hypercube setting. We provide theoretical explanations for observations from Section 3 by instantiating our lower and upper bounds for the hypercube example.

Corollary 4.2. *Consider the setting from Example 1. Suppose the classifier is $w^* = e_1 \in \mathbb{R}^k$, so the downstream*

label is $\bar{y}^(\bar{x}) = \bar{x}_1$. Furthermore, let the feature map ϕ be an identity mapping, i.e. $\phi(x) = x$. In this setting, the following statements are true:*

- (a) *Function class-agnostic guarantees will be vacuous.*
- (b) *For any $f \in \mathcal{F}_\phi$, we have $L_{clf}(f; \bar{y}^*) \leq 32k(L_{spec}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{spec}(f^*))$.*

Result (b) suggests that finding an approximate minimizer of the contrastive loss, within the class \mathcal{F}_ϕ , is sufficient to guarantee good downstream performance; this explains the good performance of linear representation in Table 1. The proof of this part is presented in Appendix A.1. Result (a) explains the presence of spurious representation in the same table and also why prior analyses fail on this example, and follows from Corollary 4.1.

5. Experiments

Our theoretical examples and analysis show that prior function class agnostic transfer bounds can be near vacuous, particularly when augmentation distributions are disjoint (Corollary 4.1) or near disjoint (Lemma 4.4). Furthermore they suggest that meaningful downstream guarantees for contrastive learning would need to depend not only on the contrastive loss but also on the representation function class and possibly training algorithm. In this section, we ask, in the context of modern contrastive learning pipelines: (a) *how sensitive to the function class is the contrastive loss \rightarrow downstream accuracy transfer in practice?*, (b) *do augmentations sufficiently overlap in standard settings* and (c) *can contrastive learning work when there is little to no overlap?*

5.1. CIFAR-10 + SimCLR experiments

We consider the setting of CIFAR-10 image classification, where the augmentation distribution for contrastive learning is derived from the popular SimCLR protocol (Chen et al., 2020). An augmentation is generated by applying a series of transformations (each with some probability) to an image, like random cropping, horizontal flipping, color jittering, grayscaling and Gaussian blurring (see Appendix D.2).

We run contrastive learning with standard function classes (architectures): convolutional networks (ResNet) (He et al., 2016), Vision Transformers (ViT) (Dosovitskiy et al., 2021) and MLP-Mixer (Tolstikhin et al., 2021). Like in the hypercube example, we compare the transfer performance of different function class and algorithmic choices, as contrastive pre-training proceeds, by plotting the trajectories through $(L_{cont}(f), 1 - L_{clf}(f))$ space of different setups. Figure 3 summarizes our findings at a glance and we list the key observations below:

- **Effect of function class.** While standard training using

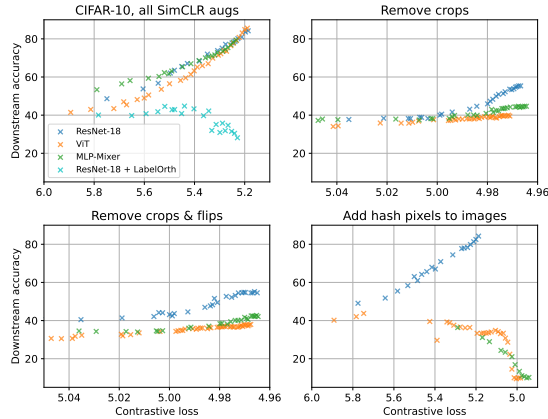


Figure 3. Contrastive loss \rightarrow accuracy transfer plots for CIFAR-10 with ResNet-18, ViT, and MLP-Mixer architectures for different augmentations. **TL**: Full pipeline of augmentations from SimCLR (Chen et al., 2020). **TR**: Remove random cropping. **BL**: Remove random cropping and horizontal flip. **BR**: Add “hash pixels” to each image, as described in Section 5.1 to ensure that there is no overlap in augmentations. Here, we observe transfer collapse for the ViT and MLP-Mixer architectures, as they overfit to these uninformative features; ResNet-18 ignores these pixels.

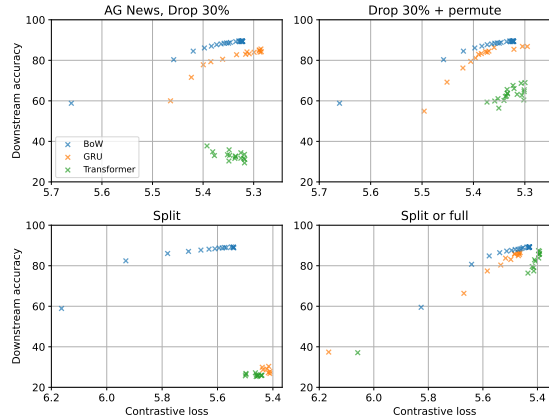


Figure 4. Contrastive loss \rightarrow accuracy transfer plots for AG News with bag-of-words (BoW), GRU and Transformer architectures. Augmentations are: **TL**: Drop random 30% of tokens. **TR**: Drop random 30% of tokens and randomly permute the rest. **BL**: Either the first half or second half of the input. **BR**: Either the first half, second half or the full input. In all cases BoW makes the contrastive loss reasonably small and does quite well downstream ($\sim 90\%$), but Transformer (and GRU) demonstrate brittleness of transfer for some augmentations.

the full pipeline of SimCLR augmentations (top left) displays very similar behaviors for different architectures, removal of certain transformations like random cropping (top right) and horizontal flipping (bottom left), leading to “weaker” augmentations, can accentuate the difference in transfer performances between different architectures.

- **Label-orthogonal training.** All architectures behaving similarly for the full SimCLR pipeline (top left) might superficially suggest that the role of inductive biases is not that significant, and that function class agnostic guarantees are good enough to explain the practical success of contrastive learning for these augmentations. However for the same augmentations, we can find pathological representations that have small contrastive loss but poor downstream performance, by introducing an adversarial modification to the training algorithm and minor tweak to ResNet architecture that is detailed in Appendix D.2.2. This suggests that guarantees depending only on the contrastive loss, but not the function class or algorithm, cannot explain the effectiveness of contrastive learning with standard architectures and augmentations.
- **Hash pixels.** The difference in architectures is even more prominent in the hash pixels setting (bottom right). Here we non-destructively⁶ modify images and augmentations in order to force the augmentation to be in the disjoint augmentation regime, as defined in Definition 4.1. In this case, ViT and MLP-Mixer representations make

⁶Only add a small number of pseudorandom pixels in random locations of a 2D image; this kind of noise can be easily removed and do not visually change images by much.

the contrastive loss much smaller than ResNet, but have close to random guessing downstream performance. ResNet training however is unaffected by this hash pixel modification, and it does well on the downstream task, despite being far from minimizing the contrastive loss. This experiment not only highlights the difference in function classes, but also concretely demonstrate a case where *contrastive learning can succeed despite the augmentation distributions being disjoint*. Details on the hash pixel augmentation are in Appendix D.2.1.

An important point to note is that the contrastive losses and downstream accuracies in Figure 3 are measured on unseen data and are thus reflective of the population versions of these metrics; thus the difference in transfer performance is not an issue of generalization. Further details of experimental setups and hyperparameters are in Appendix D.2. The next question we tackle is understanding how much overlap there is in augmentations for standard settings.

5.2. Are we in the disjoint augmentation regime?

Central to previous theory is the assumption that there exists overlap between augmentations distributions of data within a class. To get a sense of amount of overlap, we set up a classification task of predicting the image \bar{x} that could have generated an augmentation x , similar to Definition 4.2. The standard ResNet-18 architecture is modified to have 5000 output classes, one for each image in a CIFAR-10 class. We train the model to take augmentations from a fixed CIFAR-10 class and predict index of the original image generating it.

Performance is measured on *unseen* augmented data from inputs from the same class, by evaluating the accuracy of predicting the original input. The results of these experiments in Figure 7 (averaged over the 10 classes) suggest that with extremely high accuracy, the trained model is able to identify the image given an augmentation, with accuracies higher than 99.5% for different augmentation types. This suggests that we may be closer to the disjoint augmentation setting than we think. An important note is that the overlap here was only measured on images in the training set, not on the full population distribution as assumed in prior work (HaoChen et al., 2021). It is still an open question whether sufficient overlap exists in the population distribution or not. The experiment on the training set is intended as a starting point towards this question.

5.3. Experiments on text domain

In order to understand if our findings apply beyond images, we evaluate the contrastive pipeline on text domain. We use the AG News classification dataset (Zhang et al., 2015), where inputs are new articles and the 4 classes correspond to topics of the articles. Inspired by recent augmentations strategies like word/span deletion and word reordering (Wu et al., 2020b; Giorgi et al., 2021; Meng et al., 2021; Yan et al., 2021), we consider four simple augmentations for our study: (i) *Drop*: randomly drop 30% tokens but keep the order of remaining tokens, (ii) *Drop+Permute*: randomly drop 30% tokens and randomly permute the remaining tokens, (iii) *Split*: randomly return either the left half or right half of the text, and (iv) *Split+Full*: randomly return from the full text, its left half, or its right half.

We run contrastive learning with three models on this dataset. The first is a simple Bag-of-Word (BoW) model that learns word embeddings and returns the average word embedding of a text, second model is a Gated Recurrent Unit (GRU) (Chung et al., 2014) and the last is a Transformer (Vaswani et al., 2017) model. Both GRU and Transformer are unidirectional and we pick the final word representation for contrastive learning and downstream linear classification evaluation. All models are trained from scratch to minimize the contrastive loss, without the auxiliary MLM objective as in some prior work (Wu et al., 2020b; Giorgi et al., 2021; Meng et al., 2021). See Appendix D.3 for details on the experimental setup and hyperparameters.

Figure 4 visualizes the training trajectories through the $(L_{\text{cont}}(f), 1 - L_{\text{clf}}(f))$ space, i.e. the contrastive learning \rightarrow downstream accuracy transfer plots. We observe that for all augmentations, BoW has the best downstream performance, despite having worse contrastive loss. For the drop augmentation (top left), the BoW and GRU plots might suggest that the augmentation is good; however the Transformer

model leads to brittle transfer, i.e. it fails to solve the downstream task despite achieving lower contrastive loss. This kind of difference in transfer performance is unexplained by existing function class agnostic theoretical guarantees. Since the BoW representation is order invariant, we also test the augmentation that permutes tokens after dropping 30% of them (top right). This change does help the downstream accuracy of Transformer, however it does not completely bridge the gap. While the split augmentation (bottom left) works for BoW, both GRU and Transformer display brittle transfer. However a simple change of including the original text as an augmentation leads to both GRU and Transformer doing well downstream. This is particularly surprising, since including the identity augmentation only decreases the probability of overlap between augmentations, undesirable based on our current understanding of contrastive learning. In Appendix D.3.2 we visualize the learned representations and discuss reasons for the brittle transfer. The main takeaway is that a less expressive function class can succeed with weaker augmentations, while more expressive ones like GRU and Transformer might require stronger augmentations to transfer well to downstream tasks. This phenomenon is not well understood by current theory and deserves more exploration.

6. Conclusion

Contrastive learning has emerged as a unifying paradigm for building flexible learners that can adapt to many tasks. It is imperative to understand it better at a conceptual and mathematical level. The current paper lays out simple experiments and theoretical examples which suggest gaps in our current understanding. Filling these gaps will require incorporating the inductive bias of the deep nets being used, which has primarily been studied in simplistic architectures (e.g., depth 2 or 3) so far. The hypercube example from Section 3 and the behavior of simple architectures like MLPs is already an open problem. Incorporating function class bias into transfer bounds is quite non-trivial and our results show how this can be done for linear representations. Extending these results to more complex function classes, and incorporating training procedures could potentially give us new insights. Our study in this paper has been diagnostic in nature: identifying gaps in our understanding. Converting these insights into algorithmic approaches is a very promising direction. We also hope that visualizations of contrastive loss \rightarrow downstream performance can aid selection of more robust augmentations.

Acknowledgments. NS and SA are supported by NSF, ONR, Simons Foundation, DARPA and SRC. SK acknowledges funding from the Office of Naval Research under award N00014-22-1-2377 and the National Science Foundation Grant under award #CCF-1703574.

References

- Abnar, S., Dehghani, M., Neyshabur, B., and Sedghi, H. Exploring the limits of large scale pre-training. In *International Conference on Learning Representations*, 2022.
- Arora, S., Khandeparkar, H., Khodak, M., Plevrakis, O., and Saunshi, N. A theoretical analysis of contrastive unsupervised representation learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Ash, J. T., Goel, S., Krishnamurthy, A., and Misra, D. Investigating the role of negatives in contrastive representation learning. *arXiv preprint arXiv:2106.09943*, 2021.
- Bachman, P., Hjelm, R. D., and Buchwalter, W. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, 2019.
- Bansal, Y., Kaplun, G., and Barak, B. For self-supervised learning, rationality implies generalization, provably. In *International Conference on Learning Representations*, 2021.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 2020.
- Chen, T., Luo, C., and Li, L. Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 2021.
- Chen, X. and He, K. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- Chung, J., Gulcehre, C., Cho, K., and Bengio, Y. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Dosovitskiy, A., Springenberg, J. T., Riedmiller, M., and Brox, T. Discriminative unsupervised feature learning with convolutional neural networks. *Advances in neural information processing systems*, 27, 2014.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houselby, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- Gao, T., Yao, X., and Chen, D. SimCSE: Simple contrastive learning of sentence embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021.
- Giorgi, J., Nitski, O., Wang, B., and Bader, G. DeCLUTR: Deep contrastive learning for unsupervised textual representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021.
- Grill, J.-B., Strub, F., Althé, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. 2020.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010.
- HaoChen, J. Z., Wei, C., Gaidon, A., and Ma, T. Provable guarantees for self-supervised deep learning with spectral contrastive loss. In *Advances in Neural Information Processing Systems*, 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*, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019.
- Jin, C., Ge, R., Netrapalli, P., Kakade, S. M., and Jordan, M. I. How to escape saddle points efficiently. In *International Conference on Machine Learning*, 2017.
- Jing, L., Vincent, P., LeCun, Y., and Tian, Y. Understanding dimensional collapse in contrastive self-supervised learning. In *International Conference on Learning Representations*, 2022.
- 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*, 2021.

- Logeswaran, L. and Lee, H. An efficient framework for learning sentence representations. In *Proceedings of the International Conference on Learning Representations*, 2018.
- Meng, Y., Xiong, C., Bajaj, P., Bennett, P., Han, J., Song, X., et al. Coco-lm: Correcting and contrasting text sequences for language model pretraining. *Advances in Neural Information Processing Systems*, 2021.
- Mitrovic, J., McWilliams, B., Walker, J., Buesing, L., and Blundell, C. Representation learning via invariant causal mechanisms. In *International Conference on Learning Representations*, 2021.
- Nozawa, K. and Sato, I. Understanding negative samples in instance discriminative self-supervised representation learning. In *Advances in Neural Information Processing Systems*, 2021.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Robinson, J., Sun, L., Yu, K., Batmanghelich, K., Jegelka, S., and Sra, S. Can contrastive learning avoid shortcut solutions? In *Advances in Neural Information Processing Systems*, 2021.
- Saunshi, N., Malladi, S., and Arora, S. A mathematical exploration of why language models help solve downstream tasks. In *International Conference on Learning Representations*, 2021.
- Tian, Y., Krishnan, D., and Isola, P. Contrastive multiview coding. In *European conference on computer vision*, 2020a.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning? In *Advances in Neural Information Processing Systems*, 2020b.
- Tolstikhin, I. O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems*, 2021.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive learning, multi-view redundancy, and linear models. 2021a.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 2021b.
- Tsai, Y.-H. H., Wu, Y., Salakhutdinov, R., and Morency, L.-P. Self-supervised learning from a multi-view perspective. In *International Conference on Learning Representations*, 2021.
- Tschannen, M., Djolonga, J., Rubenstein, P. K., Gelly, S., and Lucic, M. On mutual information maximization for representation learning. In *International Conference on Learning Representations*, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Von Kügelgen, J., Sharma, Y., Gresele, L., Brendel, W., Schölkopf, B., Besserve, M., and Locatello, F. Self-supervised learning with data augmentations provably isolates content from style. In *Advances in Neural Information Processing Systems*, 2021.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, 2020.
- Wang, X. and Gupta, A. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- Wang, Y., Zhang, Q., Wang, Y., Yang, J., and Lin, Z. Chaos is a ladder: A new understanding of contrastive learning. In *International Conference on Learning Representations*, 2022.
- Wen, Z. and Li, Y. Toward understanding the feature learning process of self-supervised contrastive learning. 2021.
- Wu, S., Zhang, H., and Ré, C. Understanding and improving information transfer in multi-task learning. In *International Conference on Learning Representations*, 2020a.
- Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3733–3742, 2018.
- Wu, Z., Wang, S., Gu, J., Khabsa, M., Sun, F., and Ma, H. Clear: Contrastive learning for sentence representation. *arXiv preprint arXiv:2012.15466*, 2020b.
- Yan, Y., Li, R., Wang, S., Zhang, F., Wu, W., and Xu, W. ConSERT: A contrastive framework for self-supervised sentence representation transfer. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.

Zhang, X., Zhao, J., and LeCun, Y. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 2015.

Zimmermann, R. S., Sharma, Y., Schneider, S., Bethge, M., and Brendel, W. Contrastive learning inverts the data generating process. 2021.

A. Omitted Proofs

A.1. Proof of Proposition 4.2

The proof of (a) follows directly from Lemma 4.1, since all the conditions are satisfied and the augmentations are disjoint.

For the eigenvalues, we can compute the covariances by using $\tau \sim \mathcal{U}((0, 1])$

$$\begin{aligned}\Sigma(\phi) &= \text{diag} \left(\mathbf{1}_k, \mathbb{E}_{\tau}[\tau^2] \mathbf{1}_{D-k} \right) = \text{diag}(\mathbf{1}_k, 1/3 \mathbf{1}_{D-k}) \\ \Sigma(\phi_{\mathcal{A}}) &= \text{diag} \left(\mathbf{1}_k, (\mathbb{E}[\tau])^2 \mathbf{1}_{D-k} \right) = \text{diag}(\mathbf{1}_k, 1/4 \mathbf{1}_{D-k})\end{aligned}$$

Thus the matrix of interest is

$$I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}} = \text{diag}(\mathbf{0}_k, 1/4 \mathbf{1}_{D-k})$$

giving us $\lambda_i = 0$ for $i \leq k$ and $\lambda_i = 1/4$ for $i > k$. Plugging into Theorem 4.5 for $d' = k$ finishes the proof.

B. Low-rank Linear Representation Proof

Firstly we set up some notation. For sets P and set Q , where P is finite, we denote Q^P to denote the set of all functions from $P \rightarrow Q$. We abuse notation and also denote Q^P to be a subset of $Q^{|P|}$, where an element $r \in Q^P$ is a vector of $|P|$ dimensions and coordinates are indexed by elements of P . For instance, when $Q = \{\pm 1\}$ and P is finite, $Q^P = \{\pm 1\}^P$ denotes all functions mapping elements in P to either 1 or -1 . Furthermore, $r \in \{\pm 1\}^P$ denotes a vector in $\{\pm 1\}^{|P|}$ that looks like $(r(p))_{p \in P}$. Similarly we denote $Q^{P \times R}$ to denote a matrix in $Q^{|P| \times |R|}$. For a matrix $Q \in \mathbb{R}^{m \times n}$, $Q_{:,d} \in \mathbb{R}$

We now prove function class dependent guarantees for the class of linear representations. As in Section 4.3, for a feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^D$, we define the linear representation class $\mathcal{F}_{\phi} = \{f(\cdot) = W^{\top} \phi(\cdot) \mid W \in \mathbb{R}^{D \times D}\}$. We wish to show downstream guarantees for contrastive learning that depend not only on the contrastive loss of a representation $f : \mathcal{X} \rightarrow \mathbb{R}^d$ but also uses the fact that it belongs to the class \mathcal{F}_{ϕ} . In particular, we desire a bound that looks like $L_{\text{clf}}(f; \bar{y}^*) \leq \mathcal{T}(\mathcal{A}, \bar{y}^*, L_{\text{cont}}(f), \mathcal{F}_{\phi})$, as described in Equation (6).

We employ the strategy from HaoChen et al. (2021) and show guarantees for the spectral contrastive loss, defined in Equation (2) as

$$L_{\text{spec}}(f) = -2 \mathbb{E}_{(x, x^+) \sim \mathcal{D}_{\text{sim}}} [f(x)^{\top} f(x^+)] + \mathbb{E}_{x, x^- \sim \mathcal{D}_{\text{neg}}^2} \left[(f(x)^{\top} f(x^-))^2 \right] \quad (10)$$

We first provide a sketch of their proof in our notation and highlight the main steps. Our result is similar in spirit to theirs, but deviates at crucial junctions due to incorporation of the function class.

1. Rewrite as matrix factorization. Lemma 3.2 from HaoChen et al. (2021) shows that this objective can be rewritten as matrix factorization. For any two augmentations $x, x' \in \mathcal{X}$, define $w_{x, x'} = \mathcal{D}_{\text{sim}}(x, x') = \mathbb{E}_{\bar{x}} [\mathcal{A}(x \mid \bar{x}) \mathcal{A}(x' \mid \bar{x})]$ to be the probability that x and x' appear as a similar pair, i.e. two augmentations of the same input. Let $w_x = \sum_{x' \in \mathcal{X}} w_{x, x'} = \mathcal{D}_{\text{neg}}(x)$ be the marginal probability. Then the objective can be rewritten as follows:

$$\begin{aligned}L_{\text{spec}}(f) &= \mathbb{E}_{(x, x^+) \sim \mathcal{D}_{\text{sim}}} [f(x)^{\top} f(x^+)] + \mathbb{E}_{x, x^- \sim \mathcal{D}_{\text{neg}}^2} \left[(f(x)^{\top} f(x^-))^2 \right] \\ &= -2 \sum_{x, x^+ \in \mathcal{X}} w_{x, x^+} f(x)^{\top} f(x^+) + \sum_{x, x^- \in \mathcal{X}} w_x w_{x^-} (f(x)^{\top} f(x^-))^2 \\ &= \sum_{x, x' \in \mathcal{X}} \left(-2 w_{x, x'} f(x)^{\top} f(x') + w_x w_{x'} (f(x)^{\top} f(x'))^2 \right)\end{aligned}$$

Table 2. Notations

Notation	Definition	Description
<u>Distributions</u>		
$\bar{\mathcal{X}}, \mathcal{X}$		Set of inputs and augmentations
\mathcal{A}	$x \sim \mathcal{A}(\cdot \bar{x})$	Augmentation distribution
$\mathcal{D}_{\bar{\mathcal{X}}}$	$\bar{x} \sim \mathcal{D}_{\bar{\mathcal{X}}}$	Marginal distribution on inputs $\bar{\mathcal{X}}$
$\mathcal{D}_{\mathcal{X}}$	$\mathbb{E}_{\bar{x}} [\mathcal{A}(\cdot \bar{x})]$	Marginal distribution on augmentations \mathcal{X}
$\bar{D} \in \mathbb{R}^{\bar{\mathcal{X}} \times \bar{\mathcal{X}}}$	$\bar{D}[\bar{x}, \bar{x}] = \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})$	Matrix of marginal distributions on $\bar{\mathcal{X}}$
$D \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$	$D[x, x] = \mathcal{D}_{\mathcal{X}}(x)$	Matrix of marginal distributions on \mathcal{X}
$\bar{A} \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$	$\bar{A}[\bar{x}, x] = \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \mathcal{A}(x \bar{x})$	Input augmentation distribution
$\bar{A}_o \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$	$\bar{D}^{-\frac{1}{2}} \bar{A} D^{-\frac{1}{2}}$	Normalized matrix version of \bar{A}
$A \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$	$A[x, x] = \mathcal{D}_{\text{sim}}(x, x)$	Matrix of joint distribution of augmentations
$A_o \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$	$D^{-\frac{1}{2}} A D^{-\frac{1}{2}} = \bar{A}_o^\top \bar{A}_o$	Normalized matrix version of A
<u>Fixed features</u>		
$\phi : \mathcal{X} \rightarrow \mathbb{R}^D$		Fixed feature map for augmentations
$\phi_{\mathcal{A}} : \bar{\mathcal{X}} \rightarrow \mathbb{R}^D$	$\mathbb{E}_{x \sim \mathcal{A}(x \cdot)} [\phi(x)]$	Augmentation averaged feature
$\Sigma(\phi) \in \mathbb{R}^{D \times D}$	$\mathbb{E}_x [\phi(x) \phi(x)^\top]$	Covariance of feature map ϕ
$\Phi \in \mathbb{R}^{\mathcal{X} \times D}$	$\Phi[x] = \phi(x)$	Matrix version of feature map ϕ
$\Phi_o \in \mathbb{R}^{\mathcal{X} \times D}$	$D^{\frac{1}{2}} \Phi$	Normalized version of Φ
<u>Representation</u>		
$f : \mathcal{X} \rightarrow \mathbb{R}^d$		Representation function
$f_o : \mathcal{X} \rightarrow \mathbb{R}^d$	$\sqrt{\mathcal{D}_{\mathcal{X}}(\cdot)} f(\cdot)$	Normalized version of f
$F \in \mathbb{R}^{\mathcal{X} \times d}$	$F[x, i] = f(x)_i$	Matrix version of f
$F_o \in \mathbb{R}^{\mathcal{X} \times d}$	$D^{\frac{1}{2}} F$	Normalized version of F
<u>Function classes</u>		
$\mathcal{F} \subseteq \{f : \mathcal{X} \rightarrow \mathbb{R}^d\}$		Representation function class
$\mathcal{F}_\phi \subseteq \{f : \mathcal{X} \rightarrow \mathbb{R}^d\}$	$\{W^\top \phi(\cdot) \mid W \in \mathbb{R}^{D \times d}\}$	Linear representation class
$\mathcal{F}_\Phi \subseteq \mathbb{R}^{\mathcal{X} \times d}$	$\{\Phi W \mid W \in \mathbb{R}^{D \times d}\}$	Linear representation class (matrix version)

$$= C + \sum_{x, x' \in \mathcal{X}} \left(\frac{w_{x, x'}}{\sqrt{w_x w_{x'}}} - (\sqrt{w_x} f(x))^\top (\sqrt{w_{x'}} f(x')) \right)^2$$

where C depends only on w and thus only on \mathcal{A} , but not f . Thus $L_{\text{spec}}(f)$ can be interpreted as a matrix factorization objective, with the matrix being $A_\circ \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ such that $A_\circ[x, x'] = \frac{w_{x, x'}}{\sqrt{w_x w_{x'}}$ and scaled version of representation $u_x = \sqrt{w_x} f(x)$ is being used to factorize this. Note that A_\circ only depends on w 's which in turn only depend on the distributions \mathcal{A} , $\mathcal{D}_{\bar{\mathcal{X}}}$ and $\mathcal{D}_{\mathcal{X}}$. We stack the representation f into a matrix $F_\circ \in \mathbb{R}^{\mathcal{X} \times d}$, where the column corresponding to $x \in \mathcal{X}$ is $F_\circ[x] = \sqrt{w_x} f(x)$. Then the objective can be written as

$$L_{\text{spec}}(f) = C + \|A_\circ - F_\circ F_\circ^\top\|_F^2. \quad (11)$$

This helps characterize the optimal solution f^* of the contrastive objective, which corresponds to the matrix F_\circ^* learning the top d eigen-directions of the matrix A_\circ . Inspired by this analysis, we also show that the spectral loss with the function class \mathcal{F}_ϕ is a matrix factorization problem, but for a different matrix that depends on both A_\circ and ϕ .

2. ϵ -optimal solution f While the above characterization tells us something about the optimal representation f^* , in general we might have a representation that has sub-optimality of $\epsilon = L_{\text{spec}}(f) - L_{\text{spec}}(f^*)$. In this case, it can be argued that such a representation captures significant mass of the first d eigen-directions of A_\circ as long as ϵ is small and the eigen-gap is large. More specifically, if $\gamma_1, \dots, \gamma_{\mathcal{X}}$ denote the eigenvalues of A_\circ , then the suboptimal f will capture all except $\mathcal{O}\left(\frac{\epsilon}{(\gamma_{d+1} - \gamma_d)^2}\right) = \mathcal{O}\left(\frac{L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*)}{(\gamma_{d+1} - \gamma_d)^2}\right)$ mass of the first d eigen-directions of A_\circ . For our analysis, we will suffer a suboptimality only w.r.t. the function class \mathcal{F}_ϕ , i.e. the ϵ will be $L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)$ rather than $L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*)$

3. Connecting to downstream. It remains to show why approximately learning the top d directions of the augmentation matrix A_\circ can help with a downstream task \bar{y}^* . This step uses two assumptions, (1) there is sufficient overlap in augmentation distributions overall, and (2) augmentations are approximately label invariant, i.e. there is not much overlap in augmentations of inputs from different classes. These assumptions imply that the true label vector $y^* \in \{\pm 1\}^{\mathcal{X}}$ has a high component on the first d directions of A_\circ . We use similar properties but with less stringent conditions on the amount of overlap between augmentations. In addition to this we need a crucial assumption that the function class \mathcal{F}_ϕ is expressive enough to solve the classification task on augmentations.

B.1. Matrix notation

Given the backdrop of the results from [HaoChen et al. \(2021\)](#), we now present the matrix notations for various functions that will be helpful to prove our main result. All definitions and notations are summarized in [Table 2](#).

B.1.1. DISTRIBUTIONS TO MATRICES

Let $w_{\bar{x}} = \mathcal{D}_{\bar{\mathcal{X}}}(x)$ denote the marginal probabilities of input $\bar{x} \in \bar{\mathcal{X}}$ and $w_{\bar{x}, x} = \mathcal{A}(x | \bar{x}) w_{\bar{x}}$ denote the joint probability of input and augmentation. The marginal for augmentations can then be defined as $w_x = \mathcal{D}_{\mathcal{X}}(x) = \sum_{\bar{x}} w_{\bar{x}, x}$. To summarize

$$w_{\bar{x}} = \mathcal{D}_{\bar{\mathcal{X}}}(x) \quad (12)$$

$$w_{x|\bar{x}} = \mathcal{A}(x | \bar{x}) \quad (13)$$

$$w_{x, \bar{x}} = w_{\bar{x}, x} = \mathcal{A}(x | \bar{x}) w_{\bar{x}} = w_{x|\bar{x}} w_{\bar{x}} \quad (14)$$

$$w_x = \mathcal{D}_{\mathcal{X}}(x) \quad (15)$$

Let $\bar{D} \in \mathbb{R}^{\bar{\mathcal{X}} \times \bar{\mathcal{X}}}$ denote a diagonal matrix of marginal probabilities, i.e. $\bar{D} = \text{diag}((w_{\bar{x}})_{\bar{x} \in \bar{\mathcal{X}}})$. Similarly $D \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is the diagonal matrix of augmentation marginals. Thus these diagonal matrices satisfy

$$\bar{D}[\bar{x}, \bar{x}] = w_{\bar{x}}, \quad D[x, x] = w_x \quad (16)$$

We express the augmentation distributions $\mathcal{A}(\cdot | \bar{x})_{\bar{x} \in \bar{\mathcal{X}}}$ as a matrix $\bar{A} \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$, where $\bar{A}[\bar{x}, x] = w_{\bar{x}, x}$. A normalized version of \bar{A} is denoted by $\bar{A}_o \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$ and defined as $\bar{A}_o[\bar{x}, x] = \frac{w_{\bar{x}, x}}{\sqrt{w_{\bar{x}} w_x}}$. We summarize these definitions below, along with a matrix equation that follows easily from the definition

$$\bar{A}[\bar{x}, x] = w_{\bar{x}, x}, \bar{A}_o[\bar{x}, x] = \frac{w_{\bar{x}, x}}{\sqrt{w_{\bar{x}} w_x}}, \bar{A}_o = \bar{D}^{-\frac{1}{2}} \bar{A} \bar{D}^{-\frac{1}{2}} \quad (17)$$

For the similarity distribution \mathcal{D}_{sim} on pairs of augmentations, define the following

$$w_{x, x'} = \mathcal{D}_{\text{sim}}(x, x') = \mathbb{E}_{\bar{x}}[\mathcal{A}(x | \bar{x})\mathcal{A}(x' | \bar{x})] = \sum_{\bar{x}} w_{\bar{x}} w_{x|\bar{x}} w_{x'|\bar{x}}. \quad (18)$$

\mathcal{D}_{sim} is expressed as a matrix $A \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$, where $A[x, x'] = w_{x, x'}$. The normalized version of A is defined as $A_o \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$, where $A_o[x, x'] = \frac{w_{x, x'}}{\sqrt{w_x w_{x'}}}$. We summarize these definitions below, along with a matrix equation that follows easily from the definition

$$A[x, x'] = w_{x, x'}, A_o[x, x'] = \frac{w_{x, x'}}{\sqrt{w_x w_{x'}}}, A_o = D^{-\frac{1}{2}} A D^{-\frac{1}{2}} \quad (19)$$

The following lemma connects the \bar{A}_o and A_o .

Lemma B.1. For \bar{A}_o and A_o defined in Table 2, we have the following

$$A_o = \bar{A}_o^\top \bar{A}_o \quad (20)$$

Proof. Firstly from Equation (17), we get that $\bar{A}_o^\top \bar{A}_o = D^{-\frac{1}{2}} \bar{A}^\top \bar{D}^{-1} \bar{A} D^{-\frac{1}{2}}$. Given that $A_o = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$ from Equation (19), it suffices to show that $A = \bar{A}^\top \bar{D}^{-1} \bar{A}$. The (x, x') entry of the RHS is as follows

$$(\bar{A}^\top \bar{D}^{-1} \bar{A})[x, x'] = \sum_{\bar{x}} \frac{w_{\bar{x}, x} w_{\bar{x}, x'}}{w_{\bar{x}}} \stackrel{(a)}{=} \sum_{\bar{x}} \frac{w_{x|\bar{x}} w_{\bar{x}} w_{x'|\bar{x}} w_{\bar{x}}}{w_{\bar{x}}} = \sum_{\bar{x}} w_{\bar{x}} w_{x|\bar{x}} w_{x'|\bar{x}} \stackrel{(b)}{=} w_{x, x'} = A[x, x']$$

where (a) follows from Equation (14) and (b) follows from Equation (18). This completes the proof. \square

B.1.2. REPRESENTATIONS TO MATRICES

The previous section described how to convert distributions to matrices. We now do the same for representation functions. For a feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^D$, we denote $\Phi \in \mathbb{R}^{\mathcal{X} \times D}$ to be the matrix of representations, with the rows being $\Phi[x] = \phi(x)$. The distributionally normalized version of the representation $\phi_o(x) = \sqrt{\mathcal{D}_{\mathcal{X}}(x)}\phi(x) = \sqrt{w_x}\phi(x)$ is denoted by $\Phi_o \in \mathbb{R}^{\mathcal{X} \times D}$ with row for $x \in \mathcal{X}$ being $\phi_o(x)$. We similarly define the matrices for representation $f : \mathcal{X} \rightarrow \mathbb{R}^d$ to be F, F_o for the distributionally normalized version. It is easy to see the following relationship between F and F_o : $F_o = D^{\frac{1}{2}} F$. For the function class of linear representations $\mathcal{F}_\phi = \{W^\top \phi(\cdot) \mid W \in \mathbb{R}^{D \times d}\}$, the matrix version is defined as $\mathcal{F}_\Phi = \{\Phi W \mid W \in \mathbb{R}^{D \times d}\}$.

B.2. Connecting losses to matrix notations

We first define various downstream evaluation metrics for representation.

Definition B.1. We define the classification and regression error for any augmentation representation function $h : \mathcal{X} \rightarrow \mathbb{R}^d$. For any ground-truth labeling $\bar{y}^* : \bar{\mathcal{X}} \rightarrow \{\pm 1\}$ on original inputs, we define the following

$$L_{\text{clf}}(h; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x} \sim \bar{\mathcal{X}}} [\mathbb{1} \{\text{sign}(w^\top h_{\mathcal{A}}(\bar{x})) = \bar{y}^*(\bar{x})\}] \quad (21)$$

$$L_{\text{reg}}(h; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x} \sim \bar{\mathcal{X}}} [(w^\top h_{\mathcal{A}}(\bar{x}) - \bar{y}^*(\bar{x}))^2] \quad (22)$$

where $h_{\mathcal{A}}(\bar{x}) = \mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})} [h(x)]$ is the augmentation averaged representation (see Table 2). For any labeling $g : \mathcal{X} \rightarrow \{\pm 1\}$ on augmentations, we define the following

$$L_{\text{reg}}(h; g) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{x \sim \mathcal{X}} [(w^\top h(\bar{x}) - g(x))^2] \quad (23)$$

We now connect the downstream regression loss with matrix versions of feature map ϕ .

Lemma B.2. For an arbitrary predictor on augmentations $g \in \{\pm 1\}^{\mathcal{X}}$ and its normalized version $g_\circ = D^{\frac{1}{2}}g$, and an augmentation feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ and its normalized matrix Φ_\circ ,

$$L_{\text{reg}}(\phi; g) = \|P_{\Phi_\circ}^\perp g_\circ\|^2 \quad (24)$$

Proof. Note that $\Phi_\circ = D^{\frac{1}{2}}\Phi$, where $\Phi \in \mathbb{R}^{\mathcal{X} \times d}$ is the matrix version of the augmentation feature map ϕ (refer Table 2). We prove the result by rewriting L_{reg} as follows

$$L_{\text{reg}}(\phi; g) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_x (\phi(x)^\top w - g(x))^2 = \inf_{w \in \mathbb{R}^d} \sum_{x \in \mathcal{X}} D(x) (\phi(x)^\top w - g(x))^2 \quad (25)$$

$$= \inf_{w \in \mathbb{R}^d} \sum_x \left(\sqrt{D(x)} \phi(x)^\top w - \sqrt{D(x)} g(x) \right)^2 \quad (26)$$

$$= \inf_{w \in \mathbb{R}^d} \left\| D^{\frac{1}{2}} \Phi w - D^{\frac{1}{2}} g \right\|^2 = \inf_{w \in \mathbb{R}^d} \|\Phi_\circ w - g_\circ\|^2 \quad (27)$$

$$= \|P_{\Phi_\circ}^\perp g_\circ\|^2 \quad (28)$$

□

We now express the spectral contrastive loss and upper bound the downstream classification error using matrix versions of distributions and representations.

Lemma B.3. For any representation f and its corresponding normalized matrix $F_\circ \in \mathbb{R}^{\mathcal{X} \times d}$, the spectral contrastive loss (Equation (2)) and classification loss (Equation (3)) can be rewritten and upper bounded as

$$L_{\text{spec}}(f) = L_{\text{spec}}(F_\circ) = \|A_\circ - F_\circ F_\circ^\top\|_F^2 - \|A_\circ\|_F^2 = \|\bar{A}_\circ^\top \bar{A}_\circ - F_\circ F_\circ^\top\|_F^2 - \|\bar{A}_\circ^\top \bar{A}_\circ\|_F^2 \quad (29)$$

$$L_{\text{clf}}(f; \bar{y}^*) \leq L_{\text{reg}}(f; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \|\bar{A}_\circ F_\circ w - \bar{y}^*\|_2^2 = \|P_{\bar{A}_\circ F_\circ}^\perp \bar{y}^*\|_2^2 \quad (30)$$

Proof. We first prove the expression for $L_{\text{spec}}(f)$. Note that $A_\circ[x, x'] = w_{x, x'} = \mathcal{D}_{\text{sim}}(x, x')$ from Equation (18). Furthermore $F_\circ[x] = \sqrt{\mathcal{D}_{\mathcal{X}}(x)} f(x)$ from Table 2. On expanding out the contrastive loss, we get

$$\begin{aligned} L_{\text{spec}}(f) &= \mathbb{E}_{(x, x^+) \sim \mathcal{D}_{\text{sim}}} [f(x)^\top f(x^+)] + \mathbb{E}_{x, x^- \sim \mathcal{D}_{\text{neg}}} \left[(f(x)^\top f(x^-))^2 \right] \\ &= -2 \sum_{x, x^+ \in \mathcal{X}} w_{x, x^+} f(x)^\top f(x^+) + \sum_{x, x^- \in \mathcal{X}} w_x w_{x^-} (f(x)^\top f(x^-))^2 \\ &= \sum_{x, x' \in \mathcal{X}} \left(-2w_{x, x'} f(x)^\top f(x') + w_x w_{x'} (f(x)^\top f(x'))^2 \right) \\ &= - \sum_{x, x' \in \mathcal{X}} \left(\frac{w_{x, x'}}{\sqrt{w_x w_{x'}}} \right)^2 + \sum_{x, x' \in \mathcal{X}} \left(\frac{w_{x, x'}}{\sqrt{w_x w_{x'}}} - (\sqrt{w_x} f(x))^\top (\sqrt{w_{x'}} f(x')) \right)^2 \\ &= - \sum_{x, x'} A_\circ[x, x']^2 + \sum_{x, x'} (A_\circ[x, x'] - F_\circ[x]^\top F_\circ[x'])^2 \\ &= -\|A_\circ\|_F^2 + \|A_\circ - F_\circ F_\circ^\top\|_F^2 \end{aligned}$$

Now we prove the upper bound of $L_{\text{clf}}(f; \bar{y}^*)$. Firstly note that for any input representation $h : \bar{\mathcal{X}} \rightarrow \{\pm 1\}$, we have that

$$L_{\text{clf}}(h; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x}} \left[\mathbb{1} \left\{ \bar{y}^*(\bar{x}) \left(h(\bar{x})^\top w \right) < 0 \right\} \right] \stackrel{(a)}{\leq} \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x}} \left[\left(\bar{y}^*(\bar{x}) - h(\bar{x})^\top w \right)^2 \right] = L_{\text{reg}}(h; \bar{y}^*)$$

where (a) from the fact that whenever $\bar{y}^*(\bar{x}) \left(h(\bar{x})^\top w \right) < 0$, $h(\bar{x})^\top w$ has different sign compared to $\bar{y}^* \in \{\pm 1\}$, and so $(h(\bar{x})^\top w - \bar{y}^*)^2 \geq \bar{y}^{*2} = 1$. Thus for an augmentation representation $f : \mathcal{X}$, we have

$$\begin{aligned} L_{\text{reg}}(f; \bar{y}^*) &= L_{\text{reg}}(f_{\mathcal{A}}; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \mathbb{E}_{\bar{x}} \left[\left(f_{\mathcal{A}}(\bar{x})^\top w - \bar{y}^*(\bar{x}) \right)^2 \right] = \inf_{w \in \mathbb{R}^d} \sum_{\bar{x}} \left[\left(\sqrt{w_{\bar{x}}} f_{\mathcal{A}}(\bar{x})^\top w - \sqrt{w_{\bar{x}}} \bar{y}^*(\bar{x}) \right)^2 \right] \\ &= \inf_{w \in \mathbb{R}^d} \sum_{\bar{x}} \left[\left(\sqrt{w_{\bar{x}}} f_{\mathcal{A}}(\bar{x})^\top w - \bar{y}^*(\bar{x}) \right)^2 \right] \end{aligned}$$

We first observe the following about $f_{\mathcal{A}}$:

$$\begin{aligned} \sqrt{w_{\bar{x}}} f_{\mathcal{A}}(\bar{x}) &= \sqrt{w_{\bar{x}}} \mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})} [f(x)] = \sqrt{w_{\bar{x}}} \sum_{x \in \mathcal{X}} \mathcal{A}(x | \bar{x}) f(x) = \sum_{x \in \mathcal{X}} \sqrt{w_{\bar{x}}} \frac{w_{\bar{x}, x}}{w_{\bar{x}}} f(x) = \sum_{x \in \mathcal{X}} \frac{w_{\bar{x}, x}}{\sqrt{w_{\bar{x}}} w_x} \sqrt{w_x} f(x) \\ &= \sum_x A_{\circ}[\bar{x}, x] F_{\circ}[x] = (A_{\circ} F_{\circ})[\bar{x}] \end{aligned}$$

Plugging this back into the previous calculation, we get

$$L_{\text{reg}}(f; \bar{y}^*) = \inf_{w \in \mathbb{R}^d} \sum_{\bar{x}} \left[\left((A_{\circ} F_{\circ})[\bar{x}]^\top w - \bar{y}^*(\bar{x}) \right)^2 \right] = \inf_{w \in \mathbb{R}^d} \|A_{\circ} F_{\circ} w - \bar{y}^*\|_F^2$$

The final step follows from the standard expression for error of linear regression, which is the norm of the component of \bar{y}^* on the null space of $A_{\circ} F_{\circ}$, i.e. $\|P_{A_{\circ} F_{\circ}}^{\perp} \bar{y}^*\|_F^2$. \square

We now show a more specialized form of matrix factorization objective that results from the representation belonging to a particular linear function class.

Lemma B.4. *For any representation $f \in \mathcal{F}_{\phi}$ and its normalized matrix $F_{\circ} \in \mathbb{R}^{\mathcal{X} \times d}$, the spectral contrastive loss (Equation (2)) can be rewritten as*

$$L_{\text{spec}}(f) = L_{\text{spec}}(F_{\circ}) = \|P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}} - F_{\circ} F_{\circ}^{\top}\|_F^2 + C = \|P_{\Phi_{\circ}} \bar{A}_{\circ}^{\top} \bar{A}_{\circ} P_{\Phi_{\circ}} - F_{\circ} F_{\circ}^{\top}\|_F^2 + C \quad (31)$$

where C is a constant independent of f but dependent on features ϕ . Here Φ_{\circ} is the normalized matrix for the features ϕ and $P_{\Phi_{\circ}} \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is the column projection matrix of Φ_{\circ} .

Proof. From Lemma B.3, we know that $L_{\text{spec}}(f)$ can be written as a matrix factorization objective as

$$L_{\text{spec}}(f) = \|A_{\circ} - F_{\circ} F_{\circ}^{\top}\|_F^2 - \|A_{\circ}\|_F^2$$

Since f is from the class \mathcal{F}_{ϕ} , the matrix form F belongs to the class $\mathcal{F}_{\Phi} = \{\Phi W \mid W \in \mathbb{R}^{D \times d}\}$ (refer to Table 2). Thus $F_{\circ} = D^{\frac{1}{2}} F$ can be written as $F_{\circ} = D^{\frac{1}{2}} F = D^{\frac{1}{2}} \Phi W = \Phi_{\circ} W$ for some $W \in \mathbb{R}^{D \times d}$. We can conclude that $P_{\Phi_{\circ}} F_{\circ} = F_{\circ}$ and $P_{\Phi_{\circ}}^{\perp} F_{\circ} = 0$ and further simplify the contrastive loss as

$$\begin{aligned} L_{\text{spec}}(f) &\stackrel{(a)}{=} \|P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}} + P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}}^{\perp} + P_{\Phi_{\circ}}^{\perp} A_{\circ} P_{\Phi_{\circ}} + P_{\Phi_{\circ}}^{\perp} A_{\circ} P_{\Phi_{\circ}}^{\perp} - F_{\circ} F_{\circ}^{\top}\|_F^2 - \|A_{\circ}\|_F^2 \\ &\stackrel{(b)}{=} \|P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}} - F_{\circ} F_{\circ}^{\top}\|_F^2 + \|P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}}^{\perp} + P_{\Phi_{\circ}}^{\perp} A_{\circ} P_{\Phi_{\circ}}\|_F^2 + \|P_{\Phi_{\circ}}^{\perp} A_{\circ} P_{\Phi_{\circ}}^{\perp}\|_F^2 - \|A_{\circ}\|_F^2 \\ &\stackrel{(c)}{=} \|P_{\Phi_{\circ}} A_{\circ} P_{\Phi_{\circ}} - F_{\circ} F_{\circ}^{\top}\|_F^2 + C \end{aligned}$$

where (a) follows by decomposing $A_{\circ} = (P_{\Phi_{\circ}} + P_{\Phi_{\circ}}^{\perp}) A_{\circ} (P_{\Phi_{\circ}} + P_{\Phi_{\circ}}^{\perp})$, (b) follows because cross terms cancel through $P_{\Phi_{\circ}} P_{\Phi_{\circ}}^{\perp}$ multiplications, and (c) because all other terms are independent of F_{\circ} (and so f). This completes the proof. \square

We now restate the definition of Inconsistency from Section 4.3 and then relate it to some matrix form.

Definition B.2 (Inconsistency). We define inconsistency of a labeling function $g \in \{\pm 1\}^{\mathcal{X}}$ on augmentations w.r.t. some ground truth labeling $\bar{y}^* \in \{\pm 1\}^{\bar{\mathcal{X}}}$ on original inputs, as followed:

$$\Delta_{\mathcal{A}}(g, \bar{y}^*) = \mathbb{E}_{\bar{x}} \left[\mathbb{E}_{x \sim \mathcal{A}(\cdot|\bar{x})} [\mathbb{1}\{g(x) \neq \bar{y}^*(\bar{x})\}] \right] \quad (32)$$

Lemma B.5. For the normalized matrix $\bar{A}_{\circ} \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$ corresponding to augmentation distribution \mathcal{A} (refer Table 2), ground-truth labeling $\bar{y}^* \in \{\pm 1\}^{\bar{\mathcal{X}}}$ on original inputs and its normalized version $\bar{y}_{\circ}^* = D^{\frac{1}{2}} \bar{y}^*$, and an arbitrary predictor $g \in \{\pm 1\}^{\mathcal{X}}$ on augmentations and its normalized version $g_{\circ} = D^{\frac{1}{2}} g$, we have

$$\bar{y}_{\circ}^{*\top} \bar{A}_{\circ} g_{\circ} = 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) \quad (33)$$

Proof. Since $\bar{y}_{\circ}^* = \bar{D}^{\frac{1}{2}} \bar{y}^*$, $\bar{A}_{\circ} = \bar{D}^{-\frac{1}{2}} \bar{A} \bar{D}^{-\frac{1}{2}}$ and $g_{\circ} = D^{\frac{1}{2}} g$, the left hand size is equivalent to $\bar{y}^{*\top} \bar{A} g$. Expanding this further we get

$$\bar{y}^{*\top} \bar{A} g = \sum_{\bar{x}, x} \bar{A}[\bar{x}, x] \bar{y}^*(\bar{x}) g(x) \stackrel{(a)}{=} \sum_{\bar{x}, x} w_{\bar{x}, x} \bar{y}^*(\bar{x}) g(x) \stackrel{(b)}{=} \sum_{\bar{x}, x} w_{\bar{x}, x} (1 - 2\mathbb{1}\{\bar{y}^*(\bar{x}) \neq g(x)\}) \quad (34)$$

$$\stackrel{(c)}{=} \sum_{\bar{x} \in \bar{\mathcal{X}}} w_{\bar{x}} \sum_{x \in \mathcal{X}} w_{x|\bar{x}} (1 - 2\mathbb{1}\{\bar{y}^*(\bar{x}) \neq g(x)\}) \quad (35)$$

$$= 1 - 2 \mathbb{E}_{\bar{x}} \mathbb{E}_{x \sim \mathcal{A}(\cdot|\bar{x})} [\mathbb{1}\{\bar{y}^*(\bar{x}) \neq g(x)\}] = 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) \quad (36)$$

where (a) follows from Equation (19), (b) follows from Equation (14), and (c) follows from $\bar{y}^*(\bar{x}), g(x) \in \{\pm 1\}$. \square

B.3. Proof of main result

We first state the key lemmas that will be used to prove the main result.

The following lemma says that if there is a predictor on augmentations that is consistent with \bar{y}^* and also expressible enough by fixed features ϕ , then most of \bar{y}^* is retained by multiplication by $P_{\Phi_{\circ}} \bar{A}_{\circ}$.

Lemma B.6. For the normalized matrix $\bar{A}_{\circ} \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$ corresponding to augmentation distribution \mathcal{A} , an augmentation feature map ϕ and corresponding normalized matrix Φ_{\circ} (refer Table 2), ground-truth labeling $\bar{y}^* \in \{\pm 1\}^{\bar{\mathcal{X}}}$ on original inputs and its normalized version $\bar{y}_{\circ}^* = D^{\frac{1}{2}} \bar{y}^*$, and an arbitrary predictor $g \in \{\pm 1\}^{\mathcal{X}}$ on augmentations and its normalized version $g_{\circ} = D^{\frac{1}{2}} g$, we have

$$\|P_{\Phi_{\circ}} \bar{A}_{\circ}^{\top} \bar{y}_{\circ}^*\| \geq 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) - \sqrt{L_{\text{reg}}(\phi; g)} \quad (37)$$

Proof. We will use Lemma B.5 to prove this result. Note that $\|\bar{y}_{\circ}^*\| = \|g_{\circ}\| = 1$. First we lower bound $\|P_{\Phi_{\circ}} \bar{A}_{\circ}^{\top} \bar{y}_{\circ}^*\|$ by computing $\bar{y}_{\circ}^{*\top} \bar{A}_{\circ} P_{\Phi_{\circ}} g_{\circ}$

$$\begin{aligned} \|P_{\Phi_{\circ}} \bar{A}_{\circ}^{\top} \bar{y}_{\circ}^*\| &\stackrel{(a)}{\geq} \frac{\bar{y}_{\circ}^{*\top} \bar{A}_{\circ} P_{\Phi_{\circ}} g_{\circ}}{\|g_{\circ}\|} = \bar{y}_{\circ}^{*\top} \bar{A}_{\circ} P_{\Phi_{\circ}} g_{\circ} \\ &= \bar{y}_{\circ}^{*\top} \bar{A}_{\circ} g_{\circ} - \bar{y}_{\circ}^{*\top} \bar{A}_{\circ} P_{\Phi_{\circ}}^{\perp} g_{\circ} \\ &\stackrel{(b)}{=} 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) - \bar{y}_{\circ}^{*\top} \bar{A}_{\circ} P_{\Phi_{\circ}}^{\perp} g_{\circ} \\ &\stackrel{(c)}{\geq} 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) - \|\bar{y}_{\circ}^*\| \|\bar{A}_{\circ}\|_2 \|P_{\Phi_{\circ}}^{\perp} g_{\circ}\| \\ &\stackrel{(d)}{\geq} 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) - \|P_{\Phi_{\circ}}^{\perp} g_{\circ}\| \\ &\stackrel{(e)}{\geq} 1 - 2\Delta_{\mathcal{A}}(g, \bar{y}^*) - \sqrt{L_{\text{reg}}(\phi; g)} \end{aligned}$$

where (a) and (c) follow from Cauchy-Schwarz inequality, (b) follows from Lemma B.5, (d) follows from the fact that $\|\bar{A}_{\circ}\|_2 = 1$ and (e) follows from Lemma B.2 \square

The next lemma quantifies how much of the top singular directions of $P_{\Phi_o} \bar{A}_o$ are captured by an ϵ -optimal representation f (or its matrix version F_o). This is related to Lemma D.10 from HaoChen et al. (2021), however it differs in the fact that we are decomposing $P_{\Phi_o} \bar{A}_o$ instead of \bar{A}_o , and we have a better dependence on d on the right hand side. Furthermore, the sub-optimality term is w.r.t. the best representation in the class \mathcal{F}_ϕ rather than the unconstrained optimizer of L_{spec} .

Lemma B.7. *Let $f \in \mathcal{F}_\phi \in \mathcal{F}_\phi$ be an augmentation representation function. Suppose $F_o \in \mathbb{R}^{\mathcal{X} \times d}$ is the normalized representation matrix corresponding to f , \bar{A}_o is the normalized matrix corresponding to augmentation distribution \mathcal{A} and Φ_o is normalized version of Φ (refer Table 2). Let $P_{\Phi_o} \bar{A}_o^\top = USV^\top$ be the singular value decomposition, with $\sqrt{\gamma_1}, \dots, \sqrt{\gamma_D}$ being the singular values in decreasing order. Then for $d' \leq d$,*

$$\|P_{F_o}^\perp U_{:d'}\|_F^2 \leq \frac{L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)}{\gamma_{d'}^2 - \gamma_{d+1}^2} \leq \frac{L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)}{(\gamma_{d'} - \gamma_{d+1})^2} \quad (38)$$

where $U_{:d'} \in \mathbb{R}^{\mathcal{X} \times d'}$ corresponds to the first d' columns (and thus singular vectors) of $P_{\Phi_o} \bar{A}_o$.

Proof. We first note, using Lemma B.4, that the contrastive loss can be written as the following matrix factorization objective

$$L_{\text{spec}}(f) = \|P_{\Phi_o} A_o P_{\Phi_o} - F_o F_o^\top\|_F^2 = \|P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o} - F_o F_o^\top\|_F^2 = \|US^2V^\top - F_o F_o^\top\|_F^2$$

It is easy to see that $\gamma_i \leq 1$ for every i , since $\max_i \gamma_i = \|P_{\Phi_o} \bar{A}_o\|_2^2 \leq \|\bar{A}_o\|_2^2 \leq 1$. Thus we can invoke Lemma D.10 from HaoChen et al. (2021), but for matrix $P_{\Phi_o} A_o P_{\Phi_o}$ instead of A_o , to argue that

$$\|P_{F_o}^\perp u_i\|_F^2 \leq \frac{\epsilon}{(\gamma_i^2 - \gamma_{d+1}^2)} \leq \frac{\epsilon}{(\gamma_i - \gamma_{d+1})^2}$$

where ϵ is the suboptimality $L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*) = L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)$, since the optimal decomposition for $P_{\Phi_o} A_o P_{\Phi_o}$ must lie in the span of Φ_o , and thus $f^* \in \mathcal{F}_\phi$. Adding these for $i \in [d']$ we get

$$\|P_{F_o}^\perp U_{:d'}\|_F^2 \leq \sum_{i=1}^{d'} \frac{\epsilon}{(\gamma_i - \gamma_{d+1})^2} \leq \frac{\epsilon d'}{(\gamma_{d'} - \gamma_{d+1})^2}$$

This completes the proof. \square

We now show conditions under which the top d' directions of the matrix being factorized captures significant mass of the ground-truth labels.

Lemma B.8. *Let \bar{A}_o be the normalized matrix corresponding to augmentation distribution \mathcal{A} and Φ_o be the normalized version of Φ (refer Table 2). Let $P_{\Phi_o} \bar{A}_o^\top = USV^\top$ be the singular value decomposition, with $\sqrt{\gamma_1}, \dots, \sqrt{\gamma_D}$ being the singular values in decreasing order. Then we have,*

$$\|V_{d'}^\top \bar{y}_o^*\|_2^2 \leq \frac{1 - \|P_{\Phi_o} \bar{A}_o^\top \bar{y}_o^*\|_2^2}{1 - \gamma_{d'+1}} \quad (39)$$

where $V_{d'} \in \mathbb{R}^{\mathcal{X} \times (|\mathcal{X}| - d')}$ corresponds to the last $|\mathcal{X}| - d'$ columns (and thus singular vectors) of $P_{\Phi_o} \bar{A}_o$.

Proof. We expand out the term $1 - \|P_{\Phi_o} \bar{A}_o^\top \bar{y}_o^*\|_2^2$ as follows

$$\begin{aligned} 1 - \|P_{\Phi_o} \bar{A}_o^\top \bar{y}_o^*\|_2^2 &= 1 - \|USV^\top \bar{y}_o^*\|_2^2 = 1 - \|SV^\top \bar{y}_o^*\|_2^2 = 1 - \sum_{i=1}^D \gamma_i (v_i^\top \bar{y}_o^*)^2 \\ &\stackrel{(a)}{=} \|\bar{y}_o^*\|_2^2 - \sum_{i=1}^D \gamma_i (v_i^\top \bar{y}_o^*)^2 \\ &\stackrel{(b)}{\geq} \sum_{i=1}^D (v_i^\top \bar{y}_o^*)^2 - \sum_{i=1}^D \gamma_i (v_i^\top \bar{y}_o^*)^2 = \sum_{i=1}^D (1 - \gamma_i) (v_i^\top \bar{y}_o^*)^2 \end{aligned}$$

$$\geq^{(c)} (1 - \gamma_{d'+1}) \sum_{i=d'+1}^D (v_i^\top \bar{y}_o^*)^2 = (1 - \gamma_{d'+1}) \|V_{d'} \bar{y}_o^*\|^2$$

where (a) follows because $\|\bar{y}_o^*\|^2 = \sum_{\bar{x}} \mathcal{D}_{\bar{x}}(\bar{x}) \bar{y}^*(\bar{x}) = 1$, (b) follows because $\{v_i\}_{i=1}^D$ form a partial orthonormal basis and (c) is true because $\gamma_i \leq 1$ for every i and since γ_i 's are in decreasing order. Rearranging terms completes the proof \square

The following lemma connects the eigenvalues of augmentation averaged features that shows up in the final bound, to the eigenvalues of the matrix being decomposed in the spectral contrastive loss.

Lemma B.9. *Let $\Sigma(\cdot)$ be the covariance operator for features and $\phi_{\mathcal{A}}$ denote augmentation averaged representation obtained from ϕ (see Table 2). Let $\lambda_1, \dots, \lambda_D$ be the eigenvalues of $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$ in increasing order. Let \bar{A}_o be the normalized matrix corresponding to augmentation distribution \mathcal{A} and Φ_o be the normalized version of ϕ (refer Table 2). Let $P_{\Phi_o} \bar{A}_o^\top = USV^\top$ be the singular value decomposition, with $\sqrt{\gamma_1}, \dots, \sqrt{\gamma_D}$ being the singular values in decreasing order. Then we have,*

$$\lambda_i = 1 - \gamma_i, \forall i \in [D] \quad (40)$$

Proof. Let $w_x, w_{\bar{x}}, w_{x,\bar{x}}, w_{x|\bar{x}}$ be as defined in Equations (12) to (15). Using $\Phi_o = D^{\frac{1}{2}} \Phi$ and that D is diagonal with $D[x, x] = w_x$, we first simplify $\Sigma(\phi)$ as follows

$$\Sigma(\phi) = \mathbb{E}_{x \sim \mathcal{D}_{\mathcal{X}}} [\phi(x) \phi(x)^\top] = \sum_{x \in \mathcal{X}} [w_x(x) \phi(x) \phi(x)^\top] = \Phi^\top D \Phi = \Phi_o^\top \Phi_o$$

Next we find the matrix version of $\phi_{\mathcal{A}}$ using $\bar{A}[\bar{x}, x] = w_{\bar{x},x}$ and the following sequence of equalities.

$$\phi_{\mathcal{A}}(\bar{x}) = \mathbb{E}_{x \sim \mathcal{A}(\cdot|\bar{x})} [\phi(x)] = \sum_{x \in \mathcal{X}} w_{x|\bar{x}} \phi(x) = \frac{1}{w_{\bar{x}}} \sum_{x \in \mathcal{X}} w_{x,\bar{x}} \phi(x) = \frac{1}{w_{\bar{x}}} \sum_x \bar{A}[\bar{x}, x] \Phi[x]$$

Thus the matrix form of $\phi_{\mathcal{A}}$ is $\Phi_{\mathcal{A}} = \bar{D}^{-1} \bar{A} \Phi$. Similar to the argument for $\Sigma(\phi)$, we can then write $\Sigma(\phi_{\mathcal{A}})$ as follows

$$\begin{aligned} \Sigma(\phi_{\mathcal{A}}) &= \Phi_{\mathcal{A}}^\top \bar{D} \Phi_{\mathcal{A}} = (\bar{D}^{-1} \bar{A} \Phi)^\top \bar{D} (\bar{D}^{-1} \bar{A} \Phi) = \Phi^\top \bar{A}^\top \bar{D}^{-1} \bar{A} \Phi \\ &= \Phi^\top D^{\frac{1}{2}} \left(D^{-\frac{1}{2}} \bar{A}^\top \bar{D}^{-\frac{1}{2}} \right) \left(\bar{D}^{-\frac{1}{2}} \bar{A} D^{-\frac{1}{2}} \right) D^{\frac{1}{2}} \Phi \\ &\stackrel{(a)}{=} \Phi_o^\top \bar{A}_o^\top \bar{A}_o \Phi_o \end{aligned}$$

where (a) follows from Equation (17). Let $MNR^\top = \Phi_o$ be the SVD, so $P_{\Phi_o} = MM^\top$ and $\Phi_o^\top \Phi_o = R^\top N^2 R$. Using this, we simplify the matrix $\Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$.

$$\begin{aligned} \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}} &= (\Phi_o^\top \Phi_o)^{-\frac{1}{2}} \Phi_o^\top \bar{A}_o^\top \bar{A}_o \Phi_o (\Phi_o^\top \Phi_o) \\ &= (RN^2 R^\top)^{-\frac{1}{2}} (RNM^\top) \bar{A}_o^\top \bar{A}_o (MNR^\top) (RN^2 R^\top)^{-\frac{1}{2}} \\ &= (RN^{-1} R^\top) (RNM^\top) \bar{A}_o^\top \bar{A}_o (MNR^\top) (RN^{-1} R^\top) \\ &= RM^\top \bar{A}_o^\top \bar{A}_o MR^\top \end{aligned}$$

Since $\{\lambda_i\}_{i=1}^D$ are the eigenvalues of $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$, $\{1 - \lambda_i\}_{i=1}^D$ are the eigenvalues of $\Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}} = RM^\top \bar{A}_o^\top \bar{A}_o MR^\top$. Thus $\{\sqrt{1 - \lambda_i}\}_{i=1}^D$ are the singular values of $RM^\top \bar{A}_o^\top$ and thus $M^\top \bar{A}_o^\top$ and thus $MM^\top \bar{A}_o^\top = P_{\Phi_o} \bar{A}_o^\top$. The previous statements are true because multiplication by an orthogonal matrix does not change the singular values. Thus $\sqrt{\gamma_i} = \sqrt{1 - \lambda_i}$, finishing the proof. \square

Lemma B.10. *Let $\Sigma(\cdot)$ be the covariance operator for features and $\phi_{\mathcal{A}}$ denote augmentation averaged representation obtained from ϕ (see Table 2). Let $L_o = I - A_o$ be the Laplacian of the augmentation graph. If the features ϕ are full rank, then the eigenvalues of $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$ are the same as the eigenvalues of L_o .*

Proof. Note from [Lemma B.1](#) that the normalized adjacency matrix can be rewritten as $A_o = \bar{A}_o^\top \bar{A}_o$. Also from [Lemma B.9](#), we can imply that the eigenvalues of $I - P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o}$ are the same as the eigenvalues of $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$. Since ϕ is full rank, so is the matrix Φ_o , thus $P_{\Phi_o} = I$. So $I - P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o} = I - \bar{A}_o^\top \bar{A}_o = I - A_o = L_o$; this completes the proof. \square

We are now ready to present our main result.

Theorem 4.5. *Let $\Sigma(\cdot)$ be the covariance operator for features and $\phi_{\mathcal{A}}$ denote augmentation averaged representation obtained from ϕ (see [Table 2](#)). Let $\lambda_1, \dots, \lambda_d$ be the eigenvalues of $I_d - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$ in increasing order, then for $d \leq D$, any representation $f \in \mathcal{F}_\phi$, will satisfy*

$$L_{\text{clf}}(f; \bar{y}^*) \leq \min_{1 \leq d' \leq d} \left\{ \frac{\min_{g \in \{\pm 1\}^x} 4(2\Delta_{\mathcal{A}}(g, \bar{y}^*) + \sqrt{L_{\text{reg}}(\phi; g)})}{\lambda_{d'+1}} + \frac{2d' \left(L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*) \right)}{(1 - \lambda_{d'}) (\lambda_{d+1} - \lambda_{d'})^2} \right\} \quad (41)$$

where $\Delta_{\mathcal{A}}$ is defined in [Definition 4.4](#) and L_{reg} in [Definition B.1](#).

Proof. We first sketch an outline of the proof and how different lemmas will be used to prove the final result. We will use matrix versions of distributions and functions from [Table 2](#) throughout the proof. The following are the main steps:

1. (Matrix factorization) The spectral contrastive loss is shown to be equivalent to a matrix factorization objective as in [HaoChen et al. \(2021\)](#). For representations in the class \mathcal{F}_ϕ , [Lemma B.4](#) shows that the problem of contrastive learning is reduced to matrix factorization of a projected adjacency matrix $P_{\Phi_o} A_o P_{\Phi_o}$ through the objective $L_{\text{spec}}(F) = \|P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o} - F_o F_o^\top\|_F^2 + C$, where A_o is the normalized matrix corresponding to augmentation distribution \mathcal{A} and F_o is the normalized matrix for representation function f (refer [Table 2](#)). Thus the spectral contrastive loss L_{spec} is attempting to find a rank d approximation for $P_{\Phi_o} \bar{A}_o^\top$.
2. (ϵ -optimal solutions) If $P_{\Phi_o} \bar{A}_o^\top = USV^\top$ is the singular value decomposition, then any ϵ -optimal representation f (and corresponding F_o), i.e. $L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*) \leq \epsilon$, can be shown ([Lemma B.7](#)) to capture most of the signal for the top d' directions of $P_{\Phi_o} \bar{A}_o^\top$, i.e. $\|P_{F_o}^\perp U_{:d'}\|_F = \mathcal{O}(\epsilon)$ is small.
3. (Connecting to downstream) Top d' directions of $P_{\Phi_o} \bar{A}_o^\top$ can be shown to capture a lot of the mass of \bar{y}^* if the features ϕ and augmentation distribution \mathcal{A} are ‘‘nice’’ enough, as quantified by [Lemma B.8](#) that upper bounds $\|V_{d'}: \bar{y}^*\|_2^2$, in conjunction with [Lemma B.6](#) which quantifies these nice properties of \mathcal{A} and features ϕ .
4. (Wrapping up) Both of the above steps will have upper bounds that depend on the singular values of $P_{\Phi_o} \bar{A}_o$. Relating the singular values of $P_{\Phi_o} \bar{A}_o$ with the eigenvalues of $\left(I_d - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}} \right)$ through [Lemma B.9](#) completes the proof.

Step 1: We can rewrite the contrastive loss in matrix forms using [Lemma B.3](#).

$$L_{\text{spec}}(f) = L_{\text{spec}}(F_o) = \|A_o - F_o F_o^\top\|_F^2 - \|A_o\|_F^2 = \|\bar{A}_o^\top \bar{A}_o - F_o F_o^\top\|_F^2 - \|\bar{A}_o^\top \bar{A}_o\|_F^2 \quad (42)$$

For representation $F \in \mathcal{F}_\phi$, we can write it as $F = \Phi W$, thus giving $F_o = D^{\frac{1}{2}} F = D^{\frac{1}{2}} \Phi W = \Phi_o W$. Note that $P_{\Phi_o} = \Phi_o \Phi_o^\top$ is the projection matrix of column space of Φ_o ; then we have $F_o = P_{\Phi_o} F_o$. From [Lemma B.4](#), we know that $L_{\text{spec}}(F) = \|P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o} - F_o F_o^\top\|_F^2 + C$.

Thus from this we see that the contrastive learning is aiming to learn a good rank d decomposition of the matrix $P_{\Phi_o} A_o P_{\Phi_o} = P_{\Phi_o} \bar{A}_o^\top \bar{A}_o P_{\Phi_o}$. This is similar to the formulation in [HaoChen et al. \(2021\)](#) where the matrix A_o is being factorized instead. The classical result on low-rank approximation of matrices tells us that the minimizer F_o will span the top d singular of $P_{\Phi_o} \bar{A}_o$.

Step 2: Let $P_{\Phi_o} \bar{A}_o = USV^\top$ be the singular value decomposition, with $S = \text{diag}(\sqrt{\gamma_1}, \dots, \sqrt{\gamma_d})$ being the singular values in decreasing order. Then we know that the optimal solution F_o^* will be $U_{:d} S_{:d} R$ for any orthogonal matrix R . Note that $F_o^* \in \mathcal{F}_{\Phi_o}$ since the matrix $P_{\Phi_o} \bar{A}_o$ being decomposed is in the span of Φ_o . This argument can be extended to ϵ -optimal representation f (or matrix F) by invoking [Lemma B.7](#), which gives us that

$$\|P_{F_o}^\perp U_{:d'}\|_F^2 \leq \frac{\epsilon d'}{\gamma_{d'}^2 - \gamma_{d+1}^2} \leq \frac{\epsilon d'}{(\gamma_{d'} - \gamma_{d+1})^2} \quad (43)$$

This tells us that being close to optimality ensures that the representation captures most of the top d' singular directions of U and thus $P_{\Phi_o} \bar{A}_o$, whenever $d' \leq d$. Note that the gap $\gamma_{d'} - \gamma_{d+1}$ in singular values determines how the suboptimality affects the magnitude of ‘‘signal’’ captured.

Step 3: Given that ϵ -optimal solutions can capture the top directions of $P_{\Phi_o} \bar{A}_o$ (or U), we now focus our attention on what this means for downstream performance. We invoke [Lemma B.3](#) again to upper bound the downstream classification error L_{clf} (refer [Definition B.1](#)) as

$$L_{\text{clf}}(f; \bar{y}^*) \leq L_{\text{reg}}(f; \bar{y}^*) \leq L_{\text{reg}}(F_o) = \inf_{w \in \mathbb{R}^d} \|\bar{A}_o F_o w - \bar{y}_o^*\|_2^2 \quad (44)$$

$$\stackrel{(a)}{\leq} \inf_{w \in \mathbb{R}^d} \|\bar{A}_o P_{\Phi_o} F_o w - \bar{y}_o^*\|_2^2 \stackrel{(b)}{\leq} \inf_{w \in \mathbb{R}^d} \|\bar{A}_o P_{\Phi_o} P_{F_o} w - \bar{y}_o^*\|_2^2 \quad (45)$$

$$\stackrel{(c)}{\leq} \inf_{w \in \mathbb{R}^d} \|VSU^\top P_{F_o} w - \bar{y}_o^*\|_2^2 \quad (46)$$

where (a) follows from the fact that $F_o \in \mathcal{F}_{\Phi_o}$ and thus $P_{\Phi_o} F_o = F_o$, (b) is true since for any $w \in \mathbb{R}^d$, there exists $w' \in \mathbb{R}^d$ such that $F_o w = P_{F_o} w'$, and (c) uses the singular value decomposition of $P_{\Phi_o} \bar{A}_o$. Thus the downstream error is upper bounded by a quantity that depends on how much of \bar{y}_o^* is not captured by the columns of $\bar{A}_o P_{\Phi_o} P_{F_o} = VSU^\top P_{F_o}$. We show this quantity is small, by arguing that the top d' directions of V captures enough component of \bar{y}_o^* [Lemma B.6](#), and that an ϵ -optimal representation will capture a large enough portion of the top d' directions. Note that for any matrix $B \in \mathbb{R}^{n \times n}$, $B_{:m} \in \mathbb{R}^{m \times n}$ denotes the first m columns of B and $B_{m:} \in \mathbb{R}^{n-m \times n}$ denotes that last m columns of B . The calculation is as follows

$$\begin{aligned} L_{\text{clf}}(f; \bar{y}^*) &\leq \inf_{w \in \mathbb{R}^d} \|VSU^\top P_{F_o} w - \bar{y}_o^*\|_2^2 \\ &= \inf_{w \in \mathbb{R}^d} \|VSU^\top P_{F_o} w - V_{:d'} V_{:d'}^\top \bar{y}_o^* + V_{d'} V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &\stackrel{(a)}{\leq} 2 \left(\inf_{w \in \mathbb{R}^d} \|VSU^\top P_{F_o} w - V_{:d'} V_{:d'}^\top \bar{y}_o^*\|_2^2 + \|V_{d'} V_{d'}^\top \bar{y}_o^*\|_2^2 \right) \\ &= 2 \inf_{w \in \mathbb{R}^d} \|V_{:d'} S_{:d'} U_{:d'}^\top P_{F_o} w - V_{:d'} V_{:d'}^\top \bar{y}_o^*\|_2^2 + 2 \|V_{d'} V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &= 2 \inf_{w \in \mathbb{R}^d} \|S_{:d'} U_{:d'}^\top P_{F_o} w - V_{:d'}^\top \bar{y}_o^*\|_2^2 + 2 \|V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &\stackrel{(b)}{\leq} 2 \|S_{:d'} U_{:d'}^\top P_{F_o} U_{:d'} S_{:d'}^{-1} V_{:d'}^\top \bar{y}_o^* - V_{:d'}^\top \bar{y}_o^*\|_2^2 + 2 \|V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &= 2 \|S_{:d'} U_{:d'}^\top P_{F_o}^\perp U_{:d'} S_{:d'}^{-1} V_{:d'}^\top \bar{y}_o^*\|_2^2 + 2 \|V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &\stackrel{(c)}{\leq} 2 \|S_{:d'}\|_2^2 \|P_{F_o}^\perp U_{:d'}\|_2^2 \|S_{:d'}^{-1}\|_2^2 \|\bar{y}_o^*\|_2^2 + 2 \|V_{d'}^\top \bar{y}_o^*\|_2^2 \\ &\stackrel{(d)}{\leq} \frac{2 \|P_{F_o}^\perp U_{:d'}\|_F^2}{\gamma_{d'}} + 2 \|V_{d'}^\top \bar{y}_o^*\|_2^2 \end{aligned}$$

where (a) follows from the inequality $\|a + b\|_2^2 \leq 2(\|a\|_2^2 + \|b\|_2^2)$, (b) follows by a picking a specific value $w = U_{:d'} S_{:d'}^{-1} V_{:d'}^\top \bar{y}_o^*$, (c) follows from multiple applications of Cauchy-Schwarz inequality and that $\|V_{:d'}^\top \bar{y}_o^*\|_2 \leq \|\bar{y}_o^*\|_2$ and (d) follows from $\|S_{:d'}\|_2 \leq 1$ and $\|S_{:d'}^{-1}\|_2 \leq \gamma_{d'}^{-1}$.

The first term is upper bounded in step 2 already by the sub-optimality of f , while the second term is upper bounded using [Lemma B.8](#). Plugging these in, we get

$$L_{\text{clf}}(f; \bar{y}^*) \leq \frac{2\epsilon d'}{\gamma_{d'}(\gamma_{d'} - \gamma_{d+1})^2} + \frac{2(1 - \|P_{\Phi_o} \bar{A}_o^\top \bar{y}_o^*\|_2)}{1 - \gamma_{d'+1}} \leq \frac{2\epsilon d'}{\gamma_{d'}(\gamma_{d'} - \gamma_{d+1})^2} + \frac{4(1 - \|P_{\Phi_o} \bar{A}_o^\top \bar{y}_o^*\|_2)}{1 - \gamma_{d'+1}}$$

where for the last inequality we use that $1 - x^2 = (1 - x)(1 + x) \leq 2(1 - x)$ for $x \in [0, 1]$. This is further simplified using Lemma B.6 to

$$L_{\text{clf}}(f; \bar{y}^*) \leq \frac{2\epsilon d'}{\gamma_{d'}(\gamma_{d'} - \gamma_{d+1})^2} + \frac{4(2\Delta_{\mathcal{A}}(g, \bar{y}^*) + \sqrt{L_{\text{reg}}(\phi; g)})}{1 - \gamma_{d'+1}}$$

Step 4: Finally the singular values γ_i are linked the eigenvalues λ_i in the theorem statement through Lemma B.9. Specifically, we have $\gamma_i = 1 - \lambda_i$, giving us the final result

$$L_{\text{clf}}(f; \bar{y}^*) \leq \frac{4(2\Delta_{\mathcal{A}}(g, \bar{y}^*) + \sqrt{L_{\text{reg}}(\phi; g)})}{\lambda_{d'+1}} + \frac{2\epsilon d'}{(1 - \lambda_{d'})^2(\lambda_{d+1} - \lambda_{d'})^2}$$

where ϵ is the suboptimality $L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)$. The above inequality holds for every $g \in \{\pm 1\}^{\mathcal{X}}$ and for every $d' \in [d]$. Taking a min over both completes the proof. This completes the proof. \square

B.4. Discussion of upper bound

We dissect our result from Theorem 4.5, and compare it to the result from HaoChen et al. (2021), presented in Theorem 4.3. For the representation $f \in \mathcal{F}_\phi$, downstream performance is good if

- $L_{\text{spec}}(f) - \inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*)$ is small: The contrastive loss of f is close to the optimal loss in \mathcal{F}_ϕ , even if best in class is far from the absolute minimizer. The equivalent term in Theorem 4.3 was the global sub-optimality of f , i.e. $L_{\text{spec}}(f) - \inf_{f^*} L_{\text{spec}}(f^*)$.
- $2\Delta_{\mathcal{A}}(g, \bar{y}^*) + \sqrt{L_{\text{reg}}(\phi; g)}$ is small: This happens if there exists a predictor $g \in \{\pm 1\}^{\mathcal{X}}$ on augmentations that is expressible by the features ϕ and is sufficiently consistent with the ground-truth labels \bar{y}^* on inputs. Note that if augmentation distributions overlap across classes, then $\Delta_{\mathcal{A}}(g, \bar{y}^*)$ cannot be made small. In fact, $\Delta_{\mathcal{A}}(g, \bar{y}^*)$ is of the same order as α from Theorem 4.3. The extra condition we need here is that $\sqrt{L_{\text{reg}}(\phi; g)}$ is small, i.e. despite ϕ not being full rank, it can roughly express a function that is consistent with ground-truth labels.
- Eigenvalues $\lambda_{d'}$ and eigen-gaps $\lambda_{d+1} - \lambda_{d'}$ are not too small: This is very similar to Theorem 4.3, except there the eigenvalues were of the normalized Laplacian (that only depended on distributions), while here the eigenvalues also depend on ϕ and thus the function class. Intuitively these values are large if the augmentation graph is dense in the view of the features ϕ .

B.4.1. FULL RANK ϕ

When ϕ is full rank, i.e. $D = |\mathcal{X}|$, the bound from Theorem 4.5 can recover the results from HaoChen et al. (2021). In this case this case we have

- $L_{\text{reg}}(\phi; g) = 0$, since a full rank ϕ can express any function in $\mathbb{R}^{\mathcal{X}}$.
- $\inf_{f^* \in \mathcal{F}_\phi} L_{\text{spec}}(f^*) = \inf_{f^*} L_{\text{spec}}(f^*)$ since \mathcal{F}_ϕ can express all d -dimensional representations.
- $\min_g \Delta_{\mathcal{A}}(g, \bar{y}^*) = \mathcal{O}(\alpha)$, where α (from Theorem 4.3) is the minimum error in predicting labels from augmentations. Setting g to be an optimal augmentations to label predictor and plugging into Definition 4.4 proves this.
- Finally, the matrix $I_D - \Sigma(\phi)^{-\frac{1}{2}} \Sigma(\phi_{\mathcal{A}}) \Sigma(\phi)^{-\frac{1}{2}}$ is closely related to the normalized Laplacian from Section 4.2. Proof of this is presented in Lemma B.10.

C. Lower bounds for (approximately) disjoint augmentations

Here we prove that the global minimizer of the contrastive objective can achieve trivial downstream performance when the augmentation distributions do not overlap.

Theorem C.1. *Let $N \in \mathbb{N}$ be given and let $d \in \mathbb{N}$ satisfy $3 \leq d \leq cN / \log_2(N)$ for a universal constant $c > 0$. Let $\bar{\mathcal{X}}$ be a set of $|\bar{\mathcal{X}}| = N$ instances, $\bar{y}^* \in \{\pm 1\}^N$ be any labeling function with $\sum_i y_i^* = 0$, and let \mathcal{D} be the uniform distribution over $\bar{\mathcal{X}}$. Suppose that the augmentation distribution $\mathcal{A}(\cdot | \bar{x})$ is such that $\forall \bar{x}, \bar{x}' \in \bar{\mathcal{X}} : \text{supp}(\mathcal{A}(\cdot | \bar{x})) \cap \text{supp}(\mathcal{A}(\cdot | \bar{x}')) = \emptyset$. Additionally assume either*

- **Unnormalized case:** representations are unconstrained; or
- **Normalized case:** representations are constrained (to any set) and there is a fixed source of randomness $W \in \Delta(\mathcal{W})$ and mapping $T : \bar{\mathcal{X}} \times \mathcal{W} \rightarrow \mathcal{X}$ that is invertible in w for any \bar{x} such that $x \sim \mathcal{A}(\cdot | \bar{x}) \equiv w \sim W, x = T(\bar{x}, w)$.

Then for any representation $f^* : \mathcal{X} \rightarrow \mathbb{R}^d$ there exists a representation $\hat{f} : \mathcal{X} \rightarrow \mathbb{R}^d$ such that:

$$L_{\text{cont}}(\hat{f}) \leq L_{\text{cont}}(f^*), \quad \text{and} \quad L_{\text{clf}}(\hat{f}) = \min_{w \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^n \mathbf{1}\{\text{sign}(w^\top \hat{f}(\bar{x}_i)) \neq y_i^*\} \geq \frac{1}{2} - O(\sqrt{d \log(N)/N})$$

Proof. Let us start by considering the unnormalized case. Let $f^* : \mathcal{X} \rightarrow \mathbb{R}^d$ be any representation function. The proof consists of three steps:

1. Show that every instance $\bar{x} \in \bar{\mathcal{X}}$ has an embedding $v_{\bar{x}}$, such that if we embed \bar{x} and all of its augmentations to $v_{\bar{x}}$ then we obtain a new embedding function \hat{f} for which $L_{\text{cont}}(\hat{f})$ is no worse than $L_{\text{cont}}(f^*)$.
2. Let $\mathcal{V} := \{v_{\bar{x}} : \bar{x} \in \bar{\mathcal{X}}\}$. Show that for any bijection $\pi : \mathcal{V} \rightarrow \mathcal{V}$, we have $L_{\text{cont}}(\pi \circ \hat{f}) = L_{\text{cont}}(\hat{f})$. In other words, if we apply a permutation to the embeddings of \hat{f} we do not change the contrastive loss.
3. Show that there exists some permutation π such that $\pi \circ \hat{f}$ has very high downstream error rate.

Part 1. Let us first show that embeds all augmentations of an instance \bar{x} identically only lowers the contrastive loss. By convexity, we have that

$$\begin{aligned} L_{\text{cont}}(f) &:= \mathbb{E}_{\bar{x}_1, \bar{x}_2 \sim \mathcal{D}} \mathbb{E}_{(x, x^+) \sim \mathcal{A}(\cdot | \bar{x}_1), x^- \sim \mathcal{A}(\cdot | \bar{x}_2)} \left[-\log \left(\frac{\exp(f(x)^\top f(x^+))}{\exp(f(x)^\top f(x^+)) + \exp(f(x)^\top f(x^-))} \right) \right] \\ &\geq \mathbb{E}_{\bar{x}_1, \bar{x}_2 \sim \mathcal{D}} \left[-\log \left(\frac{\exp(g(\bar{x}_1)^\top g(\bar{x}_1))}{\exp(g(\bar{x}_1)^\top g(\bar{x}_1)) + \exp(g(\bar{x}_1)^\top g(\bar{x}_2))} \right) \right] \end{aligned}$$

where $g(\bar{x}) = \mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})}[f(x)]$ is the mean embedding for f . Note that this inequality is strict if f does not embed all augmentations of an instance identically. If f^* did not embed augmentations identically, we could replace f^* with the mean embedding \hat{f} and reduce the contrastive loss. Thus there exists a set $\mathcal{V} := \{v_{\bar{x}} : \bar{x} \in \bar{\mathcal{X}}\}$ such that for each \bar{x} , we can assume that \hat{f} embeds \bar{x} and all of its augmentations as $v_{\bar{x}}$.

The same argument also hold for the spectral contrastive loss defined in Equation (2), since it is also convex in the inner products.

Part 2. Let us rename the embedding vectors $\{v_{\bar{x}}\}_{\bar{x} \in \bar{\mathcal{X}}}$ to $\mathcal{V} := \{v_i\}_{i=1}^N$. Then we can rewrite the objective as

$$L_{\text{cont}}(\hat{f}) = \frac{1}{N^2} \sum_{i,j} -\log \left(\frac{\exp(v_i^\top v_i)}{\exp(v_i^\top v_i) + \exp(v_i^\top v_j)} \right)$$

Any bijection from \mathcal{V} to \mathcal{V} can be equivalently viewed as a permutation $\pi : [N] \rightarrow [N]$. We claim that the above objective is invariant to permuting the indices. This is easy to see, since all pairs (i, j) appear with equal weighting in the above expression. Thus we see that $L_{\text{cont}}(\pi \circ \hat{f}) = L_{\text{cont}}(\hat{f})$.

Part 3. In the last step of the proof, we use a combinatorial argument to show that there exists some permutation π with high error rate. First, note that the embedding function $\pi \circ \hat{f}$ embeds \bar{x}_i and all of its augmentations to $v_{\pi(i)}$. Thus, the downstream loss when using linear function w is

$$\frac{1}{N} \sum_{i=1}^n \mathbf{1}\{\text{sign}(w^\top v_{\pi(i)}) \neq y_i^*\} = \frac{1}{N} \sum_{i=1}^n \mathbf{1}\{\text{sign}(w^\top v_i) \neq y_{\pi^{-1}(i)}^*\}$$

So instead of permuting the embeddings $\{v_i\}$, we can equivalently permute the labels $\{y_i^*\}$. Define:

$$\begin{aligned} \mathcal{Y} &:= \{(y_{\pi(i)}^*)_{i=1}^N : \pi \text{ is a permutation}\} \\ \mathcal{W} &:= \{(\text{sign}(w^\top v_i))_{i=1}^N : w \in \mathbb{R}^d\} \\ \mathcal{Z}_\tau &:= \{x \in \{\pm 1\}^N : \exists b \in \mathcal{W} \text{ s.t. } \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{x_i \neq b_i\} \leq \tau\}. \end{aligned}$$

Here \mathcal{Y} are the possible labellings we can generate by permuting the indices (which as we discussed is equivalent to permuting the embedding vectors). \mathcal{W} is the labels we can generate via a linear function of the embeddings. Finally \mathcal{Z}_τ is the set of labellings that are τ close to the ones that our embeddings can generate. The statement of the theorem is equivalent to $\mathcal{Y} \setminus \mathcal{Z}_\tau \neq \emptyset$ for some large τ , which means that there is some permutation of the labels that is far from all linear functions of our embeddings.

We prove this via a combinatorial argument. First, since $\sum_i y_i^* = 0$ (meaning that the classes are balanced), we have $|\mathcal{Y}| = \binom{N}{N/2} \geq 2^{N/2}$. On the other hand, by Sauer's lemma,

$$|\mathcal{W}| \leq \sum_{i=0}^d \binom{N}{i}, \quad \text{hence} \quad |\mathcal{Z}_\tau| \leq \sum_{i=0}^d \binom{N}{i} \cdot \sum_{i=0}^{N\tau} \binom{N}{i} \quad (47)$$

Let $H(p) := p \log_2(1/p) + (1-p) \log_2(1/(1-p))$ be the binary entropy function, for $p \in [0, 1]$. Standard bounds on the volume of Hamming cubes (Cover, 1999) gives that

$$|\mathcal{Z}_\tau| \leq 2^{H(d/N) \cdot N} 2^{H(\tau) \cdot N}.$$

We also have

$$\binom{N}{N/2} \geq 2^{H(1/2) \cdot N} \cdot \frac{2}{eN} \geq 2^{N - \log_2(eN/2)}$$

Therefore, a sufficient condition is

$$H(d/N) + H(\tau) \leq 1 - \frac{\log_2(eN/2)}{N}$$

To proceed, we upper bound the entropy functional on the left hand side using the Taylor expansion. For the $H(d/N)$ term we use a first order expansions around $p = 1/N$, which, by concavity, yields an upper bound.

$$\begin{aligned} H(d/N) &\leq H(1/N) + \left. \frac{\partial H(x)}{\partial x} \right|_{x=1/N} (d/N - 1/N) \\ &= 1/N \log_2(N) + (1 - 1/N) \log_2(N/(N-1)) - \log_2\left(\frac{1/N}{1-1/N}\right) \cdot (d/N - 1/N) \\ &= \log_2(N/(N-1)) + d/N \cdot \log_2(N-1) \\ &\leq 2/N + d \log_2(N)/N. \end{aligned}$$

The last inequality holds for $N \geq 2$. For $H(\tau)$ we have the upper bound

$$H(\tau) = H(1/2) - \frac{4}{\ln(2)} \cdot \frac{1}{2} (1/2 - \tau)^2 + \frac{H^{(3)}(\xi)}{6} (\tau - 1/2)^3 \leq 1 - \frac{2}{\ln(2)} (1/2 - \tau)^2,$$

Here the first equality is Taylor’s remainder theorem where $\xi \in [\tau, 1/2]$ and the second holds because the third derivative is non-negative on the interval $[0, 1/2]$ and we will take $\tau \leq 1/2$. Putting these together, a sufficient condition is

$$\begin{aligned} \frac{2}{N} + \frac{d \log_2(N)}{N} + 1 - \frac{2}{\ln(2)}(1/2 - \tau)^2 &\leq 1 - \frac{\log_2(eN/2)}{N} \\ \Leftrightarrow \tau &< \frac{1}{2} - \sqrt{\frac{\ln(2)}{2} \left(\frac{d \log_2(N)}{N} + \frac{2}{N} + \frac{\log_2(eN/2)}{N} \right)} \end{aligned}$$

So the error rate is $1/2 - O(\sqrt{d \log(N)/N})$.

For the normalized case, the proof is structurally very similar, except that we cannot rely on the argument in part 1 to show that f^* embeds \bar{x} and all of its augmentations to the same vector $v_{\bar{x}}$. However, we only use the mean vector $v_{\bar{x}} := \mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})}[f^*(x)]$ in subsequent steps of the proof and we will see that we can remap embeddings $f^*(x)$ so that we (a) preserve the NCE loss of f^* and (b) permute all of the mean vectors $v_{\bar{x}}$.

Let us number the original inputs $\bar{x}_1, \dots, \bar{x}_N$ and let $\pi : [N] \rightarrow [N]$ be any permutation. Let \mathcal{W} be the choices for the random seed and for input \bar{x}_i let $x_{i,w} = T(\bar{x}_i, w)$ be the augmentation obtained when using seed w on input \bar{x} . The invertibility of $\mathcal{A}(\bar{x}_i, \cdot)$ implies that $x_{i,w} \neq x_{i,w'}$. This means that we can define a new predictor f_π as

$$f_\pi : x_{i,w} \mapsto f^*(x_{\pi(i),w}).$$

Since the examples are sampled uniformly at random and since the random seed is independent of the example, we can show that $L_{\text{cont}}(f_\pi) = L_{\text{cont}}(f^*)$ using a similar argument to the one we used to show permutation invariance in the unnormalized case. At the same time, we have changed the mean embeddings so that \bar{x} is now embedded as $v_{\bar{x}_{\pi(i)}}$. So now we can continue with part 3 to obtain the result. □

C.1. Prior lower bounds

Proposition 1 in [Robinson et al. \(2021\)](#) discusses a similar lower bound to [Lemma 4.2](#) when augmentations are disjoint, arguing that contrastive learning can find “shortcut solutions” that can lead to feature suppression. While the motivation is similar to ours, those results are shown specifically for contrastive loss with normalized representations, in the regime of large number of negative samples and with a specific uniform over sphere assumption on latent variables generating the data. The above results are shown in much more general settings. [Wang et al. \(2022\)](#) also show a lower bound for why the alignment and uniformity properties from [Wang & Isola \(2020\)](#) is insufficient to guarantee good downstream performance, through a non-overlapping augmentation example.

C.2. Approximately disjoint augmentations

Definition C.1. For an augmentation distribution \mathcal{A} , we define *Bayes-error*(\mathcal{A}) as the Bayes error of augmentation classification as the minimum error achievable in the input identification task, i.e. predicting the input that could have generated an augmentation. Formally we define it as follows:

$$\text{Bayes-error}(\mathcal{A}) = \inf_{g: \mathcal{X} \rightarrow \bar{\mathcal{X}}} \mathbb{E}_{\bar{x}} \left[\mathbb{E}_{x \sim \mathcal{A}(\cdot | \bar{x})} [\mathbb{1}\{g(x) \neq \bar{x}\}] \right] \quad (48)$$

Lemma C.2. For an augmentation distribution \mathcal{A} , the Bayes error from [Definition C.1](#) has the following expression

$$\text{Bayes-error}(\mathcal{A}) = 1 - \mathbb{E}_{x \sim \mathcal{D}} [\|\mathcal{A}(\cdot | x)\|_\infty] \quad (49)$$

where $\mathcal{A}(\cdot | x)$ is the posterior distribution over original inputs given an augmentation x .

Proof. In the above definition of Bayes error, we pick the optimal predictor g to be $g(x) = \arg \max_{\bar{x}} \mathcal{A}(\bar{x} | x)$, which will give us the expression for Bayes error. □

Lemma C.3. Consider the augmentation distribution \mathcal{A} and its normalized adjacency matrix A_\circ , and let $\lambda_1, \dots, \lambda_{|\mathcal{X}|}$ be the eigenvalues of the normalized Laplacian $I_{|\mathcal{X}|} - A_\circ$ in increasing order. The eigen-gap $\lambda_{d+1} - \lambda_d$ can be upper bounded as follows:

$$\lambda_{d+1} - \lambda_d \leq \lambda_{d+1} \leq \frac{2\bar{\rho} \text{Bayes-error}(\mathcal{A})}{1 - d/|\bar{\mathcal{X}}|} \quad (50)$$

where $\bar{\rho} = \frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}}$ is the ratio of max and min probabilities over inputs.

Proof. Let γ_d be the d^{th} largest eigenvalue of the normalized adjacency matrix $A_\circ = D^{-\frac{1}{2}}AD^{-\frac{1}{2}} \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$, where $A[x, x']$ is the joint probability of augmentations x and x' appearing as two augmentations of the same input. Then we know that the i^{th} smallest eigenvalue of $I_{|\mathcal{X}|} - A_\circ$ is $\lambda_i = 1 - \gamma_i$. Furthermore we note that $A_\circ \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ has rank at most $|\bar{\mathcal{X}}|$, since from Table 2 we know that $A_\circ = \bar{A}_\circ^\top \bar{A}_\circ$, where $\bar{A}_\circ \in \mathbb{R}^{\bar{\mathcal{X}} \times \mathcal{X}}$ is the normalized of the input-augmentation distribution (refer Table 2) that has entries $\bar{A}_\circ[\bar{x}, x] = \frac{\mathcal{D}_{\text{sim}}(x, \bar{x})}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)}\sqrt{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})}}$. Thus we can conclude that $\gamma_i = 0$ for $|\bar{\mathcal{X}}| < i \leq |\mathcal{X}|$. First we prove the statement $A_\circ = \bar{A}_\circ^\top \bar{A}_\circ$ below

$$\begin{aligned} (\bar{A}_\circ^\top \bar{A}_\circ)[x, x'] &= \sum_{\bar{x}} \bar{A}_\circ[\bar{x}, x] \bar{A}_\circ[\bar{x}, x'] = \sum_{\bar{x}} \frac{\mathcal{A}(x, \bar{x})}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)}\sqrt{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})}} \frac{\mathcal{A}(x', \bar{x})}{\sqrt{\mathcal{D}_{\mathcal{X}}(x')}\sqrt{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})}} \\ &= \frac{1}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)}\sqrt{\mathcal{D}_{\mathcal{X}}(x')}} \sum_{\bar{x}} \frac{\mathcal{A}(x, \bar{x})\mathcal{A}(x', \bar{x})}{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})} = \frac{1}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)}\sqrt{\mathcal{D}_{\mathcal{X}}(x')}} \sum_{\bar{x}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \mathcal{A}(x | \bar{x}) \mathcal{A}(x' | \bar{x}) \\ &= \frac{A[x, x']}{\sqrt{\mathcal{D}_{\mathcal{X}}(x)}\sqrt{\mathcal{D}_{\mathcal{X}}(x')}} = A_\circ[x, x'] \end{aligned}$$

We now connect $\text{Bayes-error}(\mathcal{A})$ to the normalized augmentation matrix A_\circ by using Lemma C.2.

$$\begin{aligned} (1 - \text{Bayes-error}(\mathcal{A}))^2 &\stackrel{(a)}{=} \left(\mathbb{E}_{x \sim \mathcal{D}} [\|\mathcal{A}(\cdot | x)\|_\infty] \right)^2 \leq \stackrel{(b)}{=} \left(\mathbb{E}_{x \sim \mathcal{D}} [\|\mathcal{A}(\cdot | x)\|_2] \right)^2 \\ &\leq \stackrel{(c)}{=} \mathbb{E}_{x \sim \mathcal{D}} [\|\mathcal{A}(\cdot | x)\|_2^2] = \sum_{x \in \mathcal{X}} \mathcal{D}_{\mathcal{X}}(x) \sum_{\bar{x} \in \bar{\mathcal{X}}} \mathcal{A}(\bar{x} | x)^2 \\ &= \sum_{\bar{x} \in \bar{\mathcal{X}}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \sum_{x \in \mathcal{X}} \mathcal{A}(x | \bar{x}) \mathcal{A}(\bar{x} | x) \\ &= \sum_{\bar{x}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \sum_x \frac{\mathcal{A}(x, \bar{x})}{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})} \frac{\mathcal{A}(x, \bar{x})}{\mathcal{D}_{\mathcal{X}}(x)} \\ &= \sum_{\bar{x}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \sum_x \frac{\mathcal{A}(x, \bar{x})}{\sqrt{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})}\sqrt{\mathcal{D}_{\mathcal{X}}(x)}} \frac{\mathcal{A}(x, \bar{x})}{\sqrt{\mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})}\sqrt{\mathcal{D}_{\mathcal{X}}(x)}} \\ &= \sum_{\bar{x}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) \sum_x \bar{A}_\circ[\bar{x}, x] \bar{A}_\circ[\bar{x}, x] = \sum_{\bar{x}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x}) (\bar{A}_\circ \bar{A}_\circ^\top)[\bar{x}, \bar{x}] \\ &= \text{tr}(\bar{D} \bar{A}_\circ \bar{A}_\circ^\top) \end{aligned}$$

where (a) follows from Lemma C.2, (b) follows from $\|\cdot\|_\infty \leq \|\cdot\|_2$, (c) follows from Jensen's inequality since $h(x) = x^2$ is convex. This upper bound can be used to lower bound the Bayes error as follows:

$$\begin{aligned} 2 \text{Bayes-error}(\mathcal{A}) &\geq 1 - (1 - \text{Bayes-error}(\mathcal{A}))^2 \\ &\stackrel{(a)}{\geq} 1 - \text{tr}(\bar{D} \bar{A}_\circ \bar{A}_\circ^\top) \stackrel{(b)}{=} \text{tr}(\bar{D}) - \text{tr}(\bar{D} \bar{A}_\circ \bar{A}_\circ^\top) \stackrel{(c)}{=} \text{tr}(\bar{D}(I_{|\bar{\mathcal{X}}|} - \bar{A}_\circ \bar{A}_\circ^\top)) \\ &\stackrel{(d)}{\geq} \|\bar{D}^{-1}\|_2^{-1} \text{tr}(I_{|\bar{\mathcal{X}}|} - \bar{A}_\circ \bar{A}_\circ^\top) = \bar{D}_{\min} \text{tr}(I_{|\bar{\mathcal{X}}|} - \bar{A}_\circ \bar{A}_\circ^\top) \end{aligned} \quad (51)$$

where $\bar{D}_{\min} = \min_{\bar{x} \in \bar{\mathcal{X}}} \mathcal{D}_{\bar{\mathcal{X}}}(\bar{x})$. In the above sequence, (a) follows the preceding calculation, (b) follows from $\text{tr}(\bar{D}) = \sum_{\bar{x}} \bar{D}(\bar{x}) = 1$ and (c) follows from linearity of the trace operator. The penultimate step (d) follows from the fact that

$\text{tr}(XY) \leq \|X\|_2 \text{tr}(Y)$ for symmetric psd matrices X, Y ; a proof for this can be found in Lemma 18 from Jin et al. (2017). We now connect this quantity to the eigenvalues of A_o as follows:

$$\begin{aligned} \text{tr}(I_{|\bar{\mathcal{X}}|} - \bar{A}_o \bar{A}_o^\top) &= |\bar{\mathcal{X}}| - \text{tr}(\bar{A}_o \bar{A}_o^\top) \stackrel{(a)}{=} |\bar{\mathcal{X}}| - \text{tr}(\bar{A}_o^\top \bar{A}_o) = |\bar{\mathcal{X}}| - \text{tr}(A_o) \\ &\stackrel{(b)}{=} |\bar{\mathcal{X}}| - \sum_{i=1}^{|\bar{\mathcal{X}}|} \gamma_i \stackrel{(c)}{=} |\bar{\mathcal{X}}| - \sum_{i=1}^{|\bar{\mathcal{X}}|} \gamma_i = |\bar{\mathcal{X}}| - \sum_{i=1}^{|\bar{\mathcal{X}}|} (1 - \lambda_i) = \sum_{i=1}^{|\bar{\mathcal{X}}|} \lambda_i \\ &\geq \sum_{i=d+1}^{|\bar{\mathcal{X}}|} \lambda_i \geq (|\bar{\mathcal{X}}| - d) \lambda_{d+1} \end{aligned} \quad (52)$$

where (a) follows from $\text{tr}(PQ) = \text{tr}(QP)$, (b) is true because γ_i 's are the eigenvalues of A_o and because trace of a symmetric matrix is the sum of its eigenvalues, (c) follows because A_o is rank $|\bar{\mathcal{X}}|$ and so $\gamma_i = 0$ for $i > |\bar{\mathcal{X}}|$. Combining Equations (51) and (52), we get $\text{Bayes-error}(\mathcal{A}) \geq 1/2 \bar{\mathcal{D}}_{\min} (|\bar{\mathcal{X}}| - d) \lambda_{d+1}$. Note that

$$\bar{\mathcal{D}}_{\min} = \frac{\bar{\mathcal{D}}_{\min}}{\bar{\mathcal{D}}_{\max}} \bar{\mathcal{D}}_{\max} \geq \frac{\bar{\mathcal{D}}_{\min}}{\bar{\mathcal{D}}_{\max}} |\bar{\mathcal{X}}|^{-1} = \bar{\rho}^{-1} |\bar{\mathcal{X}}|^{-1} \quad (53)$$

Plugging this into the bound gives $\text{Bayes-error}(\mathcal{A}) \geq \frac{1}{2\bar{\rho}} \left(1 - \frac{d}{|\bar{\mathcal{X}}|}\right) \lambda_{d+1}$, giving us

$$\lambda_{d+1} \leq \frac{2\bar{\rho} \text{Bayes-error}(\mathcal{A})}{1 - d/|\bar{\mathcal{X}}|}$$

□

D. Experiment details

In this section, we provide additional notes, tables, and figures on the experiments.

D.1. Synthetic experiments: hypercube example

Figure 5 shows the results from Section 3 in greater detail. This section completes the details omitted in the main paper.

Data and augmentations. As outlined in Section 3, the data are drawn uniformly from the hypercube in dimension $D = 50$. The downstream labels are determined by a randomly drawn linear classifier w , whose first $k = 10$ coefficients are drawn from $\mathcal{N}(0, 1)$; the rest are 0. The training set (under which L_{cont} is minimized) is of size 50000; the downstream accuracies under a linear classifier are evaluated on a holdout validation set of 12500. The augmentations are selected by i.i.d. random scaling factors $\tau \sim \text{Unif}([0, 1])$ and scaling down the last 40 coordinates.

Training and evaluation. The two-layer MLP models used a hidden layer width of $2D = 100$, and an output (i.e. representation) dimension of 20. Adam was run with a learning rate of 10^{-3} , and default parameters $\beta_1 = 0.9, \beta_2 = 0.99$. The weight decay parameter for Adam was 0.004, selected from $\{0.001, 0.002, \dots, 0.007\}$ based on best transfer performance. SGD was run with learning rate 0.01. Quantitative results are shown in Table 1; means and 95% confidence intervals are computed from 10 random seeds. 500 epochs of pre-training were run, with batch size 512.

D.2. CIFAR-10 + SimCLR experiments

For all ResNet experiments, we use the ResNet-18 architecture from PyTorch, with the standard modification for CIFAR-10 of replacing the first 7×7 convolution layer with a 3×3 convolution and removing the maxpool layer. We use the ViT implementation from <https://github.com/lucidrains/vit-pytorch> with patch size: 4, hidden dimension: 256, depth: 6 and number of heads: 8. For MLP-Mixer we use the implementation from <https://github.com/lucidrains/mlp-mixer-pytorch> with patch size: 4, hidden dimension: 256 and number of heads: 8. In each model, the representation for contrastive learning is computed by adding an extra MLP (projection layer)

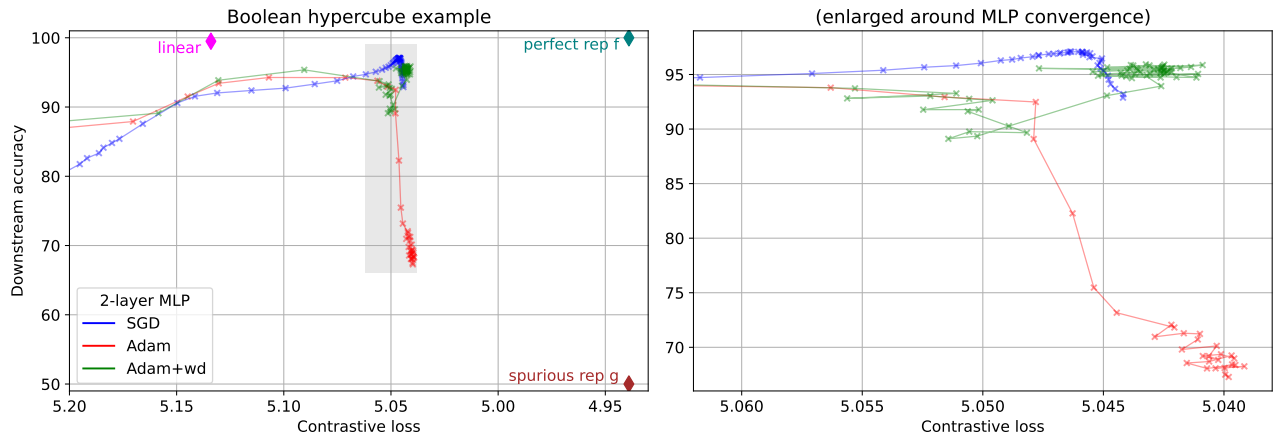


Figure 5. Full plots for the synthetic experiments, with all contrastive loss minimizers shown from various function classes (left) and enlarged plot near convergence of trajectories of solutions found by training 2-layered MLPs with various configurations of first-order optimizers (right).

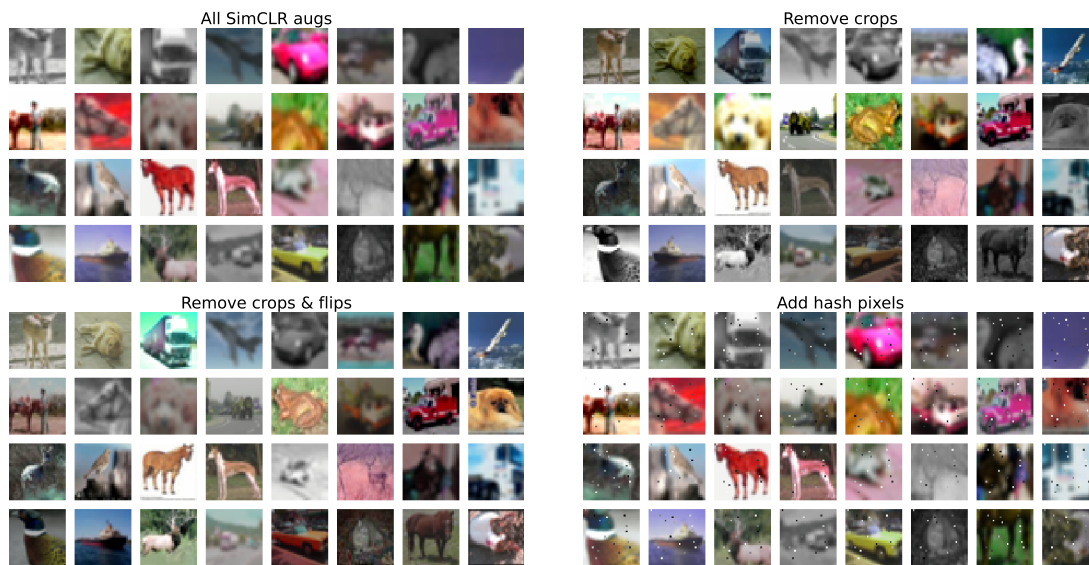


Figure 6. Examples of augmented images from CIFAR-10 used in the SimCLR experiments. **TL**: Full pipeline of augmentations from SimCLR (Chen et al., 2020). **TR**: Remove random cropping. **BL**: Remove random cropping and horizontal flip. **BR**: Add “hash pixels” to each image, which uniquely identify the particular example.

Table 3. Hyperparameter values for experiments on CIFAR-10 trained using ResNet-18.

Hyperparameters	Values	
Contrastive training	Max epoch	1000
	Learning rate	0.001
	Optimizer	Adam + weight decay (0.0005)
	Batch size	512
	Representation dimension	1024
Downstream training	Epochs	1000
	Learning rate (start)	0.01
	Optimizer	Adam + weight decay (0.000005)
	Scheduler	ExponentialLR (gamma: $10^{0.004}$)
	Batch size	1000

on top of the base model, as proposed in Chen et al. (2020). The projection layer has 1 hidden layer with 2048 dimensions, followed by a batch norm layer and ReLU non-linearity, and output dimensionality of 1024.

Augmentations. The following augmentations are used, inspired by (Chen et al., 2020):

```
transforms.Compose([
    RandomResizedCrop(32, scale=(0.3, 1.0)),
    RandomHorizontalFlip(p=0.5),
    transforms.RandomApply([transforms.ColorJitter(0.4, 0.4, 0.4, 0.1)],
        p=0.8),
    RandomGrayscale(p=0.2),
    GaussianBlur(kernel_size=3)
])
```

For experiments in Figure 3 we use the full pipeline of augmentations (top left) and sequentially remove random cropping (top right) and horizontal flipping (bottom left). Examples of augmented CIFAR-10 images are shown in Figure 6

Contrastive training. We train the model for 1000 epochs, by performing a pass over this training dataset and minimize the SimCLR contrastive learning loss. We normalize the representations f to unit norm when computing the SimCLR loss, as is common in many works. Formally, given a batch $\{(x_i, x'_i)\}_{i=1}^B$ of pairs of augmentations we perform a single update of Adam to minimize the following loss:

$$L(f) = -\frac{1}{2B} \sum_{i=1}^n \frac{w(x_i, x'_i)}{\sum_{j=1}^n w(x_i, x'_j) + \sum_{j=1, j \neq i}^n w(x_i, x_j)} - \frac{1}{2B} \sum_{i=1}^n \frac{w(x_i, x'_i)}{\sum_{j=1}^n w(x'_i, x_j) + \sum_{j=1, j \neq i}^n w(x'_i, x'_j)}, \quad (54)$$

where $w(x, x') = \exp\left(\frac{f(x)^\top f(x')}{\tau \|f(x)\|_2 \|f(x')\|_2}\right)$ and we pick the temperature parameter as $\tau = 0.5$. Training hyperparameters are presented in Table 3.

Downstream evaluation. The downstream evaluation is linear classification accuracy of the learned representation f to predict the class for an image. The linear classifier is trained for 1000 epochs using Adam; hyperparameter details are presented in Table 3.

For the plots in Figure 3 we evaluate every the contrastive loss and downstream accuracy every 5 epochs of training and stop when the average test contrastive loss (window size of 5) is minimized.

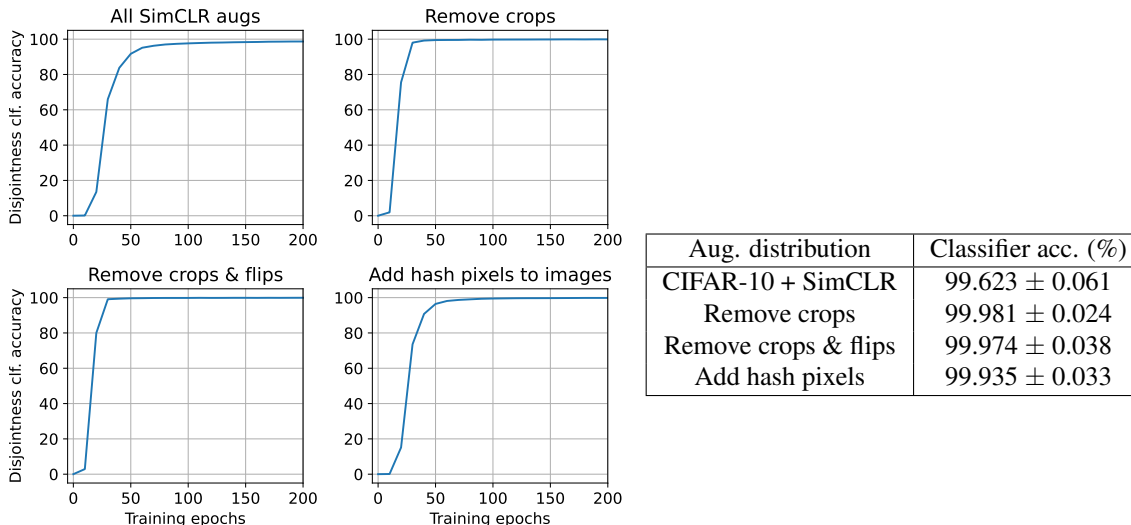


Figure 7. Demonstration of augmentation disjointness for CIFAR-10 with SimCLR augmentations. As described in Section 5.2 we train classifiers to distinguish between 5000 same-class examples, for each class. These classifiers reach $\approx 100\%$ accuracy (averaged over 10 classes) in the 5000-way classification task, in all 4 settings from Figure 3 (standard deviations shown over 10 random epoch picked close to end of training). This is evidence that these distributions are close to the disjoint regime, despite contrastive learning leading to good downstream accuracy.

D.2.1. HASH EXPERIMENT.

To enforce the disjoint augmentation regime, we select a set of 16 pixels, and modify an augmentation by replacing those 16 pixels (8 bits each) with an 128-bit MD5 hash of the image that generated the augmentation. This way the original image hash (and thus its identity) can be recovered from any of its augmentations. The result of training on this small variation of the standard pipeline is presented in Figure 3 (bottom right). Some examples of this augmentation is shown in Figure 6 (bottom right).

D.2.2. LABEL-ORTHOGONAL TRAINING.

We largely follow the same procedure as standard training, but modify the representation $f(x)$ for an augmentation before passing it to the contrastive loss. In particular for an augmentation and label pair (x, y) , compute the representation $f(x)$ as usual. Before passing it into the contrastive loss, convert apply the transformation $f'(x) = f(x) - \mu_y$, where μ_y is the mean representation for augmentations from class y . μ_y is computed at every step, using augmentations from a memory bank of 10240 pairs of (x, y) collected over training. Then $f'(x)$ is passed into the SimCLR loss in Equation (54) instead of $f(x)$ and everything else remains the same. The subtraction of the mean μ_y from the representation make the representation orthogonal to the labels, thus declining its ability to linearly classify images. The result of training with this procedure is presented in Figure 3 (top left). Note that the implicit assumption in the calculation of the contrastive loss is that different classes do not share any augmentations, i.e. the labels are almost invariant to standard augmentations.

D.3. Experiments on Text Domain

Experimental Setup. We evaluate on the AG News classification dataset (Zhang et al., 2015). This dataset contains 4 classes (“World”, “Sports”, “Business”, “Sci/Tech”) and each class contains news articles from that topic. We use the tokenizer from torchtext library. If a token sequence is of length more than 60, we then trim it to its first 60 tokens, leading to a vocabulary size of 11970.

We perform contrastive learning similar to SimCLR. We train the model in epochs and in each epoch we sample pairs of augmentation for 50,000 randomly chosen pieces of text in the training dataset. We then perform a single pass over this

Table 4. Hyperparameter values for experiments on AG News. Unless specified, the same hyperparameter value is used for both contrastive learning and the downstream classification task.

Hyperparameters	Values
Max epoch	100
Learning rate for contrastive learning	0.01 for BoW, 0.001 otherwise
Learning rate for downstream linear classification	0.01
Patience	10
Batch size	128
Representation dimension	768
Gradient clipping norm	2.5

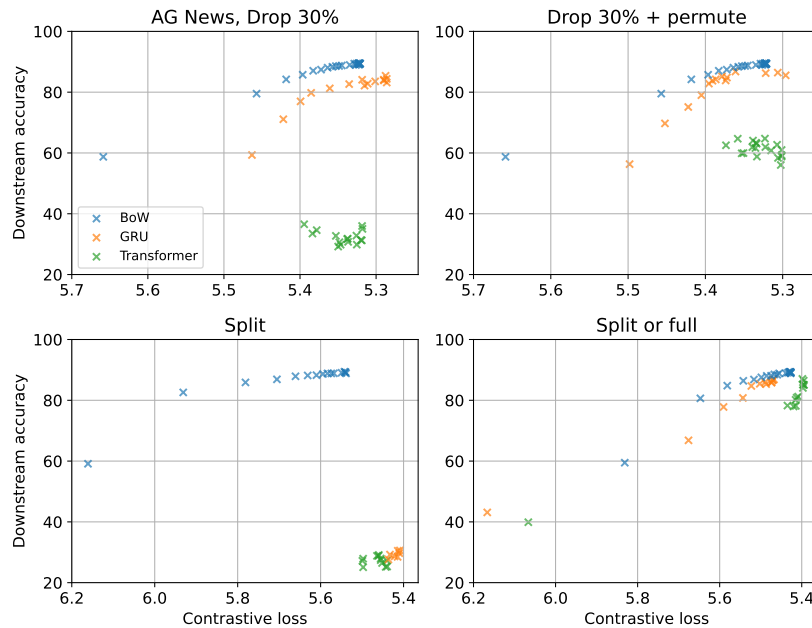


Figure 8. Contrastive loss \rightarrow accuracy transfer plots for AG News with bag-of-words (BoW), GRU and Transformer architectures with representation dimensionality $d = 128$. These plots use the average representation of augmentations $f_{\mathcal{A}}$ for downstream evaluation rather than the representation f directly. Augmentations in each case are as follows: **TL**: Drop random 30% of tokens. **TR**: Drop random 30% of tokens and randomly permute the rest. **BL**: Either the first half or second half of the input. **BR**: Either the first half, second half or the full input. The plots here are almost identical to the plots from Figure 4, suggesting that the distribution shift from augmentations to unaugmented inputs from contrastive learning to downstream evaluation does not play a big role.

dataset and minimize the SimCLR contrastive learning loss from Equation (54), with temperature $\tau = 1$. The downstream evaluation task is to simply predict the class given the text.

At the start of contrastive learning, we create a held-out validation set of pairs of augmentation sampled for 10,000 randomly chosen examples from the original validation set. At the end of each epoch of contrastive learning, we evaluate the model on this held-out validation set by computing the SimCLR loss. We also train a linear classifier on top of fixed model representations, to evaluate the model on the downstream classification task. During the downstream training, we evaluate the model at the end of epoch on the validation set and report the linear classifier with the best validation loss. We stop training if the best validation loss does not improve for κ consecutive epochs where κ is the patience hyperparameter, or if we hit a maximum number of epochs. Hyperparameter values are listed in Table 4.

Model Details. We evaluate three models on the AG News task. All models encode a given text to a d -dimensional representation. The first model is a bag of word (BoW) that trains a word embedding matrix and simply returns the average

word embedding of tokens in the text. The second model is Gated Recurrent Unit (GRU), which is a recurrent neural network (Chung et al., 2014). The GRU is uni-directional, uses a 300 dimensional input word embedding, dropout of 0.3, hidden dimension of 768, has 4 layers, and linearly maps the hidden state representation of the final token from the layer to d -dimensions. The final model is a Transformer (Vaswani et al., 2017), which is the base model for many state-of-the-art neural networks in NLP. The Transformer is uni-directional, hidden dimension of 128, has 4 layers and 4 attention heads, and linearly maps the hidden state representation of the final token from the layer to d -dimensions.

D.3.1. ROBUST EVALUATION.

Standard practice is to train a representation f on augmentations x , and use the same function to compute representations for unaugmented inputs \bar{x} . This is the strategy we employ for the plots in Figure 4. However, as discussed in Section 2, this causes an obvious distribution shift, since the representations have been trained to output something meaningful for unaugmented inputs. This could be a potential reason for the brittle transfer performance of GRU and Transformer. However we verify that this distribution shift is not the reason, by instead evaluating downstream performance using the augmentation-averaged representation $f_{\mathcal{A}}$, as defined in Equation (4). These robust evaluation transfer plots are presented in Figure 8, which look almost identical to those in Figure 4.

D.3.2. VISUALIZING 2-DIMENSIONAL REPRESENTATIONS.

We train contrastive learning models with output dimensionality $d = 2$ and visualize the contrastive loss \rightarrow accuracy in Figure 9. Firstly we note that the trends are not exactly the same as in Figure 4 that plot the same for $d = 128$. Most interestingly, for the split augmentation, GRU does not perform well on downstream accuracy for $d = 128$, but it does almost as well as BoW at $d = 2$. This kind of non-monotonic behavior w.r.t. representation dimensionality d is also unexplained by existing theory.

Next we visualize the learned representations for augmentations from different classes (normalized to unit norm) in Figure 10, for the drop augmentation. We sample 100 inputs per class and 4 augmentations per input, and encode them with the trained BoW, gru and Transformer models. For clear visualization, we plot the 4 augmentations per image with the same color, **with each of them plotted at different radii** (1.0, 1.133, 1.267, 1.4). We observe that the BoW representations look roughly linearly separable since different classes tend to roughly occupy different quadrants of the circle, corroborating its good downstream performance from Figure 9. It does so by roughly bringing augmentations of the same input (points with the same color) closer to each, although the representations not perfectly augmentation invariant. The GRU representations in every class, on the other hand, are spread out and brings augmentations very closer to each other than BoW representations, reminiscent of the uniformity and alignment properties from Wang & Isola (2020). However these representations are not linearly separable. The Transformer representations are intriguing since they are not uniformly spread out, but almost perfectly augmentation invariant. Furthermore the representation distributions for different classes are identical to each other, justifying its bad downstream performance from Figure 9. This phenomenon aligns with our lower bound Lemmas 4.1 and 4.2, whose proofs reveal how such spurious representations can be constructed.

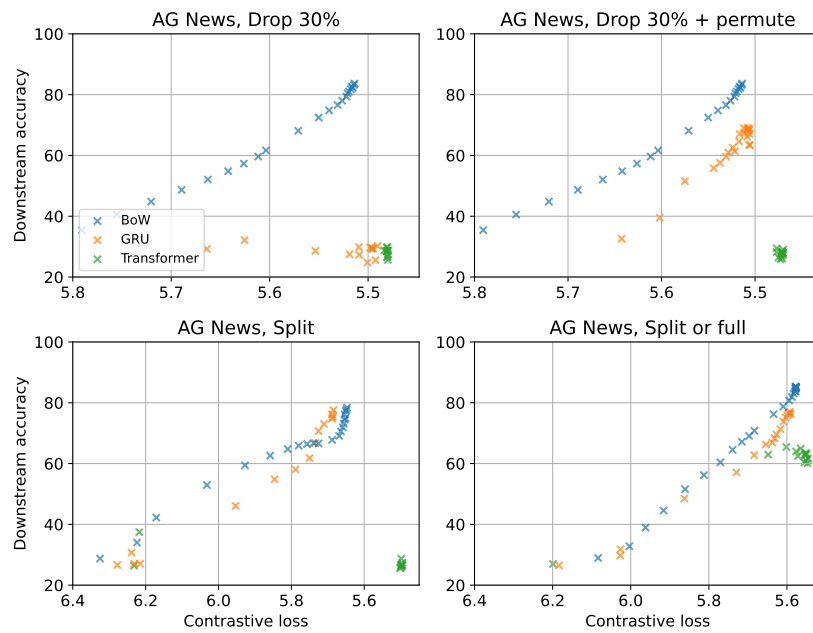


Figure 9. Contrastive loss \rightarrow accuracy transfer plots for AG News with bag-of-words (BoW), GRU and Transformer architectures with representation dimensionality $d = 2$. Augmentations in each case are as follows: **TL**: Drop random 30% of tokens. **TR**: Drop random 30% of tokens and randomly permute the rest. **BL**: Either the first half or second half of the input. **BR**: Either the first half, second half or the full input. In all cases BoW representation does quite well downstream ($\sim 80\%$), but either Transformer or both GRU and Transformer demonstrate brittleness of transfer for different augmentations.

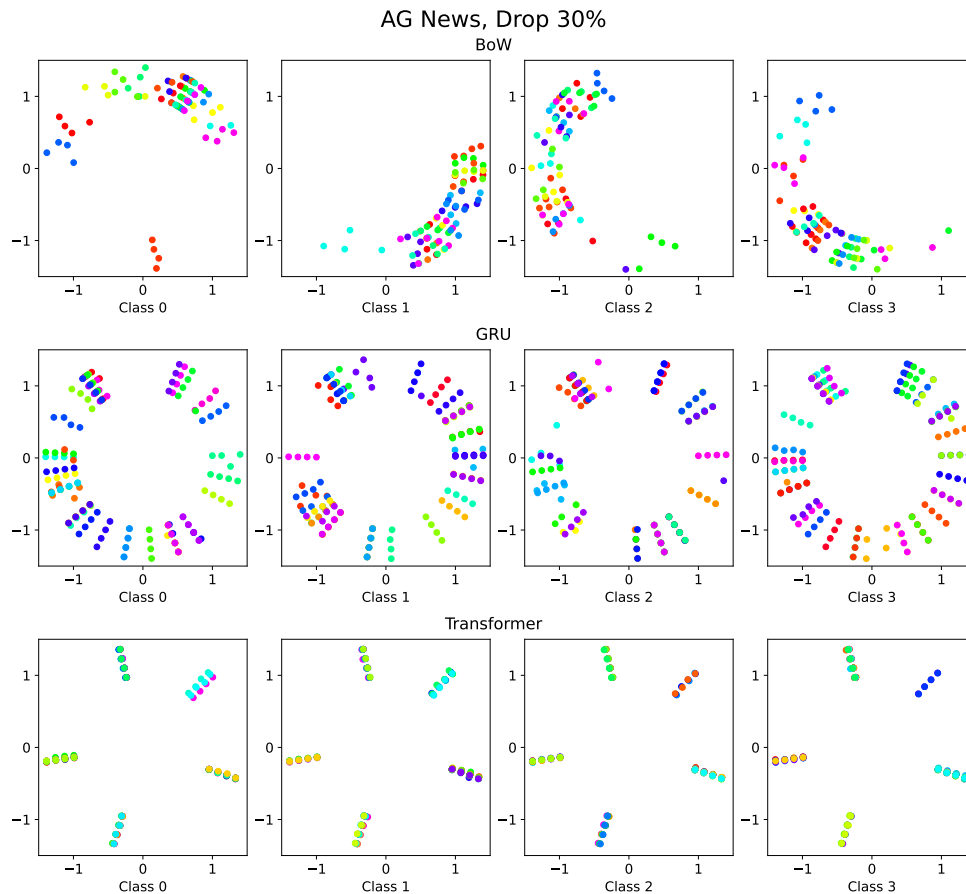


Figure 10. We plot representations of augmentations from different classes, for BoW, GRU and Transformer respectively. The 30% drop augmentation is used for these plots. While all representations are supposed to be normalized to unit norm, for clear visualization, we plot the 4 augmentations per image with the same color, **with each of them plotted at different radii** (1.0, 1.133, 1.267, 1.4). We observe that GRU and Transformer are quite augmentation invariant, but are not linearly separable. See Appendix D.3.2 for more discussion about this.

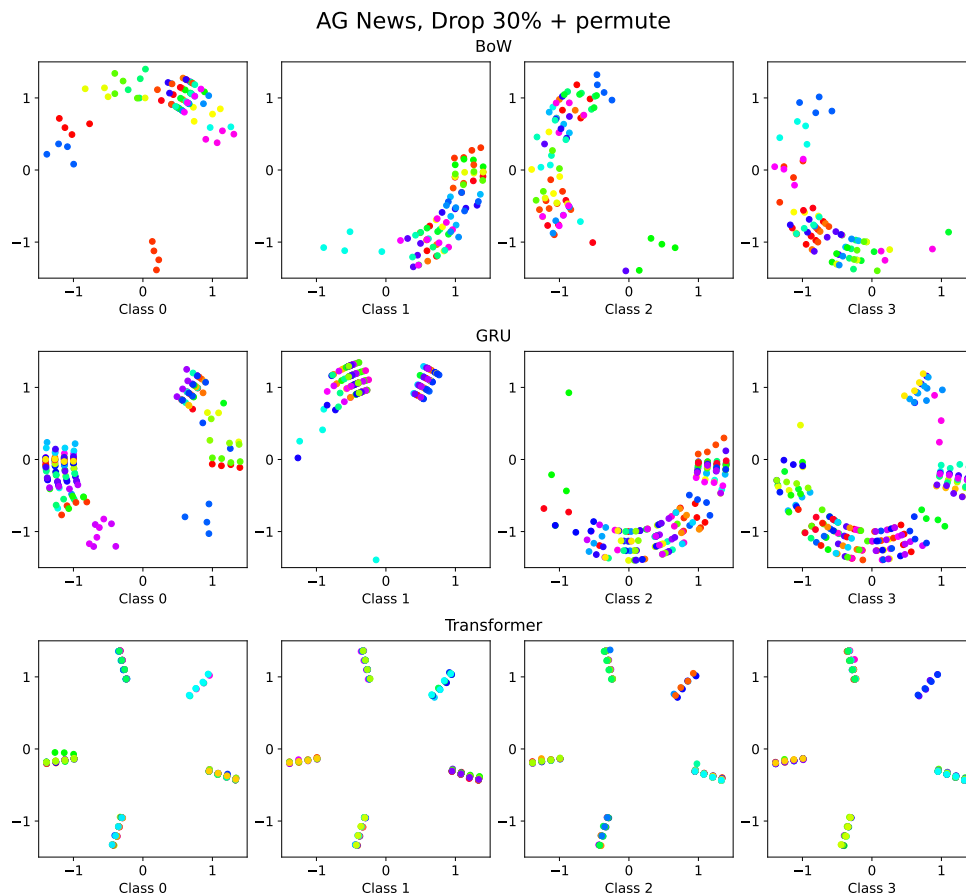


Figure 11. We plot representations of augmentations from different classes, for BoW, GRU and Transformer respectively. The 30% drop + permute augmentation is used for these plots. While all representations are supposed to be normalized to unit norm, for clear visualization, we plot the 4 augmentations per image with the same color, **with each of them plotted at different radii** (1.0, 1.133, 1.267, 1.4). We observe that BoW representations are roughly linearly classifiable, GRU representations are somewhat classifiable while Transformer are quite augmentation invariant, but not linearly separable.