

InfoNCE Loss Provably Learns Cluster-Preserving Representations

Advait Parulekar

University of Texas at Austin

ADVAITP@UTEXAS.EDU

Liam Collins

University of Texas at Austin

LIAMC@UTEXAS.EDU

Karthikeyan Shanmugam

Google Research India

KARTHIKEYANVS@GOOGLE.COM

Aryan Mokhtari

University of Texas at Austin

MOKHTARI@AUSTIN.UTEXAS.EDU

Sanjay Shakkottai

University of Texas at Austin

SANJAY.SHAKKOTTAI@UTEXAS.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

The goal of contrasting learning is to learn a representation that preserves underlying clusters by keeping samples with similar content, e.g. the “dogness” of a dog, close to each other in the space generated by the representation. A common and successful approach for tackling this unsupervised learning problem is minimizing the InfoNCE loss associated with the training samples, where each sample is associated with their augmentations (positive samples such as rotation, crop) and a batch of negative samples (unrelated samples). To the best of our knowledge, it was unanswered if the representation learned by minimizing the InfoNCE loss preserves the underlying data clusters, as it only promotes learning a representation that is faithful to augmentations, i.e., an image and its augmentations have the same representation. Our main result is to show that the representation learned by InfoNCE with a finite number of negative samples is also consistent with respect to *clusters* in the data, under the condition that the augmentation sets within clusters may be non-overlapping but are close and intertwined, relative to the complexity of the learning function class.

Keywords: Contrastive learning, Representation learning, Self-supervised learning

1. Introduction

Representations pretrained on partially or completely unlabeled data are becoming ubiquitous in machine learning applications (Peters et al., 2018; Radford et al., 2021), in large part due to the availability of large unlabeled datasets and significant computing power offline, and the effectiveness of self-supervised representation learning algorithms, especially contrastive learning (CL). CL aims to learn representations that treat natural images similarly to their augmentations, while maximizing the average distance between random pairs of images. In recent years CL has demonstrated numerous successes in pretraining representations with unlabeled data that learn meaningful relationships between data points that generalize well to downstream tasks in computer vision (Hjelm et al., 2018; Oord et al., 2018; Bachman et al., 2019; Caron et al., 2020; Chen et al., 2020a,b; He et al., 2020; Henaff, 2020; Li et al., 2020; Misra and Maaten, 2020; Tian et al., 2020a,b) and natural language processing (Brown et al., 2020; Gao et al., 2021; Su et al., 2021; Radford et al., 2019).

Despite its empirical success, it is not well-understood how CL learns meaningful relationships between data points. Since data are unlabeled, the only immediate structure in datasets leveraged

by CL are the sets of images and their augmentations. Without further assumptions, this structure is insufficient to learn relationships between images across augmentation sets. To circumvent this issue there are two approaches. The first is to assume that augmentation sets of semantically similar natural images overlap, i.e. for two images of cats, some of the augmentations of each image are equivalent (Arora et al., 2019; HaoChen et al., 2021, 2022; Shen et al., 2022; Wang et al., 2022). However, this assumption is unlikely to hold in practice, as pointed out by recent work (Saunshi et al., 2022). The second approach is to consider inductive biases of the representation function class and/or optimization algorithm and use these to argue that only certain types of representations (that capture semantic relationships) can be learned.

Prior works have initiated the study of how inductive biases of the representation class can lead to meaningful representations in CL (Saunshi et al., 2022; HaoChen and Ma, 2022), but their analysis is for the spectral contrastive loss, which is not used in practice. Instead, the vast majority of CL approaches, including the widely popular SimCLR framework (Chen et al., 2020a), optimize a loss function based on InfoNCE (Gutmann and Hyvärinen, 2010; Oord et al., 2018). A variety of works have studied properties of the InfoNCE loss, but due to its unwieldy log-sum structure have made restrictive assumptions, such as having infinite (Wang and Isola, 2020; Robinson et al., 2020; Von Kügelgen et al., 2021) or only a single (Tosh et al., 2021; Huang et al., 2021) negative sample(s).

Main Contributions. We are given a collection of clusters of natural images, with each image associated with augmentations (positive samples such as ‘rotation’) and a finite set of negative samples (unrelated images). Using the InfoNCE loss, our goal is to learn a d -dimensional representation $g \in \mathcal{G}$, where $g = (f_1, f_2, \dots, f_d)$ and $\{f_i\}$ are binary functions mapping images to $\{-1, 1\}$ (thus g maps images on the hypercube $\mathcal{H}_d = \{-1, 1\}^d$). Our setting is one where the function class has bounded expressivity with respect to the augmentation sets, meaning that the augmentation sets within clusters are intertwined, and hard to separate from the rest of the cluster using functions in \mathcal{F} .

(Realizable Setting) Suppose there exists a representation $g^* \in \mathcal{G}$ that is: (a) cluster preserving, and (b) different clusters of images are uniformly mapped over distinct vertices on the hypercube (qualitatively, class-balance in the image dataset). We show that with *any finite number of negative samples*, the representation learned by the InfoNCE loss is cluster-preserving and uniform. Furthermore, this learned representation when composed with a two-layer ReLU head, achieves zero downstream error on any cluster-preserving binary classification task. Our proof hinges on a novel Markov Chain construction showing that the InfoNCE loss of any non-uniform representation can be improved by “blurring” the representation through the Markov Chain transitions. Conversely, we show that solutions to the InfoNCE loss optimized over an arbitrarily powerful representation class \mathcal{G}_* cannot have meaningful downstream performance guarantees on such tasks.

(Agnostic Setting) In the agnostic (non-realizable) case, through sensitivity analysis, we show that for any close-to-uniform and non-cluster-preserving representation, there exists a representation that preserves one additional cluster and thus improves the InfoNCE loss. Our proof uses a novel partitioning of the image space that is of independent interest for future analysis of the InfoNCE loss.

1.1. Related Work

Several works have aimed to explain the success of contrastive learning in recent years. Wang and Liu (2021) and Wang and Isola (2020) showed empirically that CL encourages aligned and uniform representations, and improving alignment and uniformity improves downstream performance. The work in Chen et al. (2021) generalizes the InfoNCE loss to a larger family of losses with alignment

and uniformity terms weighted according to a hyperparameter. Early theoretical studies attributed the success of CL to its proclivity to maximize the mutual information between augmentations of the same image (Bachman et al., 2019), but later work cast doubt on this viewpoint by showing that optimizing a tighter bound on the mutual information leads to worse performance (McAllester and Stratos, 2020; Tschannen et al., 2019). The work in (Wang and Isola, 2020) further showed that solutions to the InfoNCE loss are aligned and uniform in the limit of infinite negative samples per batch.

A variety of works have studied CL’s ability to recover meaningful clusters or latent variables in the data (Arora et al., 2019; Tosh et al., 2021; Zimmermann et al., 2021; Ash et al., 2021; Nozawa and Sato, 2021; HaoChen et al., 2021; Shen et al., 2022; HaoChen et al., 2022; HaoChen and Ma, 2022; Wang et al., 2022; Awasthi et al., 2022; Bao et al., 2022). However, the majority of these works consider arbitrary function classes, which requires strong assumptions on the connectedness of augmentation sets within each cluster, such as assuming positive pairs are conditionally independent given their cluster identity, in order to give downstream guarantees (Saunshi et al., 2022). The work by HaoChen and Ma (2022) is the most related work to ours, as they study function classes that induce a similar bias towards preserving clusters as ours without any assumption on the connectedness of augmentation sets. However, their study is focused on minimizing a spectral contrastive loss which serves as a surrogate for the more practically used InfoNCE loss. While studying spectral contrastive loss is enlightening and provides some intuition, it cannot be extended to the InfoNCE loss because of two major reasons: First, the loss function fails to highlight the role of finite batches of negative samples, which is a well-studied and key component of the InfoNCE loss (Awasthi et al., 2022; Bao et al., 2022; Ash et al., 2021; Nozawa and Sato, 2021). Second, their analysis does not translate to our setting because the key difficulty in our proof is to show that negative samples promote uniformity; this aspect directly follows with the spectral loss due to the covariance regularizer.

Additional theoretical works have studied the feature learning process of CL with (stochastic) gradient descent on linear (Tian, 2022a; Ji et al., 2021) and two-layer ReLU neural networks (Wen and Li, 2021; Tian, 2022b), properties augmentations must satisfy in order for CL to be successful (Tian et al., 2020b), the role of the projection head in CL (Wen and Li, 2022; Gupta et al., 2022), and the behavior of contrastive losses in (semi-)supervised settings (Khosla et al., 2020; Zheng et al., 2021; Chen et al., 2022). Several other works analyze non-contrastive self-supervised learning methods (Wei et al., 2020; Balestrieri and LeCun, 2022; Garrido et al., 2022; Lee et al., 2021).

2. Problem Formulation

Our learning task consists of (i) a pretraining phase – wherein we are not provided supervised labels but rather only *associations* between images and (ii) a supervised learning phase in which we are provided (a few) labeled data points, labeled according to some specific downstream task. During the pre-training phase, we do not know what the downstream task is. However, we are provided *augmentations* of the raw data points that the learner knows should be classified the same way as the raw data for *any* downstream task. In a sense, the augmentations can be seen as modifying the data in a way that leaves the information contained in the data invariant with respect to the downstream tasks. Ideally, we aim to learn a representation that is invariant to such augmentations so that downstream learning can be statistically efficient. For interpretability, we will work in the setting of “images”.

Images and augmentations. The images consist of features that are either important for classification or which function only as irrelevant details. Inspired by (Von Kügelgen et al., 2021),

we consider an image generation model that consists of (i) content variables denoted by \mathbf{c} which capture innate qualities of the images (e.g., the ‘catness’ of a cat), and (ii) style variables denoted by \mathbf{s} which capture the appearance of the image (e.g., ‘rotation’ and ‘crop’ for creating augmentations to an image; ‘long tail’ and ‘furry’ for different natural images of dogs). More precisely, each image x is generated according to $x = \mathcal{I}(\mathbf{c}, \mathbf{s})$, where $\mathcal{I}(\cdot, \cdot)$ is a mapping from the space of content and style variables to the space of images. We assume that the natural images are generated such that their content variables \mathbf{c} belong to the set C and their style variables \mathbf{s} belong to the set S_\circ .

We further consider that there is a set of augmentations Λ , which is a set of functions mapping natural images to augmented images. An augmented image of an image x is denoted $\mathcal{A}(x)$, where $\mathcal{A} \in \Lambda$. We assume that the augmented image preserves the content of the original image, while its style may differ from the original image. More precisely, if the original image is given by $x = \mathcal{I}(\mathbf{c}, \mathbf{s})$, then its augmented image $\mathcal{A}(x)$ satisfies the following property: $\mathcal{A}(x) = \mathcal{A}(\mathcal{I}(\mathbf{c}, \mathbf{s})) = \mathcal{I}(\mathbf{c}, \mathbf{s}^+)$ for some $\mathbf{s}^+ \in S$, where the set S contains S_\circ . So the augmented images have possibly different style variables *but the same content variables* as the natural images. Further, the set of augmented images of the image $x = \mathcal{I}(\mathbf{c}, \mathbf{s})$ is called its *augmentation set* and is defined as $A(x) := A(\mathcal{I}(\mathbf{c}, \mathbf{s})) := \{\mathcal{A}(\mathcal{I}(\mathbf{c}, \mathbf{s})) \mid \mathcal{A} \in \Lambda\}$, with all images having equal-sized augmentation sets for simplicity. We typically refer to an image $\mathcal{I}(\mathbf{c}, \mathbf{s})$ as x and its augmentation $\mathcal{I}(\mathbf{c}, \mathbf{s}^+) \sim A(x)$ as x^+ , where, for all sets of images B , $\sim B$ denotes a random sample drawn uniformly from the set B . We let D denote the set of all images and their augmentations and $D_\circ \subset D$ denote the set of all natural images.

Representations and heads. We consider a function class \mathcal{F} of binary functions, $f(x') \in \{-1, +1\}$, where x' is either an image x or its augmentation x^+ . This is a function class with bounded expressivity (e.g., a class of functions that can be expressed as the thresholded output of a neuron from a neural network with bounded width and depth). The function class \mathcal{F} and the augmentations Λ define a set of *clean* functions $\mathcal{F}_c \subseteq \mathcal{F}$ that separate the data in a way that respects the augmentations, $\mathcal{F}_c = \{f : f(x) = f(\mathcal{A}(x)) \forall x \in D_\circ, \mathcal{A} \in \Lambda\}$. In other words, the binary function f is clean if it does not separate any image and its augmentations from each other.

We search over d -dimensional representations, denoted by \mathcal{G} , such that each coordinate of the representation is an element of \mathcal{F} , i.e., $g = (f_1, f_2, \dots, f_d)$. Thus a representation $g \in \mathcal{G} := \mathcal{F}^d$ is simply a concatenation of d binary classifiers, mapping an image x to the vertex of the Rademacher hypercube¹ $\mathcal{H}_d = \{-1, 1\}^d$. Note that each $g \in \mathcal{G}$ denotes only the representation (e.g., the body of a neural network). For downstream tasks, a full classifier is formed by composing g with a *head* $\omega \in \mathcal{J}$ for some class \mathcal{J} of heads (e.g., the final classification layer of a neural network).

Clusters. We consider an equivalence class on content variables \sim such that

$$\mathbf{c} \sim \mathbf{c}' \iff f(\mathcal{I}(\mathbf{c}, \mathbf{s})) = f(\mathcal{I}(\mathbf{c}', \mathbf{s}')) \forall f \in \mathcal{F}_c.$$

We refer to the images in the equivalence classes as clusters $\Gamma_{\mathbf{c}}$, so

$$\Gamma_{\mathbf{c}} = \{x : \exists \mathbf{c}' \sim \mathbf{c}, \mathbf{s} \text{ such that } x = \mathcal{I}(\mathbf{c}', \mathbf{s})\}.$$

As an example, suppose that the content \mathbf{c} captures the ‘dogness’ of an image. Then, different images of dogs would have the same content, but have different style variables (e.g., furry, skinny, long ears).

1. Representations in CL often map to the unit hypersphere (Wang and Isola, 2020). Here, we consider a discretized version of this output space for two reasons: (1) it allows us to construct a naturally restricted representation function class by extending natural properties of binary classifiers, and (2) it provides a tractable setting for us to show the first results that InfoNCE prefers cluster-preserving and uniform representations with finite samples, as it is still an open problem to determine uniform arrangements of finite points on the unit hypersphere (Thomson, 1904).

Recall that the augmentations of an image also share the same content, but the style might be chosen from a different set (e.g., rotation, cropping, blur).

Our partition of the space of images into clusters highlights the interplay between the richness of the function class \mathcal{F} and the diversity in the augmentations Λ . A more diverse set of augmentations generally results in a smaller set \mathcal{F}_c , and hence generally increases the size of the clusters. Meanwhile, a richer class \mathcal{F} generally results in a larger \mathcal{F}_c , and smaller clusters.

In this work we show conditions under which for a fixed choice of \mathcal{F} and Λ , contrastive pre-training learns the equivalence class. This is useful for solving downstream classification tasks, as described below.

Goal of pretraining. Ultimately, we aim to find a representation that allows for easily solving tasks from a set of possible downstream binary classification tasks $h \in \mathcal{T}$, where each task h maps an image to a binary label $\{-1, 1\}$. These tasks are assumed to be faithful to the clusters, meaning that for any pair of images x, x' belonging to the same cluster, $h(x) = h(x')$.

Note that during pretraining, the learner does not have any knowledge about which task will be assigned among the solvable ones. After pretraining, the learner fixes the representation but can learn a task-specific head when it encounters a downstream task. We define the error a representation g on the downstream task $h \in \mathcal{T}$ with respect to the class $\mathcal{J} \subseteq \{\omega : \mathbb{R}^d \rightarrow \mathbb{R}\}$ of allowed heads as

$$\mathcal{L}_{h, \mathcal{J}}(g) := \inf_{\omega \in \mathcal{J}} \mathbb{P}_{x \sim \mathcal{D}}[\omega \circ g(x) \neq h(x)]. \quad (1)$$

The error of g on a family of downstream tasks $\mathcal{T} \subseteq \mathcal{F}_c$ is the worst case error among tasks in \mathcal{T} :

$$\mathcal{L}_{\mathcal{T}, \mathcal{J}}(g) := \sup_{h \in \mathcal{T}} \mathcal{L}_{h, \mathcal{J}}(g). \quad (2)$$

To summarize, for a task that is realizable *with supervision* using function class \mathcal{F} , we would like to learn a representation entirely from unlabelled data such that the task on the embedded images is still realizable for \mathcal{J} . The overall motivation is that learning $\omega \in \mathcal{J}$ can generally require fewer *labeled* samples than learning the joint model $\omega \circ g$.

2.1. InfoNCE loss

We denote $\mathbb{E}_{x, x^+} := \mathbb{E}_{x \sim D_0, x^+ \sim A(x)}$ and $\mathbb{E}_{x, x^+, \{x_i^-\}_\ell} := \mathbb{E}_{x \sim D_0, x^+ \sim A(x), \{x_i^-\}_\ell \sim D_0^\ell}$ for simplicity. The InfoNCE loss we consider is given by²

$$\mathcal{L}(g) = - \underbrace{\beta \mathbb{E}_{x, x^+} [g(x)^\top g(x^+)]}_{\text{alignment}} + \underbrace{\mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right]}_{\text{uniformity}} \quad (3)$$

Following Wang and Isola (2020), we refer to the first term as the *alignment* term, or the *positive* term, and we refer to the second term as the *uniformity* term or the *negative* term. By minimizing the first term, we are maximizing the alignment between the representation of an image and its augmentation, and by minimizing the second term we are enforcing the representation of different images to be as different as possible.

2. For ease of exposition we consider the case wherein negative samples are drawn from the set of natural images, as in (Wen and Li, 2021). Although this may not hold in practice, it greatly simplifies the presentation of our results.

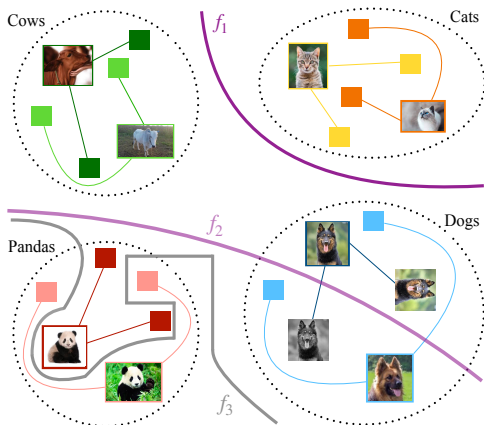


Figure 1: We illustrate a setting with four clusters, and two augmentation sets within each cluster indicated by linked rectangles in distinct colors. The clean function f_1 does not split any augmentation or cluster (meaning, it maps all images from the same augmentation set alike, and likewise for clusters). f_2 splits the cluster of dogs, and in accordance with Assumption 2, also splits augmentation sets within that cluster (and is therefore non-clean). Finally, f_3 violates Assumption 2, because it does not split any augmentation of pandas, yet it splits the pandas cluster. Images are from the Animals V2 dataset: DeepNets (2022).

The above formulation suggests that the representation learned by minimizing the above loss forces images and their augmentations to have a similar representation. What we show in the following sections is a stronger result which guarantees by minimizing the InfoNCE loss, all images that belong to the same cluster (share the same content) will have a similar representation.

3. Bounded Function Class

The goal of contrastive learning is to learn a representation from unlabeled samples that is useful for downstream tasks. Recall that the representations we consider map images to vertices on the Rademacher hypercube \mathcal{H}_d . A “good” representation should map images from the same cluster to the same vertex, and images from distinct clusters to distinct vertices.

Intuitively, this seems possible if images having the same content (i.e., from the same cluster) along with their set of augmentations are “close and intertwined” (see Figure 1), such that any function $f \in \mathcal{F}$ cannot split the cluster without also splitting an image from its augmentation. Note that we do *not* need connected clusters with overlapping augmentations (meaning two images have the same augmentation, which is an unrealistic assumption); merely that the cluster has a complex geometry relative to the function class.

3.1. Complexity of \mathcal{F} Relative to Augmentations

We formalize the notion of bounded expressivity of \mathcal{F} relative to the geometry of clusters. We use this assumption to show in Section 4 that solutions to the InfoNCE loss optimized over \mathcal{G} satisfy useful uniformity and alignment properties that lead to downstream performance guarantees on tasks that adhere to the clusters. Formally, the function class \mathcal{F} and the augmentations Λ define a set of *clean* functions $\mathcal{F}_c \subseteq \mathcal{F}$ that separate the data in a way that respects the augmentations.

Definition 1 (Clean Function) $f \in \mathcal{F}_c$ is clean $\iff f(x) = f(\mathcal{A}(x)) \forall x \in D_o, \mathcal{A} \in \Lambda$.

In other words, the binary function f is clean if it does not separate any image and its augmentations from each other. Our main assumption is that if a classifier in \mathcal{F} splits a cluster, then it is not clean.

Assumption 2 (Intertwined Augmentations) For all $f \in \mathcal{F}$, if $f(x) \neq f(x')$ for some $x, x' \in \Gamma_c$, then $f(x'') \neq f(\mathcal{A}(x''))$ for some $x'' \in \Gamma_c \cap D_o$, where $\mathcal{A}(x'') \in A(x'')$.

Note that if a classifier does not split any cluster, then it must be clean, since augmentation sets are contained within clusters. Thus, Assumption 2 implies that $f \in \mathcal{F}_c$ if and only if f labels all images with the same content (belonging to the same cluster) alike, in other words it is *cluster-preserving*. This assumption holds if the augmentation sets within clusters are close and intertwined (they cannot be easily split from the rest of the cluster), relative to the complexity of \mathcal{F} . Importantly, the augmentation sets need not overlap, meaning a single image need not be an augmentation to multiple natural images, consistent with practice (Saunshi et al., 2022). As prior works have pointed out (Saunshi et al., 2022; HaoChen and Ma, 2022), Assumption 2 or variants on the bounded complexity of the function class are necessary for the success of CL in the realistic setting in which the augmentation sets do not overlap.

However, while some condition like Assumption 2 is necessary, it is not clear if this suffices to show that CL learns useful representations. Consider the example in Figure 1. It may be the case, for instance, that CL on \mathcal{F} does not learn the cluster-preserving classifiers, as in addition to trying to maximize the similarity between images and their augmentations, CL also tries to minimize the similarity between negative pairs of images. Thus, it may choose a non-cluster-preserving classifier such as f_2 in an effort to minimize similarity of negative pairs. This would lead to poor downstream generalization on tasks involving classifying dogs, since f_2 separates images of dogs. It thus becomes critical to quantify the extent to which non-cluster-preserving classifiers must intersect augmentation sets such that CL will not learn them, as we do in Section 5. Before this, we must show that even if CL learns a representation consisting of cluster-preserving classifiers, this representation generalizes well, which may not happen if it maps two or more clusters to the same vertex. For instance, if CL simply learned d copies of the cluster-preserving classifier f_1 in Figure 1, this representation would not be able to distinguish cows from pandas from dogs on downstream tasks. We thus desire representations to be *both* cluster-preserving and uniform such that their mapping is a bijection from clusters to vertices. Next, we show that when a cluster-preserving and uniform representation is realizable, CL with the InfoNCE loss learns it, even with finite negative samples per batch.

4. Results for the Realizable Setting

Our first result shows that when the dataset D and representation class \mathcal{G} allow for mapping the data *uniformly* on the hypercube in a *cluster-preserving* manner, then the representation learned by minimizing the InfoNCE loss over \mathcal{G} results in such a mapping. We first formally define the terms *uniform* and *cluster-preserving* below.

Definition 3 (Cluster-Preserving) A cluster-preserving representation $g \in \mathcal{G}$ is one that for all $c \in C$ and all $x, x' \in \Gamma_c$, $g(x) = g(x')$.

Definition 4 (Uniform) A uniform representation $g \in \mathcal{G}$ satisfies $\mathbb{P}_{x \sim D_\circ}[g(x) = v] = 2^{-d}$ for all $v \in \mathcal{H}_d$.

Next, our results in this section assume a cluster-preserving and uniform representation exists in \mathcal{G} .

Assumption 5 (Realizability) There exists a $g \in \mathcal{G}$ that is both cluster-preserving and uniform.

In order for there to exist a representation that is both cluster-preserving and uniform, there must be an integral multiple of 2^d clusters in the dataset and they must be balanced. Before stating our main result, we must prove a key lemma that shows that among all “clean” representations, those

that minimize the InfoNCE loss are uniform. We define $\mathcal{G}_c \subseteq \mathcal{G}$ as the set of clean representations in \mathcal{G} consisting of d concatenated clean classifiers from \mathcal{F}_c .

Lemma 6 *If Assumptions 2 and 5 hold, $\beta > c \log d$ for an absolute constant c , $d > 3$, and $\ell \geq 1$, then $g^* \in \arg \min_{g \in \mathcal{G}_c} \mathcal{L}(g)$ if and only if g^* is uniform.*

Proof [Proof sketch] Since the optimization problem is over representations composed of clean functions, we know that for all $g \in \mathcal{G}_c$, the term $g(x)^\top g(x^+)$ in the InfoNCE loss is exactly equal to d . Hence, by regrouping the terms in (3), the optimization problem simplifies to:

$$\min_{g \in \mathcal{G}_c} \mathcal{L}(g) = \min_{g \in \mathcal{G}_c} \left\{ \hat{\mathcal{L}}(g) := \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-) - \beta d} \right) \right] \right\} \quad (4)$$

By Assumption 5, at least one uniform representation belongs to the set \mathcal{G}_c . We show that it minimizes the loss $\hat{\mathcal{L}}(g)$. To do so, we observe that we can think of minimizing $\hat{\mathcal{L}}(g)$ as an optimization with respect to distributions over the hypercube induced by g . To better understand this connection, consider the random variable $g(x)$ for $x \sim D_o$. Further, denote the corresponding induced distribution over $g(x)$ as \mathcal{D}_g , i.e., \mathcal{D}_g is a distribution over the vertices of the hypercube \mathcal{H}_d . Letting $y = g(x)$, the objective above can now be rewritten in terms of these distributions:

$$\min_{\{\mathcal{D}_g: g \in \mathcal{G}_c\}} \left\{ \tilde{\mathcal{L}}(\mathcal{D}_g) := \mathbb{E}_{y, \{y_i^-\}_\ell \sim \mathcal{D}_g} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta y^\top y_i^- - \beta d} \right) \right] \right\} \quad (5)$$

Suppose the the minimizing distribution was not uniform over the hypercube, i.e. for $\mathcal{D}_g^* \in \arg \min_{\{\mathcal{D}_g: g \in \mathcal{G}_c\}} \tilde{\mathcal{L}}(\mathcal{D}_g)$, $\mathcal{D}_g^* \neq \mathcal{U}$, where \mathcal{U} is the uniform distribution over the hypercube \mathcal{H}_d . For any sample $y, \{y_i^-\}_\ell \sim \mathcal{D}_g$, consider a random walk that starts from this sample and evolves over time. For this random walk, denote the variables at time t by $y^t, \{y_i^t\}_\ell$ where y^t (and similarly y_i^t for all i), with $y^0 = y$ (correspondingly $y_i^0 = y_i$). The random walk evolves from y^{t-1} to y^t by flipping a uniformly random bit of y^{t-1} with probability $\frac{1}{2}$, and with probability $\frac{1}{2}$, not changing anything; this construction is independent across all samples. We now observe that this construction induces an irreducible, aperiodic Markov chain with uniform stationary distribution over the hypercube.

With this construction, the critical step in our proof is a surprising ‘‘monotonicity’’ property over time: we show in Appendix A that each transition over time decreases the function value as long as \mathcal{D}_g is not uniform. Intuitively, ‘‘blurring’’ the distribution \mathcal{D}_g decreases the objective.

This result implies that g is a minimizer of the loss $\hat{\mathcal{L}}(g)$ if and only if g is a uniform representation. Consequently, we obtain that among all the representations in \mathcal{G}_c , the ones that are uniform minimize the loss in (5) and the statement of Lemma 6 follows. See Appendix A.1 for details. ■

Using Lemma 6, we show our main result that all minimizers of the InfoNCE loss are uniform and cluster-preserving. To the best of our knowledge, this is the first result characterizing the minimizers of the InfoNCE loss with a finite batch of negative samples. The proof is provided in Appendix A.2.

Theorem 7 *If Assumptions 2 and 5 hold, and we have $d > 3$, $\ell \geq 1$, and $\beta > c \log d$ for an absolute constant c , then a representation $g^* \in \mathcal{G}$ is a global minimizer of the loss $\mathcal{L}(g)$ optimized over \mathcal{G} if and only if it is uniform and cluster-preserving.*

4.1. Downstream Guarantees

Next we translate the aforementioned representation learning results for \mathcal{G} into downstream performance guarantees. We consider the class of heads consisting of single-layer ReLU neural networks with m neurons. Formally,

$$\mathcal{J}_{\text{ReLU}} := \{\omega_{a, \mathbf{W}, b} : \mathbb{R}^d \rightarrow \mathbb{R} \text{ s.t. } \omega_{a, \mathbf{W}, b}(g(x)) = a^\top \text{ReLU}(\mathbf{W}g(x) - b), a, b \in \mathbb{R}^m, \mathbf{W} \in \mathbb{R}^{m \times d}\},$$

where $\text{ReLU}(h) = \max(h, 0)$.

Theorem 8 *Suppose there are exactly 2^d clusters, and the representation $g^* \in \arg \min_{g \in \mathcal{G}} \mathcal{L}(g)$ under Assumptions 2 and 5, $\beta > c \log d$ for an absolute constant c and $d > 3$. Then for any set of cluster-preserving downstream tasks \mathcal{T} , $\mathcal{L}_{\mathcal{T}, \mathcal{J}_{\text{ReLU}}}(g^*) = 0$.*

Theorem 8 shows that any representation learned by minimizing the InfoNCE loss achieves zero downstream error on any task from \mathcal{F}_c with a sufficiently wide two-layer ReLU head.

Next, we show that controlling the expressivity of \mathcal{G} is necessary to achieve meaningful downstream performance guarantees. Suppose that instead of optimizing the InfoNCE loss over \mathcal{G} , we instead optimized it over a representation class $\mathcal{G}_* := \mathcal{F}_*^d$ where $\mathcal{F}_* := \{f : D \rightarrow \{-1, 1\}\}$ consists of all classifiers mapping from images to binary labels.

Theorem 9 *Let $\beta > c \log d$ for an absolute constant c and $d > 3$. There exists a dataset D that satisfies Assumptions 2 and 5 for \mathcal{G} , representation $g \in \arg \min_{g' \in \mathcal{G}_*} \mathcal{L}(g')$, and a downstream task $h \in \mathcal{F}_c$ such that $\mathcal{L}_{h, \mathcal{J}_*}(g) \geq 0.5$, where $\mathcal{J}_* = \{\omega : \mathcal{H}_d \rightarrow \{-1, 1\}\}$ is the set of all mappings from $\mathcal{H}_d \rightarrow \{-1, 1\}$.*

5. Results for the Agnostic Setting

In this section, we consider the setting in which there may not exist any cluster-preserving and uniform representation (that is, Assumption 5 is violated). We show that even in this setting, the InfoNCE loss prioritizes cluster-preserving representations. Specifically, we show that if an optimal solution of the InfoNCE loss on \mathcal{G} is close to uniform, then it must also be cluster-preserving. This result requires two new assumptions that we describe below.

First, the function class \mathcal{F} must be closed under operations that make classifiers cluster-preserving, in the sense that if $f \in \mathcal{F}$ and f does not preserve the cluster Γ_c , then the two perturbations of f that preserve Γ_c (by assigning ± 1 to all images within it) and do not change f otherwise are also in \mathcal{F} .

Assumption 10 (Expressivity of \mathcal{F}) *For any cluster Γ_c , if any $f \in \mathcal{F}$ is such that $f(x) \neq f(x')$ for some $x, x' \in \Gamma_c$, then $f' \in \mathcal{F}$ and $f'' \in \mathcal{F}$, where $f'(x) = f''(x) = f(x) \forall x \notin \Gamma_c$, and $f'(x) = 1, f''(x) = -1 \forall x \in \Gamma_c$.*

Remark 11 *This assumption is used in the analysis to perturb candidate representations that are not cluster-preserving towards improved representations that are cluster-preserving but have a lower loss. If \mathcal{F} is expressive enough to isolate each cluster Γ_c (that is, if $\chi_{\{x \in \Gamma_c\}} \in \mathcal{F}$ for all c), then for Assumption 10 to be satisfied, it is sufficient for \mathcal{F} to be closed under negations and the 'OR' operations, that is, if $f \in \mathcal{F} \implies \neg f \in \mathcal{F}$ and $f_1, f_2 \in \mathcal{F} \implies f_1 \wedge f_2 \in \mathcal{F}$.*

Next we define a regularity condition of a function class and augmentation set that captures the extent to which non-cluster-preserving classifiers classify images in positive pairs differently within clusters that they intersect. So far, we have only assumed that non-cluster-preserving classifiers misclassify at least *one* positive pair differently within any cluster they intersect (Assumption 2). However, for regular classes of binary classifiers and intertwined augmentation sets within clusters, we can expect that the number of positive pairs split in a cluster that are split by any binary classifier scales with the number of negative pairs in the same cluster that are split by the classifier. For a set of images $B \subseteq D$, we employ the notations $\|B\|_o := \mathbb{P}_{x \sim D_o}[x \in B]$ and $\|B\| := \mathbb{P}_{x \sim D \setminus D_o}[x \in B]$.

Definition 12 (δ -Regularity) For any $f \in \mathcal{F}$, let $\Sigma_f := \{\mathbf{c} \in C : \exists x, x' \in \Gamma_{\mathbf{c}} \text{ s.t. } f(x) \neq f(x')\}$ be the set of content variables corresponding to clusters split by f . For all $\mathbf{c} \in \Sigma_f$ and $\sigma \in \{-1, 1\}$, define

$$f^{(\mathbf{c}, \sigma)}(x) := \begin{cases} f(x) & x \notin \Gamma_{\mathbf{c}} \\ \sigma & x \in \Gamma_{\mathbf{c}} \end{cases}$$

as the classifier that outputs the same label as f on all images not in $\Gamma_{\mathbf{c}}$ and σ on $\Gamma_{\mathbf{c}}$. Further define

$$\Delta_{f, \mathbf{c}} := \min_{\sigma \in \{-1, 1\}} \|\{x \in \Gamma_{\mathbf{c}, o} \mid f(x) \neq f^{(\mathbf{c}, \sigma)}(x)\}\|_o$$

as the minimum measure of the set on which $f^{(\mathbf{c}, \sigma)}$ and f differ among all possible choices of $\sigma \in \{-1, 1\}$. Then (\mathcal{F}, Λ) is δ -regular if for all $\mathbf{c} \in \Sigma_f$,

$$\|\{\mathcal{A}(x) : \mathcal{A} \in \Lambda, x \in \Gamma_{\mathbf{c}, o}, f(x) \neq f(x^+)\}\| \geq \delta \Delta_{f, \mathbf{c}}.$$

Next, we state our regularity assumption and the result for the agnostic case.

Assumption 13 (δ -Regularity of (\mathcal{F}, Λ)) The pair (\mathcal{F}, Λ) is δ -regular with $\delta \geq 0.4$.

Remark 14 [Discussion of Assumption 13] Assumption 13 can be interpreted as a relationship between mislabelling of data and generalization error in the supervised learning problem associated with \mathcal{F} . Specifically, suppose $f^{(\mathbf{c}, \sigma)}(x) := f'(x)$ is a perfect classifier. Consider an entity that mislabels images before giving them to a supervised learner using class \mathcal{F} . For some f , the quantity $\|\{x \in \Gamma_{\mathbf{c}, o} \mid f(x) \neq f'(x)\}\|_o$ for some cluster can be interpreted as a “mislabelling budget”. The classifier f can be thought of solving a dataset in which the $\{x \in \Gamma_{\mathbf{c}, o} \mid f(x) \neq f'(x)\}$ has been mislabelled. Further, suppose generalization loss is given by the measure of augmented images (recall that these are not provided during training in the supervised problem considered in this remark) that are misclassified by f . This is precisely the quantity $\|\{\mathcal{A}(x) : \mathcal{A} \in \Lambda, x \in \Gamma_{\mathbf{c}, o}, f(x) \neq f(x^+)\}\|$. To summarize, this assumption says that training on the “true” labels, results in good generalization, but classifying a mislabelled dataset results in a generalization error that scales in the size of the mislabelling.

Note that this is related to Assumption 2, which simply states that $\delta_f > 0$ whenever $\Delta_{f, \mathbf{c}} > 0$. Assumption 2 is sufficient in the realizable setting.

Theorem 15 Suppose Assumptions 10 and 13 hold and $g = [f_1, \dots, f_d]$ is **not** cluster-preserving with $\min_{j \in [d]} \min_{\mathbf{c} \in \Sigma_{f_j}} \mathbb{P}_{x, x' \sim D}[x, x' \in \Gamma_{\mathbf{c}}, f_j(x) \neq f_j(x')] \geq \epsilon > 0$. Let $\ell \geq \frac{\epsilon}{c} d 2^d$, $\beta \geq c \log(\frac{\epsilon}{c}) 2^d$ for

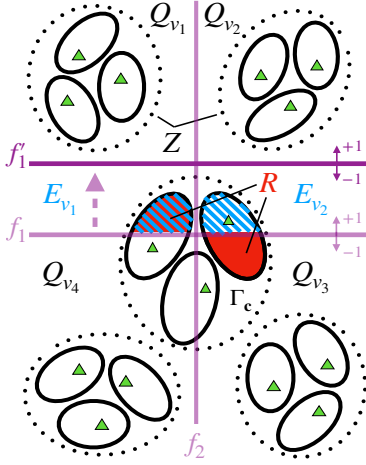


Figure 2: Example partitioning of images with $d = 2$ and $D \subset \mathbb{R}^2$. Green triangles denote natural images and solid black ellipses denote their corresponding augmentation sets (here we have drawn the augmentation sets as compact convex sets for ease of presentation, but in reality they may be non-simply connected and non-smooth). Clusters are indicated by dotted black ellipses. The non-cluster-preserving representation $g = (f_1, f_2)$, and we construct $g' = (f'_1, f_2)$ by making f_1 preserve the cluster Γ_c . The region R consisting of augmentations in Γ_c misclassified by f_1 is shaded red, and the set E of images which are classified differently by f_1 and f'_1 is indicated by blue diagonal lines. By Assumption 13, $\|R\| \geq \delta \|E\|_o$, where $\|B\| := \mathbb{P}_{x \sim D \setminus D_o}[x \in B]$ and $\|B\|_o := \mathbb{P}_{x \sim D_o}[x \in B]$ for any set of images $B \subseteq D$.

a sufficiently large constant c . Moreover, suppose g is close to a uniform representation in the sense that $\mathbb{P}_{x \sim D_o}[g(x) = v] \geq \frac{10}{cd2^d}$ or $\mathbb{P}_{x \sim D_o}[g(x) = v] \leq \frac{\epsilon}{100cd2^{2d}}$ for all $v \in \mathcal{H}_d$. Then g is **not** a minimizer of the InfoNCE loss.

Proof [Proof sketch of Theorem 15] For a non-cluster-preserving representation g that is “close” to a uniform representation, we construct a nearby representation g' by changing one coordinate of g such that it preserves one additional cluster, and show that the resulting g' achieves smaller InfoNCE loss than g . Suppose WLOG that f_1 does not preserve the cluster Γ_c . Further, let $f_1^{(c,\sigma)}$ be the smallest perturbation of f_1 that preserves Γ_c , as defined in Definition 12. Denote $f'_1 = f_1^{(c,\sigma)}$. By Assumption 10, $f'_1 \in \mathcal{F}$. Construct $g' = [f'_1, f_2, \dots, f_d] \in \mathcal{G}$. Note that g' is equivalent to g on all but one coordinate, and the one differing coordinate differs only on one cluster.

To characterize the variation in the InfoNCE loss when moving from g to g' , we first consider a specific partition of the space of images defined based on the representations g and g' . In particular, for a given vertex $v \in \mathcal{H}_d$, consider the set $Q_v := \{x \in D : g(x) = v, g'(x) = v\}$ which denotes the set of images that both g and g' map to vertex v , and the set $E_v := \{x \in D : g(x) = v, g'(x) \neq v\}$ which denotes the set of images that g maps to v and g' maps to another vertex. Considering these definitions, the set $Q := \cup_{v \in \mathcal{H}_d} Q_v$ corresponds to the set of all images that g and g' map to the same vertex, while $E := \cup_{v \in \mathcal{H}_d} E_v$ denotes the set of all images which g and g' map to different vertices. Based on this construction, it is not hard to observe that for any $v \neq v'$ the sets $Q_v, Q_{v'}, E_v,$ and $E_{v'}$ are disjoint, and each image belongs to either some Q_v or E_v . Hence, the concatenation of these sets partitions the space of images. Figure 2 illustrates this partition for a special case with $d = 2$. The above partition is critical as we divide our sensitivity analysis into multiple cases based on the location of the positive and negative images in this partition.

Let us define \mathcal{L}^+ and \mathcal{L}^- as the alignment and uniformity losses in (3), respectively. We refer to \mathcal{L}^+ as the positive part of the loss as it deals with positive samples (augmented images), and we refer to \mathcal{L}^- as the negative part of the loss as it contains negative samples. To prove that moving from g to g' decreases the loss, i.e., $\mathcal{L}(g) - \mathcal{L}(g') > 0$, we show that the amount that

3. Note that this near-uniformity condition allows for representations that for each vertex put mass at least a constant factor of $\frac{1}{d}$ times 2^{-d} , or essentially treat the vertex as inactive, which allows for the case wherein the number of clusters is less than 2^d and some vertices are inactive for cluster-preserving representations.

the positive part of the loss decreases is more than the amount the negative part might increase: $\mathcal{L}^+(g) - \mathcal{L}^+(g') > \mathcal{L}^-(g') - \mathcal{L}^-(g)$. To do so, first, note that the variation in the positive part is

$$\mathcal{L}^+(g) - \mathcal{L}^+(g') = 2\beta (\mathbb{P}[x \in Q, x^+ \in E] + \mathbb{P}[x \in Q, x^+ \in E]). \quad (6)$$

This holds as $\beta g(x)^\top g(x^+) = \beta g(x)^\top g'(x^+)$ except for the cases that $x \in Q, x^+ \in E$ or $x \in Q, x^+ \in E$. In these two cases, they differ by 2β . Note that the augmentations that belong to either of these two cases lie in the area shaded red in Figure 2. We refer to the set of augmentations in this region as R , in other words, R is the set of augmentations in Γ_c that are classified differently than their natural image by f_1 . Thus, we can write $\mathcal{L}^+(g) - \mathcal{L}^+(g') = 2\beta \|R\|$.

Next, we consider the difference in negative parts of the loss. To bound this difference, we leverage the partitioning of the space of images defined above to decompose the variation of the losses based on the set that image x belongs to. In particular, if we define the function

$$\mathcal{L}_B^-(g) := \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi(B) \log (e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)})]$$

for any event B , where $\chi(B)$ is the indicator random variable for the event B , then using the fact that each image x either belongs to one of the Q_v 's or E_v 's we can write

$$\mathcal{L}^-(g') - \mathcal{L}^-(g) = \sum_{v \in \mathcal{H}_d} [\mathcal{L}_{\{x \in Q_v\}}^-(g') - \mathcal{L}_{\{x \in Q_v\}}^-(g)] + [\mathcal{L}_{\{x \in E_v\}}^-(g') - \mathcal{L}_{\{x \in E_v\}}^-(g)], \quad (7)$$

Since the cases with $x \in E_v$ utilize similar analysis for those with $x \in Q_v$, we focus on the $x \in Q_v$ cases here and defer the $x \in E_v$ cases to Appendix B.

To analyze $\mathcal{L}_{\{x \in Q_v\}}^-(g') - \mathcal{L}_{\{x \in Q_v\}}^-(g)$, we first observe that this difference is non-positive for a subset of the Q_v 's. Note in Fig. 2 that if x belongs to Q_{v_1} or Q_{v_2} , then moving from f_1 to f'_1 decreases the representation similarity for some pairs of negative samples (those with $x_i^- \in E$) while keeping the rest the same. So, the negative part of the loss cannot increase going from g to g' if x lies in either Q_{v_1} or Q_{v_2} . We formally define this set of Q_v 's as $Z := \{Q_v : f'_1(x) \neq f'_1(x^-) \forall x \in Q_v, x^- \in E_v\}$. At a high level, the reason this definition implies the negative part of the loss does not increase if $x \in Q_v \in Z$ is because f_1 and f'_1 must agree on $Q_v \in Z$ and disagree on E , so since f'_1 differs on $Q_v \in Z$ and E , f_1 must agree on these sets. Thus, the similarity between negative pairs consisting of $x \in Q_v \in Z$ and $x_i^- \in E$ diminishes when moving from g to g' . Thus, we have

$$\sum_{v \in \mathcal{H}_d} \mathcal{L}_{\{x \in Q_v\}}^-(g') - \mathcal{L}_{\{x \in Q_v\}}^-(g) \leq \sum_{v \in \mathcal{H}_d} \mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g') - \mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g) \quad (8)$$

Now, for each event $\{x \in Q_v \notin Z\}$, we consider two cases depending on the number of negative samples in Q_v . (1) If there is at least one negative sample $x_i^- \in Q_v$, then $g'(x) = g'(x_i^-) = g(x) = g(x_i^-)$, so both the log-sums in $\mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g')$ and $\mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g)$ are dominated by $e^{\beta d}$ terms and the losses do not significantly differ (using that log-sum is approximately a max operation). (2) If no negative samples lie in Q_v , then the dominant terms in the log-sum for $\mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g')$ may be a factor of $e^{2\beta}$ larger than the dominant terms for $\mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g)$, requiring a sharp analysis to control the probability these events occur. Letting $n_{1,v}$ denote the number of negative samples in Q_v , we define

these two cases above as $B_{v,1} := \{x \in Q_v \notin Z, n_{1,v} > 0\}$ and $B_{v,2} := \{x \in Q_v \notin Z, n_{1,v} = 0\}$, respectively. Note that they form a partition of $\{x \in Q_v \notin Z\}$, so we have

$$\mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g') - \mathcal{L}_{\{x \in Q_v \notin Z\}}^-(g) = \sum_{j=1}^2 \mathcal{L}_{B_{v,j}}^-(g') - \mathcal{L}_{B_{v,j}}^-(g).$$

We detail each case below, where n_2 is the number of negative samples in E .

Case 1: $B_{1,v} := \{x \in Q_v \notin Z, n_{1,v} > 0\}$. In this case the dominant terms in the log sums for $\mathcal{L}_{B_{1,v}}^-(g')$ and $\mathcal{L}_{B_{1,v}}^-(g)$ are both $e^{\beta d}$, although the losses may differ in the number of such terms, which can be, in the worst case, $n_{1,v} + n_2 + 1$ for g' and $n_{1,v}$ for g . This is because g and g' can disagree on at most n_2 negative samples, and they can also disagree on the positive sample. Thus,

$$\mathcal{L}_{B_{1,v}}^-(g') - \mathcal{L}_{B_{1,v}}^-(g) \leq \mathbb{E}[\chi(B_{1,v}) \log\left(\frac{n_{1,v} + n_2 + 1}{n_{1,v}}\right)] \leq \mathbb{E}[\chi(B_{1,v}) \frac{2(n_2 + 1)}{n_{1,v} + 1}],$$

where the last inequality follows using $\log(1+x) \leq x$. We bound $\mathbb{E}[\chi(B_{1,v}) \frac{2(n_2 + 1)}{n_{1,v} + 1}]$ by writing the trinomial expansion of the expectation (note that the joint distribution of $(n_{1,v}, n_2)$ is trinomial with parameters $(\|Q_v\|_\circ, \|E\|_\circ)$), and further simplifying to result in an upper bound of $\|E\|_\circ$. Importantly, this bound is $O(\|R\|)$ by Assumption 13 and independent of β , so we control it by making β large enough.

Case 2: $B_{2,v} := \{x \in Q_v \notin Z, n_{1,v} = 0\}$. Since here there is no shared dominant $e^{\beta d}$ term in the log-sums for $\mathcal{L}_{B_{2,v}}^-(g')$ and $\mathcal{L}_{B_{2,v}}^-(g)$, the dominant terms for g' may involve strictly larger similarities than those for g , corresponding to $x_i^- \in E$ and $x^+ \in E$ (the only samples on which g' and g can disagree). These events are bounded depending on whether $n_2 = 0$. If $n_2 = 0$, the loss of g' exceeds that of g iff $g'(x)^\top g'(x^+) = g(x)^\top g(x^+) + 2$, which occurs iff $x^+ \in E$. If $n_2 > 0$, the loss can increase by 2β regardless of the value of x^+ . Combining these sub-cases yields

$$\begin{aligned} \mathcal{L}_{B_{2,v}}^-(g') - \mathcal{L}_{B_{2,v}}^-(g) &\leq 2\beta \mathbb{P}(x \in Q_v \notin Z, x^+ \in E) \mathbb{P}(n_{1,v} = 0, n_2 = 0) \\ &\quad + 2\beta \mathbb{P}(x \in Q_v \notin Z) \mathbb{P}(n_{1,v} = 0, n_2 > 0). \end{aligned}$$

For each term above, we need to show that the coefficient of 2β is $o(\|R\|)$ even after it is summed over v . Note that both terms scale with the probability that $x \in Q_v \notin Z$ and no negative samples are in Q_v . To control this probability we leverage that the distribution induced by g is close to uniform in the sense that every ‘‘active’’ vertex v has mass $\mathbb{P}[g(x) = v] = \tilde{\Omega}(\frac{1}{d^{2d}})$. We use this fact to bound $\mathbb{P}[x \in Q_v]$. Note that the set of images that g maps to v is $Q_v \cup E_v$, yet for all $Q_z \notin Z$, $E_v = \emptyset$ since, at a high level, f_1 must separate these Q_v from E . So, $\mathbb{P}[g(x) = v] = \mathbb{P}[x \in Q_v]$ for all $Q_v \notin Z$. Therefore, we can show that with large ℓ it is highly unlikely that $x \in Q_v \notin Z$ and none of the negative samples are in Q_v . To complete the bounds, we leverage the facts that $\mathbb{P}[x \in Q_v \notin Z, x^+ \in E]$ scales with $\|R\|$ for the first term, and $\mathbb{P}[n_2 > 0]$ scales with $\|E\|_\circ$ for the second term, where $\|E\|_\circ = O(\|R\|)$ by Assumption 13.

After performing a similar analysis for $\{x \in E_v\}$ and summing the resulting bounds over $\{v \in \mathcal{H}_d\}$, as in (6), we obtain $\mathcal{L}^-(g') - \mathcal{L}^-(g) < 2\beta \|R\| = \mathcal{L}^+(g) - \mathcal{L}^+(g')$. ■

Theorem 15 shows that for large ℓ and β , all minimizers of the InfoNCE loss that are near-uniform must be cluster-preserving regardless of the sizes of each cluster or the number of clusters. However,

it does not rule out that there could be a highly non-uniform and non-cluster-preserving optimal representation. In Appendix B.1, we show that if we re-weight the alignment and uniformity losses in the InfoNCE loss, we can ensure that *all* minimizers of the InfoNCE loss are cluster-preserving.

6. Conclusion

We study properties of minimizers of the InfoNCE loss optimized over function classes with restricted complexity relative to the complexity of augmentations in the dataset, in realistic settings with disjoint augmentation sets and finite negative samples. Our results show that such representations are uniform and cluster-preserving in the realizable setting, and must be cluster-preserving if they are close to uniform in the agnostic setting. We believe that our novel analytical tools, namely our stochastic argument for the optimality of representations and our inverse partitioning of the space of images, may be of use for future studies of the InfoNCE loss.

Acknowledgments

This research is supported in part by NSF Grants 2019844, 2107037 and 2112471, ONR Grant N00014-19-1-2566, the Machine Learning Lab (MLL) at UT Austin, and the Wireless Networking and Communications Group (WNCG) Industrial Affiliates Program.

References

- Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. In *36th International Conference on Machine Learning, ICML 2019*, pages 9904–9923. International Machine Learning Society (IMLS), 2019.
- Jordan T Ash, Surbhi Goel, Akshay Krishnamurthy, and Dipendra Misra. Investigating the role of negatives in contrastive representation learning. *arXiv preprint arXiv:2106.09943*, 2021.
- Pranjal Awasthi, Nishanth Dikkala, and Pritish Kamath. Do more negative samples necessarily hurt in contrastive learning? In *International Conference on Machine Learning*, pages 1101–1116. PMLR, 2022.
- Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32, 2019.
- Randall Balestriero and Yann LeCun. Contrastive and non-contrastive self-supervised learning recover global and local spectral embedding methods. *arXiv preprint arXiv:2205.11508*, 2022.
- Han Bao, Yoshihiro Nagano, and Kento Nozawa. On the surrogate gap between contrastive and supervised losses. In *International Conference on Machine Learning*, pages 1585–1606. PMLR, 2022.
- Pierre Bremaud. *Markov chains: Gibbs fields, Monte Carlo simulation, and queues; 1st ed.* Texts in applied mathematics. Springer, Berlin, 2001.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020.
- Mayee Chen, Daniel Y Fu, Avaniika Narayan, Michael Zhang, Zhao Song, Kayvon Fatahalian, and Christopher Ré. Perfectly balanced: Improving transfer and robustness of supervised contrastive learning. In *International Conference on Machine Learning*, pages 3090–3122. PMLR, 2022.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020a.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020b.
- Ting Chen, Calvin Luo, and Lala Li. Intriguing properties of contrastive losses. *Advances in Neural Information Processing Systems*, 34:11834–11845, 2021.
- DeepNets. Animals - v2: Image classification dataset, Nov 2022. URL <https://www.kaggle.com/datasets/utkarshsaxenadn/animal-image-classification-dataset>.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*, 2021.
- Quentin Garrido, Yubei Chen, Adrien Bardes, Laurent Najman, and Yann Lecun. On the duality between contrastive and non-contrastive self-supervised learning. *arXiv preprint arXiv:2206.02574*, 2022.
- Kartik Gupta, Thalaisyasingam Ajanthan, Anton van den Hengel, and Stephen Gould. Understanding and improving the role of projection head in self-supervised learning. *arXiv preprint arXiv:2212.11491*, 2022.
- Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010.
- Jeff Z HaoChen and Tengyu Ma. A theoretical study of inductive biases in contrastive learning. *arXiv preprint arXiv:2211.14699*, 2022.
- Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34:5000–5011, 2021.

- Jeff Z HaoChen, Colin Wei, Ananya Kumar, and Tengyu Ma. Beyond separability: Analyzing the linear transferability of contrastive representations to related subpopulations. *arXiv preprint arXiv:2204.02683*, 2022.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *International conference on machine learning*, pages 4182–4192. PMLR, 2020.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- Weiran Huang, Mingyang Yi, and Xuyang Zhao. Towards the generalization of contrastive self-supervised learning. *arXiv preprint arXiv:2111.00743*, 2021.
- Wenlong Ji, Zhun Deng, Ryumei Nakada, James Zou, and Linjun Zhang. The power of contrast for feature learning: A theoretical analysis. *arXiv preprint arXiv:2110.02473*, 2021.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34: 309–323, 2021.
- Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020.
- David McAllester and Karl Stratos. Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pages 875–884. PMLR, 2020.
- Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6707–6717, 2020.
- Kento Nozawa and Issei Sato. Understanding negative samples in instance discriminative self-supervised representation learning. *Advances in Neural Information Processing Systems*, 34: 5784–5797, 2021.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *CoRR*, abs/1802.05365, 2018. URL <http://arxiv.org/abs/1802.05365>.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. *arXiv preprint arXiv:2010.04592*, 2020.
- Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. *arXiv preprint arXiv:2202.14037*, 2022.
- Kendrick Shen, Robbie M Jones, Ananya Kumar, Sang Michael Xie, Jeff Z HaoChen, Tengyu Ma, and Percy Liang. Connect, not collapse: Explaining contrastive learning for unsupervised domain adaptation. In *International Conference on Machine Learning*, pages 19847–19878. PMLR, 2022.
- Yixuan Su, Fangyu Liu, Zaiqiao Meng, Tian Lan, Lei Shu, Ehsan Shareghi, and Nigel Collier. Tacl: Improving bert pre-training with token-aware contrastive learning. *arXiv preprint arXiv:2111.04198*, 2021.
- Joseph John Thomson. Xxiv. on the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 7(39):237–265, 1904.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European conference on computer vision*, pages 776–794. Springer, 2020a.
- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020b.
- Yuandong Tian. Deep contrastive learning is provably (almost) principal component analysis. *arXiv preprint arXiv:2201.12680*, 2022a.
- Yuandong Tian. Understanding the role of nonlinearity in training dynamics of contrastive learning. *arXiv preprint arXiv:2206.01342*, 2022b.
- Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pages 1179–1206. PMLR, 2021.
- Michael Tschannen, Josip Djolonga, Paul K Rubenstein, Sylvain Gelly, and Mario Lucic. On mutual information maximization for representation learning. *arXiv preprint arXiv:1907.13625*, 2019.
- Julius Von Kügelgen, Yash Sharma, Luigi Gresele, Wieland Brendel, Bernhard Schölkopf, Michel Besserve, and Francesco Locatello. Self-supervised learning with data augmentations provably isolates content from style. *Advances in neural information processing systems*, 34:16451–16467, 2021.

- Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504, 2021.
- Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020.
- Yifei Wang, Qi Zhang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. Chaos is a ladder: A new theoretical understanding of contrastive learning via augmentation overlap. *arXiv preprint arXiv:2203.13457*, 2022.
- Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data. *arXiv preprint arXiv:2010.03622*, 2020.
- Zixin Wen and Yuanzhi Li. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pages 11112–11122. PMLR, 2021.
- Zixin Wen and Yuanzhi Li. The mechanism of prediction head in non-contrastive self-supervised learning. *arXiv preprint arXiv:2205.06226*, 2022.
- Mingkai Zheng, Fei Wang, Shan You, Chen Qian, Changshui Zhang, Xiaogang Wang, and Chang Xu. Weakly supervised contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10042–10051, 2021.
- Roland S Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In *International Conference on Machine Learning*, pages 12979–12990. PMLR, 2021.

Appendix A. Proof of Theorem 7

We prove Theorem 7. To do so, we first prove Lemma 6, and use this to prove our main result.

A.1. Proof of Lemma 6

Lemma 16 (Lemma 6 Restated) *If Assumptions 2 and 5 hold, $\beta > c \log d$ for an absolute constant c , $d > 3$, and $\ell \geq 1$, then $g^* \in \arg \min_{g \in \mathcal{G}_c} \mathcal{L}(g)$ if and only if g^* is uniform.*

Proof To prove this claim we first note that since we are optimizing over \mathcal{G}_c and $g(x) = g(x^+)$ for all $x \in D_\circ$, $x^+ \in A(x)$ and $g \in \mathcal{G}_c$, the optimization problem $\min_{g \in \mathcal{G}_c} \mathcal{L}(g)$ is equivalent to

$$\min_{g \in \mathcal{G}_c} \hat{\mathcal{L}}(g) = \mathbb{E}_{x, \{x_i^-\}_\ell} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-) - \beta d} \right) \right]. \quad (9)$$

Note that we can think of this as an optimization over distributions over the hypercube induced by g . That is, consider the random variable Y supported on the hypercube that is given by $Y = g(X)$ for $X \sim D_\circ$, and denote its distribution as \mathcal{D}_g . Using this notation, the optimization above can be rewritten in terms of distributions

$$\min_{\mathcal{D} \in \{\mathcal{D}_g : g \in \mathcal{G}\}} \tilde{\mathcal{L}}(\mathcal{D}) = \mathbb{E}_{y, \{\bar{y}_i\}_\ell \sim \mathcal{D}} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta y^\top \bar{y}_i - \beta d} \right) \right].$$

where \bar{y}_i corresponds to the representation of the i -th negative sample, and here we overload notation by using $\sim \mathcal{D}$ to denote an i.i.d. draw from the distribution \mathcal{D} .

Next, we define a Markov chain as follows. We begin with a fresh set of samples denoted by $y^0, \bar{y}_1^0, \dots, \bar{y}_\ell^0$ that are drawn i.i.d. from the distribution $\mathcal{D}^0 = \mathcal{D}$. At each step, for each sample, we either with probability $\frac{1}{2}$ flip one bit uniformly at random, or with probability $\frac{1}{2}$ we do not change it. Concretely, we take for all i a random variable $j_{i,t} \in [d]$ (both uniformly random and independent of each other and every other such sample) and set $(\bar{y}_i^t)_{j_{i,t}} = -(\bar{y}_i^{t-1})_{j_{i,t}}$. After this operation, each y_i^t (and y) can be considered to be an i.i.d. drawn from \mathcal{D}^t , where \mathcal{D}^t is another distribution over \mathcal{H}_d . We show that $\mathcal{L}(\mathcal{D}^{t-1}) > \mathcal{L}(\mathcal{D}^t)$ if \mathcal{D}^{t-1} is not uniform. Since $\{\mathcal{D}^t\}_t$ converges to the uniform distribution by Lemma 19, these two arguments together imply the claim of Lemma 6.

To show $\mathcal{L}(\mathcal{D}^{t-1}) > \mathcal{L}(\mathcal{D}^t)$ for the case that \mathcal{D}^{t-1} is not the uniform distribution, considering the definition of \mathcal{L} we need to study the variation in the inner products between the vectors (y^t, \bar{y}_i^t) when we move from one distribution to another. Note that as these vectors are binary vectors, their inner product can be written as a function of their Hamming distances. More precisely, for any pair (y, y') we have $y^\top y' = d - 2h(y, y')$, where the Hamming distance between them is defined as $h(y, y') := \sum_{j=1}^{\ell} \chi\{y_j \neq y'_j\}$ or the number of bits that are different in the two points y, y' (note that $\chi\{U\}$ is the indicator variable for the event U).

For ease of notation we let $h_i^t := h(y^t, \bar{y}_i^t)$ for all i, t . Due to the fact that each of the \bar{y}_i^t 's are independent and identically distributed and are evolving according to a Markov chain, the h_i^t 's also evolve according to a Markov chain. In particular, for every distribution \mathcal{D}^t over \mathcal{H}_d that describes the distribution of each \bar{y}_i^t , there is induced a distribution $\tilde{\mathcal{D}}_h^t$ over $[d]$ that specifies the distribution for h_i^t . By direct computation, one can check that h_i^t has the following transition kernels which differ for different values of h_i :

For $2 \leq h_i^{t-1} \leq d-2$:

$$h_i^t \rightarrow \begin{cases} h_i^{t-1} - 2 & \text{w.p. } \frac{(h_i^{t-1})(h_i^{t-1}-1)}{4d^2} \\ h_i^{t-1} - 1 & \text{w.p. } \frac{h_i^{t-1}}{2d} \\ h_i^{t-1} & \text{w.p. } \frac{1}{4} + \frac{(h_i^{t-1})(d-h_i^{t-1}+1)+(d-h_i^{t-1})(h_i^{t-1}+1)}{4d^2} \\ h_i^{t-1} + 1 & \text{w.p. } \frac{d-h_i^{t-1}}{2d} \\ h_i^{t-1} + 2 & \text{w.p. } \frac{(d-h_i^{t-1})(d-h_i^{t-1}+1)}{4d^2} \end{cases}$$

For $h_i^{t-1} = 1$:

$$1 \rightarrow \begin{cases} 0 & \text{w.p. } \frac{1}{2d} \\ 1 & \text{w.p. } \frac{1}{4} + \frac{1}{4d} + \frac{2(d-1)}{4d^2} \\ 2 & \text{w.p. } \frac{d-1}{2d} \\ 3 & \text{w.p. } \frac{(d-1)(d-2)}{4d^2} \end{cases} \quad (10)$$

For $h_i^{t-1} = d-1$:

$$1 \rightarrow \begin{cases} d & \text{w.p. } \frac{1}{2d} \\ d-1 & \text{w.p. } \frac{1}{4} + \frac{1}{4d} + \frac{2(d-1)}{4d^2} \\ d-2 & \text{w.p. } \frac{d-1}{2d} \\ d-3 & \text{w.p. } \frac{(d-1)(d-2)}{4d^2} \end{cases}$$

For $h_i^{t-1} = 0$:

$$0 \rightarrow \begin{cases} 0 & \text{w.p. } \frac{1}{4} + \frac{1}{4d} \\ 1 & \text{w.p. } \frac{1}{2} \\ 2 & \text{w.p. } \frac{d-1}{4d} \end{cases} \quad (11)$$

For $h_i^{t-1} = d$:

$$d \rightarrow \begin{cases} d & \text{w.p. } \frac{1}{4} + \frac{1}{4d} \\ d-1 & \text{w.p. } \frac{1}{2} \\ d-2 & \text{w.p. } \frac{d-1}{4d} \end{cases}$$

For ease of notation we drop the t superscripts and refer to quantities at time $t-1$ without any superscript, and quantities at times t with a $'$ superscript, e.g. \mathcal{D}^{t-1} as \mathcal{D} and \mathcal{D}^t as \mathcal{D}' , for the remainder of the proof.

Next, let $\mathbf{h} = [h_1, \dots, h_\ell]$ denote the vector of concatenated Hamming distances between y and $\{\bar{y}_i\}_{i=1}^\ell$. Using the above definitions, and the fact that $d - y^\top \bar{y}_i$ is twice of the hamming distance between y and \bar{y}_i the loss $\tilde{\mathcal{L}}$ can be written as

$$\tilde{\mathcal{L}}(\mathcal{D}) = \mathbb{E}_{y, \{\bar{y}_i\}_{i=1}^\ell \sim \mathcal{D}} \left[\log \left(1 + \sum_{i=1}^\ell e^{\beta(y^\top \bar{y}_i - d)} \right) \right] = \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h} \left[\log \left(1 + \sum_{i=1}^\ell e^{-2\beta h_i} \right) \right]$$

Now to characterize the difference between $\tilde{\mathcal{L}}(\mathcal{D})$ and $\tilde{\mathcal{L}}(\mathcal{D}')$ we need to study the evolution of the distribution of the Hamming distance \mathbf{h} from $\tilde{\mathcal{D}}_h$ to $\tilde{\mathcal{D}}'_h$, i.e.,

$$\tilde{\mathcal{L}}(\mathcal{D}') - \tilde{\mathcal{L}}(\mathcal{D}) = \mathbb{E}_{\mathbf{h}' \sim \tilde{\mathcal{D}}'_h} \left[\log \left(1 + \sum_{i=1}^\ell e^{-2\beta h'_i} \right) \right] - \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h} \left[\log \left(1 + \sum_{i=1}^\ell e^{-2\beta h_i} \right) \right]$$

For each i , consider the random variable $s_i = h'_i - h_i$ that indicates which of the transitions is undertaken by h_i , and let $\mathbf{s} := [s_1, \dots, s_\ell]$ be its concatenation. Note that each s_i takes values in $\{-2, -1, 0, 1, 2\}$ and its distribution depends on the value of h_i , defined according to the transition kernel of h_i defined above. Given this, we can express $\tilde{\mathcal{L}}(\mathcal{D}')$ as

$$\tilde{\mathcal{L}}(\mathcal{D}') = \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h, \mathbf{s}} \left[\log 1 + \sum_{i=1}^{\ell} e^{-2\beta(h_i + s_i)} \right]$$

Now we consider the difference $\tilde{\mathcal{L}}(\mathcal{D}') - \tilde{\mathcal{L}}(\mathcal{D})$. According to the above definitions, this difference can be written as

$$\begin{aligned} & \tilde{\mathcal{L}}(\mathcal{D}') - \tilde{\mathcal{L}}(\mathcal{D}) \\ &= \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h, \mathbf{s}} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{-2\beta(h_i + s_i)} \right) - \log \left(1 + \sum_{i=1}^{\ell} e^{-2\beta h_i} \right) \right] \\ &= \sum_{i=1}^{\ell} \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h, \mathbf{s}} \left[\log \left(1 + \sum_{j=1}^i e^{-\beta(h_j + s_j)} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j} \right) \right. \\ & \quad \left. - \log \left(1 + \sum_{j=1}^{i-1} e^{-2\beta(h_j + s_j)} + \sum_{j=i}^{\ell} e^{-2\beta h_j} \right) \right] \end{aligned} \quad (12)$$

where (12) follows from telescoping over each negative sample indexed by i . Now if we take out the i -th term of each of the above two expressions the difference can be written as

$$\begin{aligned} & f(\mathcal{D}') - f(\mathcal{D}) \\ &= \sum_{i=1}^{\ell} \mathbb{E}_{\mathbf{h} \sim \tilde{\mathcal{D}}_h, \mathbf{s}} \left[\log \left(1 + \sum_{j=1}^{i-1} e^{-2\beta(h_j + s_j)} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j} + e^{-2\beta(h_i + s_i)} \right) \right. \\ & \quad \left. - \log \left(1 + \sum_{j=1}^{i-1} e^{-2\beta h_j} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j} + e^{-2\beta h_i} \right) \right] \\ &= \sum_{i=1}^{\ell} \mathbb{E}_{h_{-i}, s_i} \mathbb{E}_{h_i \sim \tilde{\mathcal{D}}_h, s_i} \left[\log \left(1 + \sum_{j=1}^{i-1} e^{-2\beta(h_j + s_j)} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j} + e^{-2\beta(h_i + s_i)} \right) \right. \\ & \quad \left. - \log \left(1 + \sum_{j=1}^{i-1} e^{-2\beta(h_j + s_j)} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j} + e^{-2\beta h_i} \right) \right] \end{aligned}$$

To show that the RHS above is strictly less than 0 when \mathcal{D} is not uniform, it is sufficient to show that each term of the sum is strictly less than zero. To do this, we show that the inner expectation is strictly negative. In other words, all that remains to prove is that for all instances of $C_{-i} := 1 + \sum_{j=1}^{i-1} e^{-2\beta(h_j + s_j)} + \sum_{j=i+1}^{\ell} e^{-2\beta h_j}$, each of the terms satisfies

$$\mathbb{E}_{h_i \sim \tilde{\mathcal{D}}_h, s_i} \left[\log \left(C_{-i} + e^{-2\beta(h_i + s_i)} \right) - \log \left(C_{-i} + e^{-2\beta h_i} \right) \right] < 0 \quad (13)$$

when \mathcal{D} is not uniform. Once we have this result, the claim that $f(\mathcal{D}') - f(\mathcal{D}) < 0$ holds.

From now on, for ease of notation, we replace C_{-i} by C . To prove the claim in (13), we first introduce the function $\Delta_C(h)$ defined as

$$\begin{aligned}\Delta_C(h) &:= \mathbb{E}_{s_i} \left[\log(C + e^{-2\beta(h+s_i)}) - \log(C + e^{-2\beta h}) \right] \\ &= \mathbb{E}_{s_i} \left[\log\left(1 + \frac{e^{-2\beta(h+s_i)}}{C}\right) - \log\left(1 + \frac{e^{-2\beta h}}{C}\right) \right].\end{aligned}$$

Considering this definition the claim in (13) can be translated into

$$\sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = k] \Delta_C(h) < 0. \quad (14)$$

To show this, it is easiest to compare this expression with the analogous expression in the case that \mathcal{D} is uniform. That is $y, \{\bar{y}_i\}_{i=1}^\ell$ are drawn from the uniform distribution \mathcal{U} . We let $\tilde{\mathcal{U}}_h$ denote the distribution on $[d]$ of Hamming distances induced by \mathcal{U} . By the stationarity of the uniform distribution (Lemma 19), the distribution of $h'_i = h_i + s_i$ is identical to that of h_i if h_i is drawn from a uniform distribution. Thus we have the following result:

$$\begin{aligned}& \mathbb{E}_{h_i \sim \tilde{\mathcal{U}}_h, s_i} \left[\log\left(C_{-i} + e^{-2\beta(h_i+s_i)}\right) - \log\left(C_{-i} + e^{-2\beta h_i}\right) \right] \\ &= \mathbb{E}_{h'_i \sim \tilde{\mathcal{U}}'_h} \left[\log\left(C_{-i} + e^{-2\beta(h'_i)}\right) \right] - \mathbb{E}_{h_i \sim \tilde{\mathcal{U}}_h} \left[\log\left(C_{-i} + e^{-2\beta h_i}\right) \right] \\ &= 0,\end{aligned} \quad (15)$$

Next, we show that $\sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = k] \Delta_C(h) - \sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = k] \Delta_C(h) < 0$, which by (15) immediately implies $\sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = k] \Delta_C(h) < 0$. To achieve this we invoke Lemmas 17 and 18, which describe the behavior of $\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = k] \Delta_C(h) - \sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = k] \Delta_C(h)$ and $\Delta_C(k)$, respectively. Using these lemmas we obtain:

$$\begin{aligned}& \sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = k] \Delta_C(k) - \sum_{k=0}^d \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = k] \Delta_C(k) \\ &\leq (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = 0]) \Delta_C(0) \\ &\quad + \sum_{k>0} \binom{d}{k} (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = 0]) \Delta_C(k)\end{aligned} \quad (16)$$

$$= (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = 0]) \sum_{k \geq 0} \binom{d}{k} \Delta_C(k) \quad (17)$$

$$< 0 \quad (18)$$

where (16) holds by Lemma 17-3, (17) follows by simply combining terms, and (18) holds by Lemma 17-2, which states that $\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h} [h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h} [h = 0] > 0$, and Lemma 18, which states that $\sum_{k \geq 0} \binom{d}{k} \Delta_C(k) < 0$. This completes the proof. \blacksquare

Lemma 17 For any the distribution $\tilde{\mathcal{D}}_h$ on $[d]$ induced by any non-uniform distribution \mathcal{D} on \mathcal{H}_d , the following are true:

1. $\mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = k] = \binom{d}{k} \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = 0]$
2. $\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = 0] > 0$
3. $\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h = k] - \mathbb{P}_{Y \sim \mathcal{U}}[h = k] < \binom{d}{k} (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h = 0] - \mathbb{P}_{Y \sim \mathcal{U}}[h(y, Y) = k])$

Proof For any vertex $v \in \mathcal{H}_d$, we let $p_v := \mathbb{P}_{y \sim \mathcal{D}}[y = v]$ for ease of notation.

1. Note that for any $k \in [d]$,

$$\mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = k] = \mathbb{P}_{y, y' \sim \mathcal{U}}[h(y, y') = k] = \mathbb{P}_{y \sim \mathcal{U}}[h(w, y) = k] \quad (19)$$

for any fixed vertex $w \in \mathcal{H}_d$ by the symmetry of the uniform distribution. Then

$$\begin{aligned} \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = k] &= \mathbb{P}_{y \sim \mathcal{U}}[h(w, y) = k] \\ &= \sum_{v \in \mathcal{H}_d: h(w, v) = k} \mathbb{P}_{y \sim \mathcal{U}}[y = v] \\ &= \sum_{v \in \mathcal{H}_d: h(w, v) = k} \frac{1}{2^d} \\ &= \binom{d}{k} \frac{1}{2^d}, \end{aligned}$$

which implies that $\mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = 0] = \frac{1}{2^d}$.

2. Using the above observation that $\mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = 0] = \frac{1}{2^d}$, we have

$$\begin{aligned} \mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h = 0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = 0] &= \sum_{v \in \mathcal{H}_d} (p_v^2 - \frac{1}{2^{2d}}) \\ &= \sum_{v \in \mathcal{H}_d} (p_v^2 - \frac{2p_v}{2^d} + \frac{1}{2^{2d}}) + \sum_{v \in \mathcal{H}_d} (-\frac{2}{2^{2d}} + \frac{2p_v}{2^d}) \\ &= \sum_{v \in \mathcal{H}_d} (p_v - \frac{1}{2^d})^2 \\ &> 0 \end{aligned}$$

where the strict inequality holds since \mathcal{D} is not uniform.

3. We argue similarly as in the proofs of the previous two statements. We have

$$\begin{aligned} &2 \left(\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h = k] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h = k] \right) \\ &= 2 \sum_{v, u \in \mathcal{H}_d, h(v, u) = k} (p_v p_u - \frac{1}{2^{2d}}) \end{aligned} \quad (20)$$

$$= \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (-p_v^2 + 2p_v p_u - p_u^2) + \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (p_v^2 + p_u^2 - \frac{2}{2^{2d}}) \quad (21)$$

$$= - \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (p_v - p_u)^2 + \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (p_v^2 - \frac{1}{2^{2d}} + p_u^2 - \frac{1}{2^{2d}})$$

$$= - \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (p_v - p_u)^2 + 2 \binom{d}{k} \sum_{v \in \mathcal{H}_d} (p_v^2 - \frac{1}{2^{2d}}) \quad (22)$$

$$= - \sum_{v,u \in \mathcal{H}_d, h(v,u)=k} (p_v - p_u)^2 + 2 \binom{d}{k} (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h=0] - \mathbb{P}_{h \sim \tilde{\mathcal{U}}_h}[h=0])$$

$$< 2 \binom{d}{k} (\mathbb{P}_{h \sim \tilde{\mathcal{D}}_h}[h=0] - \mathbb{P}_{Y \sim \mathcal{U}}[h(y, Y)=0]) \quad (23)$$

where (21) is obtained by adding and subtracting $p_v^2 + p_u^2$, (22) follows by the symmetry of the hypercube and the fact that for every $v \in \mathcal{H}_d$, there are $\binom{d}{k}$ vertices $u \in \mathcal{H}_d$ satisfying $h(v, u) = k$, and (23) follows by the fact that $p_v \neq p_u$ for some $v, u \in \mathcal{H}_d$ for all non-uniform distributions \mathcal{D} . ■

Lemma 18 *If $\beta > c \log d$ for an absolute constant c , $d > 3$ and $C > 1$, then $\sum_{k \geq 0} \binom{d}{k} \Delta_C(k) < 0$.*

Proof According to the transition matrix of h for the case that $h = 0$ we know that s could be either 0, 1, or 2, with probabilities denoted in (11). Hence, we can simplify the expression for $\Delta_C(0)$ as

$$\begin{aligned} \Delta_C(0) &= \log \left(1 + \frac{1}{C} \right) \frac{d+1}{4d} + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \frac{1}{2} + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \frac{d-1}{4d} - \log \left(1 + \frac{1}{C} \right) \\ &= -\frac{3d-1}{4d} \log \left(1 + \frac{1}{C} \right) + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \frac{1}{2} + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \frac{d-1}{4d} \end{aligned}$$

We can similarly compute $\Delta_C(1), \Delta_C(2), \dots$ to obtain:

$$\begin{aligned} &\sum_{k \geq 0} \binom{d}{k} \Delta_C(k) \\ &= \log \left(1 + \frac{1}{C} \right) \frac{-3d+1}{4d} + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \frac{1}{2} + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \frac{d-1}{4d} \\ &\quad + \binom{d}{1} \left(\log \left(1 + \frac{1}{C} \right) \frac{1}{2d} + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \frac{-3d^2+3d-2}{4d^2} \right. \\ &\quad \quad \left. + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \frac{d-1}{2d} + \log \left(1 + \frac{e^{-6\beta}}{C} \right) \frac{(d-1)(d-2)}{4d^2} \right) \\ &\quad + \binom{d}{2} \left(\log \left(1 + \frac{1}{C} \right) \frac{1}{2d^2} \right) \end{aligned}$$

$$\begin{aligned}
 & + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \frac{1}{d} + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \left(-\frac{3}{4} + \frac{2(d-1) + 3(d-2)}{4d^2} \right) \\
 & + \log \left(1 + \frac{e^{-6\beta}}{C} \right) \frac{(d-2)}{2d} + \log \left(1 + \frac{e^{-8\beta}}{C} \right) \frac{(d-2)(d-3)}{4d^2} \\
 & + \sum_{k>2} \binom{d}{k} \Delta(k) \\
 = & \log \left(1 + \frac{1}{C} \right) \left(\frac{-3d+1}{4d} + \frac{\binom{d}{1}}{2d} + \frac{\binom{d}{2}}{2d^2} \right) + \log \left(1 + \frac{e^{-2\beta}}{C} \right) \left(\frac{1}{2} + \binom{d}{1} \frac{-3d^2+3d-2}{4d^2} + \frac{\binom{d}{2}}{d} \right) \\
 & + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \left(\frac{d-1}{4d} + \binom{d}{1} \frac{d-1}{2d} + \binom{d}{2} \left(-\frac{3}{4} + \frac{2(d-1) + 3(d-2)}{4d^2} \right) \right) \\
 & + \log \left(1 + \frac{e^{-6\beta}}{C} \right) \left(\binom{d}{1} \frac{(d-1)(d-2)}{4d^2} + \binom{d}{2} \frac{(d-2)}{2d} \right) \\
 & + \log \left(1 + \frac{e^{-8\beta}}{C} \right) \binom{d}{2} \frac{(d-2)(d-3)}{4d^2} + \sum_{k>2} \binom{d}{k} \Delta(k) \\
 = & -\log \left(1 + \frac{e^{-2\beta}}{C} \right) \left(\frac{d^2-3d+2}{4d} \right) \\
 & + \log \left(1 + \frac{e^{-4\beta}}{C} \right) \left(\frac{d-1}{4d} + \binom{d}{1} \frac{d-1}{2d} + \binom{d}{2} \left(-\frac{3}{4} + \frac{2(d-1) + 3(d-2)}{4d^2} \right) \right) \\
 & + \log \left(1 + \frac{e^{-6\beta}}{C} \right) \left(\binom{d}{1} \frac{(d-1)(d-2)}{4d^2} + \binom{d}{2} \frac{(d-2)}{2d} \right) \\
 & + \log \left(1 + \frac{e^{-8\beta}}{C} \right) \binom{d}{2} \frac{(d-2)(d-3)}{4d^2} + \sum_{k>2} \binom{d}{k} \Delta(k) \\
 \leq & -\frac{e^{-2\beta}}{C} \frac{d^2-3d+2}{8d} + \frac{c'}{C} d^2 e^{-4\beta} + \sum_{k>2} \binom{d}{k} \Delta(k) \tag{24}
 \end{aligned}$$

where in (24) we have used the numerical inequalities $-\log(1+x) \leq -\frac{x}{2}$ for $x \in [0, 1]$ and $\log(1+x) \leq x$, and $\beta > c \log d$, and c' is a sufficiently large constant.

For $k > 2$, we again use $\log(1+x) \leq x$ to obtain

$$\begin{aligned}
 \Delta(k) & \leq \frac{e^{-2\beta \cdot (k-2)}}{C} \frac{(k-1)k}{4d^2} + \frac{e^{-2\beta(k-1)}}{C} \frac{k}{2d} - \frac{1}{2} \log \left(1 + \frac{e^{-2\beta k}}{C} \right) + \frac{e^{-2\beta(k+1)}}{2C} + \frac{e^{-2\beta(k+2)}}{4C} \\
 & \leq \frac{e^{-2\beta \cdot (k-2)}}{C} \frac{(k-1)k}{4d^2} + \frac{c''}{C} e^{-2\beta(k-1)} \tag{25}
 \end{aligned}$$

for an absolute constant c'' . Combining this bound with (24) yields

$$\begin{aligned}
 \sum_{k \geq 0} \binom{d}{k} \Delta_C(k) & \leq -\frac{e^{-2\beta}}{C} \frac{d^2-3d+2}{8d} + \frac{c'}{C} d^2 e^{-4\beta} + \sum_{k=3}^d \frac{e^{-2\beta \cdot (k-2)}}{C} \frac{(k-1)k}{4d^2} + \frac{c''}{C} e^{-2\beta(k-1)} \\
 & \leq \frac{e^{-2\beta}}{C} \left(-\frac{d^3-3d^2+2d}{8d^2} + \frac{12}{8d^2} \right) + \frac{c'''}{C} d^2 e^{-4\beta} \tag{26} \\
 & < 0 \tag{27}
 \end{aligned}$$

where (26) holds for an absolute constant c''' , and (27) holds for a sufficiently large constant c and $d > 3$, where throughout we have used $\beta > c \log d$. \blacksquare

Lemma 19 $\{\mathcal{D}_t\}_t$ converges to \mathcal{U}_d .

Proof The transition kernel of \mathcal{D}_t is aperiodic and irreducible over a finite state space, and has a symmetric transition kernel, so it must converge to the uniform distribution (Bremaud, 2001). \blacksquare

A.2. Proof of Theorem 7

Now using the above results, we prove the main claim of Theorem 7.

Proof Note that the InfoNCE loss can be written as

$$\mathcal{L}(g) = \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta g(x)^\top (g(x_i^-) - g(x^+))} \right) \right] \quad (28)$$

Considering that we search over representations composed of clean functions, we know that for all $g \in \mathcal{G}_c$, the term $g(x)^\top g(x^+)$ is exactly equal to d . Hence, the optimizing $\mathcal{L}(g)$ over \mathcal{G}_c simplifies to minimizing

$$\hat{\mathcal{L}}(g) := \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(1 + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-) - \beta d} \right) \right] \quad (29)$$

Below, we use the term ‘clean representation’ to indicate that the the representation is composed of clean functions, and a non-clean representation if at least one of the functions in the representation is not clean. Recall the definitions of the functions $\mathcal{L}(g)$ in (28) and $\hat{\mathcal{L}}(g)$ in (29). Note that since we always have $g(x)^\top g(x^+) \leq d$, we can argue that for any g we have $\mathcal{L}(g) \geq \hat{\mathcal{L}}(g)$. Indeed, the equality holds when g is a clean representation. Moreover, for any non-clean representation g , we know that there exists at least one image x for which its representation $g(x)$ is not exactly aligned with the representation of one of its augmented images x^+ . Therefore, for that pair (x, x^+) , $-\beta g(x)^\top g(x^+) > -\beta d$. Hence, for some sample $x, x^+, \{x_i^-\}_\ell$ with positive mass, we have: $\beta g(x)^\top (g(x_i^-) - g(x^+)) > \beta g(x)^\top g(x_i^-) - \beta d$. Therefore, for non-clean g we have $\mathcal{L}(g) > \hat{\mathcal{L}}(g)$ from (28) and (29).

Moreover, according to the result of Lemma 6, we know that the minimizer of the loss function $\hat{\mathcal{L}}$ is a uniform representation, thus for any non-uniform representation g' and uniform representation g'' we have $\hat{\mathcal{L}}(g') > \hat{\mathcal{L}}(g'')$.

Considering these two observations, we show that a uniform representation composed of clean functions, denoted by g^* , is an optimal solution of the loss \mathcal{L} . We consider the following two cases:

Case 1: If the representation g' is not composed of clean functions, then we have

$$\mathcal{L}(g') \stackrel{(a)}{>} \hat{\mathcal{L}}(g') \stackrel{(b)}{\geq} \hat{\mathcal{L}}(g^*) \stackrel{(c)}{=} \mathcal{L}(g^*)$$

where (a) holds with strict inequality since g' is not clean (discussion above), (b) holds as g^* with a uniform distribution is an optimal solution of $\hat{\mathcal{L}}$ (Lemma 6), and (c) holds because g^* is composed of clean functions.

Case 2: If the representation g' is composed of clean functions, but is not uniform, then we have

$$\mathcal{L}(g') \stackrel{(a)}{\geq} \hat{\mathcal{L}}(g') \stackrel{(b)}{>} \hat{\mathcal{L}}(g^*) \stackrel{(c)}{=} \mathcal{L}(g^*)$$

where (a) holds based on the definitions of \mathcal{L} and $\hat{\mathcal{L}}$, (b) holds since g' is not uniform and g^* is uniform (discussion above), and (c) holds because g^* is composed of clean functions.

Combining these two cases, we obtain that the representation g^* minimizes $\mathcal{L}(\cdot)$ if and only if g^* is composed of clean functions and uniform. Furthermore, by Assumption 2, g^* being composed of clean functions implies that it is cluster-preserving. \blacksquare

Appendix B. Agnostic Case

In this section we prove Theorem 15.

Theorem 20 (Theorem 15 Restated) *Suppose Assumptions 10 and 13 hold and $g = [f_1, \dots, f_d]$ is **not** cluster-preserving with $\min_{j \in [d]} \min_{\mathbf{c} \in \Sigma_{f_j}} \mathbb{P}_{x, x' \sim D}[x, x' \in \Gamma_{\mathbf{c}}, f_j(x) \neq f_j(x')] \geq \epsilon > 0$. Let $\ell \geq \frac{c}{\epsilon} d 2^d$, $\beta \geq c \log(\frac{c}{\epsilon}) 2^d$ for a sufficiently large constant $c > 1$. Moreover, suppose g is close to a uniform representation in the sense that $\mathbb{P}_{x \sim D_{\circ}}[g(x) = v] \geq \frac{10}{cd 2^d}$ or $\mathbb{P}_{x \sim D_{\circ}}[g(x) = v] \leq \frac{\epsilon}{100cd 2^{2d}}$ for all $v \in \mathcal{H}_d$. Then g is **not** a minimizer of the InfoNCE loss.*

Proof First we recall notations: for a set of images $B \subseteq D$, we employ the notations $\|B\|_{\circ} := \mathbb{P}_{x \sim D_{\circ}}[x \in B]$ and $\|B\| := \mathbb{P}_{x \sim D \setminus D_{\circ}}[x \in B]$. Also, we let $\Gamma_{\mathbf{c}, \circ} := \Gamma_{\mathbf{c}} \cap D_{\circ}$.

As discussed in the proof sketch, we construct a representation g' that is close to g by changing one coordinate of g such that it preserves one additional cluster, and show that the resulting g' achieves smaller InfoNCE loss than g .

Suppose WLOG that f_1 does not preserve the cluster $\Gamma_{\mathbf{c}}$ for some $\mathbf{c} \in C$. That is, $\exists x, x' \in \Gamma_{\mathbf{c}}$ such that $f_1(x) \neq f_1(x')$. Further, let $f_1^{(\mathbf{c}, \sigma)}$ be the smallest perturbation of f_1 that preserves $\Gamma_{\mathbf{c}}$.

Specifically, $f^{(\mathbf{c}, \sigma)}(x) := \begin{cases} f(x) & x \notin \Gamma_{\mathbf{c}} \\ \sigma & x \in \Gamma_{\mathbf{c}} \end{cases}$, where $\sigma \in \arg \min_{\sigma' \in \{-1, 1\}} \|\{x \in \Gamma_{\mathbf{c}, \circ} \mid f_1(x) \neq f_1^{(\mathbf{c}, \sigma')}(x)\}\|_{\circ}$.

Denote $f'_1 = f_1^{(\mathbf{c}, \sigma)}$. By Assumption 10, $f'_1 \in \mathcal{F}$. Construct $g' = [f'_1, f_2, \dots, f_d] \in \mathcal{G}$, and note that g' is equivalent to g on all but one coordinate, and the one differing coordinate differs only on one cluster, for which this coordinate preserves the cluster in g' but does not preserve it in g .

We show that $\mathcal{L}(g) - \mathcal{L}(g') > 0$, where

$$\mathcal{L}(g) - \mathcal{L}(g') = \mathcal{L}^+(g) - \mathcal{L}^+(g') + \mathcal{L}^-(g) - \mathcal{L}^-(g'). \quad (30)$$

where

$$\begin{aligned} \mathcal{L}^+(g) &:= -\beta \mathbb{E}_{x, x^+} \left[g(x)^{\top} g(x^+) \right] \\ \mathcal{L}^-(g) &:= \mathbb{E}_{x, x^+, \{x_i^-\}_{i=1}^{\ell}} \left[\log \left(e^{\beta g(x)^{\top} g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^{\top} g(x_i^-)} \right) \right] \end{aligned}$$

where \mathcal{L}^+ and \mathcal{L}^- respectively correspond to the alignment and uniformity losses in 3. We will refer to these losses as the positive and negative losses, respectively. We show $\mathcal{L}(g) - \mathcal{L}(g') > 0$ by showing

$$\mathcal{L}^+(g) - \mathcal{L}^+(g') > \mathcal{L}^-(g') - \mathcal{L}^-(g) \quad (31)$$

by first computing the LHS explicitly, then upper bounding the RHS. To do this, we define a partitioning of the set of images D as follows. For all $v \in \mathcal{H}_d$, define the sets $Q_v := \{x \in D : g(x) = v, g'(x) = v\}$ and $E_v := \{x \in D : g(x) = v, g'(x) \neq v\}$. In other words $Q_v \cup E_v$ is the set of images that g maps to v , and Q_v is the subset of this set which g' also maps to v , while E_v is the subset which g' does not map to v . Let $Q := \cup_{v \in \mathcal{H}_d} Q_v$ and $E := \cup_{v \in \mathcal{H}_d} E_v$. Observe that $Q \cup E = D$, and each of the Q_v and E_v 's are disjoint, so they form a partition of D . Further note that $\|E\|_o = \Delta_{f_1, \mathbf{c}} := \min_{\sigma' \in \{-1, 1\}} \|\{x \in \Gamma_{\mathbf{c}, o} \mid f_1(x) \neq f_1^{(\mathbf{c}, \sigma')}(x)\}\|_o$, as defined in Definition 12.

Now we consider the difference in the positive losses. For all augmentations $x^+ \in D \setminus D_o$, define $\tilde{A}^{-1}(x^+)$ as the natural image from which the augmentation was derived, i.e. $\tilde{A}^{-1}(x^+) = x \iff \mathcal{A}(x) = x^+$ for some $\mathcal{A} \in \Lambda$. Moreover, define $R := \{x^+ \in E \setminus D_o : \tilde{A}^{-1}(x^+) \in Q\} \cup \{x^+ \in Q \setminus D_o : \tilde{A}^{-1}(x^+) \in E\} = \{x^+ \in \Gamma_{\mathbf{c}} : f_1(\tilde{A}^{-1}(x^+)) \neq f_1(x^+)\}$ as the set of augmentations in $\Gamma_{\mathbf{c}}$ that f_1 classifies incorrectly. For any event B , let χB denote the indicator random variable for B , i.e. $\chi B = 1$ if B occurs and $\chi B = 0$ otherwise. Using this notation and the construction of g' we can write the difference in positive losses as:

$$\begin{aligned} \mathcal{L}^+(g) - \mathcal{L}^+(g') &= \beta \mathbb{E}_{x, x^+} \left[g'(x)^\top g'(x^+) - g(x)^\top g(x^+) \right] \\ &= \beta \mathbb{E}_{x, x^+} \left[f'_1(x) f'_1(x^+) - f_1(x) f_1(x^+) \right] \\ &= \beta \sum_{v \in \mathcal{H}_d} \mathbb{E}_{x, x^+} \left[\chi \{x \in Q_v\} (f'_1(x) f'_1(x^+) - f_1(x) f_1(x^+)) \right. \\ &\quad \left. + \chi \{x \in E_v\} (f'_1(x) f'_1(x^+) - f_1(x) f_1(x^+)) \right] \\ &= 2\beta \sum_{v \in \mathcal{H}_d} \mathbb{E}_{x, x^+} \left[\chi \{x \in Q_v\} \chi \{x^+ \in E\} + \chi \{x \in E_v\} \chi \{x^+ \in Q\} \right] \\ &= 2\beta \sum_{v \in \mathcal{H}_d} \left(\mathbb{P}_{x \sim D_o, x^+ \sim A(x)} [x \in Q_v, x^+ \in E] \right. \\ &\quad \left. + \mathbb{P}_{x \sim D_o, x^+ \sim A(x)} [x \in E_v, x^+ \in Q] \right) \\ &= 2\beta \|R\| \end{aligned}$$

where in the last equality we have used that all augmentation sets are of equal size.

Now we consider the negative losses. We first decompose the negative loss as

$$\begin{aligned} \mathcal{L}^-(g) &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \\ &= \sum_{v \in \mathcal{H}_d} \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in Q_v\} \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \\ &\quad + \sum_{v \in \mathcal{H}_d} \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v\} \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \end{aligned}$$

$$= \sum_{v \in \mathcal{H}_d} \mathcal{L}_{Q_v}^-(g) + \mathcal{L}_{E_v}^-(g)$$

where $\mathcal{L}_{Q_v}^-(g) := \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \log \left(e^{\beta g(x)^\top g(x^+) + \sum_{i=1}^\ell e^{\beta g(x)^\top g(x_i^-)} \right) \right]$ and $\mathcal{L}_{E_v}^-(g) := \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \log \left(e^{\beta g(x)^\top g(x^+) + \sum_{i=1}^\ell e^{\beta g(x)^\top g(x_i^-)} \right) \right]$. Note that we need to upper bound

$$\mathcal{L}^-(g') - \mathcal{L}^-(g) = \sum_{v \in \mathcal{H}_d} \mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g) + \mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g). \quad (32)$$

We analyze $\mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g)$ and $\mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g)$ separately for every $v \in \mathcal{H}_d$. To do so, we define additional notations. For a batch of negative samples $\{x_i^-\}_{i=1}^\ell$ and a vertex $v \in \mathcal{H}_d$, let $n_{1,v} := \sum_{i=1}^\ell \chi\{x_i^- \in Q_v\}$, $n_2 := \sum_{i=1}^\ell \chi\{x_i^- \in E\}$, $n_{3,v} := \sum_{i=1}^\ell \chi\{x_i^- \in E_v\}$ and $n_4 := \sum_{i=1}^\ell \chi\{x_i^- \in Q\}$. Using the fact that $g(x)^\top g(x_i^-) = d$ for all $x, x_i^- \in Q_v$, we have

$$\begin{aligned} \mathcal{L}_{Q_v}^-(g) &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \log \left(e^{\beta g(x)^\top g(x^+) + n_{1,v} e^{\beta d} + \sum_{x_i^- \notin Q_v} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \\ &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \right. \\ &\quad \left. \log \left(e^{\beta g(x)^\top g(x^+) + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \end{aligned}$$

Next, using the fact that $g'(x) = g(x)$ for all $x \notin E$, we have

$$\begin{aligned} \mathcal{L}_{Q_v}^-(g') &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \log \left(e^{\beta g'(x)^\top g'(x^+) + n_{1,v} e^{\beta d} + \sum_{x_i^- \notin Q_v} e^{\beta g'(x)^\top g'(x_i^-)} \right) \right] \\ &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \right. \\ &\quad \left. \log \left(e^{\beta g'(x)^\top g'(x^+) + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \end{aligned}$$

Before analyzing $\mathcal{L}_{E_v}^-$, we first prove the following claims.

Claim B.1 For all $x, x^- \in E_v$, $g'(x) = g'(x^-)$.

Proof By construction of g' , $g'(x)$ agrees with $g(x)$ on all but the first coordinate. Thus, $g(x) \neq g'(x) \implies f_1(x) = -f_1'(x)$. Consider any $x, x^- \in E_v$. By definition of E_v , $g(x) = g(x^-) = v$, and $g'(x) \neq g(x)$. Let $v = [v_1, v_2, \dots, v_d]$, then we have $g'(x) = g'(x^-) = [-v_1, v_2, \dots, v_d]$. ■

Claim B.2 For all $v \in \mathcal{H}_d$ and all $x \in E_v$, $x^- \in E$, $g'(x)^\top g'(x^-) = g(x)^\top g(x^-)$.

Proof Consider any $x \in E_v, x^- \in E$. As above, observe that the j -th coordinates of $g(x)$ and $g'(x)$ are the same for all $j > 1$ (and likewise for $g(x^-)$ and $g'(x^-)$) by construction of g . Moreover, $f_1(x) = -f'_1(x)$ and $f_1(x^-) = -f'_1(x^-)$ by definition of E . Thus $g'(x)^\top g'(x^-) - g(x)^\top g(x^-) = f'_1(x)f'_1(x^-) - (-f'_1(x))(-f'_1(x^-)) = 0$. ■

Claim B.3 For all $v \in \mathcal{H}_d$ and all $x \in E_v, x^- \in Q$, $g'(x)^\top g'(x^-) \neq g(x)^\top g(x^-)$.

Proof By definition of $E_v, g'(x) \neq g(x)$ for all $x \in E_v$, and by definition of $Q, g'(x^-) = g(x^-)$ for all $x^- \in Q$. Thus $g'(x)^\top g'(x^-) \neq g(x)^\top g(x^-)$. ■

Next we can decompose $\mathcal{L}_{E_v}^-(g)$ as follows, using the fact that $g(x)^\top g(x_i^-) = d$ for all $x, x_i^- \in E_v$.

$$\begin{aligned} \mathcal{L}_{E_v}^-(g) &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \notin E_v} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \\ &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \right. \\ &\quad \left. \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \end{aligned}$$

Next we use Claims B.1, B.2 and B.3 to obtain

$$\begin{aligned} \mathcal{L}_{E_v}^-(g') &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \notin E_v} e^{\beta g'(x)^\top g'(x_i^-)} \right) \right] \\ &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \right. \\ &\quad \left. \log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \end{aligned}$$

Next, define $Z := \{x \in Q : f'_1(x) \neq f'_1(x^-), \forall x^- \in E\} \subseteq Q$. as the set of images that f'_1 labels differently than it does the samples in E (note that $f'_1(x) = f'_1(x^-)$ for all $x, x^- \in E$, and $f'_1(x) = f'_1(x^-)$ for all $x, x^- \in Q_v$, so the \forall condition in the definition of Z could be replaced with 'for some'). Also note that the definition of Z here differs slightly from the one in the proof sketch in Section 5 for ease of notation.

Next we prove two claims regarding properties of Z .

Claim B.4 For all $v \in \mathcal{H}_d$ exactly one of the following holds: (i) $Q_v \cap Z = Q_v$ or (ii) $Q_v \setminus Z = Q_v$.

Proof Suppose $x \in Q_v \cap Z$. Then, for all $x' \in Q_v, f'_1(x') = f'_1(x)$ by definition of Q_v . Thus $f'_1(x') = f'_1(x) \neq f'_1(x'')$ for any $x'' \in E$ since $x \in Z$. This implies $x' \in Z$, therefore $Q_v \cap Z = Q_v$.

Likewise, suppose $x \in Q_v \setminus Z$. Then, for all $x' \in Q_v, f'_1(x') = f'_1(x)$ by definition of Q_v . Thus $f'_1(x') = f'_1(x) = f'_1(x'')$ for any $x'' \in E$ since $x \notin Z$. This implies $x' \notin Z$, therefore $Q_v \setminus Z = Q_v$. ■

Claim B.5 For all $v \in \mathcal{H}_d$, $x \in Q_v \cap Z$, $x^- \in E$, $g'(x)^\top g'(x^-) = g(x)^\top g(x^-) - 2$.

Proof Suppose $x \in Q_v \cap Z$. For any $x^- \in E$, then $f_1(x_i^-) = -f'_1(x_i^-)$ by definition of E , $f_1(x) = f'_1(x)$ by definition of Q_v and $f'_1(x) = -f'_1(x_i^-)$ by definition of Z . Therefore, $f_1(x) = f_1(x_i^-)$ and $g'(x)^\top g'(x^-) = g(x)^\top g(x^-) - 2$, noting that g and g' agree on all but the first coordinate. \blacksquare

Claim B.6 For all $v \in \mathcal{H}_d$, $Q_v \cap Z = \emptyset$ and $Q_v \neq \emptyset \implies E_v = \emptyset$.

Proof From Claim B.4, $Q_v \cap Z = \emptyset \implies Q_v \setminus Z = Q_v$. Suppose $x \in Q_v \setminus Z$ and $x' \in E_v$. Then $f_1(x) = f_1(x')$ by definition of Q_v and E_v . Also, $f'_1(x) = f'_1(x')$ since $x \notin Z$ and $x' \in E$, and $f_1(x') \neq f'_1(x')$ by definition of E . Therefore $f_1(x) \neq f'_1(x)$, but this contradicts the definition of Q_v . \blacksquare

We use Claim B.6 later in the proof. For now we use Claims B.4 and B.5 to bound $\mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g)$ for all v such that $Q_v \subseteq Z$:

$$\begin{aligned}
 & \mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g) \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \right. \\
 & \quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v\} \right. \\
 & \quad \left(\log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + e^{-2\beta} \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \quad (33)
 \end{aligned}$$

$$\leq 0 \quad (34)$$

Thus, we have

$$\begin{aligned}
 \mathcal{L}^-(g') - \mathcal{L}^-(g) &= \sum_{v \in \mathcal{H}_d} \mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g) + \mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g) \\
 &\leq \sum_{v \in \mathcal{H}_d: Q_v \setminus Z = Q_v} \mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g) + \sum_{v \in \mathcal{H}_d} \mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g) \quad (35)
 \end{aligned}$$

For each $v \in \mathcal{H}_d : Q_v \setminus Z = Q_v$, we consider three cases: (1) $n_2 = 0$, (2) $n_2 > 0, n_{1,v} = 0$, and (3) $n_2 > 0, n_{1,v} > 0$. In particular we decompose $\mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g)$ as follows:

$$\mathcal{L}_{Q_v}^-(g') - \mathcal{L}_{Q_v}^-(g)$$

$$\begin{aligned}
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z\} \chi\{n_2 = 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \\
 &+ \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z\} \chi\{n_2 > 0, n_{1,v} = 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \\
 &+ \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z\} \chi\{n_2 > 0, n_{1,v} > 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right]
 \end{aligned}$$

Likewise, for each $v \in \mathcal{H}_d$, we decompose $\mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g)$ as:

$$\begin{aligned}
 &\mathcal{L}_{E_v}^-(g') - \mathcal{L}_{E_v}^-(g) \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \chi\{n_4 = 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \\
 &+ \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \chi\{n_4 > 0, n_{3,v} = 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \\
 &+ \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \chi\{n_4 > 0, n_{3,v} > 0\} \right.
 \end{aligned}$$

$$\left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right)$$

Thus, for each v in (32), we need to upper bound a total of six terms. We consider each of these six terms individually, starting with the three terms with $\chi\{x \in Q_v \setminus Z\}$ factors.

1. $x \in Q_v \setminus Z, n_2 = 0$.

In this case, we have

$$\begin{aligned} ((1)) &:= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z\} \chi\{n_2 = 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\ &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \end{aligned} \quad (36)$$

$$\begin{aligned} &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z\} \chi\{n_2 = 0\} \right. \\ &\quad \left. \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} \right) - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} \right) \right) \right] \end{aligned} \quad (37)$$

$$= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0\} \log \left(\frac{e^{2\beta} e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d}}{e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d}} \right) \right] \quad (38)$$

$$\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0\} \log \left(\frac{e^{\beta d} + n_{1,v} e^{\beta d}}{e^{\beta(d-2)} + n_{1,v} e^{\beta d}} \right) \right] \quad (39)$$

$$\begin{aligned} &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0\} \log \left(\frac{1 + n_{1,v}}{e^{-2\beta} + n_{1,v}} \right) \right] \\ &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} = 0\} \log \left(\frac{1 + n_{1,v}}{e^{-2\beta} + n_{1,v}} \right) \right] \\ &\quad + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} > 0\} \log \left(\frac{1 + n_{1,v}}{e^{-2\beta} + n_{1,v}} \right) \right] \\ &\leq 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} = 0\} \right] \\ &\quad + \log(2) \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} > 0\} \right] \end{aligned} \quad (40)$$

where (37) follows by the submodularity of the $\log(\cdot)$ function and the fact that $g'(x)^\top g'(x^+) \geq g(x)^\top g(x^+)$ by construction of g' , (54) follows by the fact that if $x \in Q_v$, then $g'(x)^\top g'(x^+) = g(x)^\top g(x^+)$ for all $x^+ \notin B$, and (39) follows since $h(x) := \frac{ax+c}{x+c}$ is monotonically increasing

for $a > 1$. Next, by the independence of x_i^- from x and x^+ ,

$$\begin{aligned} & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} = 0\}] \\ &= \mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) \mathbb{P}(n_{1,v} = 0, n_2 = 0) \\ &= \mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) (1 - \|Q_v \setminus Z\|_o - \|E\|_o)^\ell \end{aligned} \quad (41)$$

Similarly, for the second term in (40), we have

$$\begin{aligned} & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in Q_v \setminus Z, x^+ \in E, n_2 = 0, n_{1,v} > 0\}] \\ &= \mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) \mathbb{P}(n_{1,v} > 0 | n_2 = 0) \mathbb{P}(n_2 = 0) \end{aligned} \quad (42)$$

$$\leq \mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) \min\left(1, \frac{\ell \|Q_v \setminus Z\|_o}{1 - \|E\|_o}\right) (1 - \|E\|_o)^\ell \quad (43)$$

By combining (43), (41), and (40), we obtain the following upper bound on (36):

$$\begin{aligned} ((1)) &\leq \mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) \\ &\quad \left(2\beta(1 - \|Q_v \setminus Z\|_o - \|E\|_o)^\ell + \log(2) \min\left(1, \frac{\ell \|Q_v \setminus Z\|_o}{1 - \|E\|_o}\right) (1 - \|E\|_o)^\ell\right). \end{aligned} \quad (44)$$

2. $x \in Q_v \setminus Z, n_{1,v} = 0, n_2 > 0$

In this case we have:

$$\begin{aligned} ((2)) &:= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} = 0, n_2 > 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\ &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \end{aligned} \quad (45)$$

$$\begin{aligned} &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} = 0, n_2 > 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\ &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \end{aligned} \quad (46)$$

$$\begin{aligned} &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} = 0, n_2 > 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} \right) \right) \end{aligned}$$

$$-\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} \right) \Big] \quad (47)$$

where (47) follows by the submodularity of the $\log(\cdot)$ function and the facts that $g'(x)^\top g'(x^+) \geq g(x)^\top g(x^+)$ and $g'(x)^\top g'(x_i^-) \geq g(x)^\top g(x_i^-)$ for all $x \in Q_v \setminus Z, x_i^- \in E$ by definition of g', C and E . Next,

$$((2)) \leq 2\beta \mathbb{P}(x \in Q_v \setminus Z) \mathbb{E}_{\{x_i^-\}_\ell} [\chi\{n_{1,v} = 0\} \chi\{n_2 > 0\}] \quad (48)$$

$$\begin{aligned} &= 2\beta \mathbb{P}(x \in Q_v \setminus Z) \mathbb{E}_{\{x_i^-\}_\ell} [\chi\{\cap_i \{x_i^- \notin Q_v \setminus Z\}\} \chi\{\cup_i \{x_i^- \in E\}\}] \\ &= 2\beta \mathbb{P}(x \in Q_v \setminus Z) \mathbb{P}(\cup_i \{x_i^- \in E\} | \cap_i \{x_i^- \notin Q_v \setminus Z\}) \mathbb{P}(\cap_i \{x_i^- \notin Q_v \setminus Z\}) \\ &\leq 2\beta \mathbb{P}(x \in Q_v \setminus Z) \mathbb{P}(\cap_i \{x_i^- \notin Q_v \setminus Z\}) \min \left(1, \sum_{i=1}^{\ell} \mathbb{P}(x_i^- \in E | x_i^- \notin Q_v \setminus Z) \right) \end{aligned} \quad (49)$$

$$= 2\beta \|Q_v \setminus Z\|_0 (1 - \|Q_v \setminus Z\|_0)^\ell \min \left(1, \frac{\ell \|E\|_0}{1 - \|Q_v \setminus Z\|_0} \right)$$

where (48) follows since $g'(x)^\top g'(x^+) - g(x)^\top g(x^+) \leq 2$ and $g'(x)^\top g'(x_i^-) - g(x)^\top g(x_i^-) \leq 2$ for all x, x_i^- , and (49) follows by a union bound.

3. $x \in Q_v \setminus Z, n_1 > 0, n_2 > 0$.

In this case, we have

$$\begin{aligned} ((3)) &:= \mathbb{E} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\ &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin Q_v \cup E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \end{aligned} \quad (50)$$

$$\begin{aligned} &\leq \mathbb{E} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0\} \right. \\ &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{1,v} e^{\beta d} + e^{2\beta} \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\ &\quad \left. \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v} e^{\beta d} + \sum_{x_i^- \in E} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right] \end{aligned} \quad (51)$$

where (51) follows by the submodularity of $\log(\cdot)$ and the facts that $g'(x)^\top g'(x^+) \geq g(x)^\top g(x^+)$ and $g'(x)^\top g'(x_i^-) = g(x)^\top g(x_i^-) + 2$ for all $x \in Q_v \setminus Z, x_i^- \in E$ by definition of g', Z and E . Continuing, we obtain

$$((3)) \leq \mathbb{E} \left[\chi\{x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0\} \right.$$

$$\begin{aligned}
 & \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + (n_{1,v} + n_2)e^{\beta d} \right) \right. \\
 & \quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v}e^{\beta d} + n_2e^{\beta(d-2)} \right) \right) \Bigg] \quad (52) \\
 & = \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \notin E \} \right. \\
 & \quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + (n_{1,v} + n_2)e^{\beta d} \right) \right. \\
 & \quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v}e^{\beta d} + n_2e^{\beta(d-2)} \right) \right) \Bigg] \\
 & + \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in E \} \right. \\
 & \quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + (n_{1,v} + n_2)e^{\beta d} \right) \right. \\
 & \quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{1,v}e^{\beta d} + n_2e^{\beta(d-2)} \right) \right) \Bigg] \quad (53)
 \end{aligned}$$

where (52) follows since $h(x) := \frac{a+cx}{b+x}$ is an increasing function of x for $x > 0, c > \frac{a}{b}$ (here, $\frac{a}{b} \leq \frac{2}{1+e^{-2\beta}} \leq 2$ and $c = e^{2\beta} > 2$). Note that $x \in Q_v, x^+ \in E \implies g'(x)^\top g'(x^+) = g(x)^\top g(x^+) + 2$, and $x \in Q_v, x^+ \notin E \implies g'(x)^\top g'(x^+) = g(x)^\top g(x^+)$. Using this we find

$$\begin{aligned}
 ((3)) & \leq \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \notin E \} \log \left(\frac{n_{1,v} + n_2}{n_{1,v}} \right) \right] \\
 & \quad + \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in B \} \log \left(\frac{n_{1,v} + n_2 + 1}{n_{1,v}} \right) \right] \\
 & \leq \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \notin B \} \frac{2n_2}{n_{1,v} + 1} \right] \\
 & \quad + \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in E \} \frac{2n_2 + 2}{n_{1,v} + 1} \right] \quad (54) \\
 & = \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0 \} \frac{2n_2}{n_{1,v} + 1} \right] \\
 & \quad + \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in B \} \frac{2}{n_{1,v} + 1} \right] \\
 & = \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0 \} \frac{2n_2}{n_{1,v} + 1} \right] \\
 & \quad + \mathbb{E} \left[\chi \{ x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in E \} \frac{2}{n_{1,v} + 1} \right] \quad (55)
 \end{aligned}$$

where (54) follows using the inequality $\log(1+x) \leq x$. Thus we are left with two terms in (55). For the first term we have (ignoring notation overload, as after the first line, $n_{1,v}$ and n_2

change from random variables to dummy variables):

$$\begin{aligned} & \mathbb{E} \left[\chi \{x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, Q_v \cap E = \emptyset\} \frac{2n_2}{n_{1,v} + 1} \right] \\ &= \mathbb{P}(x \in Q_v \setminus Z) \sum_{n_{1,v}, n_2: n_{1,v} + n_2 \leq \ell} \chi \{n_{1,v} > 0, n_2 > 0\} \binom{\ell}{n_{1,v} \ n_2 \ \ell - n_{1,v} - n_2} \\ & \quad \|Q_v \setminus Z\|_{\circ}^{n_{1,v}} \|E\|_{\circ}^{n_2} (1 - \|Q_v \setminus Z\|_{\circ} - \|E\|_{\circ})^{\ell - n_{1,v} - n_2} \frac{2n_2}{n_{1,v} + 1} \end{aligned} \quad (56)$$

$$\begin{aligned} &= 2 \|Q_v \setminus Z\|_{\circ} \sum_{n_{1,v}, n_2: n_{1,v}, n_2 > 0, n_{1,v} + n_2 \leq \ell} \binom{\ell}{n_{1,v} \ n_2 \ \ell - n_{1,v} - n_2} \\ & \quad \|Q_v \setminus Z\|_{\circ}^{n_{1,v}} \|E\|_{\circ}^{n_2} (1 - \|Q_v \setminus Z\|_{\circ} - \|E\|_{\circ})^{\ell - n_{1,v} - n_2} \frac{n_2}{n_{1,v} + 1} \end{aligned} \quad (57)$$

$$\begin{aligned} &= 2 \|Q_v \setminus Z\|_{\circ} \frac{\|E\|_{\circ}}{\|Q_v \setminus Z\|_{\circ}} \sum_{n_{1,v}, n_2: n_{1,v}, n_2 > 0, n_{1,v} + n_2 \leq \ell} \binom{\ell}{n_{1,v} + 1 \ n_2 - 1 \ \ell - n_{1,v} - n_2} \\ & \quad \|Q_v \setminus Z\|_{\circ}^{n_{1,v} + 1} \|E\|_{\circ}^{n_2 - 1} (1 - \|Q_v \setminus Z\|_{\circ} - \|E\|_{\circ})^{\ell - n_{1,v} - n_2} \end{aligned} \quad (58)$$

$$\begin{aligned} &\leq 2 \|Q_v \setminus Z\|_{\circ} \frac{\|E\|_{\circ}}{\|Q_v \setminus Z\|_{\circ}} \\ & \quad \sum_{n_{1,v} + 1, n_2 - 1: n_{1,v} + 1, n_2 - 1 \geq 0, n_{1,v} + 1 + n_2 - 1 \leq \ell} \binom{\ell}{n_{1,v} + 1 \ n_2 - 1 \ \ell - (n_{1,v} + 1) - (n_2 - 1)} \\ & \quad \|Q_v \setminus Z\|_{\circ}^{n_{1,v} + 1} \|E\|_{\circ}^{n_2 - 1} (1 - \|Q_v \setminus Z\|_{\circ} - \|E\|_{\circ})^{\ell - (n_{1,v} + 1) - (n_2 - 1)} \end{aligned} \quad (59)$$

$$\begin{aligned} &= 2 \|E\|_{\circ} (\|Q_v \setminus Z\|_{\circ} + \|E\|_{\circ} + 1 - \|Q_v \setminus Z\|_{\circ} - \|E\|_{\circ})^{\ell} \\ &= 2 \|E\|_{\circ} \end{aligned}$$

where in (59) we have added terms to the sum to complete the trinomial expansion, and the last equality follows since Q_v and E are disjoint.

Now we need to consider the last term in (55), which corresponds to the case wherein the positive inner products are not equal for g and g' . For this term, we simply have

$$\begin{aligned} & \mathbb{E} \left[\chi \{x \in Q_v \setminus Z, n_{1,v} > 0, n_2 > 0, x^+ \in E\} \frac{2}{n_{1,v} + 1} \right] \\ & \leq \mathbb{P} [x \in Q_v \setminus Z, x^+ \in E] \end{aligned} \quad (60)$$

In total, for the case $n_{1,v} > 0, n_2 > 0$ and $x \in Q_v \setminus Z$, we have

$$((3)) \leq 2 \|E\| + \mathbb{P} [x \in Q_v \setminus Z, x^+ \in E] \quad (61)$$

4. $x \in E_v, n_4 = 0$.

This case is symmetric to Case 1, so we argue similarly.

$$((4)) := \mathbb{E}_{x, x^+, \{x_i^-\}} \left[\chi \{x \in E_v\} \chi \{n_4 = 0\} \right]$$

$$\begin{aligned}
 & \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \chi\{n_4 = 0\} \right. \\
 & \quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \left. \right] \\
 &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v\} \chi\{n_4 = 0\} \log \left(\frac{e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d}}{e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d}} \right) \right] \tag{62}
 \end{aligned}$$

where (62) follows since $g'(x)^\top g'(x^+) \geq g(x)^\top g(x^+)$ and $\log(\cdot)$ is submodular. Next we intersect with the events $\{x^+ \in Q\}$ and $\{x^+ \notin Q\}$, obtaining

$$\begin{aligned}
 ((4)) &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 = 0, x^+ \in Q\} \log \left(\frac{e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d}}{e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d}} \right) \right] \\
 &\quad + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 = 0, x^+ \notin Q\} \log \left(\frac{e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d}}{e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d}} \right) \right] \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 = 0, x^+ \in Q\} \log \left(\frac{e^{2\beta} e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d}}{e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d}} \right) \right] \\
 &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 = 0, x^+ \in Q\} \log \left(\frac{e^{\beta d} + n_{3,v} e^{\beta d}}{e^{\beta(d-2)} + n_{3,v} e^{\beta d}} \right) \right] \\
 &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 = 0, x^+ \in Q\} \log \left(\frac{1 + n_{3,v}}{e^{-2\beta} + n_{3,v}} \right) \right] \\
 &\leq 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in E_v, n_4 = 0, x^+ \in Q, n_{3,v} = 0\}] \\
 &\quad + \log(2) \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in E_v, n_4 = 0, x^+ \in Q, n_{3,v} > 0\}] \tag{63}
 \end{aligned}$$

where

$$\begin{aligned}
 &\mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in E_v, n_4 = 0, x^+ \in Q, n_{3,v} = 0\}] \\
 &= \mathbb{P}(x \in E_v \cap x^+ \in Q) \mathbb{P}(n_{3,v} = 0, n_4 = 0) \tag{64}
 \end{aligned}$$

$$= \mathbb{P}(x \in E_v \cap x^+ \in Q) (1 - \|E_v\|_o - \|Q\|_o)^\ell \tag{65}$$

$$\begin{aligned}
 &\mathbb{E}_{x, x^+, \{x_i^-\}_\ell} [\chi\{x \in E_v, n_4 = 0, x^+ \in Q, n_{3,v} > 0\}] \\
 &= \mathbb{P}(x \in E_v \cap x^+ \in Q) \mathbb{P}(n_{3,v} > 0, n_4 = 0) \\
 &\leq \mathbb{P}(x \in E_v \cap x^+ \in Q) \min \left(1, \frac{\ell \|E_v\|_o}{1 - \|Q\|_o} \right) (1 - \|Q\|_o)^\ell \tag{66}
 \end{aligned}$$

so in total for this case we have

$$((4)) \leq \mathbb{P}(x \in E_v \cap x^+ \in Q) \left(2\beta(1 - \|E_v\|_0 - \|Q\|_0)^\ell + \log(2) \min(1, \frac{\ell \|E_v\|_0}{1 - \|Q\|_0}) (1 - \|Q\|_0)^\ell \right)$$

5. $x \in E_v, n_4 > 0, n_{3,v} = 0$.

Define $n_5 := \sum_{i=1}^\ell \chi\{x_i^- \in Q \setminus Z\}$. We have

$$((5)) = \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 > 0, n_{3,v} = 0\} \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) - \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right]$$

since we are intersecting with the event $\{n_{3,v} = 0\}$. Next we split the negative samples in Q into those in $Q \setminus Z$ and those in Z , noting that $g'(x)^\top g'(x_i^-) = g(x)^\top g(x_i^-) - 2$ for $x_i^- \in Z$ and $g'(x)^\top g'(x_i^-) = g(x)^\top g(x_i^-) + 2$ for $x_i^- \in Q \setminus Z$.

$$((5)) = \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 > 0, n_{3,v} = 0\} \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + e^{2\beta} \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + e^{-2\beta} \sum_{x_i^- \in C} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) - \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \in C} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right]$$

$$\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 > 0, n_{3,v} = 0\} \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + e^{2\beta} \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) - \log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \right]$$

$$\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi\{x \in E_v, n_4 > 0, n_{3,v} = 0\} \right]$$

$$\begin{aligned}
 & \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_5 e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_5 e^{\beta(d-2)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \Bigg] \quad (67)
 \end{aligned}$$

where (67) follows since $h(x) := \frac{a+e^{2\beta}x}{b+x}$ is an increasing function of x for $a \leq e^{2\beta}b$. Next we intersect with $\{x^+ \in Q\}$ and $\{x^+ \notin Q\}$ to obtain

$$\begin{aligned}
 ((5)) & \leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \in Q\} \right. \\
 & \quad \left(\log \left((n_5 + 1)e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. - \log \left((n_5 + 1)e^{\beta(d-2)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \Bigg] \\
 & + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \notin Q\} \right. \\
 & \quad \left(\log \left(n_5 e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. - \log \left(n_5 e^{\beta(d-2)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \Bigg] \quad (68)
 \end{aligned}$$

We have two terms above. For the first term,

$$\begin{aligned}
 & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \in Q\} \right. \\
 & \quad \left(\log \left((n_5 + 1)e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 & \quad \left. - \log \left((n_5 + 1)e^{\beta(d-2)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \Bigg] \\
 & \leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \in Q\} \left(\log \left(\frac{(n_5 + 1)e^{\beta d}}{(n_5 + 1)e^{\beta(d-2)}} \right) \right) \right] \quad (69) \\
 & \leq 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \in Q\} \right]
 \end{aligned}$$

Similarly, for the second term in (68), we have

$$\begin{aligned}
 & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \notin Q\} \right. \\
 & \quad \left(\log \left(n_5 e^{\beta d} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) - \log \left(n_5 e^{\beta(d-2)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \Bigg]
 \end{aligned}$$

$$\begin{aligned}
 &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_5 > 0, n_{3,v} = 0, x^+ \notin Q\} \log \left(\frac{n_5 e^{\beta d}}{n_5 e^{\beta(d-2)}} \right) \right] \\
 &= 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_5 > 0, n_{3,v} = 0, x^+ \notin Q\} \right]
 \end{aligned}$$

By summing the upper bounds on the two terms in (68), we obtain

$$\begin{aligned}
 ((5)) &\leq 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} = 0, x^+ \in Q\} \right] \\
 &\quad + 2\beta \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_5 > 0, n_{3,v} = 0, x^+ \notin Q\} \right] \\
 &\leq 2\beta \mathbb{P}(x \in E_v) \mathbb{P}(n_4 > 0, n_{3,v} = 0) \\
 &= 2\beta \|E_v\|_\circ (1 - \|E_v\|_\circ)^\ell \min \left(1, \frac{\ell \|Q\|_\circ}{1 - \|E_v\|_\circ} \right). \tag{70}
 \end{aligned}$$

6. $x \in E_v, n_4 > 0, n_{3,v} > 0$.

In this case, we argue similarly as in Case 3 to obtain

$$\begin{aligned}
 ((6)) &= \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g'(x)^\top g'(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \left. \right] \\
 &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + n_{3,v} e^{\beta d} + e^{2\beta} \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right. \\
 &\quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + \sum_{x_i^- \in Q \setminus Z} e^{\beta g(x)^\top g(x_i^-)} + \sum_{x_i^- \notin E_v \cup Q} e^{\beta g(x)^\top g(x_i^-)} \right) \right) \left. \right] \\
 &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + (n_{3,v} + n_5) e^{\beta d} \right) \right. \\
 &\quad \left. - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + n_5 e^{\beta(d-2)} \right) \right) \left. \right] \tag{71}
 \end{aligned}$$

where (71) follows by the analogous argument as in (52). Next we intersect with $\{x^+ \in Q\}$ and $\{x^+ \notin Q\}$. We have

$$\begin{aligned}
 ((6)) &\leq \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0\} \right. \\
 &\quad \left(\log \left(e^{\beta g'(x)^\top g'(x^+)} + (n_{3,v} + n_5) e^{\beta d} \right) \right.
 \end{aligned}$$

$$\begin{aligned}
 & - \log \left(e^{\beta g(x)^\top g(x^+)} + n_{3,v} e^{\beta d} + n_5 e^{\beta(d-2)} \right) \Big] \\
 \leq & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0, x^+ \in Q\} \log \left(\frac{n_{3,v} + n_5 + 1}{n_{3,v}} \right) \right] \\
 & + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_4 > 0, n_{3,v} > 0, x^+ \notin Q\} \log \left(\frac{n_{3,v} + n_5}{n_{3,v}} \right) \right] \quad (72) \\
 \leq & \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_{3,v} > 0, x^+ \in Q\} \log(n_5 + 2) \right] \\
 & + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{x \in E_v, n_{3,v} > 0, x^+ \notin Q\} \log(n_5 + 1) \right] \\
 = & \mathbb{P}(x \in E_v, x^+ \in Q) \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{n_{3,v} > 0\} \log(n_5 + 2) \right] \\
 & + \mathbb{P}(x \in E_v, x^+ \notin Q) \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\chi \{n_{3,v} > 0\} \log(n_5 + 1) \right] \\
 \leq & \mathbb{P}(x \in E_v, x^+ \in Q) \log(\mathbb{E}[n_5 + 2]) + \mathbb{P}(x \in E_v, x^+ \notin Q) \log(\mathbb{E}[n_5 + 1]) \quad (73) \\
 \leq & \mathbb{P}(x \in E_v, x^+ \in Q) \log(\ell \|Q \setminus Z\|_\circ + 2) + \|E_v\|_\circ \log(\ell \|Q \setminus Z\|_\circ + 1)
 \end{aligned}$$

where (72) follows since if $x \in E_v$, then $x^+ \in Q \iff g'(x)^\top g'(x^+) = g(x)^\top g(x^+) + 2$ and by the submodularity of $\log(\cdot)$, and (73) follows by Jensen's Inequality and upper bounding $\chi \{n_{3,v} > 0\} \leq 1$.

Now we combine all six cases and sum over $v \in \mathcal{H}_d$. We obtain

$$\begin{aligned}
 & \mathcal{L}^-(g') - \mathcal{L}^-(g) \\
 \leq & \sum_{v \in \mathcal{H}_d} \left(\mathbb{P}(x \in Q_v \setminus Z \cap x^+ \in E) \right. \\
 & \quad \left(2\beta(1 - \|Q_v \setminus Z\|_\circ - \|E\|_\circ)^\ell + \log(2) \min \left(1, \frac{\ell \|Q_v \setminus Z\|_\circ}{1 - \|E\|_\circ} \right) (1 - \|E\|_\circ)^\ell \right) \\
 & + 2\beta \|Q_v \setminus Z\|_\circ (1 - \|Q_v \setminus Z\|_\circ)^\ell \min \left(1, \frac{\ell \|E\|_\circ}{1 - \|Q_v \setminus Z\|_\circ} \right) \\
 & + 2\|E\|_\circ + \mathbb{P}[x \in Q_v \setminus Z, x^+ \in E] \\
 & + \mathbb{P}(x \in E_v \cap x^+ \in Q) (2\beta(1 - \|E_v\|_\circ - \|Q \setminus Z\|_\circ)^\ell \\
 & \quad + \log(2) \min \left(1, \frac{\ell \|E_v\|_\circ}{1 - \|Q \setminus Z\|_\circ} \right) (1 - \|Q \setminus Z\|_\circ)^\ell) \\
 & + 2\beta \|E_v\|_\circ (1 - \|E_v\|_\circ)^\ell \min \left(1, \frac{\ell \|Q \setminus Z\|_\circ}{1 - \|E_v\|_\circ} \right) \\
 & \left. + \mathbb{P}(x \in E_v, x^+ \in Q) \log(\ell \|Q \setminus Z\|_\circ + 2) + \|E_v\|_\circ \log(\ell \|Q \setminus Z\|_\circ + 1) \right) \quad (74)
 \end{aligned}$$

$$\begin{aligned}
 \leq & \|R\| \log(\ell + 2) + \|E\|_\circ (2^{d+1} + \log(\ell + 1)) \\
 & + \sum_{v \in \mathcal{H}_d} \left(\mathbb{P}(x \in Q_v \cap x^+ \in E) \left(2\beta(1 - \|Q_v\|_\circ - \|E\|_\circ)^\ell + \log(2)(1 - \|E\|_\circ)^\ell \right) \right. \\
 & + 2\beta \ell \|E\|_\circ \|Q_v \setminus Z\|_\circ (1 - \|Q_v \setminus Z\|_\circ)^{\ell-1} \\
 & + \mathbb{P}(x \in E_v \cap x^+ \in Q) \left(2\beta(1 - \|E_v\|_\circ - \|Q\|_\circ)^\ell + \log(2)(1 - \|Q\|_\circ)^\ell \right) \\
 & \left. + 2\beta \|E_v\|_\circ (1 - \|E_v\|_\circ)^\ell \right) \quad (75)
 \end{aligned}$$

$$\begin{aligned}
 &\leq \|R\|(\log(\ell+2) + \log(2)) + \|E\|_{\circ}(2^{d+1} + \log(\ell+1)) \\
 &\quad + 2\beta \sum_{v \in \mathcal{H}_d} \left(\mathbb{P}(x \in Q_v \cap x^+ \in E)(\|Q\|_{\circ} - \|Q_v\|_{\circ})^{\ell} + \mathbb{P}(x \in E_v \cap x^+ \in Q)(\|E\|_{\circ} - \|E_v\|_{\circ})^{\ell} \right. \\
 &\quad \left. + \ell \|E\|_{\circ} \|Q_v \setminus Z\|_{\circ} (1 - \|Q_v \setminus Z\|_{\circ})^{\ell-1} + \|E_v\|_{\circ} (1 - \|E_v\|_{\circ})^{\ell} \right) \\
 &\leq \|R\|(\log(\ell+2) + \log(2)) + \|E\|_{\circ}(2^{d+1} + \log(\ell+1)) + 2^{-(\ell-1)}\beta \|R\| \\
 &\quad + 2\beta \sum_{v \in \mathcal{H}_d} \left(\mathbb{P}(x \in Q_v \cap x^+ \in E)(\|Q\|_{\circ} - \|Q_v\|_{\circ})^{\ell} \right. \\
 &\quad \left. + \ell \|E\|_{\circ} \|Q_v \setminus Z\|_{\circ} (1 - \|Q_v \setminus Z\|_{\circ})^{\ell-1} + \|E_v\|_{\circ} (1 - \|E_v\|_{\circ})^{\ell} \right) \tag{76}
 \end{aligned}$$

where (75) follows since $\min(x, y) \leq x$, $\sum_{v \in \mathcal{H}_d} \|E_v\|_{\circ} = \|E\|_{\circ}$, and $\|R\| = \sum_{v \in \mathcal{H}_d} \mathbb{P}(x \in Q_v \cap x^+ \in E_v) + \mathbb{P}(x \in E_v \cap x^+ \in Q)$, and (76) follows since $\|E\|_{\circ} \leq \frac{1}{2}$ by construction of g' (since for $f'_1 = f_1^{(c, \sigma)}$, σ is chosen such that the induced $\|E\|_{\circ}$ cannot be larger than $\frac{1}{2}$). It remains to bound the three terms in the sum in (76).

To do so, we first define $\tilde{E}_v := \{x^+ \in B : \tilde{\mathcal{A}}^{-1}(x^+) \in Q_v\}$ as the set of partial augmentation sets that are in B , corresponding to sets whose natural image is in Q_v (where $\tilde{\mathcal{A}}^{-1}(x^+)$ is the natural image from which the augmented image x^+ was generated, i.e. $\tilde{\mathcal{A}}^{-1}(x^+) = x \iff \mathcal{A}(x) = x^+$ for some $\mathcal{A} \in \Lambda$). We have

$$\sum_{v \in \mathcal{H}_d} \mathbb{P}(x \in Q_v \cap x^+ \in E)(\|Q\|_{\circ} - \|Q_v\|_{\circ})^{\ell} = \sum_{v \in \mathcal{H}_d} \|\tilde{E}_v\| (\|Q\|_{\circ} - \|Q_v\|_{\circ})^{\ell} \tag{77}$$

$$\begin{aligned}
 &\leq \sum_{v \in \mathcal{H}_d} \|\tilde{E}_v\| (1 - \|Q_v\|_{\circ})^{\ell} \\
 &\leq \sum_{v \in \mathcal{H}_d} \|\tilde{E}_v\| \left(1 - \|\tilde{E}_v\|\right)^{\ell} \tag{78}
 \end{aligned}$$

where (77) and (78) follow since all augmentation sets are equal size.

Note $\sum_v \|\tilde{E}_v\| = \|R\|$, and $h(x_v) := x_v(1 - x_v)^{\ell}$ is maximized on $x_v \in [0, 1]$ at $x_v = \frac{1}{\ell+1}$. Thus, $\sum_v \|\tilde{E}_v\| \left(1 - \|\tilde{E}_v\|\right)^{\ell}$ is upper bounded by setting $\|\tilde{E}_v\| = \frac{1}{\ell+1}$ for all v . Thus we obtain

$$\begin{aligned}
 \sum_{v \in \mathcal{H}_d} \mathbb{P}(x \in Q_v \cap x^+ \in B)(\|Q\|_{\circ} - \|Q_v\|_{\circ})^{\ell} &\leq \sum_{v \in \mathcal{H}_d} \frac{1}{\ell+1} \left(1 - \frac{1}{\ell+1}\right)^{\ell} \\
 &= \frac{2^d}{\ell+1} \left(1 - \frac{1}{\ell+1}\right)^{\ell} \\
 &\leq \frac{\epsilon}{cd} e^{-\frac{\ell}{\ell+1}}
 \end{aligned}$$

where we have used $\ell \geq \frac{c}{\epsilon} d 2^d$ for a constant c in the last line. Next we consider $\sum_{v \in \mathcal{H}_d} \|E_v\|_{\circ} (1 - \|E_v\|_{\circ})^{\ell}$, and use a tighter method of bounding this sum than above. Note that $\sum_{v \in \mathcal{H}_d} \|E_v\|_{\circ} = \|E\|_{\circ}$. If $\frac{\|E\|_{\circ}}{2^d} \leq \frac{1}{\ell+1}$, then by the concavity of $h(x_v) := x_v(1 - x_v)^{\ell}$ on the interval $x_v \in [0, \frac{1}{\ell+1}]$,

the sum is maximized by setting $\|E_v\|_\circ = \frac{\|E\|_\circ}{2^d}$ for all v . Otherwise, the sum is upper bounded by setting $\|E_v\|_\circ = \frac{1}{\ell+1}$ for all v . Thus we have

$$\begin{aligned}
 \sum_{v \in \mathcal{H}_d} \|E_v\|_\circ (1 - \|E_v\|_\circ)^\ell &\leq \chi \left\{ \|E\|_\circ \leq \frac{2^d}{\ell+1} \right\} \|E\|_\circ \left(1 - \frac{\|E\|_\circ}{2^d} \right)^\ell \\
 &\quad + \chi \left\{ \|E\|_\circ > \frac{2^d}{\ell+1} \right\} \frac{2^d}{\ell+1} \left(1 - \frac{1}{\ell+1} \right)^\ell \\
 &\leq \chi \left\{ \|E\|_\circ \leq \frac{2^d}{\ell+1} \right\} \|E\|_\circ e^{-\frac{\|E\|_\circ \ell}{2^d}} \\
 &\quad + \chi \left\{ \|E\|_\circ > \frac{2^d}{\ell+1} \right\} \|E\|_\circ e^{-\frac{\ell}{\ell+1}}
 \end{aligned} \tag{79}$$

Finally, note that $\mathcal{D}_g(v) := \mathbb{P}_{x \sim D_\circ}[g(x) = v] = \|Q_v\|_\circ + \|E_v\|_\circ$. We have that for all v , $Q_v \setminus Z \neq \emptyset \implies Q_v \setminus Z = Q_v, E_v = \emptyset$ by Claims B.4 and B.6. Thus for all $v : Q_v \setminus Z \neq \emptyset$, $\mathcal{D}_g(v) = \|Q_v\|_\circ = \|Q_v \setminus Z\|_\circ$. This allows us to use that g is near uniform, i.e. $\mathcal{D}_g(v) > \frac{1}{c_1 d 2^d}$ or $\mathcal{D}_g(v) \leq \frac{\epsilon}{c_2 d 2^{2d}}$ for all $v \in \mathcal{H}_d$ for some constants c_1, c_2 . We have $\ell \geq \frac{c}{\epsilon} d 2^d$ and choose $c_1 < c$, such that $\frac{1}{c_1 d 2^d} > \frac{1}{\ell+1}$. Since $h(x_v) := x_v(1 - x_v)^\ell$ is a decreasing function of x_v for $x_v \geq \frac{1}{\ell+1}$, we can bound the last sum in (76) as

$$\begin{aligned}
 \ell \sum_{v \in \mathcal{H}_d} \|Q_v \setminus Z\|_\circ (1 - \|Q_v \setminus Z\|_\circ)^{\ell-1} &\leq \ell 2^d \max \left(\frac{\epsilon}{c_2 d 2^{2d}}, \frac{1}{c_1 d 2^d} \left(1 - \frac{1}{c_1 d 2^d} \right)^\ell \right) \\
 &\leq \ell \max \left(\frac{\epsilon}{c_2 d 2^{2d}}, \frac{1}{c_1 d} e^{-\frac{\ell}{c_1 d 2^d}} \right) \\
 &\leq \max \left(\frac{c}{c_2}, \frac{c}{c_1 \epsilon} e^{-(\frac{c}{c_1 \epsilon} - 1)} \right)
 \end{aligned} \tag{80}$$

Before we combining these bounds with (76), we first show that $\|R\| = \Omega(\epsilon)$.

Claim B.7 *Let $\epsilon \leq \min_{j \in [d]} \min_{\mathbf{c}' \in \Sigma_{f_j}} \mathbb{P}_{x, x' \sim D}[x, x' \in \Gamma_{\mathbf{c}'}, f_j(x) \neq f_j(x')]$ as defined in the statement of Theorem 15. Then $\epsilon \leq \frac{6}{8} \|R\|$.*

Proof From our choice of f'_1 , we have

$$\begin{aligned}
 \epsilon &\leq \min_{j \in [d]} \min_{\mathbf{c}' \in \Sigma_{f_j}} \mathbb{P}_{x, x' \sim D}[x, x' \in \Gamma_{\mathbf{c}'}, f_j(x) \neq f_j(x')] \\
 &\leq 2 \mathbb{P}_{x, x' \sim D}[x \in \Gamma_{\mathbf{c}} \setminus E, x' \in E] \\
 &= 2 \mathbb{P}_{x \sim D}[x \in \Gamma_{\mathbf{c}} \setminus E] \mathbb{P}_{x \sim D}[x \in E] \\
 &\leq 2 \mathbb{P}_{x \sim D}[x \in E] \\
 &\leq 2(\mathbb{P}_{x \sim D_\circ}[x \in E] + \mathbb{P}_{x^+ \sim D \setminus D_\circ}[x^+ \in E]) \\
 &= 2(\|E\|_\circ + \|E\|).
 \end{aligned} \tag{81}$$

Note that $\|E\|_\circ \leq \frac{1}{8} \|R\|$ by Assumption 13. Observe that

$$\|E\| = \mathbb{P}_{x^+ \sim D \setminus D_\circ}[x^+ \in E]$$

$$\begin{aligned}
 &= \mathbb{P}_{x^+ \sim D \setminus D_\circ} [x^+ \in E, \mathcal{A}^{-1}(x^+) \in E] + \mathbb{P}_{x^+ \sim D \setminus D_\circ} [x^+ \in E, \mathcal{A}^{-1}(x^+) \notin E] \\
 &\leq \mathbb{P}_{x \sim D_\circ, x^+ \sim \mathcal{A}(x)} [x \in E, x^+ \in E] + \|R\| \\
 &\leq \|E\|_\circ + \|R\|
 \end{aligned} \tag{82}$$

Note that $\|E\|_\circ \leq \frac{1}{\delta} \|R\|$ by Assumption 13. Combining this with (81) and (82) yields

$$\epsilon \leq 2(2\|E\|_\circ + \|R\|) \leq \frac{6}{\delta} \|R\|$$

■

Finally, using Claim B.7 with $\beta \geq \frac{c'}{\delta} \log\left(\frac{c}{\epsilon}\right) 2^d$, we obtain from (76) that

$$\begin{aligned}
 &\mathcal{L}^-(g') - \mathcal{L}^-(g) \\
 &\leq \|R\|(\log(\ell+2) + \log(2)) + \frac{\|R\|}{\delta}(2^{d+1} + \log(\ell+1)) + 2^{-(\ell-1)}\beta\|R\| \\
 &\quad + \frac{2\beta\epsilon}{cd}e^{-\frac{\ell}{\ell+1}} + 2\beta \max\left(\frac{c}{c_2}, \frac{c}{c_1}e^{-\left(\frac{c}{c_1\epsilon}-1\right)}\right) \\
 &\quad + 2\beta\chi \left\{ \|E\|_\circ \leq \frac{2^d}{\ell+1} \right\} \|E\|_\circ e^{-\frac{\|E\|_\circ \ell}{2^d}} + 2\beta\chi \left\{ \|E\|_\circ > \frac{2^d}{\ell+1} \right\} \|E\|_\circ e^{-\frac{\ell}{\ell+1}} \\
 &\leq \beta\|R\| \left(\frac{\log(2cd2^d/\epsilon+4) + \frac{1}{\delta} \log(cd2^d/\epsilon+1)}{c' \log(\frac{c}{\epsilon})2^d/\delta} + \frac{2}{c' \log(\frac{c}{\epsilon})2^d} + \frac{2}{2cd2^d/\epsilon} \right. \\
 &\quad \left. + \frac{12}{cd\delta}e^{-\frac{\ell}{\ell+1}} + 2 \max\left(\frac{c}{c_2}, \frac{c}{c_1}e^{-\left(\frac{c}{c_1\epsilon}-1\right)}\right) \right) \\
 &\quad + 2\beta\|E\|_\circ \left(\chi \left\{ \|E\|_\circ \leq \frac{\epsilon}{cd} \right\} e^{-\frac{\|E\|_\circ \ell}{2^d}} + \chi \left\{ \|E\|_\circ > \frac{2^d}{\ell+1} \right\} e^{-\frac{\ell}{\ell+1}} \right) \\
 &\leq \beta\|R\| \left(\frac{\log(4c/\epsilon+4) + \frac{1}{\delta} \log(2c/\epsilon+1)}{2c' \log(\frac{c}{\epsilon})/\delta} + \frac{1}{c' \log(\frac{c}{\epsilon})} + \frac{1}{2c/\epsilon} \right. \\
 &\quad \left. + \frac{12}{cd\delta}e^{-\frac{\ell}{\ell+1}} + 2 \max\left(\frac{c}{c_2}, \frac{c}{c_1}e^{-\left(\frac{c}{c_1\epsilon}-1\right)}\right) \right) \\
 &\quad + \beta\|R\| \left(\frac{12}{cd\delta} \chi \left\{ \|E\|_\circ \leq \frac{6\|R\|}{cd\delta} \right\} + \frac{1}{\delta} \chi \left\{ \|E\|_\circ > \frac{2^d}{\ell+1} \right\} e^{-\frac{\ell}{\ell+1}} \right) \\
 &\leq 2\beta\|R\| \left(\frac{2c}{c_2} + \frac{1}{100} \right) + 2\beta\|R\| \frac{1.01}{\delta e} \tag{83} \\
 &< 2\beta\|R\| = \mathcal{L}^+(g) - \mathcal{L}^+(g') \tag{84}
 \end{aligned}$$

where (83) follows for a sufficiently large constants c' and $c > c_1$. Use $\delta \geq 0.4$ and set $c' = c$, $c_1 = \frac{c}{10}$, $c_2 = 100c$ and make c sufficiently large to obtain

$$\mathcal{L}^-(g') - \mathcal{L}^-(g) < 2\beta\|R\| = \mathcal{L}^+(g) - \mathcal{L}^+(g').$$

■

B.1. Remark: Modified version of Theorem 15

Theorem 15 shows that if a minimizer of the InfoNCE loss is close to uniform, then it must be clean. Here we show that the InfoNCE loss can be interpreted as the Lagrangian which, under appropriate choice of hyperparameters, we can formally show to be minimized only by clean representations.

Weighted InfoNCE loss. Consider the following constrained optimization problem that tries to maximize uniformity while preserving alignment:

$$\begin{aligned} \min_{g \in \mathcal{G}} \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \\ \text{s.t. } \mathbb{E}_{x, x^+} [g(x)^\top g(x^+)] = d \end{aligned} \quad (85)$$

The unconstrained penalized version of this problem is

$$\min_{g \in \mathcal{G}} -\lambda \mathbb{E}_{x, x^+} [g(x)^\top g(x^+)] + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \quad (86)$$

Note that the above objective with penalty coefficient $\lambda = \beta$ is equal to the InfoNCE objective. This formulation motivates alternate choices of λ depending on how strictly we would like to enforce alignment. We refer to the loss

$$\mathcal{L}_\rho(g) = -\rho\beta \mathbb{E}_{x, x^+} [g(x)^\top g(x^+)] + \mathbb{E}_{x, x^+, \{x_i^-\}_\ell} \left[\log \left(e^{\beta g(x)^\top g(x^+)} + \sum_{i=1}^{\ell} e^{\beta g(x)^\top g(x_i^-)} \right) \right] \quad (87)$$

as the Weighted InfoNCE loss which is equivalent to the penalized version of (85) with $\lambda = \rho\beta$. We note that this loss is similar to the generalized InfoNCE loss proposed by [Chen et al. \(2021\)](#).

Corollary 21 *Consider the same setting as Theorem 15 but with any $\delta > 0$ and with the Weighted InfoNCE loss with $\rho \geq 2^{d+1} + \frac{1}{\delta}$. Then for sufficiently large ℓ and β , all solutions of the Weighted InfoNCE objective are cluster-preserving.*

Proof The result follows from the analysis in the proof of Theorem 15. The analysis for the difference in positive terms is identical except that they are scaled by ρ , so we have

$$\mathcal{L}^+(g) - \mathcal{L}^+(g') = 2\rho\beta \|R\| \quad (88)$$

For the difference in negative terms, the analysis is again identical except that we can no longer use that g is close to uniform. The only place we have used this is to bound $\ell \sum_{v \in \mathcal{H}_d} \|Q_v \setminus Z\|_o (1 - \|Q_v \setminus Z\|_o)^\ell$ in (80). Here, we bound this term using the fact that $h(x) := x(1-x)^\ell$ is maximized on the interval $x \in [0, 1]$ at $x = \frac{1}{\ell+1}$.

$$\begin{aligned} \ell \sum_{v \in \mathcal{H}_d} \|Q_v \setminus Z\|_o (1 - \|Q_v \setminus Z\|_o)^\ell &\leq \ell \sum_{v \in \mathcal{H}_d} \frac{1}{\ell+1} \left(1 - \frac{1}{\ell+1}\right)^\ell \\ &= 2^d \left(\frac{\ell}{\ell+1}\right)^{\ell+1} \end{aligned}$$

$$\leq \frac{2^d}{e} \quad (89)$$

Replacing this bound and executing the same analysis as in (83) yields

$$\begin{aligned} & \mathcal{L}^-(g') - \mathcal{L}^-(g) \\ & \leq 2\beta\|R\| \left(2^d + \frac{1}{100} \right) + 2\beta\|R\| \frac{1.01}{\delta e} \\ & < 2\rho\beta\|R\| \end{aligned} \quad (90)$$

completing the proof. \blacksquare

Appendix C. Proofs of Downstream Guarantees

C.1. Proof of Theorem 8

Proof Because there are exactly 2^d clusters, from Theorem 7, we know that since g^* is uniform, it maps each cluster to a unique vertex on the d -dimensional hypercube. Let $\psi_f := \{\Gamma_{\mathbf{c}} : f(x) = 1 \ \forall x \in \Gamma_{\mathbf{c}}\}$ be the set of clusters which f labels 1, and let $N = |\psi_f|$. Similarly let $\psi_f^c := \{\Gamma_{\mathbf{c}} : f(x) = -1 \ \forall x \in \Gamma_{\mathbf{c}}\}$. For all $j \in [N]$, let the j -th row of $\mathbf{W} \in \mathbb{R}^{m \times d}$ equal the vertex corresponding to the mapping of the j -th cluster in ψ_f by g^* . For all $j \in [m - N]$, let the $j + N$ -th row of \mathbf{W} equal the vertex corresponding to the mapping of the j -th cluster in ψ_f^c by g^* . Then for any x such that $f(x) = 1$, $\mathbf{W}g^*(x)$ has exactly one element with value d among the first N elements, and all other elements are at most $d - 2$. On the other hand, for any x such that $f(x) = -1$, $\mathbf{W}g^*(x)$ has exactly one element with value d among the last $m - N$ elements, and all other elements are at most $d - 2$. Set $b = (d - 2) \times 1_m$, that is, $d - 2$ times the m -dimensional vector of ones, and $a = [1_N^\top, -1_{m-N}^\top]^\top$, that is, the m -dimensional vector whose first N elements are 1 and whose last $m - N$ elements are -1 . Then $a^\top \text{ReLU}(\mathbf{W}g(x) - b) = f(x)$ for all x . \blacksquare

C.2. Proof of Theorem 9

Proof Since we are in the realizable setting, D has $m := 2^d$, equal-size clusters, where $d > 3$. Moreover, since \mathcal{F}_* is arbitrarily powerful and the augmentation sets are disjoint, for every pair of augmentation sets $(A(x), A(\bar{x}))$, there exists an $f \in \mathcal{F}_*$ such that $f(x^+) \neq f(\bar{x}^+)$ for all $x^+ \in A(x)$ and $\bar{x}^+ \in A(\bar{x})$, and f does not intersect any other augmentation set. Further, \mathcal{G}_* can map augmentation sets to arbitrarily different vertices, even if these sets lie in the same cluster. In other words, there are clean representations in \mathcal{G}_* (meaning they are faithful to all augmentation sets) that split clusters by augmentation sets.

Suppose the number of augmentation sets in each cluster is $k \times m \times 2^d$ for some $k \in \mathbb{N}^+$, and all augmentation sets are of equal size M . Then there exists a clean and uniform representation $g \in \mathcal{G}_*$ such that for each vertex $v \in \mathcal{H}_d$, kM of the images in the set $\{x \in D_\circ : g(x) = v\}$ are in each cluster. In other words, for all $x \in D_\circ$ and $x^+ \sim A(x)$, $g(x) = g(x^+)$. Thus, we can apply Theorem 7 to obtain $g \in \arg \min_{g' \in \mathcal{G}_*} \mathcal{L}(g')$ (note that in Theorem 7, cluster-preserving is equivalent to clean since we are optimizing over the restricted class \mathcal{G} , and the same proof can be applied exactly as is,

with the word “cluster-preserving” replaced by “clean”, to show that $g \in \arg \min_{g' \in \mathcal{G}_*} \mathcal{L}(g')$ if and only if g is clean and uniform).

Since any head $\omega \in \mathcal{J}$ composed with g must yield the same prediction for all images mapped to the same vertex on \mathcal{H}_d , and all vertices have the same number of images from each cluster mapped to them, the number of images with predicted label 1 must be the same for all clusters, and likewise for -1 . Thus, for any downstream binary classification task h that satisfies $h(x) = h(x')$ for all $x, x' \in \Gamma_{\mathbf{c}}$ for all $\mathbf{c} \in C$ and $\mathbb{P}_{x \sim D_0}[h(x) = 1] = 0.5$, any ω must have $\mathcal{L}_f(\omega \circ g) \geq 0.5$. ■