Measuring Self-Supervised Representation Quality for Downstream Classification Using Discriminative Features

Neha Kalibhat¹, Kanika Narang², Hamed Firooz², Maziar Sanjabi², Soheil Feizi¹

¹University of Maryland, College Park ²Meta AI nehamk@umd.edu

Abstract

Self-supervised learning (SSL) has shown impressive results in downstream classification tasks. However, there is limited work in understanding their failure modes and interpreting their learned representations. In this paper, we study the representation space of state-of-the-art self-supervised models including SimCLR, SwaV, MoCo, BYOL, DINO, SimSiam, VICReg and Barlow Twins. Without the use of class label information, we discover discriminative features that correspond to unique physical attributes in images, present mostly in correctly-classified representations. Using these features, we can compress the representation space by up to 40% without significantly affecting linear classification performance. We then propose Self-Supervised Representation Quality Score (or Q-Score), an unsupervised score that can reliably predict if a given sample is likely to be mis-classified during linear evaluation, achieving AUPRC of 91.45 on ImageNet-100 and 78.78 on ImageNet-1K. Q-Score can also be used as a regularization term on pre-trained encoders to remedy lowquality representations. Fine-tuning with Q-Score regularization can boost the linear probing accuracy of SSL models by up to 5.8% on ImageNet-100 and 3.7% on ImageNet-1K compared to their baselines. Finally, using gradient heatmaps and Salient ImageNet masks, we define a metric to quantify the interpretability of each representation. We show that discriminative features are strongly correlated to core attributes and, enhancing these features through Q-score regularization makes SSL representations more interpretable.

Introduction

Self-supervised models (Chen et al. 2020a; Caron et al. 2020; Chen et al. 2020b; Grill et al. 2020; Chen and He 2021; Caron et al. 2018a; Khosla et al. 2020; Caron et al. 2021; Bardes, Ponce, and LeCun 2021; Zbontar et al. 2021) learn to extract useful representations from data without relying on human supervision, and perform comparably to supervised models in downstream classification tasks. Pre-training these models can be highly resource-intensive and time-consuming. It is therefore crucial that the learned representations are of high quality such that they are explainable and generalizable. However, in practice, these representations are often quite noisy and un-interpretable, causing difficulties in understanding and debugging their failure

modes (Jing et al. 2022; Huang, Yi, and Zhao 2021; Ericsson, Gouk, and Hospedales 2021).

In this paper, our goal is to study the representation space of pre-trained self-supervised encoders (SSL) such as Sim-CLR (Chen et al. 2020a), SwaV (Caron et al. 2020), MoCo (Chen et al. 2020b), BYOL (Grill et al. 2020), SimSiam (Chen and He 2021), DINO (Caron et al. 2021), VICReg (Bardes, Ponce, and LeCun 2021) and Barlow Twins (Zbontar et al. 2021) and discover their informative features in an unsupervised manner. We observe that representations are mostly sparse, containing a small number of *highly activating features*. These features can strongly activate a small, moderate or large number of samples in the population. We refer to the moderate category of features as *discriminative features*.

We observe some intriguing properties of discriminative features: (i) Although discovered without any class label information, they can be strongly correlated to a particular class or group of classes (See Figure 1); (ii) They highlight informative concepts in the activating samples which are often related to the ground truth of those samples; (iii) They activate strongly in correctly classified representations rather than mis-classified representations (as shown in Figure 3) and finally, iv) Representations can be compressed by up to 40% using discriminative features without significantly affecting linear evaluation performance.

Building on these observations, we propose an unsupervised, sample-wise **Self-Supervised Representation Quality Score (Q-Score)**. A high Q-Score for a sample implies that its representation contains highly activating discriminative coordinates which is a favorable representation property. We empirically observe that Q-Score can be used as a zero-shot predictor in distinguishing between correct and incorrect classifications for any SSL model achieving AUPRC of 91.45 on ImageNet-100 and 78.78 AUPRC on ImageNet-1K.

We next apply Q-Score as a regularizer and further-train pre-trained SSL models at a low rate to improve low-quality representations. This improves the linear probing performance across all our baselines, highest on BYOL (5.8% on ImageNet-100 and 3.7% on ImageNet-1K). The representations, after regularization, show increased activation for discriminative features (Figure 3) due to which several previously mis-classified samples get correctly classified with

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

higher confidence.

Finally, we define a metric for quantifying representation interpretability using Salient ImageNet (Singla and Feizi 2021) masks as ground truth. Discriminative features are strongly correlated to *core* features of Salient ImageNet. We can potentially explain these features by correlating their meanings with the feature annotations provided for core features in Salient ImageNet. We also observe that discriminative features in mis-classified representations are less correlated with core features compared to correct classifications. Q-score regularization improves this correlation for both correct and mis-classified representations, thereby making representations more explainable.

Related Work

Unsupervised methods for classification has been a longstanding area of research, traditionally involving the use of clustering techniques (Bojanowski and Joulin 2017; Dosovitskiy et al. 2014; YM., C., and A. 2020; Bautista et al. 2016; Caron et al. 2018b, 2019; Huang et al. 2019). Selfsupervised learning, is a powerful approach that enables learning by preparing own labels for every sample (Bojanowski and Joulin 2017; Dosovitskiy et al. 2014; Wu et al. 2018; Dosovitskiy et al. 2016) usually with the help of a contrastive loss (Arora et al. 2019; Tosh, Krishnamurthy, and Hsu 2021; Bachman, Hjelm, and Buchwalter 2019). Positive views in SSL losses are multiple transformations (Tian et al. 2020) of a given sample using stochastic data augmentation. Through this approach, several state-of-the-art SSL techniques (Chen et al. 2020a; Caron et al. 2020; Chen and He 2021; Grill et al. 2020; Chen et al. 2020b; Khosla et al. 2020) have produced representations that show competitive linear classification accuracy to that of supervised approaches.

Understanding these learned representations is relatively less explored. Several feature interpretability techniques exist (Bau et al. 2017; Kalibhat et al. 2023; Hernandez et al. 2022), that aim to explain individual neurons with natural language. However, our goal is to study representations through the lens of failure modes and generalization. (Jing et al. 2022), observes that self-supervised representations collapse to a lower dimensional space instead of the entire embedding space. Other methods (von Kügelgen* et al. 2021; Xiao et al. 2021), propose to separate the representation space into variant and invariant information so that augmentations are not task-specific. (Grigg et al. 2021) observes representations across layers of the encoder and compare it to supervised setups. Clustering-based or prototypical-based methods have also been proposed where the representation space is collapsed into a low-rank space (Dwibedi et al. 2021; Koohpayegani, Tejankar, and Pirsiavash 2021). (Bordes, Balestriero, and Vincent 2021) uses an RCDM model to understand representation invariance to augmentations. (Garrido et al. 2022; Li, Efros, and Pathak 2022) propose a score based on the rank of all post-projector embeddings that can be used to judge and compare various self-supervised models.

In this work, we focus more on studying the properties of representations across correct and incorrect classi-



Figure 1: Discriminative Features in Self-Supervised (SSL) Models: We plot the percentage of highly activating samples for each feature in the SimCLR (ResNet-50) representation space. The features that show very low or very high percentage activations are *non-discriminative* as they likely correspond to very uncommon (lower tail) or very general attributes (upper tail). The features that activate a moderate number of samples (middle portion) are called *discriminative features*. As shown in the gradient heatmaps, these features encode important physical attributes shared among specific classes. These features play a key role in assessing the quality of SSL representations for downstream linear classification tasks.

fications in downstream linear probing (without using any labels). We investigate the connection between these unsupervised properties in the representation space and misclassifications. Unlike (Garrido et al. 2022; Li, Efros, and Pathak 2022) which requires computing rank over the entire dataset, our analysis leads to the development of an unsupervised *sample-wise* quality score which can be used as a regularizer and effectively improve downstream classification performance.

Understanding Representations and Their Failure Modes

Let us consider a pre-trained self-supervised model with a ResNet (He et al. 2016) backbone encoder f(.). Given an input sample, $\mathbf{x}_i \in \mathbb{R}^n$ its representation is denoted by,



Figure 2: Linear probing discriminative features: We train linear classifiers after selecting subsets of discriminative features of various sizes (middle portion of Figure 1) and plot their top-1 accuracy for various SSL baselines. Classifiers trained using discriminative features consistently outperform those of randomly selected features (averaged over 4 random seeds). We can achieve up to 40% reduction in representations size using discriminative features without significantly affecting the top-1 accuracy.

 $f(\mathbf{x}_i) = \mathbf{h}_i \in \mathbb{R}^r$, where r is the size of the representation space.

Upon visual analysis (See Appendix for more details), we observe that each representation is *nearly* sparse, i.e., most feature activations are close to zero (Jing et al. 2022). There exists a select few features that are strongly deviated from the remaining features in that representation. For any given representation $\mathbf{h}_i \in \mathbb{R}^r$, we formally define the set of highly activating features (L_i) as $L_i := \{j : h_{ij} > \mu_i + \epsilon \sigma_i\},\$ where μ_i and σ_i denote the mean and standard deviation of \mathbf{h}_i respectively and ϵ is a hyperparameter that is empirically selected. We use $\epsilon = 4$ in our experiments. In all our analysis, we perform L2 normalization over every h to ensure fair comparison of features. For every feature j, the percentage of highly activating samples is denoted by, $A_j = \frac{100}{N} \sum_{i=1}^{N} \mathbb{1}_{j \in L_i}$ where N is the size of the population. In the top panel of Figure 1, we plot A_j for all features j in the SimCLR representation space of the ImageNet-1K train set. The x-axis is ordered in ascending order of A_j . In the next section, we dissect this plot to group features based on A_i .

Our observations do not directly extend to ViT-based SSL encoders since, unlike ResNet encoders, their representations can also contain negatively activated features. They need not be sparse in nature (no ReLU before representation layer), rather, features can be both positively (highly activating) or negatively (lowly activating) correlated to important class-specific concepts. In our work, we observe several unique properties of highly activating features in ResNet representations which are beneficial to detect failures. We see ViT encoders as an important direction for future work and focus on ResNet-based encoders for our study.

Discriminative Features

Based on Figure 1, we can define three broad categories of highly activating features: (i) Features that are highly acti-

vating across a very small fraction of the population, corresponding to the lower tail features in Figure 1. We take the example of features 621, 251 and 1021 and visualize their highly activating samples and gradient heatmaps (using GradCAM (Selvaraju et al. 2019)). Since these features activate very few samples, they likely correspond to imagespecific or uncommon concepts. Such features would also not be useful in classification tasks as these are not shared, class-relevant attributes. (ii) Features that highly activate a very large number of samples in the population i.e, the upper tail features in Figure 1. Like feature 2021 and 401, such features are likely to encode very broad and general characteristics (like texture, color etc.) common to most samples (spanning various classes) and therefore, are not classdiscriminative. The third category includes, (iii) Features that highly activate a moderate number of samples in the population (i.e. the middle part in Figure 1). These features are most likely to encode unique physical attributes associated with particular classes. For example, feature 98 corresponds to the "stripe" pattern which is an important property of the zebra class. Similarly, feature 970 corresponds to the style of the daisy class, and feature 1392 corresponds to lionfish in different scenes. We refer to this subset of highly activating features as discriminative features. Note that we did not use any label information for this analysis. We can identify discriminative and non-discriminative features in a fully unsupervised manner by simply observing their percentage activations (A). The bar plots in Figure 1, show that these features activate more than 80% of particular classes which confirms that these features are strongly class-correlated.

Discriminative features can be regarded as a summarization of the top concepts related to each class of the dataset the encoder is trained on. We justify the described method of selection in Figure 2, where we plot the top-1 accuracy of a linear classifier trained on ImageNet-1K using subsets of discriminative features of varying sizes as chosen from Figure 1 (middle portion). We compute the percentile for each point in the distribution A and gradually increase the lower limits (from 0^{th} percentile), and decreasing the upper limits (from 100th percentile) to get multiple sets of discriminative features of varying sizes. We also plot the top-1 accuracy when random subsets of features are selected. We observe that discriminative features perform significantly better compared to randomly selected features. We also observe that we can reduce the representation size up to 40% using the discriminative features, with minimal reduction in performance. In practice, we select the discriminative features between the 50^{th} and the 95^{th} percentile of A (as shown in Figure 1). This range can be discovered empirically and can be further tuned for each model-dataset pair. In the Appendix, we also show that selecting features from either the lower or upper tail of A also under-perform compared to discriminative features from the middle portion.

While we analyze features independently in our work, it has been shown (Kalibhat et al. 2023; Elhage et al. 2022) that not all neural features are axis-aligned. Meaningful class-related concepts can also be encoded by multiple features (See Appendix for examples). In such cases, the whole *group* of features can be considered as discriminative. These



Indices of Discriminative Features in Representations

Figure 3: Comparing correct and mis-classified representations: In these heatmaps, we visualize the discriminative features of average SimCLR representations of several ImageNet-1K classes - correct (left) and incorrect (right) classifications. In the baseline, we observe that discriminative features are strongly activated only in correctly classified representations. Q-Score regularization improves discriminative features' activations, even in mis-classified representations.

Model	AUROC	AUPRC
SimCLR	73.26	91.11
SwaV	62.30	81.18
MoCo	63.09	84.59
BYOL	72.78	91.45
SimSiam	66.97	86.05
DINO	62.57	83.03
VICReg	69.68	89.06
Barlow	69.57	90.06

Table 1: AUPRC and AUROC scores of Q-Score: We measure the effectiveness of Q-Score when used as a predictor in distinguishing between correct and mis-classified representations on ImageNet-100 and ImageNet-1K (See Appendix) on each SSL model. Q-Score shows an AUPRC of up to 91.45 on ImageNet-100, 78.78 on ImageNet-1K and AU-ROC of 73.26 on ImageNet-100, 65.44 on ImageNet-1K.

groups can be highly activating for specific classes and lie in the middle portion of A. We also perform a PCA analysis on the representation space (see Appendix) to partially validate our selection method.

Mis-classified Representations

We now study how discriminative features play a key role in detecting potential mis-classifications in a fully unsupervised manner. In Figure 3, we take SimCLR ImageNet-1K representations and visualize the discriminative features. On the left, we show the average representations of correctly classified samples (after linear probing) in a subset of classes, while on the right, we show the same for the misclassified samples in those classes. The subset of features we display is the same for correct and incorrect classifications.

As we can see, in Figure 3, in the first panel, there is a clear difference between representations of correctly and incorrectly classified examples. Both correct and misclassified representations are *nearly* sparse, however, the discriminative features are significantly more activated in correct classifications. This is especially interesting because we can visually distinguish between correct and incorrect classifications, just by observing the discriminative features, without using any label information. Note that this observation does not depend on the actual ground truths or predicted labels of the linear classifier rather, just a binary outcome or whether or not a sample was correctly classified.

The correlation of discriminative features to unique physical attributes as studied in the previous section, suggests that their presence may be useful in correctly classifying representations. In Figure 3, our claim is confirmed as we observe that mis-classified representations do not show high activations on these features. Therefore, for any given sample, we can consider discriminative features as strong signals indicating classification outcome without requiring to train a linear classification head. We would like to emphasize that our results only indicate an *association* between these structural properties and classification accuracy and we do not claim any causal relationship between the two.

Self-Supervised Representation Q-Score

Our study of learned representation patterns helps us discover discriminative features in an unsupervised manner. These features encode class-specific attributes and help us visually distinguish between correct and incorrect classifications. We combine these observations to design a samplewise quality score for SSL representations. Let us define D, such that |D| < r, as the set of discriminative features for a given SSL model trained on a given dataset. For the i^{th} sample, we have \mathbf{h}_i (representation), μ_i (mean of \mathbf{h}_i), σ_i (standard deviation of \mathbf{h}_i) and the set of highly activating features $L_i = \{j : h_{ij} > \mu_i + \epsilon \sigma_i\}, |L_i| < r$. We define our Self-Supervised Quality Score for sample *i* as,

$$Q_{i} := \frac{1}{|L_{i} \cap D|} \sum_{j \in L_{i} \cap D} (h_{ij} - \mu_{i})$$
(1)

where, $L_i \cap D$ is the set of discriminative features specific to the i^{th} sample. Intuitively, higher Q_i implies that the representation contains highly activated discriminative features which are strongly deviated from the mean. Our objective with this metric is to compute a sample-specific score in an unsupervised manner indicating the quality of its representations. Ideally, we would like to argue that samples with higher Q-score have improved representations and thus are more likely to be classified correctly in the downstream task. This is a general score that can be applied to any ResNetbased SSL model trained on any dataset. See Appendix for a discussion on Q-Score in supervised models.

Next, we measure how effective our score is in differentiating between correctly and incorrectly classified representations in an unsupervised manner. We plot the Precision-Recall (PR) curve and the Receiver Operating Characteristic (ROC) curve (See Appendix) of Q-Score when used as a predictor of classification outcome (correct or incorrect). We show this for SimCLR, SwaV, MoCo, BYOL, DINO and SimSiam for ImageNet-100 and ImageNet-1K. We also compute the AUROC (area under receiver operating characteristic curve) and AUPRC (area under precision-recall curve) of these curves in Table 1. We observe AUPRC up to 91.45 on ImageNet-100 and 78.78 on ImageNet-1K on BYOL. On SimCLR, we observe AUROC up to 73.26 on ImageNet-100 and 65.44 on ImageNet-1K. Based on these results we can conclude that, Q-Score is a reliable metric in assessing the quality of representations, meaning that representations with lower Q-Score (quality), are more likely to be mis-classified.

We now check if promoting Q-Score on pre-trained representations is helpful. To do so, we take state-of-the-art pretrained SSL models and further train them for a small number of iterations with Q-Score as a regularizer. For example, we can apply this regularizer to the SimCLR optimization as follows,

$$\max_{\theta} \frac{1}{2N} \sum_{i=1}^{2N} \left[\log \frac{sim(\mathbf{z}_i, \tilde{\mathbf{z}}_i)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} sim(\mathbf{z}_i, \mathbf{z}_j)} + \lambda_1 \mathbb{1}_{Q_i < \alpha}(Q_i) \right]$$
(2)

where z is the latent vector computed by passing h through a projector network and sim(.) denotes the exponentiated cosine similarity of the normalized latent vector. α is a threshold with which we select the low-score samples whose Q-Scores should be maximized and λ_1 is the regularization coefficient. In other words the goal of this regularization is to improve low-quality representations, similar to the ones shown in Figure 3, by maximizing their discriminative features for downstream classification.

In practice, directly applying this regularization could lead to a trivial solution where a small set of features get activated for all samples. This is not a favorable situation because these representations become harder to classify accurately and more importantly, the discriminative features are no longer *informative* because they are activated for all samples (similar to the upper tail in Figure 1). Such features have significantly large L1 norms *across* samples compared to the remaining features. Therefore, in our revised optimization, we penalize features that have large L1 norms across samples. Let us denote the representation matrix of a given batch by $\mathbf{H} \in \mathbb{R}^{2N \times r}$ and $\|\mathbf{H}_{*,k}\|_1$ represents the L1 norm of the k^{th} column (corresponding to the k^{th} feature). Our regularized objective would then be,

$$\max_{\theta} \frac{1}{2N} \sum_{i=1}^{2N} \left[\log \frac{sim(\mathbf{z}_i, \tilde{\mathbf{z}}_i)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} sim(\mathbf{z}_i, \mathbf{z}_j)} + \lambda_1 \mathbb{1}_{Q_i < \alpha}(Q_i) \right] \\ - \lambda_2 \sum_{k=1}^r \mathbb{1}_{\|\mathbf{H}_{*,k}\|_1 > \beta}(\|\mathbf{H}_{*,k}\|_1)$$
(3)

where the threshold β helps us select the uninformative features whose L1 norms should be minimized. In practice, we choose α and β for each batch as the mean values of Q_i and $\|\mathbf{H}_{*,k}\|_1$ respectively.

Experimental Setup

Our setup consists of state-of-the-art self-supervised ResNet encoders (f(.)) - SimCLR, SwaV, MoCo, BYOL, DINO

Model	ImageNet-100		ImageNet-1K	
	Base	Q-Score	Base	Q-Score
SimCLR	78.64	80.79	63.80	66.18
SwaV	74.36	78.90	69.95	71.05
MoCo	79.62	85.16	67.03	69.31
BYOL	80.88	86.72	69.14	72.81
DINO	75.41	76.39	75.52	75.78
SimSiam	78.80	81.41	68.62	70.47
VICReg	79.77	81.56	73.63	74.72
Barlow	80.63	81.03	67.85	69.58

Table 2: Boosting linear classification performance with Q-Score regularization: We tabulate the top-1 accuracy of linear evaluation on SimCLR, SwaV, MoCo, BYOL, DINO, SimSiam, VICReg and Barlow Twins with and without Q-Score regularized fine-tuning. We observe that Q-Score regularization consistently improves each SSL state-of-theart baseline achieving up to 5.8% relative improvement on ImageNet-100 and 3.7% on ImageNet-1K.

(ResNet-based), SimSiam, VICReg and Barlow Twins that are pre-trained on datasets - ImageNet-1K, ImageNet-100 (Russakovsky et al. 2015). We use a ResNet-50 encoder for our ImageNet-1K experiments and ResNet-18 encoder for all other datasets. We discover discriminative features for each pre-trained model using the train set of each dataset. For Q-Score regularization, maintaining the same encoder architecture as the respective papers, we use the LARS (You, Gitman, and Ginsburg 2017) optimizer with warmup-anneal scheduling. We further-train each pre-trained model with and without Q-Score regularization (controlled by λ_1 and λ_2) using a low learning rate of 10^{-5} for 50 epochs. We find that $\lambda_1 = \lambda_2 = 10^{-4}$ generally works well for fine tuning. We use a maximum of 4 NVIDIA RTX A4000 GPUs (16GB memory) for all our experiments. Using the implementations from solo-learn (da Costa et al. 2022), we have tried to match our baseline numbers as much as possible within the error bars reported in the papers using the available resources. We follow the standard evaluation by training a linear classifier on frozen pre-trained representations for 100 epochs. For all our gradient heatmap visualizations, we utilize GradCAM (Selvaraju et al. 2019).

Q-Score Regularization

We tabulate our linear evaluation results of various SSL baselines before and after Q-Score regularization in Table 2. We also include results on lasso (L1) regularization (Tib-shirani 1996) on pre-trained models (See Appendix). Lasso promotes sparsity by minimizing the L1 norm of representations. Q-Score regularization improves the linear probing top-1 accuracy on all of the SSL state-of-the-art models. We observe the most improvement on BYOL showing 5.8% increase in accuracy on ImageNet-100 and 3.7% on ImageNet-1K. Lasso regularization shows degraded performance across most models since naively sparsifying representations can lead to loss of information. In contrast, Q-Score regularization promotes highly activating discrim-



Figure 4: Discriminative features in mis-classified samples The discriminative features' heatmaps on the Sim-CLR (baseline) activate portions that may not be relevant to the image ground truth, leading to incorrect predictions. After Q-Score regularization on these representations, the heatmaps become more localized and less noisy, whilst improving predictions and confidence.

inative coordinates which we have shown to be essential for downstream classification. We include more results on CIFAR-10 (Krizhevsky, Nair, and Hinton 2009a), STL-10 (Coates, Lee, and Ng 2011) and CIFAR-100 (Krizhevsky, Nair, and Hinton 2009b) in the Appendix. We also include the results on the transfer performance of discriminative features and Q-Score regularized ImageNet-1K models on unseen datasets in the Appendix. Q-Score is therefore a powerful regularizer that can boost the performance of state-ofthe-art SSL baselines.

In addition to top-1 accuracy, Q-Score also shows significant improvement in representation quality. In Figure 3, we compare the discriminative features of representations before and after Q-Score regularization. We observe that the magnitude of discriminative features and consequently the Q-Score, increases for both correct and mis-classified representations after regularization, making it harder to differentiate between them. For example on SimCLR ImageNet-1K, the AUROC reduces to 59.81 and AUPRC to 71.28. We also observe improved classification confidence as representations become more disentangled (see Appendix for a discussion on this). Our regularization produces better quality representations with clear discriminative features making them more distinguishable across classes and therefore, easier to classify. Due to this, we can attribute the improvement in performance to improved representation quality. Although Q-Score improves accuracy, it does not entirely prevent mis-classifications as mis-classifications may occur due to a variety of reasons such as, training augmentations, hardness of samples, encoder complexity, dataset imbalance etc.

Our motivation for using discriminative features as discussed in Section is because - a) they are at clear contrast between correct and incorrect classifications, and b) they show strong correlation to ground truth. We observed in Figure 3 in the baseline, that the discriminative features in correctly classified samples are not strongly activated in mis-classified samples. We now study some mis-classified samples and observe how their features may improve with Q-Score regularization. In Figure 4, we visualize the gradient heatmaps of

the discriminative features of some mis-classified examples in SimCLR. In the baseline, we observe that discriminative features do highlight portions of the image relevant to the ground truth, however, they may also activate other portions that are not necessarily important (see rock crab and green mamba). These heatmaps reflect low quality representations where the discriminative features are not strongly deviated from the mean. After Q-Score regularization, the maximization of discriminative features also leads to better gradient heatmaps that are more localized and cover almost all important portions of the image relevant to the ground truth. Therefore, these samples get classified correctly with higher confidence after regularization.



Figure 5: Comparing discriminative features with Salient ImageNet *core* and *spurious* features: We compare the gradient heatmaps of discriminative features correct and incorrect classifications of SimCLR on ImageNet-1K with the core and spurious masks of the same images in Salient ImageNet. We observe that discriminative features generally overlap more with core features in Salient ImageNet.

Quantifying Representation Interpretability with Salient ImageNet

We have observed that discriminative features in representations correspond to meaningful physical attributes through gradient heatmaps and they play a key role in deciding the downstream classification outcome. In this section, we quantify the interpretability of these features between correct and incorrect classifications. We utilize Salient ImageNet (Singla and Feizi 2021) as the ground truth baseline to compare our gradient heatmaps with. The Salient ImageNet dataset contains annotated masks for both core and spurious features extracted from a supervised robust ResNet-50 model for 6858 images spanning 327 ImageNet classes. It also contains some natural language keywords, provided by workers to explain each feature. Core features are those that are highly correlated with the ground truth of the image, whereas, spurious features are those that activate portions irrelevant to the ground truth. In Figure 5, we study some correct and mis-classified samples in the SimCLR baseline. We plot the gradient heatmaps of the discriminative features (combining each individual feature heatmap) of Sim-CLR for each respective image. We also plot the core and spurious masks of the same images from the Salient ImageNet dataset. We observe that discriminative SimCLR features mostly capture relevant and defining characteristics of the images, therefore are highly correlated with the groundtruth. Moreover, for every correctly classified image, these heatmaps overlap more with core features than spurious features in Salient ImageNet. Discriminative features in misclassified images also overlap with core features in most cases. Since discriminative features are very closely related (in terms of overlap) to *core* features, we can potentially explain these features better with the help of worker annotations in Salient ImageNet. Therefore, these features can be considered as *interpretable*.

We quantitatively measure the interpretability of a given representation of a given model by computing the Intersection over Union (mIoU) between the heatmap of discriminative features and the core or spurious mask of that image in Salient ImageNet. We can extend this to measure the overall interpretability of a given model by computing the mean Intersection over Union (mIoU) over the population. For the i^{th} image , we define Ar_i as the area of the heatmap of its discriminative features. Let Ar_i^{core} and Ar_i^{sp} be the area of the core and spurious masks respectively. The mIoU scores are defined as follows,

$$mIoU^{core} = \frac{1}{N} \sum_{i} \frac{s(Ar_i \cap Ar_i^{core})}{s(Ar_i \cup Ar_i^{core})}$$
$$mIoU^{sp} = \frac{1}{N} \sum_{i} \frac{s(Ar_i \cap Ar_i^{sp})}{s(Ar_i \cup Ar_i^{sp})}$$

where s(.) calculates the sum of the pixel values of the discriminative features' heatmap in the given area. Higher mIoU^{core}% indicates that, on an average higher percentage of the feature heatmap overlaps with the annotated core region, meaning that the model features are more interpretable.

In Figure 6, we show that for all SSL baselines, mIoU^{core} > mIoU^{sp} for both correct and incorrect classifications which confirms that discriminative features generally encode important and core attributes over the whole population. Among correct and mis-classified samples in the baselines, we observe that the mIoU^{core} of correct classifications is higher than mis-classifications. This aligns with our observations in Figure 4, which shows that discriminative features in mis-classified samples may not be strongly deviated from the mean and therefore, may correspond to less important portions of the image. After Q-Score regularization, we observe an increase in mIoU^{core} for both correct and mis-classified samples compared to the baseline. This shows that our regularization which enhances discriminative features produces better gradient heatmaps which are more overlapped with core portions of images and therefore, improves the overall model interpretability. Note that, spurious feature heatmaps (Figure 5) cover almost all the image content. As shown in Figure 4, discriminative feature heatmaps, after regularization, become larger or smaller to capture all core characteristics of the image. This can cause mIoU^{sp} to be higher or lower after regularization. We therefore only use mIoU^{sp} to compare with mIoU^{core} and not to analyze the regularization effect.



Figure 6: mIoU scores with Salient ImageNet features: We compute the mean mIoU^{core} and mIoU^{sp} scores of SSL baselines (using their discriminative features) before and after Q-Score regularization. We observe that discriminative features for all models generally show higher % IoU with core features than spurious features. Mis-classified representations show relatively lower % IoU with core features. After Q-score regularization, we observe that mIoU^{core} generally improves for both correct and mis-classified representations.

Conclusion

We studied the representation space of SSL models to identify discriminative features in a fully unsupervised manner. Using discriminative features, we compress the representation space by up to 40% without largely affecting the downstream performance. We defined an unsupervised samplewise score, Q-Score, that uses discriminative features and is effective in determining how likely samples are to be correctly or incorrectly classified. We regularized with Q-Score and remedied low-quality samples, thereby, improving the overall accuracy of state-of-the-art SSL models on ImageNet-1K by up to 3.7%, also producing more explainable representations. Our work poses several questions and directions for future work: (i) What are other causes for failures and poor generalization in SSL models apart from representation quality?; (ii) Studying properties for ViT-based representations in a similar fashion is a crucial next step to our work as ViT is a widely used SSL encoder; (iii) How can representations be better structured for non-classificationbased downstream tasks such as object detection and semantic segmentation, where discriminative features should correlate to the object/segment categories (which can be several per-image).

Acknowledgements

This project was supported in part by a grant from an NSF CAREER AWARD 1942230, ONR YIP award N00014-22-1-2271, ARO's Early Career Program Award 310902-00001, Meta grant 23010098, HR00112090132 (DARPA/RED), HR001119S0026 (DARPA/GARD), Army Grant No. W911NF2120076, NIST 60NANB20D134, the NSF award CCF2212458, an Amazon Research Award and an award from Capital One.

References

Arora, S.; Khandeparkar, H.; Khodak, M.; Plevrakis, O.; and Saunshi, N. 2019. A Theoretical Analysis of Contrastive Unsupervised Representation Learning. arXiv:1902.09229.

Bachman, P.; Hjelm, R. D.; and Buchwalter, W. 2019. Learning Representations by Maximizing Mutual Information Across Views. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Bardes, A.; Ponce, J.; and LeCun, Y. 2021. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. arXiv:2105.04906.

Bau, D.; Zhou, B.; Khosla, A.; Oliva, A.; and Torralba, A. 2017. Network Dissection: Quantifying Interpretability of Deep Visual Representations. arXiv:1704.05796.

Bautista, M. A.; Sanakoyeu, A.; Tikhoncheva, E.; and Ommer, B. 2016. CliqueCNN: Deep Unsupervised Exemplar Learning. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

Bojanowski, P.; and Joulin, A. 2017. Unsupervised Learning by Predicting Noise. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 517–526. PMLR.

Bordes, F.; Balestriero, R.; and Vincent, P. 2021. High Fidelity Visualization of What Your Self-Supervised Representation Knows About. arXiv:2112.09164.

Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*.

Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018a. Deep Clustering for Unsupervised Learning of Visual Features. *Lecture Notes in Computer Science*, 139–156.

Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018b. Deep Clustering for Unsupervised Learning of Visual Features. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

Caron, M.; Bojanowski, P.; Mairal, J.; and Joulin, A. 2019. Unsupervised Pre-Training of Image Features on Non-Curated Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 9912–9924. Curran Associates, Inc.

Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9650–9660.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020a. A Simple Framework for Contrastive Learning of Visual Representations. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1597–1607. PMLR.

Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020b. Improved Baselines with Momentum Contrastive Learning. arXiv:2003.04297.

Chen, X.; and He, K. 2021. Exploring Simple Siamese Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (*CVPR*), 15750–15758.

Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; ; and Vedaldi, A. 2014. Describing Textures in the Wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.

Coates, A.; Lee, H.; and Ng, A. Y. 2011. Stanford STL-10 Image Dataset. *AISTATS*.

da Costa, V. G. T.; Fini, E.; Nabi, M.; Sebe, N.; and Ricci, E. 2022. solo-learn: A Library of Self-supervised Methods for Visual Representation Learning. *Journal of Machine Learning Research*, 23(56): 1–6.

Dosovitskiy, A.; Fischer, P.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2016. Discriminative Unsupervised Feature Learning with Exemplar Convolutional Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9): 1734–1747.

Dosovitskiy, A.; Springenberg, J. T.; Riedmiller, M.; and Brox, T. 2014. Discriminative Unsupervised Feature Learning with Convolutional Neural Networks. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.

Dwibedi, D.; Aytar, Y.; Tompson, J.; Sermanet, P.; and Zisserman, A. 2021. With a Little Help from My Friends: Nearest-Neighbor Contrastive Learning of Visual Representations. 2021 IEEE/CVF International Conference on Computer Vision (ICCV).

Elhage, N.; Hume, T.; Olsson, C.; Schiefer, N.; Henighan, T.; Kravec, S.; Hatfield-Dodds, Z.; Lasenby, R.; Drain, D.; Chen, C.; Grosse, R.; McCandlish, S.; Kaplan, J.; Amodei, D.; Wattenberg, M.; and Olah, C. 2022. Toy Models of Superposition. arXiv:2209.10652.

Engstrom, L.; Ilyas, A.; Santurkar, S.; Tsipras, D.; Tran, B.; and Madry, A. 2020. Adversarial Robustness as a Prior for Learned Representations.

Ericsson, L.; Gouk, H.; and Hospedales, T. M. 2021. Why Do Self-Supervised Models Transfer? Investigating the Impact of Invariance on Downstream Tasks. arXiv:2111.11398.

Garrido, Q.; Balestriero, R.; Najman, L.; and Lecun, Y. 2022. RankMe: Assessing the downstream performance of pretrained self-supervised representations by their rank. arXiv:2210.02885.

Grigg, T. G.; Busbridge, D.; Ramapuram, J.; and Webb, R. 2021. Do Self-Supervised and Supervised Methods Learn Similar Visual Representations? arXiv:2110.00528.

Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; Piot, B.; kavukcuoglu, k.; Munos, R.; and Valko, M. 2020. Bootstrap Your Own Latent - A New Approach to Self-Supervised Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 21271–21284. Curran Associates, Inc.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778.

Hernandez, E.; Schwettmann, S.; Bau, D.; Bagashvili, T.; Torralba, A.; and Andreas, J. 2022. Natural Language Descriptions of Deep Visual Features. arXiv:2201.11114.

Huang, J.; Dong, Q.; Gong, S.; and Zhu, X. 2019. Unsupervised Deep Learning by Neighbourhood Discovery. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 2849–2858. PMLR.

Huang, W.; Yi, M.; and Zhao, X. 2021. Towards the Generalization of Contrastive Self-Supervised Learning. arXiv:2111.00743.

Jing, L.; Vincent, P.; LeCun, Y.; and Tian, Y. 2022. Understanding Dimensional Collapse in Contrastive Selfsupervised Learning. In *International Conference on Learning Representations*.

Kalibhat, N.; Bhardwaj, S.; Bruss, B.; Firooz, H.; Sanjabi, M.; and Feizi, S. 2023. Identifying Interpretable Subspaces in Image Representations. arXiv:2307.10504.

Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. arXiv:2004.11362.

Koohpayegani, S. A.; Tejankar, A.; and Pirsiavash, H. 2021. Mean Shift for Self-Supervised Learning. 2021 IEEE/CVF International Conference on Computer Vision (ICCV).

Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3D Object Representations for Fine-Grained Categorization. In 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13). Sydney, Australia.

Krizhevsky, A.; Nair, V.; and Hinton, G. 2009a. CIFAR-10 (Canadian Institute for Advanced Research).

Krizhevsky, A.; Nair, V.; and Hinton, G. 2009b. CIFAR-100 (Canadian Institute for Advanced Research).

Li, A. C.; Efros, A. A.; and Pathak, D. 2022. Understanding Collapse in Non-Contrastive Siamese Representation Learning. arXiv:2209.15007.

Maji, S.; Kannala, J.; Rahtu, E.; Blaschko, M.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. Technical report.

Nilsback, M.-E.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In Indian Conference on Computer Vision, Graphics and Image Processing. Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3): 211–252.

Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2019. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision*, 128(2): 336–359.

Singla, S.; and Feizi, S. 2021. Salient ImageNet: How to discover spurious features in Deep Learning? arXiv:2110.04301.

Tian, Y.; Sun, C.; Poole, B.; Krishnan, D.; Schmid, C.; and Isola, P. 2020. What Makes for Good Views for Contrastive Learning? In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6827–6839. Curran Associates, Inc.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1): 267–288.

Tosh, C.; Krishnamurthy, A.; and Hsu, D. 2021. Contrastive learning, multi-view redundancy, and linear models. arXiv:2008.10150.

von Kügelgen*, J.; Sharma*, Y.; Gresele*, L.; Brendel, W.; Schölkopf, B.; Besserve, M.; and Locatello, F. 2021. Selfsupervised learning with data augmentations provably isolates content from style. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*. *equal contribution.

Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised Feature Learning via Non-Parametric Instance Discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Xiao, T.; Wang, X.; Efros, A. A.; and Darrell, T. 2021. What Should Not Be Contrastive in Contrastive Learning. In *International Conference on Learning Representations*.

YM., A.; C., R.; and A., V. 2020. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*.

You, Y.; Gitman, I.; and Ginsburg, B. 2017. Large Batch Training of Convolutional Networks. *arXiv: Computer Vision and Pattern Recognition*.

Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. arXiv:2103.03230.