
Intriguing Properties of Contrastive Losses

Ting Chen
Google Research
iamtingchen@google.com

Calvin Luo
Google Research
calvinluo@google.com

Lala Li
Google Research
lala@google.com

Abstract

We study three intriguing properties of contrastive learning. First, we generalize the standard contrastive loss to a broader family of losses, and we find that various instantiations of the generalized loss perform similarly under the presence of a multi-layer non-linear projection head. Second, we study if instance-based contrastive learning (with a global image representation) can learn well on images with multiple objects present. We find that meaningful hierarchical local features can be learned despite the fact that these objectives operate on global instance-level features. Finally, we study the phenomenon of *feature suppression* among competing features shared across augmented views, such as “color distribution” vs “object class”. We construct datasets with explicit and controllable competing features and show that, for contrastive learning, a few bits of easy-to-learn shared features can suppress, and even fully prevent, the learning of other sets of competing features. In scenarios where there are multiple objects in an image, the dominant object would suppress the learning of smaller objects. Existing contrastive learning methods critically rely on data augmentation to favor certain sets of features over others, and could suffer from learning saturation for scenarios where existing augmentations cannot fully address the feature suppression. This poses open challenges to existing contrastive learning techniques¹.

1 Introduction

Contrastive learning [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14] has recently achieved great successes in learning visual representations without supervision. As shown in [13, 14], contrastive learning can learn representations that rival supervised learning, and significantly improve the state-of-the-art in semi-supervised learning on ImageNet. One successful use case of contrastive loss for self-supervised learning is to make *augmented* views of the same example agree [1, 2, 13]. A widely used contrastive loss to encourage agreement is based on cross entropy [15, 3, 4, 13]. Given an augmented view of an example, the contrastive prediction task aims to classify a set of candidates into the positive example (i.e. the other augmented view of the same example) and negative ones via the cross entropy loss.

In this work, to understand the effectiveness and limitation of existing contrastive learning methods, we study three intriguing aspects. First, we propose a generalization of the standard contrastive loss, and systematically study their performance differences. Second, we study if the instance-based contrastive learning, for which the contrastive loss operates on global representation of an input image, can learn well on images with multiple objects present, and whether or not it leads to meaningful local features. Finally, we systematically study the feature suppression phenomenon in contrastive learning. The suppression effect occurs among competing features shared across augmented views. For example, with random cropping as the augmentation, “color distribution” and “object class” are often competing features as they are likely shared between two augmented views. The suppression effect among competing features can significantly degenerate the representation quality, or even

¹Code and visualization at <https://contrastive-learning.github.io/intriguing>.

completely disable the learning of certain features, as shown in our experiments. Existing methods critically rely on hand-crafted data augmentation to favor certain sets of competing features than others.

Our main findings and contributions are summarized below.

- We propose a generalized contrastive loss, and show that differences between contrastive losses are small with a deep projection head.
- We show that the instance-based objective widely used in existing contrastive learning methods can learn on images with multiple objects, and also learn meaningful local features despite operating on global image representation.
- We construct three datasets with explicit and controllable competing features to systematically study the feature suppression effect in contrastive learning.
- We show that a few bits of easy-to-learn shared features can suppress, and even fully prevent, the learning of other sets of competing features. In scenarios where there are multiple objects in an image, the dominant object would suppress the learning of smaller objects. This poses open challenges to existing contrastive learning.

2 Generalized contrastive loss and differences among its instantiations

The common contrastive loss used in most recent work is based on cross entropy [15, 3, 4]. Following the notation in [13], the contrastive loss can be defined between two augmented views (i, j) of the same example for a mini-batch of size of n , and can be written as the following.

$$\mathcal{L}^{\text{NT-Xent}} = -\frac{1}{n} \sum_{i,j \in \mathcal{MB}} \log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\sum_{k=1}^{2n} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)} \quad (1)$$

where $\mathbf{z}_i, \mathbf{z}_j$ are hidden representations of two augmented views of the same example; $\text{sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \mathbf{v} / (\|\mathbf{u}\| \|\mathbf{v}\|)$ is the cosine similarity between two vectors; τ is a temperature scalar and \mathcal{MB} is a randomly sampled mini-batch consisting of augmented pairs of images. In [13], a MLP *projection head* is introduced between intermediate layer \mathbf{h} (e.g. output of ResNet encoder) and final output \mathbf{z} . It is shown that the projection head is very beneficial and \mathbf{h} is a much better feature representation than \mathbf{z} .

In this work, we generalize the standard contrastive loss to the following form.

$$\boxed{\mathcal{L}^{\text{generalized contrastive}} = \mathcal{L}^{\text{alignment}} + \lambda \mathcal{L}^{\text{distribution}}} \quad (2)$$

Both terms are defined on hidden representations. $\mathcal{L}^{\text{alignment}}$ encourages representations of augmented views to be consistent, while $\mathcal{L}^{\text{distribution}}$ encourages representations (or a random subset of them) to match a prior distribution (of high entropy). It is not difficult to see that the standard contrastive loss in Eq. 1 is a special case as it can be re-written as follows (scaled by a constant τ).

$$\tau \mathcal{L}^{\text{NT-Xent}} = \underbrace{-\frac{1}{n} \sum_{i,j} \text{sim}(\mathbf{z}_i, \mathbf{z}_j)}_{\mathcal{L}^{\text{alignment}}} + \underbrace{\frac{\tau}{n} \sum_i \log \sum_{k=1}^{2n} \mathbb{1}_{[k \neq i]} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)}_{\mathcal{L}^{\text{distribution}}} \quad (3)$$

This form of factorization in Eq. 3 has been proposed in [16], where the second LogSumExp term is referred to as *uniformity* since it encourages representation to uniformly distributed in the hypersphere. Different from [16], here we generalize the hypersphere uniform distribution and study a wider set of prior distributions for their effectiveness in learning representations.

SWD for supporting diverse prior distributions. One issue of using more diverse set of priors is we *cannot* rely on LogSumExp for matching the distribution. To this end, we resort to the theory of optimal transport, via Sliced Wasserstein Distance (SWD) [17, 18, 19]. For two sets of equal-sized samples from two 1-D distributions, the optimal transport can be obtained by computing two

Algorithm 1 Sliced Wasserstein Distance (SWD) loss.

input: activation vectors $\mathbf{H} \in \mathbb{R}^{b \times d}$, a prior distribution (e.g. Gaussian) sampler \mathcal{S}
draw prior vectors $\mathbf{P} \in \mathbb{R}^{b \times d}$ using \mathcal{S}
generate random orthogonal matrix $\mathbf{W} \in \mathbb{R}^{d \times d}$
make projections: $\mathbf{H}^\perp = \mathbf{H}\mathbf{W}$; $\mathbf{P}^\perp = \mathbf{P}\mathbf{W}$
initialize SWD loss $\ell = 0$
for $j \in \{1, 2, \dots, d'\}$ **do**
 $\ell = \ell + \|\text{sort}(\mathbf{H}_{:,j}^\perp) - \text{sort}(\mathbf{P}_{:,j}^\perp)\|^2$
end for
return $\ell/(dd')$

Table 1: Instantiations of the generalized contrastive loss, i.e. $\mathcal{L}_{\text{align}} + \lambda\mathcal{L}_{\text{distribution}}$, that we use in this work. $\tilde{\mathbf{z}}$ denotes ℓ_2 -normalized $\mathbf{z} \in \mathbb{R}^d$, and is only used for uniform hypersphere prior.

$\mathcal{L}_{\text{align}}$	Prior distribution	$\mathcal{L}_{\text{distribution}}$
$\frac{1}{nd} \sum_{i,j} \ \tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j\ ^2$	Uniform hypersphere	$\frac{1}{n} \sum_i \log \sum_j \exp(\tilde{\mathbf{z}}_i^T \tilde{\mathbf{z}}_j / \tau)$
$\frac{1}{nd} \sum_{i,j} \ \tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_j\ ^2$	Uniform hypersphere	SWD($\tilde{\mathbf{Z}}, Z^{\text{prior}}$)
$\frac{1}{nd} \sum_{i,j} \ \mathbf{z}_i - \mathbf{z}_j\ ^2$	Uniform hypercube	SWD(Z, Z^{prior})
$\frac{1}{nd} \sum_{i,j} \ \mathbf{z}_i - \mathbf{z}_j\ ^2$	Normal distribution	SWD(Z, Z^{prior})

permutations that order the values of both sets of samples respectively. The 1-D Wasserstein distance can then be computed with ℓ_2 distance between the ordered values. For n-D distributions, we first project the samples to n randomly-generated orthogonal 1-D subspaces, and then compute the sum of 1-D Wasserstein distance across all 1-D subspaces. By adjusting the network weights to minimize the SWD, we are able to reduce the mismatch between the distribution of hidden vectors and a known prior distribution. The detailed algorithm can be found in Algorithm 1. With SWD loss, we are able to use a wider set of priors, and Table 1 summarizes instantiations of the generalized contrastive loss with different prior distributions and distribution matching loss.

Connection with mutual information. The connection between the standard contrastive loss and mutual information has been shown before [3, 20], where the contrastive loss (a.k.a. InfoNCE loss [3]) is shown to be a lower bound of the mutual information. To connect the generalized contrastive loss to mutual information, we start by the definition of mutual information between two latent variables U, V , which is $I(U; V) = H(U) - H(U|V)$. Comparing this factorization of mutual information with generalized contrastive loss, it is not difficult to see that: 1) the alignment term $\mathcal{L}_{\text{align}}$ is directly related to $H(U|V)$ which aims to reduce uncertainty of the other views given one view of the example; and 2) the distribution matching term $\mathcal{L}_{\text{distribution}}$ can be considered as a proxy to $H(u)$ for maximizing the entropy in the representation. It is perhaps worth noting that different from mutual information, the generalized contrastive loss (Eq. 2) allows a tunable weight (λ) between the alignment and distribution matching term. The weighting scalar λ is (inversely) related to the temperature τ (details in Appendix A.2).

Comparing different instantiations of generalized contrastive loss. Here we ask: Is it essential to use a uniform hypersphere prior for the effectiveness of contrastive loss? How much difference does it make when distinct generalized contrastive losses are used? To answer this question, we conduct experiments following SimCLR settings [13, 14], and use the linear evaluation protocol. Detailed experimental setup can be found in Appendix A.1.

Figure 1 shows linear evaluation results of models trained with different losses under different training epochs. On CIFAR-10, we see little difference in terms of linear evaluation for variants of the generalized contrastive losses, especially when trained longer than 200 epochs. As for ImageNet, there are some discrepancies between different losses, but they disappear when a deeper 3-layer non-linear projection head is used.

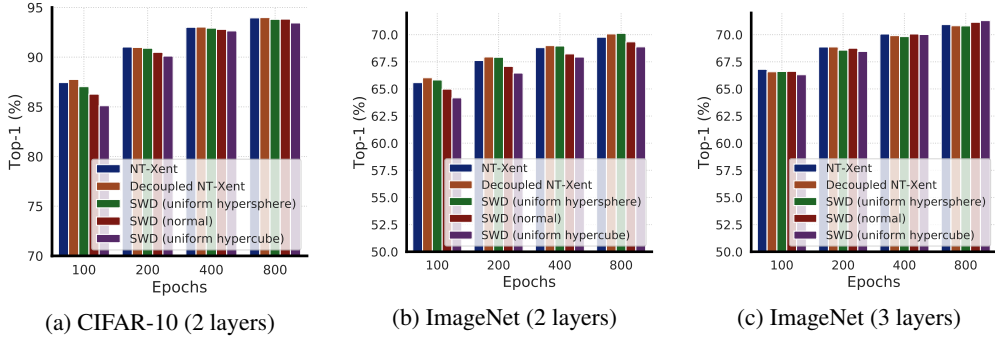


Figure 1: Linear evaluation accuracy of ResNet-50 trained with different losses on CIFAR-10 and ImageNet datasets. Numbers of projection head layers are in parentheses. Differences between variants of generalized contrastive loss are small with a deep projection head. Decoupled NT-Xent loss is introduced in A.2. Numerical results can be found in Appendix A.3.

Furthermore, we find that deep projection head not only reduces the differences among different generalized contrastive losses, but has a similar effect for batch size. With proper learning rate scaling across batch sizes (e.g. square root scaling with LARS optimizer [21]), the impact of batch size on representation quality is small. Table 2 demonstrate this phenomenon for the standard contrastive loss, and more results on other losses can be found in Appendix A.3.

Table 2: Linear eval accuracy of ResNet-50 on ImageNet.

Projection head	Batch size	Epoch			
		100	200	400	800
2 layers	512	65.4	67.3	68.7	69.3
	1024	65.6	67.6	68.8	69.8
	2048	65.3	67.6	69.0	70.1
3 layers	512	66.6	68.4	70.0	71.0
	1024	66.8	68.9	70.1	70.9
	2048	66.8	69.1	70.4	71.3
4 layers	512	66.8	68.8	70.0	70.7
	1024	67.0	69.0	70.4	70.9
	2048	67.0	69.3	70.4	71.3

3 Instance-based objective can learn on images with multiple objects and learn good local features

Most existing contrastive learning methods [13, 10, 22, 4] define their objectives at the instance level where each image is encoded into a single vector representation (e.g. representations of two random crops of the same image instance are treated as a positive pair). In other words, the objective operates on a global representation of its input rather than on some local regions (of its input). We pose two questions regarding instance-based global objective: 1) when there is only a single (dominant) object in the image, the objective seems reasonable as it encourages the model to learn features relevant to object class, but when there are multiple objects present in the image, can instance-based objective still learn well? 2) Since the instance-based objective uses a global summary of its input, can it still learn good local features (e.g. parts of an object, or multiple objects in the same scheme)? To answer these questions, we use SimCLR as representative for the instance-based objective.

3.1 SimCLR can learn on images with multiple objects

Commonly used self-supervised learning datasets, such as MNIST, CIFAR-10, ImageNet, are object centered, i.e. the image is mainly occupied by a single (dominant) object. To experiment with multiple objects in a controllable setting, we propose a new dataset setting by composing multiple digits as follows.

MultiDigits dataset. We place MNIST digits (28×28 size) on a shared canvas (112×112 size). We vary the number of digits placed on the canvas. One factor that could interfere with learning of multiple digits is overlapping digits, therefore we use two placement strategies: random vs in-grid (Figure 2). Random placement of digits incurs no constraint on where digits can be placed on the




Figure 2: MultiDigit dataset. More digits lead to more overlapping in random placement.

canvas, whereas in-grid placement puts each digit in one of the 4×4 grid cells the canvas is divided into, and no two digits can fall in the same cell. In-grid placement ensures no overlapping of digits.

We first pretrain a ResNet-18 with SimCLR or supervised learning with the same augmentation policy (random cropping and resize) on MultiDigits dataset. To access the representation quality, we then train linear classifiers for images with a single digit of size 28×28 on the canvas. Similarly during evaluation, we place only one digit of size 28×28 on the canvas.

As shown in Table 3, representations learned using supervised loss maintains its quality when up to 8 digits are placed in the image. After that the representation becomes worse as the canvas gets more crowded. Notably, representations learned using SimCLR display a similar phenomenon. Regardless of placement strategy, top-1 accuracy stays at the same level up to 8 digits, demonstrating that SimCLR can learn from images with multiple objects. In addition, the increased performance gap between the two placement strategies with increased number of digits shows that object overlapping makes it harder for contrastive losses to learn from multiple objects.

Table 3: Top-1 linear evaluation accuracy (%) for pretrained ResNet-18 on the MultiDigits dataset. We vary the number of digits placed on the canvas during training from 1 to 16. During evaluation only 1 digit is present. As a baseline, a network with random weights gives 18% top-1 accuracy.

		Placing of digits		Number of digits (size 28×28)				
				1	2	4	8	12
Supervised	Random	99.5	99.5	99.3	99.4	98.9	98.3	
	In-grid	99.5	99.6	99.5	99.3	98.6	92.4	
SimCLR	Random	98.9	98.9	99.0	98.9	98.2	96.4	
	In-grid	98.3	98.6	99.1	99.2	99.1	98.3	

3.2 SimCLR learns local features that exhibit hierarchical properties

To understand the local features learned by SimCLR, we apply K-means on intermediate features of the pretrained ResNet with SimCLR, and see how local regions of an image are grouped together. For good representations, we expect that regions of similar objects or object parts should be grouped together.

Specifically, we take a pretrained Resnet-50 $2 \times$ on ImageNet, and run inference on images (from ImageNet validation set and COCO [23]) of size 448×448 . We run K-means with various numbers of clusters on the l2-normalized hidden features from middle layers of the network (e.g. block group 2,3,4 of the ResNet). We also compare SimCLR learned features with supervised learned features, as well as the raw pixel (RGB) features extracted from each 14×14 patch.

Figure 3a shows that as the number of clusters increases, the learned representations tend to group image regions based on parts of the object (i.e. facial components of the dog). This phenomenon appears in both SimCLR and supervised learned features, but not with raw pixel features, indicating meaningful local features learned by SimCLR and supervised learning. In Figure 3b, we compare ResNet intermediate features at different layers, and it suggests that earlier layers contain more edge-related features, while later layers contain more object/part features.

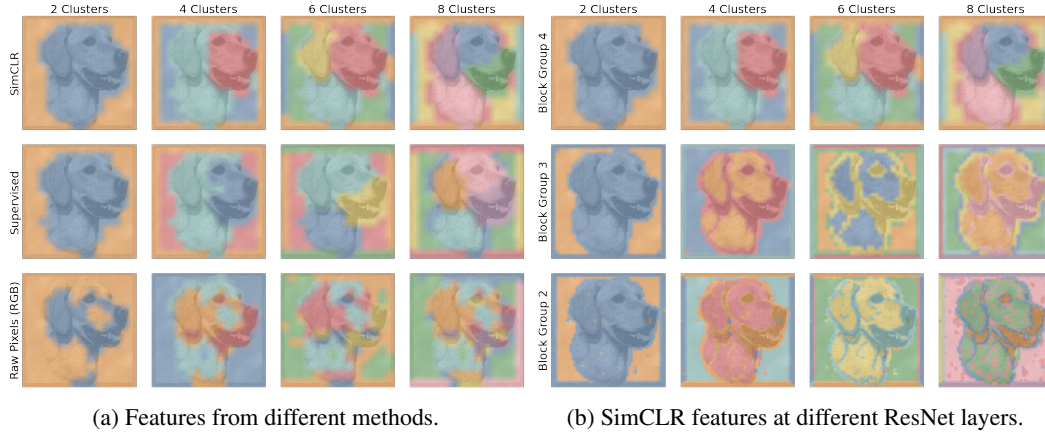


Figure 3: Visualizing features on a ImageNet validation image with K-means clustering. Each row denotes a type of local features used, and each column denotes the number of K-means clusters. Later layers of SimCLR/supervised ResNet tend to group by object parts. More visualization examples can be found in <https://contrastive-learning.github.io/intriguing>.

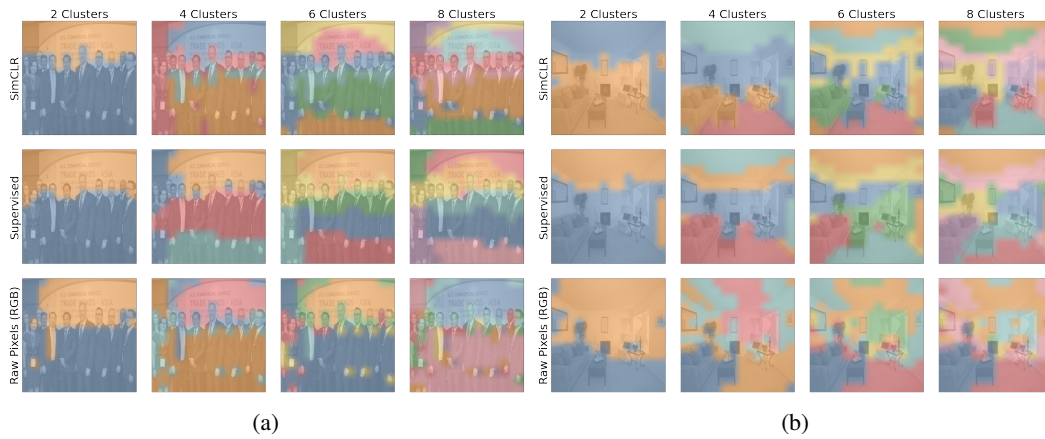


Figure 4: Visualizing features on two images from COCO. Each row denotes a type of local features (SimCLR, Supervised, and raw pixels; both SimCLR and Supervised are trained on ImageNet), and each column denotes the number of K-means clusters. Region grouping by SimCLR/supervised features tend to overlap with object class.

Region grouping results on two COCO images for SimCLR and supervised learning (trained on ImageNet) are shown in Figure 4. Again, region grouping by local features tend to overlap with object class, indicating good local features learned.

4 Feature suppression limits the potential of contrastive learning

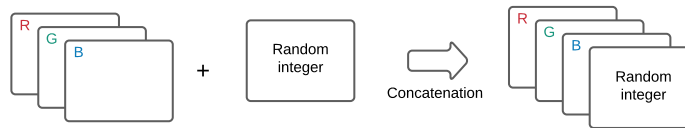
Contrastive learning requires good design of data augmentation to work well. As shown in [13], without color augmentations that randomly shift color distribution (while maintaining information regarding object class), the quality of learned representations are significantly worse. In other words, the presence of “color distribution” features suppresses their competing feature of “object class”, and is addressed by color augmentation. However, there may be scenarios where the known augmentations cannot fully address this feature suppression effect, and it can thus limit the potential of contrastive learning. Here we quantitatively study the feature suppression phenomenon by constructing datasets with explicit and controllable competing features, and see how well contrastive learning method could learn.



(a) ImageNet images overlaid with MNIST digits. The left most column is original image, and others are augmented views via random crop and color distortion. MNIST digits and ImageNet classes are competing features. We vary the number of unique MNIST digits to control the competing features.



(b) Two MNIST digits randomly placed on a shared canvas (of size 112×112). The two digits can have the same size (upper row) or different sizes (lower row), and digits of different sizes can be considered as competing features. We fix the size of one digit and vary the other.



(c) Images (of RGB channels) are concatenated with additional channels of random integer sampled from range of $[1, \log_2(n)]$. The integer, shared between two views, is replicated for spatial dimension and represented as n binary channels. RGB channels and random bits are competing features.

Figure 5: Probing datasets with explicit and controllable competing features.

4.1 Datasets with explicit and controllable competing features

To construct datasets with controllable competing features, we leverage two strategies: *channel addition* that adds different feature information in a shared canvas, and *channel concatenation* that expand the RGB channels to include additional features. With these strategies, we construct three datasets below.

DigitOnImageNet dataset. We overlay MNIST digits on ImageNet images via channel addition/summation (Figure 5a). For each ImageNet image, we assign a unique MNIST digit and replicate it in nine fixed locations before the standard SimCLR augmentations [13] are applied to create augmented views. Therefore the original ImageNet images and added MNIST digits are competing features. Although it is difficult to quantify information in MNIST digits, we can manually control the number of unique MNIST digits used. Ideally, we want the model to learn both set of features so that it could perform well for both MNIST digit and ImageNet object recognition.

MultiDigits dataset (varying the size of one digit). This dataset is modified from MultiDigits introduced above. Here we only consider two digits, and vary the size of one of them (Figure 5b). In this work, we place two digits on a canvas of size 112×112 . We fix the size of one of the digits to be 20×20 while varying the other from 20×20 to 80×80 . Digits of different sizes can be considered as competing features. Ideally, we want the model to learn features for digits of all sizes appeared during training.

RandBit dataset. We concatenate a real image with an image of a random integer in the channel dimension (Figure 5c). The random integer is randomly sampled from range of $[1, \log_2(n)]$ where n is a parameter to control. It is replicated across spatial dimension (i.e. all pixel location shares the same value), and it is also represented as n binary bits/channels instead of an integer or floating

number to make it easily learnable. Furthermore, unlike RGB channels, these additional channels of random bits will *not* be altered by augmentation, so they are identical for both augmented views of the same image. The RGB channels and the added channels of random bits are competing features, and this construction allows us to control the amount of information in the added competing feature, which is n bits. Also, we know that the mutual information between two views given this construction is at least $\log_2(n)$.

4.2 Easy-to-learn features (MNIST digit) suppress the learning of other features (ImageNet object class)

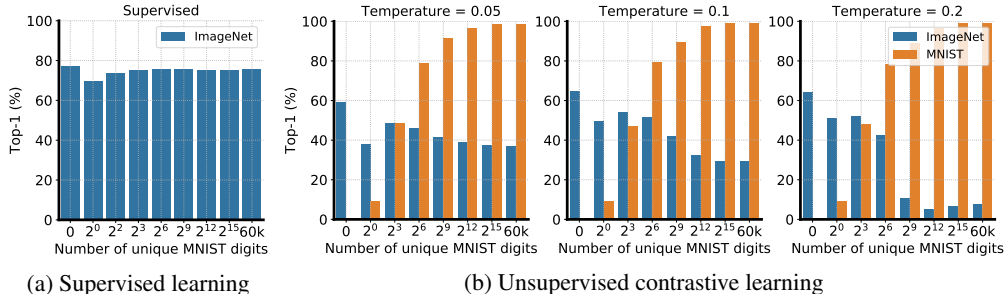


Figure 6: (a) Supervised learning accuracy on ImageNet classification. (b) Linear evaluation of learned features for both MNIST classification and ImageNet classification on the DigitOnImageNet dataset. Batch size of 1024 and 2-layer projection head is used. Different batch sizes and projection head layers have negligible influence on the trade-off between ImageNet vs MNIST accuracy.

On DigitOnImageNet datasets, we vary the number of unique MNIST digits used in the training set, and all MNIST digits are used in the validation/test set. As a baseline, we train supervised ResNet-50 on the created datasets with ImageNet labels, and the number of unique MNIST digits has little impact on the top-1 ImageNet classification accuracy (Figure 6a).

We then train SimCLR on the datasets with different temperatures. As shown in Figure 6b, when we increase the number of unique MNIST digits, the linear evaluation performance of the learned features for MNIST classes increases accordingly, while the accuracy for ImageNet classes decreases dramatically. The trade-off between digit recognition ability and object recognition ability shows that simple features suppress the learning of difficult features, when both are shared between two augmented views. Different batch sizes and projection head depths have negligible influence to the outcome we observe here. Therefore, it is difficult to learn both of the competing features using existing contrastive losses (e.g. SimCLR).

4.3 The presence of dominant object suppresses the learning of features of smaller objects

On the MultiDigits dataset, as mentioned, we fix one digit to be size of 20×20 while varying the other from 20×20 to 80×80 , on a canvas of 112×112 . We first pretrain a ResNet-18 with SimCLR or supervised learning with the same augmentation policy (random cropping and resize) and batch size of 1024. To access the representation quality, we then train linear classifiers for each of the digit sizes that appeared during pretraining. For training of the linear classifier, we only place a single digit at a time on the canvas of the same size as during pretraining.

The results are summarized in Table 4. For supervised learning, the learned representations for the smaller digit do not change much as the other digit increases its size, and the model perform well for both small and large digits (accuracy $> 99\%$). However, for SimCLR, the learned representations of the smaller digit degenerate significantly when the size of the other digit increases, almost to the level of a random untrained network. The dominant object can be learned very well (accuracy $> 99\%$) while suppressing the learning of the smaller object. Although tuning temperature has some effects on reducing the feature suppression, the trend stays unchanged.

Table 4: Top-1 linear evaluation accuracy (%) for pretrained ResNet-18 on the MultiDigits dataset. We fix the size of 1st digit while increasing the size of the 2nd digit. For SimCLR, results are presented for two temperatures. Accuracies suffered from a significant drop when increasing 2nd digit size are red colored.

		2nd digit size (1st digit is kept the same size of 20×20)						
		20×20	30×30	40×40	50×50	60×60	70×70	80×80
Supervised	1st digit	99.1	99.2	99.2	99.2	99.1	99.1	99.0
	2nd digit	99.1	99.5	99.5	99.6	99.5	99.5	99.6
SimCLR ($\tau = 0.05$)	1st digit	97.8	97.6	96.2	96.5	88.5	74.5	39.9
	2nd digit	97.8	97.9	97.8	98.3	98.2	97.7	98.2
SimCLR ($\tau = 0.2$)	1st digit	98.7	98.8	98.3	87.5	24.9	19.8	20.3
	2nd digit	98.7	99.2	99.2	99.0	99.1	98.9	99.4
Random net (untrained)	1st digit	16.5	16.7	16.6	16.6	16.6	16.9	16.5
	2nd digit	16.5	19.1	21.9	24.1	26.5	28.1	29.0

4.4 Extra channels with a few bits of easy-to-learn mutual information suppress the learning of all features in RGB channels

In the RandBit datasets, we add additional channels (identical across pixels) of random bits to MNIST and ImageNet. As mentioned above, SimCLR augmentation is only applied to RGB channels so extra added channels will be shared among two view.

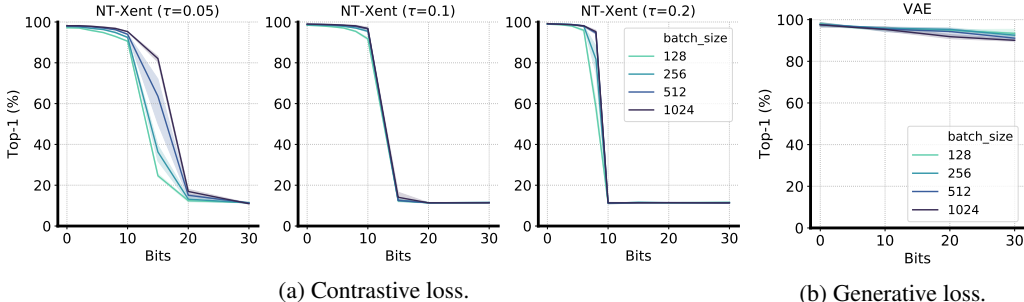


Figure 7: Linear evaluation of learned features when a few bits of competing features added (on MNIST). Adding a few bits completely disables contrastive learning (across various batch size or losses). Interestingly, it has little effects on a generative model (VAE). The detrimental effects are just as strong for larger datasets such as CIFAR-10 and ImageNet (Appendix B.1).

Figure 7 shows the linear evaluation accuracy of models trained on MNIST (with additional random bits added). We observe that the linear evaluation accuracy quickly drops with a few bits of competing feature added. This detrimental effect on the representation quality persists on bigger datasets like CIFAR-10 and ImageNet as well, and cannot be avoided by using different contrastive losses, batch sizes, or memory mechanism based on momentum contrast (details in Appendix B.1). We believe the fact that just a few bits of easy-to-learn features can completely disable the good representation learning is related to the saturation of the distribution matching loss. As shown in Appendix B.2, the linear increase in bits requires an exponential increase in batch size, which is not sustainable as the required batch size can quickly go beyond the size of the dataset size. In practice, we rely on using data augmentation to remove those uninformative easy-to-learn features so that contrastive learning can learn useful representations. Interestingly, the extra bits do not affect a generative model, variational autoencoder [24, 25], nearly as much, despite other settings such as model size are held the same, prompting a potential direction of addressing the issue.

5 Related Work

Our work studies the contrastive loss based on cross entropy loss [15, 3, 4, 13]. This loss is widely used in recent successful contrastive learning methods [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14]. In

terms of the contrastive loss, our work is perhaps most related to [16], which shows that formulating contrastive loss as alignment and uniformity in the hypersphere gives similar performance as the standard contrastive loss. We further generalize this factorization, and show other distribution matching losses can be used, and they could achieve similar results. Other than standard contrastive loss that directly utilize negative examples, BYOL [22] demonstrates another way to maintain representation distribution/entropy without directly relying on distribution matching, and SWAV [26] shows clustering-based method equipped with proper data augmentations could also achieve similar performance. We conduct preliminary experiments of BYOL on RandBit and found that it also suffers from feature suppression as generalized contrastive loss. It is expected that SWAV would exhibit similar behaviors on RandBit as those random bits could fuel representations for perfect clustering.

The connection between contrastive loss and mutual information has been studied before [3, 20]. We show that for the generalized contrastive loss, it can also be related to mutual information. Despite the connection between contrastive loss and mutual information, it has been pointed out that mutual information estimation may suffer from certain limitations [27, 28]. Moreover, [29, 12] show that higher mutual information learned by the network does not warrant better representation quality. In our work, we find adding mutual information bits between two views which are irrelevant to downstream tasks can be harmful for the quality of learned representations. Data augmentation plays an important role at favoring certain bits of mutual information than others.

There is a growing number of recent work on the topic of understanding contrastive learning, both theoretically [30, 31, 32, 33, 34] and empirically [16, 12, 35, 36]. However, little work has been done to study the phenomenon of feature suppression. To our knowledge, we are the first one to quantitatively and systematically study this problem. We believe this is still a very open question and could benefit from more future investigation. Finally, the feature suppression effect in unsupervised contrastive learning that we study in this work may also exist in standard supervised learning (“contrastive loss” between examples and class labels), as suggested by [37, 38], though the specific form would be different.

6 Conclusion

In this work, we study three intriguing properties of contrastive losses. In particular, our results highlight that feature suppression is still an open challenge in contrastive learning. While there is a plethora of work on improving contrastive learning, few of them directly aim to address feature suppression. This limitation of contrastive learning becomes a bottleneck for scenarios where existing augmentation cannot fully address the feature suppression phenomenon, and learning would saturate at a level of dissatisfaction.

We would also like to point out some limitations of our study. Firstly, we focus mostly on contrastive learning with explicit negatives (e.g. SimCLR and MoCo). We believe other methods based on clustering and/or without negative pairs would exhibit similar phenomenon but we leave that as future work. Secondly, many of our proposed image datasets are not fully realistic despite being composed from some (challenging) natural image datasets such as ImageNet. We admit it is very hard to explore competing features or multiple objects in a controllable fashion on realistic large scale image datasets.

Acknowledgements

We specially thank Geoffrey Hinton for many inspiring discussions and helpful advice. We would also like to thank David Fleet, Simon Kornblith, Mohammad Norouzi, Kevin Swersky and Katherine Hermann for insightful discussions. In addition, we are thankful to William Chan and Sara Sabour for ideas on implementation of sorting on TPUs. We also thank the anonymous reviewers for their constructive feedback.

References

- [1] Suzanna Becker and Geoffrey E Hinton. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(6356):161–163, 1992.
- [2] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in neural information processing systems*, pages 766–774, 2014.
- [3] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [4] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3733–3742, 2018.
- [5] 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.
- [6] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems*, pages 15509–15519, 2019.
- [7] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019.
- [8] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- [9] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. *arXiv preprint arXiv:1912.01991*, 2019.
- [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- [11] Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020.
- [12] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. *arXiv preprint arXiv:2005.10243*, 2020.
- [13] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [14] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [15] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Advances in neural information processing systems*, pages 1857–1865, 2016.
- [16] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. *arXiv preprint arXiv:2005.10242*, 2020.
- [17] Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.
- [18] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1):22–45, 2015.
- [19] Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. Generalized sliced wasserstein distances. In *Advances in Neural Information Processing Systems*, pages 261–272, 2019.
- [20] Ben Poole, Sherjil Ozair, Aaron van den Oord, Alexander A Alemi, and George Tucker. On variational bounds of mutual information. *arXiv preprint arXiv:1905.06922*, 2019.
- [21] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- [22] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning, 2020.
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

- [24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [25] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [26] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020.
- [27] David McAllester and Karl Stratos. Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pages 875–884, 2020.
- [28] Jiaming Song and Stefano Ermon. Understanding the limitations of variational mutual information estimators. *arXiv preprint arXiv:1910.06222*, 2019.
- [29] 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.
- [30] Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv preprint arXiv:1902.09229*, 2019.
- [31] Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Demystifying self-supervised learning: An information-theoretical framework. *arXiv preprint arXiv:2006.05576*, 2020.
- [32] Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. *arXiv preprint arXiv:2008.10150*, 2020.
- [33] Yuandong Tian, Lantao Yu, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning with dual deep networks. *arXiv preprint arXiv:2010.00578*, 2020.
- [34] Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *arXiv preprint arXiv:2008.01064*, 2020.
- [35] Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. What makes instance discrimination good for transfer learning? *arXiv preprint arXiv:2006.06606*, 2020.
- [36] Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. *arXiv preprint arXiv:2007.13916*, 2020.
- [37] Katherine L Hermann and Andrew K Lampinen. What shapes feature representations? exploring datasets, architectures, and training. *arXiv preprint arXiv:2006.12433*, 2020.
- [38] Harshay Shah, Kaustav Tamuly, Aditi Raghunathan, Prateek Jain, and Praneeth Netrapalli. The pitfalls of simplicity bias in neural networks. *arXiv preprint arXiv:2006.07710*, 2020.
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.