

Divide and Contrast: Self-supervised Learning from Uncurated Data

Yonglong Tian
MIT

Olivier J. Hénaff
DeepMind

Aaron van den Oord
DeepMind

Abstract

Self-supervised learning holds promise in leveraging large amounts of unlabeled data, however much of its progress has thus far been limited to highly curated pre-training data such as ImageNet. We explore the effects of contrastive learning from larger, less-curated image datasets such as YFCC, and find there is indeed a large difference in the resulting representation quality. We hypothesize that this curation gap is due to a shift in the distribution of image classes—which is more diverse and heavy-tailed—resulting in less relevant negative samples to learn from. We test this hypothesis with a new approach, Divide and Contrast (DnC), which alternates between contrastive learning and clustering-based hard negative mining. When pretrained on less curated datasets, DnC greatly improves the performance of self-supervised learning on downstream tasks, while remaining competitive with the current state-of-the-art on curated datasets.

1. Introduction

Recent developments in self-supervised learning have shown that it is possible to learn high-level representations of object categories from unlabeled images [39, 43, 88, 12, 97], phonetic information from speech [69, 79] and language understanding from raw text [21, 101]. The most studied benchmark in self-supervised learning is ImageNet [20], where representations learned from unlabeled images can surpass supervised representations, both in terms of their data-efficiency and transfer-learning performance [13, 34].

One of the caveats with self-supervised learning on ImageNet is that it is not completely “self-supervised”. The training set of ImageNet, on which the representations are learned, is heavily curated and required extensive human effort to create [20]. In particular, ImageNet contains many fine-grained classes (such as subtly different dog breeds), each one containing roughly the same number of images. While this *consistency* may facilitate the learning of high-level visual representations, limiting self-supervised learning to such curated datasets risks biasing their development

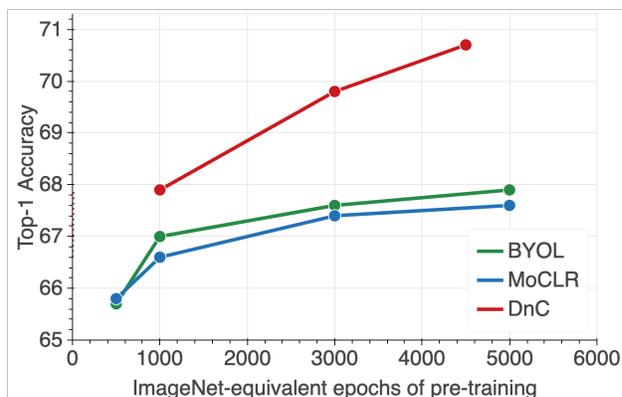


Figure 1. Linear evaluation on ImageNet of representations learned from a large-scale uncurated dataset using ResNet-50. Divide and Contrast (DnC) is better able to handle the diverse and long-tailed distribution of images and improves more with longer training. X-axis represents total computation, in ImageNet-equivalent epochs.

towards methods which require this consistency, limiting their applicability to more diverse downstream tasks and larger datasets for pre-training.

In this paper we assess how well recent self-supervised learning methods perform on downstream tasks (including ImageNet) when they are pre-trained on significantly less curated datasets, such as YFCC100M [87]. We observe a notable drop in performance of over 9% Top-1 accuracy (from 74.3% to 65.3%) for a ResNet50 model trained with the current state-of-the-art in self-supervised learning.

We hypothesize that this *curation gap* is due to the heavy-tailed nature of images collected in the wild, which present much more diverse content, breaking the *global consistency* exploited in previous datasets. We test this hypothesis with a new method, Divide and Contrast (DnC), which attempts to recover *local consistency* in subsets of the larger, uncurated dataset, such that self-supervised learning methods can learn high-level features that are specific to each subset. We find that such semantically coherent subsets can be straightforwardly obtained by clustering the representations of standard self-supervised models.

Divide and Contrast (DnC) proceeds by training individual “expert” models on each subset and distilling them into

a single model. As a result, DnC can be used in combination with any self-supervised learning technique, and requires the same amount of computation, as each expert is trained for significantly less time. Finally, this computation is trivially parallelized, allowing it to be scaled to massive datasets.

The remainder of this paper is structured as follows. We first review related work in self-supervised learning. We then present a new stronger baseline (MoCLR) which improves over current contrastive methods, matching the performance of the current state-of-the-art (BYOL [34]). Next we present the main method, Divide and Contrast, and how this model can be used together with any SSL method. In the experiments we evaluate the different hypotheses that support DnC, and compare its ability to learn from uncurated dataset with existing methods.

2. Related Work

Recent self-supervised representation learning generally includes three types of methods: *generative models* that directly model the data distribution, *pretext tasks* that are manually designed according to the data, and *contrastive learning* that contrasts positive pairs with negative pairs.

Generative models. While the primary goal of generative models such as GAN [30, 14] or VAE [50] is to model the data distribution (e.g., sample new data or estimate likelihood), the encoder network can also extract good representations [76]. Recent state of the art generative models for representation learning include BiGAN [24] and BigBiGAN [25], which learn a bidirectional mapping between the latent codes and the images, and iGPT [11] which trains an autoregressive model on raw pixels.

Pretext tasks. Good representations may also be learned by solving various pretext tasks. Examples include denoising [93], relative patch prediction [22], image inpainting [73], noise prediction [5], colorization [107, 108, 94], Jigsaw [68], exemplar modeling [26], motion segmentation [72], image transformation prediction [29, 106], tracking [96], or even the combination of multiple tasks [23]. Another line of methods generates pseudo labels by clustering features [8, 9, 46, 105, 1]. Most recently, SeLa [102] jointly clusters images and balances the clusters. SwAV [10] learns representations by having different views of the same image assigned to the same cluster. Another work [45] directly optimizes the transferability of representation by integrating clustering with meta-learning.

Contrastive learning. Contrastive learning is a widely-used generic method. The loss function for contrastive learning has evolved from early margin-based binary classification [36], to triplet loss [80], and to recent k-pair loss [82, 69]. The core idea lying at the heart of the recent series of self-supervised contrastive learning methods [97, 69, 43, 88, 113, 3, 39, 64, 12, 15, 90, 13, 55, 7]

is to maximize the agreement between two “views” of the same image while repulsing “views” from different images. Such views can be created by color decomposition [88], patch cropping [69, 43, 3], data augmentation [12, 13, 83], or image segmentation [42, 92, 109]. Indeed, contrastive learning is very general such that it can be easily adapted to different data types. Examples include different frames of video [69, 112, 81, 37, 31, 38], point clouds [100], multiple sensory data [65, 18, 74], text and its context [63, 101, 59, 51], or video and language [84, 62, 56]. A set of other work [2, 90, 110, 99, 91, 75, 95] focuses on providing empirical and theoretical understanding of contrastive learning. Recently a non-contrastive method BYOL [34] applies a momentum-encoder to one view and predicts its output from the other, inspired by bootstrapping RL [35]. Finally, contrastive learning has also been applied to supervised image classification [48], image translation [70], knowledge distillation [89, 77], and adversarial learning [49].

This paper is also related to knowledge distillation [44]. In [44], several expert models were also trained in parallel on a large scale dataset, and then distilled into a single model. While labels are assumed available in [44] to partition the dataset and distill into a single model, we are dealing with self-supervised learning without supervision. Our distillation procedure is also inspired by FitNet [78].

Lastly, while self-supervised representation learning on uncurated datasets is largely unexplored, there are a few prior attempts [9, 33]. In [9], clustering is applied to generate training targets, and in order to capture the long-tailed distribution of images in the uncurated YFCC100m [87], a hierarchical formulation is proposed. The work of [33] benchmarked pretext-based self-supervised methods in a large scale setting, e.g., jigsaw, colorization and rotation prediction, and found that these pretext tasks are not ‘hard’ enough to take full advantage of large scale data. Concurrent work SEER [32] directly scales up SwAV with larger models and datasets.

3. Divide and Contrast

Though Divide and Contrast can be used in combination with any self-supervised learning technique, in this paper we will combine it with recent state of the art techniques (BYOL, SimCLR, MoCo), such that the model can be compared to a strong baseline and make the experiments relevant with respect to recent developments in the literature. We will start by introducing our baseline, MoCLR, which is a simple hybrid based on BYOL [34], SimCLR [12] and MoCo [39], and as a contrastive method outperforms SimCLR v2 [13], achieving similar performance to BYOL (by using a momentum encoder similar to BYOL and MoCo). Even though DnC can be coupled with BYOL, empirically we have found it to work better with methods that use a contrastive loss.

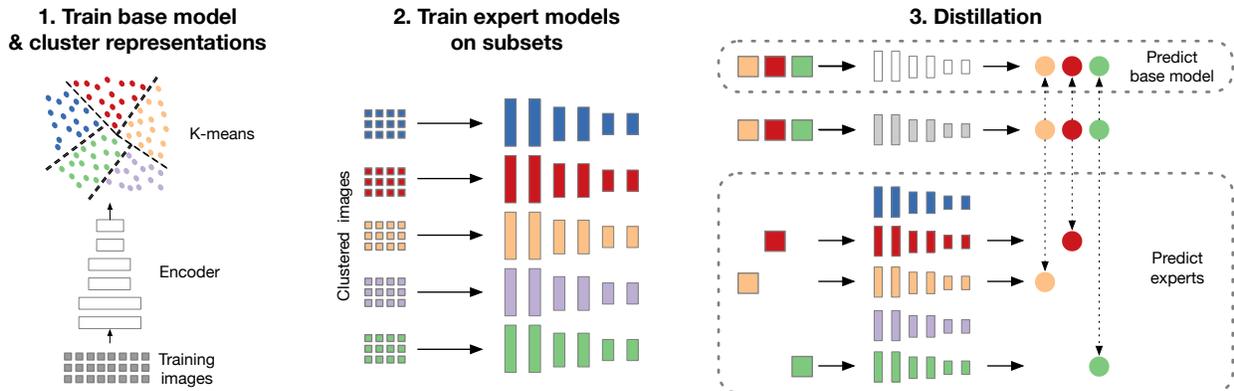


Figure 2. Overview of Divide and Contrast (DnC). DnC can be used in conjunction with any self-supervised learning method (we use MoCLR, an improvement to SimCLR). In the first step a self-supervised learning method is trained on the whole dataset, which we call the base model. The image representations of the base model are then clustered with k-means into 5, 10 or more groups. In the second step, the clustered dataset is then used to train an expert model on each of the image clusters. In the third step the experts and base model are distilled into a single model by predicting their representations. By splitting the dataset into semantically-similar subsets, contrastive methods need to pay more attention to the differences between the images in those clusters and learn more specific representations.

3.1. An Improved Contrastive Baseline: MoCLR

MoCLR roughly uses a similar setup to SimCLR and BYOL, we will briefly describe the main components of this setup and highlight the differences.

Augmenting two views. Given an image x and two distributions of image augmentations \mathcal{T} and \mathcal{T}' , two *views* are created $v \triangleq t(x)$ and $v' \triangleq t'(x)$ by respectively applying random image augmentations $t \sim \mathcal{T}$ and $t' \sim \mathcal{T}'$ from these distributions. The augmentations \mathcal{T} and \mathcal{T}' here are exactly the same as in BYOL [34].

Architecture. The first augmented view v is fed into an online encoder $f(\cdot)$, followed up with an MLP projection head $g(\cdot)$ to produce a projection z . Similarly, an exponential moving average of the online encoder and projection head, also known as momentum encoder [39] or mean teacher [86], is applied on the second view v' to generate z' . The MLP head consists of two layers with a hidden dimension of 4096 and output size of 256, similar to BYOL [34].

Loss function. Given a batch \mathcal{B} , we follow the InfoNCE loss [69] with a cosine similarity function $s(\cdot, \cdot)$ and a scalar temperature value τ :

$$\mathcal{L}_{\text{NCE}} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{e^{s(z_i, z'_i)/\tau}}{e^{s(z_i, z'_i)/\tau} + \sum_{j \in \mathcal{B}/i} e^{s(z_i, z'_j)/\tau}} \quad (1)$$

We symmetrize the loss \mathcal{L}_{NCE} by separately feeding v' to the online network and v to the momentum encoder, resulting in $\tilde{\mathcal{L}}_{\text{NCE}}$. The final loss is $\mathcal{L} = \mathcal{L}_{\text{NCE}} + \tilde{\mathcal{L}}_{\text{NCE}}$.

The difference between MoCLR and other standard methods are as follows. Compared with SimCLR [12], we use a momentum encoder, and double the size of the projection head (from 2048 to 4096 for the hidden layer, and from 128 to 256 for the output layer). In comparison with BYOL [34], we remove the predictor head and use the con-

Table 1. Evaluating the MoCLR baseline which is used in the proposed Divide And Contrast (DnC) model. This comparison is on the ImageNet linear classification benchmark. MoCLR is a contrastive method which achieves similar performance as BYOL, while requiring only a small change to SimCLR (see Section 3.1).

Method	Epochs	Top-1	Top-5
SimCLR [12]	1000	69.3	89.0
SimCLR v2 [13]	1000	71.7	90.4
MoCo v3 [16]	800	73.8	-
BYOL [34]	1000	74.3	91.6
MoCLR (ours)	1000	74.3	92.2

trastive loss instead of the mean squared prediction loss. While concurrent work MoCo v3 [16] inherits from BYOL the asymmetric “projector & predictor” design (the online encoder has an additional predictor compared to the momentum network), our MoCLR removes the predictor for simplicity.

In our experiments we set the batch size to 4096 and do not use a memory buffer [97, 88, 39]. As we will show in our experiments, with these simple changes our *MoCLR* baseline trained for 1,000 epochs outperforms SimCLR v1/v2 [13] and is on par with BYOL [34], see Table 1.

3.2. Divide and Contrast

The motivation behind Divide and Contrast is that, when training on diverse, large-scale datasets, the density of informative negatives will be sparse if we sample randomly from the whole dataset. Instead, if we contrast locally between semantically-similar classes, the sampled negatives will be more informative and the learned model will capture a more discriminative representation.

As visualized in Figure 2, the training of our DnC model

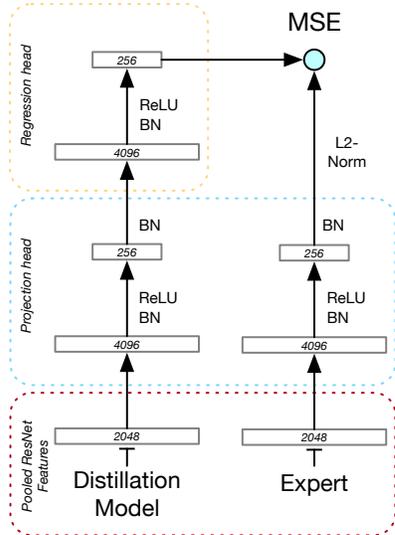


Figure 3. Detailed architectural overview of the distillation of the expert’s features. The distillation model and experts are applied to the same augmented image view.

consists of three stages:

(1) We first train a MoCLR model on the given dataset for N_1 epochs (though other self-supervised learning methods can be used as well). We will call it the **base** model. We use the base model to extract representations for a set of samples in the training set, and cluster them in to K clusters. With these clusters we partition the dataset into K subsets.

(2) For each subset, we train a separate MoCLR model from scratch, which we call **expert** models. In this stage, we distribute a *total* computational budget of N_2 epochs (measured on the whole dataset) to these expert models, proportionally to their corresponding cluster sizes.

(3) Finally, given a base model that captures the general knowledge of the dataset and expert models focusing on locally similar categories, we distill knowledge from these models into a **distillation** model. In this stage, we train for N_3 epochs.

The encoder-architecture for the base model, expert models, and the distillation model are all identical. Therefore, the computational footprint can roughly be measured by summing up the training epochs across all three stages, resulting in a total training of $N_1 + N_2 + N_3$ epochs (except for the clustering overhead and extra forward pass during distillation, which we discuss later).

3.3. Distillation

To leverage the information learned by each of the different experts and more general information from the base model, we distill their representation into one model in the last stage of training. During the distillation we use a single augmented image (instead of the 2-view setup) in combination with a simple regression loss to predict the representa-

tions in these models (no contrastive loss).

The distillation model’s architecture is mostly identical to the other models and is visualized in Figure 3. All have a backbone encoder $f(\cdot)$ and an MLP projection head $g(\cdot)$. On top of the projection head in the distillation model there are $K + 1$ regression networks: $r_k(\cdot)$, $k = 1, \dots, K$, one to predict each of the K expert models and another regression network $r_b(\cdot)$ to predict the base model. The architecture of these regressors is the same as the projection head, except that we remove the final global BatchNorm after the last output layer.

For distillation we use the same augmentation as during self-supervised learning. Given an augmented input image x with clustering id k , we feed it into the distillation model to produce the projection-head output z . Similarly we get z_b and z_k from the base model and the k -th expert model respectively. We also ℓ_2 -normalize z_b and z_k to be unit-norm. The distillation objective is then the average of the two mean squared errors:

$$\mathcal{L}(x) = \frac{1}{2} \|r_b(z) - z_b\|_2^2 + \frac{1}{2} \|r_k(z) - z_k\|_2^2 \quad (2)$$

Note that the outputs of r_b and r_k are *not* ℓ_2 normalized.

To make it possible to compare to our baseline methods in terms of the number of epochs trained, we use two augmented views from the same input image and average their losses. This is otherwise not necessary and alternatively one could also increase the batch-size.

The computational cost in this stage is slightly higher than the self-supervised learning stage (e.g., BYOL and MoCLR), as there are now **two** forward passes (for the expert and base model) for each view (not backward pass and gradient computation). In contrast, BYOL and MoCLR only need **one** forward pass from the momentum encoder. However, we found that always feeding a center crop to expert and base models only leads to very marginal drop in performance (instead of an augmented view). This strategy offers the possibility of first doing a single forward pass over the dataset and storing the activations offline.

4. Experiments

In this section, we compare DnC to BYOL and MoCLR by pre-training on two large-scale uncurated datasets and evaluating transfer performance on different downstream tasks.

Datasets. We consider two large-scale uncurated datasets. The first is a private dataset of roughly 300 million images (JFT-300M [85]). For the second dataset we use YFCC100M [87], a public dataset of 95M Flickr images available under the Creative Commons license. Figures 4 and 5 show a visual comparison between images from ImageNet and YFCC100M. ImageNet images often contain the object or animal of interest in the center of the image. ImageNet also does not have a long-tailed distribution (e.g.,

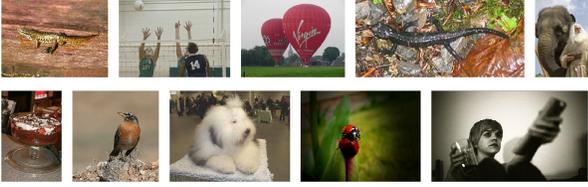


Figure 4. Example ImageNet images



Figure 5. Example YFCC100M images

power law) over object-classes but only considers a specific set of 1000 different classes, which are (roughly) equally represented in the dataset. As a result, specific objects or animals (e.g., common tench, Bedlington terrier, . . .) are over-represented compared to more typically occurring scenes such as human faces and landscapes (which are better represented in YFCC100M).

Settings. ResNet-50 [41] is used in all experiments, unless noted otherwise. For ease of comparison, we report the computational footprint of all experiments in ImageNet-epoch equivalents (e.g. 1 “epoch” = 1281167/batch_size iterations). More implementation and optimization details are included in Appendix.

Table 2. We consider three different training schedules for DnC. The number of total training epochs in each stage, as well as the number of clusters, is specified.

Schedules	Base epochs	Experts total epochs	Distillation epochs
1,000 epochs	200	600 (5 clusters)	200
3,000 epochs	1,000	1,500 (5 clusters)	500
4,500 epochs	1,000	3,000 (10 clusters)	500

DnC Schedules. Table 2 shows three training schedules with different number of epochs. For example, in the schedule of 3,000 epochs, we first train the base model for 1,000 epochs, after which we cluster the samples into 5 groups. The 5 experts are trained in parallel on these subsets. We use 1,500 epochs in total, spread out over the experts according to the number of images in each cluster (300 on average per expert). The distillation model is then trained for 500 epochs. See Section ?? for analysis of run time.

4.1. Linear Evaluation on ImageNet and Places-365

Table 3 shows the results of models pre-trained on YFCC100M and JFT-300M and tested on ImageNet and Places-365 [111] with linear evaluation, i.e., features are frozen and a linear classifier is trained. For JFT-300M the

Table 3. Comparison of self-supervised learning methods pre-trained on uncurated datasets YFCC100M and JFT-300M. For evaluation a linear classifier is trained on ImageNet and Places-365. Computation is measured as ImageNet-equivalent epochs.

Method	Arch	pre-training # epochs	ImageNet Top-1 Acc	Places 365 Top-1 Acc
<i>Concurrent work trained on IG 1B images:</i>				
SEER [32]	R-50	≈1,000	61.6	-
	R-101	≈1,000	65.8	-
<i>Pre-training on YFCC100M:</i>				
MoCLR	R-50	1,000	65.1	53.2
BYOL		1,000	65.3	52.9
MoCLR		3,000	65.7	53.2
BYOL	R-50	3,000	66.6	52.9
DnC		3,000	67.8	54.1
MoCLR		5,000	66.1	53.5
BYOL	R-50	5,000	67.0	53.2
DnC		4,500	68.5	54.4
<i>Pre-training on JFT-300M:</i>				
MoCLR		1,000	66.6	52.1
BYOL	R-50	1,000	67.0	51.9
DnC		1,000	67.9	52.5
MoCLR		3,000	67.4	52.5
BYOL	R-50	3,000	67.6	52.4
DnC		3,000	69.8	53.3
MoCLR		5,000	67.6	52.4
BYOL	R-50	5,000	67.9	52.4
DnC		4,500	70.7	53.5
<i>With larger ResNet:</i>				
MoCLR	R-200x2	3,000	74.2	54.6
DnC		3,000	77.3	56.2

results are also visualized in Figure 1. On the ImageNet linear benchmark we see a large drop in performance compared to pre-training on ImageNet (Table 1): **-9.0%**, **-7.3%** for BYOL-1k and **-9.2%**, **-7.7%** for MoCLR-1k, showing the difficulty of learning representations from uncurated (and more diverse) data.

In these experiments DnC always uses MoCLR-1k for the clustering, and uses the remaining pre-training epochs for the expert training and distillation. Therefore a good comparison is with MoCLR trained for longer (from scratch). For 3,000 epochs of training, MoCLR-3k improves over MoCLR-1k by **+0.6** and **+0.8** on YFCC100M and JFT-300M respectively, while DnC-3k improves by **+2.7** and **+3.2**. On Places-365 we see similar relative improvements. We also include BYOL-3k for completeness and again see small improvements with longer training for BYOL. For further longer schedules (4,500-5,000 epochs), we notice similar behavior on both YFCC100M and JFT-300M. Besides, we see DnC significantly outperforms concurrent efforts SEER [32] when using ResNet-50. We further test DnC with a larger model (i.e., ResNet-200 with a

Table 4. Transfer learning experiments. We evaluate models pre-trained on ImageNet, YFCC100M and JFT-300M with a linear classifier on 12 downstream classification tasks: Food-101 [6], CIFAR-10/100 [54], Birdsnap [4], SUN397 [98], Stanford Cars [53], FGVC Aircraft [61], PASCAL VOC 2007 [27], Describable Textures (DTD) [19], Oxford-IIIT Pets [71], Caltech-101 [28] and Oxford 102 Flowers [67].

		Food-101	CIFAR10	CIFAR100	Birdsnap	SUN397	Cars	Aircraft	VOC2007	DTD	Pets	Caltech-101	Flowers	Average
YFCC	BYOL-5k	69.1	85.8	66.8	35.5	64.1	50.1	51.9	82.5	74.5	74.0	87.6	95.8	69.8
	MoCLR-5k	68.4	87.6	69.7	30.5	63.9	41.0	46.7	82.4	76.2	68.5	86.0	93.0	67.8
	DnC-4.5k	72.1	88.0	71.1	35.5	67.2	52.6	49.2	83.7	76.5	75.9	87.0	97.8	71.4
JFT-300M	BYOL-5k	73.3	89.8	72.4	38.2	61.8	64.4	54.4	81.3	75.5	77.0	90.1	94.3	72.7
	MoCLR-5k	72.8	90.7	72.5	33.8	62.2	60.6	50.9	81.9	75.3	75.8	89.5	93.8	71.7
	DnC-4.5k	78.7	91.7	74.9	42.1	65.0	75.3	54.1	83.1	76.6	86.1	90.2	98.2	76.3

Table 5. Fine-tuning pre-trained model for transfer learning experiments, including object detection on COCO dataset, semantic segmentation on Pascal VOC 2012, and depth estimation on NYU v2 dataset. For the evaluation metrics of *rms* and *rel*, lower is better.

		COCO detection			COCO instance seg.			PASCAL seg.	NYU v2 depth estimation				
		AP ^{bb}	AP ^{bb} ₅₀	AP ^{bb} ₇₅	AP ^{mk}	AP ^{mk} ₅₀	AP ^{mk} ₇₅	mIoU	<1.25	<1.25 ²	<1.25 ³	rms↓	rel↓
ImageNet Super.		39.5	60.1	43.3	35.4	56.9	38.1	74.4	81.1	95.3	98.8	0.573	0.127
YFCC	BYOL-5k	41.1	62.0	45.1	36.6	58.6	38.9	75.5	83.5	96.4	99.0	0.558	0.130
	MoCLR-5k	40.8	61.7	44.8	36.6	58.5	39.0	75.1	86.7	97.4	99.3	0.503	0.117
	DnC-4.5k	41.5	62.5	45.6	37.0	59.3	39.6	76.6	86.2	97.2	99.3	0.512	0.121
JFT-300M	BYOL-5k	40.6	61.2	44.3	36.2	58.1	38.8	75.8	84.4	96.5	99.0	0.544	0.129
	MoCLR-5k	41.1	62.0	45.4	36.9	58.9	39.5	76.1	86.3	97.2	99.3	0.513	0.120
	DnC-4.5k	41.7	62.5	45.9	37.2	59.3	39.8	76.9	86.1	97.2	99.4	0.509	0.119

width multiplier of 2) and observe that DnC outperforms MoCLR by **+3.1**.

4.2. Transfer Learning

In this section, we consider both using frozen representations for fine-grained linear classification and fine-tuning for different downstream tasks.

Fine-grained linear classification. Following SimCLR [12] and BYOL [34], we further perform linear classification evaluation on 12 classification datasets (introduced by [52]), to assess whether the learned representation is generic across different image domains (see more details in Section ??). As shown in Table 4, when pre-training on YFCC100M or JFT-300M, DnC significantly and consistently outperforms BYOL and MoCLR.

Detection, segmentation, and depth estimation. In Table 5, we evaluate the representation on three different fine-tuning tasks: (1) for object detection and instance segmentation on COCO [58], we train a standard Mask-RCNN [40] using FPN [57] with a 1x schedule, *i.e.*, 12 epochs; (2) for semantic segmentation on VOC2012, we used FCN [60] as in [39]; (3) for depth estimation on NYU-v2 dataset [66], the setup is the same as [34]. In all three tasks, DnC significantly outperforms ImageNet supervised pre-training, *e.g.*, **+2.2** in AP^{bb} and **+1.8** in AP^{mk} for detection, **+2.5** in mIoU for segmentation, and **+5.0** in <1.25 metric for

depth prediction. DnC also significantly outperforms both self-supervised baselines when transferring to PASCAL and COCO tasks, and performs on-par with MoCLR while outperforming BYOL for depth estimation.

For more implementation details, please refer to Section ?? in the Appendix; Complete results for transfer learning are included in Section ??.

5. Hypothesis and Analysis

The Divide and Contrast (DnC) method hinges on two main hypotheses. The first hypothesis is that clustering activations of powerful self-supervised learning models should provide us with *locally consistent* clusters of images (*e.g.* having similar class labels). The second is that contrasting against similar (but different) object categories allows self-supervised methods to learn more fine-grained, discriminative representations.

We empirically assess these hypotheses in isolation. Next we compare DnC with current state of the art methods on ImageNet to see how well it performs on standard (curated) datasets, and analyze the design choices of DnC.

5.1. Clustered Representations are Object Categories

Our first hypothesis is that clusters of self-supervised representations are semantically meaningful. To this effect

Table 6. Evaluating clustered representations from self-supervised learning methods using Top-1 accuracy and Mutual Information (MI) with the class labels. The representations of each method are clustered using k-means with 1000 centroids. To compute Top-1 accuracy, every cluster is mapped to the most frequent class label for the images in that cluster.

Method	Layer	Dimension	Top-1 Acc	MI
SimCLR [12]	pool	2048	30.1	5.64
	hidden	2048	33.3	5.83
	output	128	31.8	5.68
BYOL [34]	pool	2048	39.9	6.23
	hidden	4096	51.0	6.99
	output	256	50.0	6.78
MoCLR (ours)	pool	2048	40.1	6.26
	hidden	4096	51.6	7.19
	output	256	49.8	7.11

we cluster the representations of various self-supervised learning methods trained on ImageNet with k-means. Specifically, we consider representations from three different layers: the `pool` layer right after the mean pooling, the `hidden` layer of the projection head, and the final projection.

We start with 1000-way clustering, and assign every cluster to a single ImageNet class with a simple majority vote to measure the Top-1 Accuracy. We also measure the mutual information between the clustering assignments and the class labels. Table 6 gives an overview of these results. In particular, these methods can group images from the same category surprisingly well, with some representations achieving over 50% Top-1 clustering accuracy. We also notice that the `hidden` layer performs the best for all methods, and thus use this layer for clustering in the DnC method.

To give an orthogonal view with a smaller number of clusters, Figure 6 plots the 5 clusters used in the DnC model for ImageNet (based on the clustering of a MoCLR ResNet-50). Qualitatively, it appears that groups of classes are jointly assigned to the same cluster. Indeed, the fraction of images in each class that belong to the same cluster is 87.4%, lending further evidence that individual clusters are semantically coherent.

5.2. Training on Semantically Similar Data Subsets

DnC is based on the second hypothesis that training self-supervised learning methods on a subset of images from similar object categories should improve performance on those object classes. Contrastive methods in particular stand to benefit from this procedure, as distinguishing positive samples from negatives from nearby classes might require learning more fine-grained features (similarly to hard-negative mining [80, 47]). On the other hand, it might hin-

Table 7. Linear classification evaluated on the Canine subset of ImageNet (130 classes). The feature extraction models were pre-trained either on ImageNet (Full) or the subset of Canine images without using any labels. All computation is reported in terms of “full imagenet” epochs. We report the difference in performance relative to training on the full dataset for 1,000 epochs. Even though the canine-only models were trained with $5\times$ fewer gradient updates, they largely outperform the self-supervised learning models that were trained on the full dataset. We also observe that contrastive methods (SimCLR, MoCLR) benefit the most.

Method	Pre-training Dataset	Epochs	Top-1 Acc
SimCLR	Full	200	67.4
BYOL	Full	200	70.7
MoCLR	Full	200	68.7
SimCLR	Full	1,000	69.4
BYOL	Full	1,000	76.5
MoCLR	Full	1,000	75.3
SimCLR	Canine	100	72.8 (+3.4)
BYOL	Canine	100	76.0 (-0.5)
MoCLR	Canine	100	76.1 (+0.8)
SimCLR	Canine	200	74.1 (+4.7)
BYOL	Canine	200	77.3 (+0.8)
MoCLR	Canine	200	77.5 (+2.2)

der performance by drawing negative samples that are too similar, including more false-negatives [17].

To test this hypothesis in isolation and gain a better intuitive understanding of our method, we train various self-supervised learning models on the subset of ImageNet classes that belong to the canine family (including dogs, wolves and foxes, 130 classes in total) and compare them to models trained on the full dataset. For all models, we train a linear classifier on the canine-only subset, and evaluate on validation images from the canine subset.

From Table 7 it can be seen that models pre-trained on the canine-only subset perform significantly better than those trained on the entire ImageNet dataset, even though they have significantly less images to learn from and were trained with $5\times$ less computation.

5.3. ImageNet Results

Even though our main goal is to improve self-supervised learning on uncurated datasets, we asked whether it remains competitive on heavily studied datasets such as ImageNet.

From Table 8 we see that training the baseline MoCLR for 2,000 more epochs does not improve the results by much (+0.2). DnC on the other hand convincingly outperforms the baseline (+1.3). Interestingly, DnC even slightly outperforms MoCLR or BYOL when giving a computational budget of 1,000 epochs. Though DnC aims at uncurated datasets, the previous results from Section 5.2 on the canine subset have shown that even on ImageNet it might be

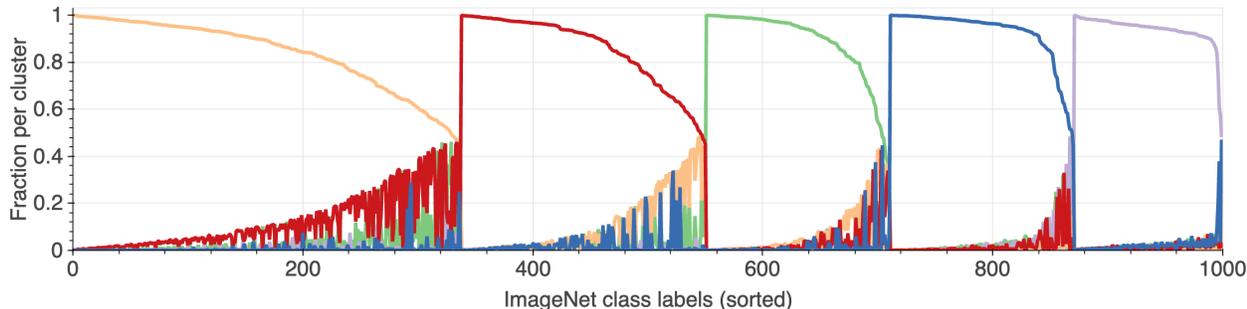


Figure 6. Visualization of the 5-way ImageNet clustering used in DnC. For each ImageNet class we compute the fraction of images that belong to each cluster. For better visualization the x-axis was sorted per cluster. From the figure it is clear that most images in each class belong to a single cluster. The fraction of images in each class that belong to the same cluster is **87.4%**.

beneficial to draw negatives from more similar categories.

Table 8. Comparison with BYOL and MoCLR on ImageNet linear evaluation benchmark with training budgets of 1000 and 3000 epochs. Top-1 accuracy is reported using ResNet-50.

Method	1000 epochs	3000 epochs	Δ
BYOL	74.3	73.9	-0.4
MoCLR	74.3	74.5	0.2
DnC (ours)	74.5	75.8	1.3

5.3.1 Semi-supervised learning

We evaluate the performance of DnC when fine-tuning on a subset of ImageNet’s train set. Following the semi-supervised protocol [52, 104, 12, 34], we use the same splits of 1% and 10% ImageNet data as in [12, 34]. As shown in Table 9, DnC consistently outperforms BYOL, SwAV, and MoCLR and Barlow Twins.

Table 9. Semi-supervised results with a fraction of ImageNet labels following the protocol of [12, 34]. The encoder is ResNet-50.

method	Top-1 Label fraction			Top-5 Label fraction		
	1%	10%	100%	1%	10%	100%
SimCLR [12]	48.3	65.6	76.0	75.5	87.8	93.1
BYOL [34]	53.2	68.8	77.7	78.4	89.0	93.9
SwAV [10]	53.9	70.2	-	78.5	89.9	-
MoCLR	53.0	68.8	77.4	79.1	89.6	94.0
Barlow Tw. [103]	55.0	69.7	-	79.2	89.3	-
DnC	59.9	71.1	78.2	83.0	90.4	94.2

5.3.2 Ablations

We provide further experiments for isolating the factors that make DnC work, shown in Table 10. If we train the expert models on the full dataset instead of subsets (similar to an ensemble), but with the same computational budget, the resulting model achieves the same performance as the base model (no improvement). Alternatively, splitting the dataset

Table 10. Ablating the experts used in DnC: We notice a big drop in performance if the experts models are trained on the whole dataset (ensemble), or on random subsets.

	Partitioning	Experts trained on	Top-1 Acc
DnC	Clustering	local partition	75.8
- local experts	-	full dataset	74.3
- clustering	Randomly	local partition	73.1

Table 11. Evaluation of what models to predict during distillation.

base model	local experts	use center-crop	Top-1 Acc
✓			74.5
	✓		75.2
✓	✓		75.8
✓	✓	✓	75.6

into random subsets hurts the final performance, showing the importance of the clustering used.

In Table 11 we study the distillation process, showing it is important to regress to both the base model and experts. And as discussed in Section 3.3, using center-crops instead of augmented views does not hurt the performance by much.

6. Conclusion

In this paper we have studied how state of the art self-supervised learning methods perform when they are pre-trained on uncurated data – datasets that did not require human annotations or labels to create – as a step towards fully self-supervised learning. We have observed that current methods suffer from a large drop in performance of up to **-9%** when pre-trained on these uncurated datasets. To alleviate this issue, we have proposed Divide and Contrast (DnC) that requires a few simple changes to existing self-supervised learning methods, and which largely outperforms state of the art SSL methods on uncurated datasets, as well as achieving similar or better performance on ImageNet. We hope this work draws more attention to uncurated datasets as a benchmark for self-supervised learning.

References

- [1] Humam Alwassel, Dhruv Mahajan, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *arXiv:1911.12667*, 2019. 2
- [2] Sanjeev Arora, Hrishikesh Khandeparkar, Mikhail Khodak, Orestis Plevrakis, and Nikunj Saunshi. A theoretical analysis of contrastive unsupervised representation learning. *arXiv:1902.09229*, 2019. 2
- [3] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv:1906.00910*, 2019. 2
- [4] Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L Alexander, David W Jacobs, and Peter N Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *CVPR*, 2014. 6
- [5] Piotr Bojanowski and Armand Joulin. Unsupervised learning by predicting noise. *arXiv:1704.05310*, 2017. 2
- [6] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *ECCV*, 2014. 6
- [7] Yue Cao, Zhenda Xie, Bin Liu, Yutong Lin, Zheng Zhang, and Han Hu. Parametric instance classification for unsupervised visual feature learning. *NeurIPS*, 2020. 2
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 2
- [9] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *ICCV*, 2019. 2
- [10] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2, 8
- [11] Mark Chen, Alec Radford, Rewon Child, Jeff Wu, Heewoo Jun, Prafulla Dhariwal, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *ICML*, 2020. 2
- [12] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv:2002.05709*, 2020. 1, 2, 3, 6, 7, 8
- [13] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. In *NeurIPS*, 2020. 1, 2, 3
- [14] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, 2016. 2
- [15] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv:2003.04297*, 2020. 2
- [16] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised visual transformers. *arXiv preprint arXiv:2104.02057*, 2021. 3
- [17] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debiased contrastive learning. *NeurIPS*, 2020. 7
- [18] Soo-Whan Chung, Joon Son Chung, and Hong-Goo Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. In *ICASSP*, 2019. 2
- [19] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *CVPR*, 2014. 6
- [20] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv:1810.04805*, 2018. 1
- [22] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 2
- [23] Carl Doersch and Andrew Zisserman. Multi-task self-supervised visual learning. In *ICCV*, 2017. 2
- [24] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *ICLR*, 2017. 2
- [25] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019. 2
- [26] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NIPS*, 2014. 2
- [27] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 6
- [28] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPR*, 2004. 6
- [29] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv:1803.07728*, 2018. 2
- [30] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [31] Daniel Gordon, Kiana Ehsani, Dieter Fox, and Ali Farhadi. Watching the world go by: Representation learning from unlabeled videos. *arXiv:2003.07990*, 2020. 2
- [32] Priya Goyal, Mathilde Caron, Benjamin Lefaudeux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021. 2, 5
- [33] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019. 2
- [34] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad

- Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. In *NeurIPS*, 2020. 1, 2, 3, 6, 7, 8
- [35] Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-bastien Grill, Florent Alché, Rémi Munos, and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. *arXiv:2004.14646*, 2020. 2
- [36] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 2
- [37] Tengda Han, Weidi Xie, and Andrew Zisserman. Video representation learning by dense predictive coding. In *ICCV Workshop*, 2019. 2
- [38] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. *arXiv:2008.01065*, 2020. 2
- [39] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3, 6
- [40] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 6
- [41] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [42] Olivier J Hénaff, Skanda Koppula, Jean-Baptiste Alayrac, Aaron van den Oord, Oriol Vinyals, and João Carreira. Efficient visual pretraining with contrastive detection. *arXiv preprint arXiv:2103.10957*, 2021. 2
- [43] Olivier J Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv:1905.09272*, 2019. 1, 2
- [44] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv:1503.02531*, 2015. 2
- [45] Kyle Hsu, Sergey Levine, and Chelsea Finn. Unsupervised learning via meta-learning. *arXiv:1810.02334*, 2018. 2
- [46] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, 2019. 2
- [47] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In *NeurIPS*, 2018. 7
- [48] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 2
- [49] Minseon Kim, Jihoon Tack, and Sung Ju Hwang. Adversarial self-supervised contrastive learning. *NeurIPS*, 2020. 2
- [50] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013. 2
- [51] Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. A mutual information maximization perspective of language representation learning. In *ICLR*, 2020. 2
- [52] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *CVPR*, 2019. 6, 8
- [53] Jonathan Krause, Jia Deng, Michael Stark, and Li Fei-Fei. Collecting a large-scale dataset of fine-grained cars. 2013. 6
- [54] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 6
- [55] Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv:2005.04966*, 2020. 2
- [56] Tianhao Li and Limin Wang. Learning spatiotemporal features via video and text pair discrimination. *arXiv:2001.05691*, 2020. 2
- [57] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 6
- [58] 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 *ECCV*, 2014. 6
- [59] Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations. In *ICLR*, 2018. 2
- [60] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 6
- [61] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv:1306.5151*, 2013. 6
- [62] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. *arXiv:1912.06430*, 2019. 2
- [63] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, 2013. 2
- [64] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. *arXiv:1912.01991*, 2019. 2
- [65] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. *arXiv:2004.12943*, 2020. 2
- [66] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 6
- [67] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, 2008. 6
- [68] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 2
- [69] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. 1, 2, 3

- [70] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, 2020. 2
- [71] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *CVPR*, 2012. 6
- [72] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017. 2
- [73] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 2
- [74] Mandela Patrick, Yuki M Asano, Ruth Fong, João F Henriques, Geoffrey Zweig, and Andrea Vedaldi. Multi-modal self-supervision from generalized data transformations. *arXiv:2003.04298*, 2020. 2
- [75] Senthil Purushwalkam and Abhinav Gupta. Demystifying contrastive self-supervised learning: Invariances, augmentations and dataset biases. *arXiv:2007.13916*, 2020. 2
- [76] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv:1511.06434*, 2015. 2
- [77] Anyi Rao, Jiase Wang, Linning Xu, Xuekun Jiang, Qingqiu Huang, Bolei Zhou, and Dahua Lin. A unified framework for shot type classification based on subject centric lens. In *ECCV*, 2020. 2
- [78] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv:1412.6550*, 2014. 2
- [79] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised pre-training for speech recognition. *arXiv:1904.05862*, 2019. 1
- [80] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. 2, 7
- [81] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *ICRA*, 2018. 2
- [82] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NIPS*, 2016. 2
- [83] Aravind Srinivas, Michael Laskin, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv:2004.04136*, 2020. 2
- [84] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Contrastive bidirectional transformer for temporal representation learning. *arXiv:1906.05743*, 2019. 2
- [85] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*, 2017. 4
- [86] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 3
- [87] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 2016. 1, 2, 4
- [88] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv:1906.05849*, 2019. 1, 2, 3
- [89] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 2
- [90] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. In *NeurIPS*, 2020. 2
- [91] Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive learning, multi-view redundancy, and linear models. *arXiv:2008.10150*, 2020. 2
- [92] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. *arXiv preprint arXiv:2102.06191*, 2021. 2
- [93] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008. 2
- [94] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *ECCV*, 2018. 2
- [95] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. *arXiv:2005.10242*, 2020. 2
- [96] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 2
- [97] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 1, 2, 3
- [98] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010. 6
- [99] Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. *arXiv:2008.05659*, 2020. 2
- [100] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas J Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. *arXiv:2007.10985*, 2020. 2
- [101] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, 2019. 1, 2
- [102] Asano YM., Rupprecht C., and Vedaldi A. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. 2
- [103] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021. 8
- [104] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *ICCV*, 2019. 8
- [105] Xiaohang Zhan, Jiahao Xie, Ziwei Liu, Yew-Soon Ong, and Chen Change Loy. Online deep clustering for unsupervised representation learning. In *CVPR*, 2020. 2

- [106] Liheng Zhang, Guo-Jun Qi, Liqiang Wang, and Jiebo Luo. Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In *CVPR*, 2019. [2](#)
- [107] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. [2](#)
- [108] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. [2](#)
- [109] Xiao Zhang and Michael Maire. Self-supervised visual representation learning from hierarchical grouping. *arXiv preprint arXiv:2012.03044*, 2020. [2](#)
- [110] Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. What makes instance discrimination good for transfer learning? *arXiv:2006.06606*, 2020. [2](#)
- [111] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *TPAMI*, 2017. [5](#)
- [112] Chengxu Zhuang, Alex Andonian, and Daniel Yamins. Unsupervised learning from video with deep neural embeddings. *arXiv:1905.11954*, 2019. [2](#)
- [113] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. *arXiv:1903.12355*, 2019. [2](#)