

Unsupervised Visual Representation Learning by Synchronous Momentum Grouping

Bo Pang¹, Yifan Zhang¹, Yaoyi Li², Jia Cai², and Cewu Lu^{1*}

¹ Shanghai Jiao Tong University
{pangbo, zhangyf_sjtu, lucewu}@sjtu.edu.cn
² HuaWei Technologies Co., Ltd.
{liyaoqi, caijiaai1}@huawei.com

Abstract. In this paper, we propose a genuine group-level contrastive visual representation learning method whose linear evaluation performance on ImageNet surpasses the vanilla supervised learning. Two mainstream unsupervised learning schemes are the instance-level contrastive framework and clustering-based schemes. The former adopts the extremely fine-grained instance-level discrimination whose supervisory signal is not efficient due to the false negatives. Though the latter solves this, they commonly come with some restrictions affecting the performance. To integrate their advantages, we design the SMoG method. SMoG follows the framework of contrastive learning but replaces the contrastive unit from instance to group, mimicking clustering-based methods. To achieve this, we propose the momentum grouping scheme which synchronously conducts feature grouping with representation learning. In this way, SMoG solves the problem of supervisory signal hysteresis which the clustering-based method usually faces, and reduces the false negatives of instance contrastive methods. We conduct exhaustive experiments to show that SMoG works well on both CNN and Transformer backbones. Results prove that SMoG has surpassed the current SOTA unsupervised representation learning methods. Moreover, its linear evaluation results surpass the performances obtained by vanilla supervised learning and the representation can be well transferred to downstream tasks.

1 Introduction

In the era of adopting deep learning and data-driving as the mainstream framework [33], the quality of the learned representation largely determines the performance of models on a majority of tasks [75]. For a long time, people utilize supervised tasks and adopt a large amount of annotations to train models and get good representations. Nevertheless, this simple and efficient scheme faces the problems of expensive and time-consuming annotating costs [2, 17, 18, 57], and non-ideal generalization performance caused by bias from the annotation information [4, 51]. To solve these problems, people extensively study the unsupervised representation learning method [58], including the generative models

* Cewu Lu is the corresponding author.

providing expressive latent features [30, 55] and self-supervised methods with different pretext tasks [62, 62, 53, 48]. However, there are still performance gaps between these methods and the supervised scheme. In recent years, unsupervised representation learning achieves great improvements, SOTA contrastive learning [8, 23, 6, 21] based unsupervised methods reduce the performance gap to less than 3%. We continue the study of the unsupervised representation learning, stand on the shoulder of previous contrastive and clustering-based methods, and finally, push the performance of unsupervised methods above the level of the vanilla supervised scheme.

Traditional contrastive learning methods adopt instance discrimination as the pretext task. This kind of refinement of classification treats every sample as a category to conduct discrimination. Thus, it will introduce many false-negative pairs, leading to inefficient supervisory signals. And another problem is that the accurate instance discrimination requires pair-wise comparison on the entire dataset. However, with limited resources, most implementations compromise to adopt a subset of the comparison with a large batch size or memory bank [23, 67], which further decreases the efficiency of learning.

Previous clustering-based methods relax instance discrimination into group discrimination to solve these problems. The core task of clustering is how to split instances into groups. They try to adopt K-means clustering method [5, 71] or optimal transport [3, 6] method to achieve this, but due to the limitations they introduce (asynchronous two-stage training or local equipartition respectively), they don't show an advantage in terms of performance.

Here, we integrate the advantages of instance contrastive and clustering-based methods to the newly proposed SMoG method which follows the instance contrastive framework, and extend it to group-level (see Fig. 1). One major design is the momentum grouping scheme that allows gradient propagating through groups to synchronously conduct the instance grouping process together with the representation learning. It is the first method to apply contrasting directly on groups instead of instances, since no gradient can propagate through group fea-

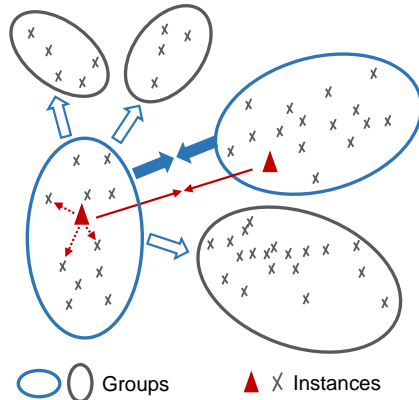


Fig. 1: Group contrastive *vs.* instance contrastive learning. Compared with instance contrasting (red parts), group contrasting (blue part) learns representations through higher-level semantics, which can reduce the chance that the already learned similar instances are still treated as negative pairs (false negatives). The significant reduction of contrast elements after the grouping process also makes global contrasting easier to calculate. The colored instances or groups denote the positive pairs and the grey ones are negatives.

tures in previous methods, With it, SMOG avoids the false negatives in instance contrasting and wipes off limiting factors in previous clustering-based methods.

The proposed SMOG method is effective and pretty simple. We evaluate it on several standard self-supervised benchmarks and multiple downstream tasks. In particular, under the linear protocol, it achieves 76.4% top-1 accuracy on ImageNet with a standard ResNet-50 [25], which surpasses the vanilla supervised-level performance for the first time! Extensive experiments show that SMOG works robustly and we observe consistent SOTA performance under linear (+0.8%), semi-supervised classification (+2%), detection, and segmentation tasks. Besides, SMOG works well on both CNN and Transformer backbones.

2 Related Work

2.1 Handcraft Pretext Task

Adopting pretext tasks to provide training supervisory signals is an effective unsupervised representation learning scheme. By now, well studied pretext tasks include jigsaw solving [48, 49], colorization [62, 32, 73], denoising [62], inpainting [53], super-resolution [34], and patch position prediction [13, 14]. In addition, for video visual inputs with temporal dimension, sequential-related pretext tasks are proved useful, such as ordering frames [19, 46, 66], motion estimation [1, 29, 61, 39], and further frame estimation [43, 44, 56, 63, 64]. Due to the specificity of these tasks, the learned representation often carry a certain bias, leading to relatively limited performances on downstream tasks.

2.2 Instance Discrimination Method

Present mainstream contrastive learning methods [23, 8, 21, 6, 74, 37, 47, 20, 59, 52] learn representations by instance discrimination which treats each image (pixel) in a training set as its own category. This scheme achieves state-of-the-art performances on downstream tasks. Its common training direction provided by InfoNCE [50, 8, 22] or its variations is to maximize the mutual information [50, 27], which requires a large number of negative pairs to get good performances [8]. Adopting a large batch size is a straightforward method, but it consumes lots of resources. To solve this, MoCo [23] and [67] propose to utilize memory structures. Some latest designs [21, 9, 7] conduct contrasting without negative samples by an asymmetry Siamese structure or normalization techniques. Recent work [37] proposes self-supervised learning with the Hilbert-Schmidt Independence Criterion, which yields a new understanding of InfoNCE. While RELIC [47] improves the generalization performance through an invariance regularizer under the causal framework. NNCLR [15] adopts the nearest neighbour from the dataset as the positives. UniGrad [59] provides a unified contrastive formula through gradient analysis. Instance discrimination, as an excessive fine-grained classification setting implies some problems like false-negative pairs [76].

2.3 Group Discrimination Method

Our work follows the group discrimination scheme. DeepClustering [5] adopts the K-means clustering method to get the groups and learn features from them in an iterative manner. However, the two-step training scheme (clustering, learning) leads to delayed supervisory signal, against effective representation learning. ODC [71] and CoKe [54] shortens the two-step circulation period but does not solve the delay problem. SeLa [3] and SwAV [6] treat the grouping problem as pseudo-labelling and solve it as an optimal transport task. But, to avoid degeneracy, they add an equipartition constraint, decreasing the validity of grouping. Our work stands on the shoulders of these works, solves the mentioned problems, and finally achieves the level of vanilla supervised learning for the first time.

3 Approach

Present SOTA methods commonly adopt the instance contrastive learning, apply a two-stream Siamese network structure, and take InfoNCE or its variants as the loss function to conduct the contrasting. This is an eminent and robust structure that has been proved by many models. However, there are two potential problems: 1) In terms of the problem setting, instance discrimination is too fine-grained and the learned representation goes against the downstream tasks that rely on the high-level semantics to some extent, because for high-level semantics, instance-level contrasting introduces many false-negative pairs which damage the quality of representation learning. 2) In terms of technical practice, contrastive learning based on InfoNCE needs a large number of negative pairs to improve the upper bound of the theoretical performance [8, 23], which poses a challenge on computing resources or needs specific model designs.

To deal with the above two problems, previous works [6, 5, 65] propose the clustering-based method: adopting groups as negatives to reduce false-negatives and the total number of negative pairs. However, these methods introduce some limitations and lose the advantages of contrastive methods: the always-updating optimization signal and totally unrestricted feature distribution. To integrate the advantages of contrastive and clustering-based method, we propose to contrast in the group unit and design the first group-level contrastive learning algorithm Synchronous Momentum Grouping (SMoG) which inherits and develops current contrastive learning techniques, and pushes the performance of unsupervised methods to the level of vanilla supervised scheme.

3.1 Group-Level Contrastive Learning

In general, given a dataset $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ with n instances and a visual model f_θ with parameter set θ which maps each instance to a vector representation $f_\theta(\mathbf{x})$, the pipeline of the instance contrastive learning framework can be expressed as:

$$L_i = -\log \frac{\exp(\text{sim}(f_\theta(\mathbf{x}_i^a), \hat{f}_\eta(\mathbf{x}_i^b))/\tau)}{\sum_{\mathbf{x}_j \neq \mathbf{x}_i^a, \mathbf{x}_j \in \mathbf{Y} \subseteq \mathbf{X}} \exp(\text{sim}(f_\theta(\mathbf{x}_i^a), \hat{f}_\eta(\mathbf{x}_j))/\tau)}, \quad (1)$$

where $\text{sim}(\mathbf{u}, \mathbf{v})$ is commonly instantiated as the normalized inner product, \hat{f}_η is just f_θ or a variant of f_θ such as its momentum updated version or the version without the last few layers (predictor). \mathbf{x}_i^a and \mathbf{x}_i^b are two different augmented views of \mathbf{x}_i . In theory, \mathbf{x}_j should come from the training set \mathbf{X} , but in practical, taking computing complexity into account, \mathbf{x}_j is commonly selected from a much smaller subset $\mathbf{Y} \subseteq \mathbf{X}$.

In order to extend this framework to the group level, we need to first define and generate a certain number of groups with group feature \mathbf{g} , attach each instance to a certain group, allow gradients propagating through groups, and update them synchronously during training: $\mathbf{c}_i^a, \{\mathbf{g}\} \leftarrow \phi(f_\theta(\mathbf{x}_i^a) | \{\mathbf{g}\})$, where ϕ is the group assigning function that takes in instance and group features, generates instances' corresponding group \mathbf{c}_i^a , and updates the groups $\{\mathbf{g}\}$. Thus, $\mathbf{c}_i = \mathbf{g}_k$ means that instance \mathbf{x}_i is attached to group k . With the group-level features, we can derive the group-level contrastive learning framework as:

$$L_i = -\log \frac{\exp(\text{sim}(\mathbf{c}_i^a, \mathbf{c}_i^b)/\tau)}{\sum_{\mathbf{g}_j \in \mathbf{G}} \exp(\text{sim}(\mathbf{c}_i^a, \mathbf{g}_j)/\tau)}, \quad (2)$$

where \mathbf{G} is the set of group features. Since the number of groups is a certain small value, we can conduct group contrasting globally with all the group pairs.

The algorithm of gaining group features (ϕ) will be detailed introduced in the next section. Here, we emphasize that besides well expressing a group of similar instances without extra limitation, a qualified ϕ also needs to make sure every group feature \mathbf{c}_i is updated synchronously with the instance feature $f_\theta(\mathbf{x}_i)$ since the gradients need to be back-propagated through \mathbf{c}_i to $f_\theta(\mathbf{x}_i)$ for training the parameters which means \mathbf{c}_i can replace the latest $f_\theta(\mathbf{x}_i)$ to participate in contrasting. This is the core difference from the previous clustering-based method, since their group features gained by clustering cannot back-propagate

the gradient, and directly contrasting group features cannot train the network.

Intuitively, the newly designed group contrastive learning method aims at directly adjusting the distribution of group features, and since gradients can propagate through group features to instance features, this algorithm can grad-




Fig. 2: The left part illustrates the group contrastive target described by Eq. 2. As the training progresses, preliminary formed meaningful representations have a large chance to make \mathbf{x}_i^a and \mathbf{x}_i^b belong to the same group. In this case, the gathering force (the green arrow) cannot make them get closer. Thus, we modify it to the version in Eq. 3 illustrated in the right part, where the push and pull force required by contrastive learning can always take effect.

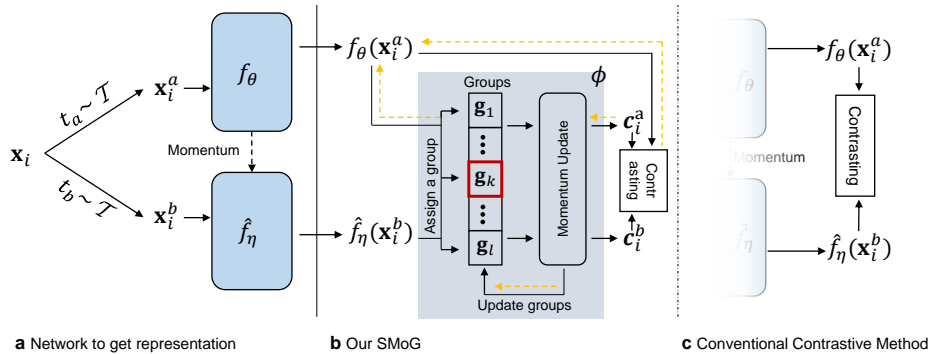


Fig. 3: Pipeline of our SMOG *vs.* conventional contrastive method. Like typical contrastive learning methods, SMOG gets instance features from two augmentation views of the same image by a Siamese structure. Rather than directly contrasting the instance features, SMOG works at the group level. Specifically, it dynamically splits the already seen instances into groups, assigns the new instance to the closest group (the red one in the figure), and adopts the group feature which replaces the instance one to conduct the contrasting. Synchronous with the representation learning, group features are updated by the new instances in a momentum manner. The yellow dashed lines show the back-propagation paths, and we can see that the group feature can also back-propagate gradients.

ually learn the instance representation. However, unfortunately, it has flaws in the second half of the training process. With the meaningful representations gradually formed, \mathbf{c}_i^a and \mathbf{c}_i^b tend to be the same group feature. This will make the supervisory signal fail to gather similar instances and get them closer (see Fig. 2). To solve this problem, considering that \mathbf{c}_i^a is the combination of a group of $f_\theta(\mathbf{x})$, we change the group contrastive learning loss to:

$$L_i = -\log \frac{\exp(\text{sim}(f_\theta(\mathbf{x}_i^a), \mathbf{c}_i^b)/\tau)}{\sum_{\mathbf{g}_j \in \mathbf{G}} \exp(\text{sim}(f_\theta(\mathbf{x}_i^a), \mathbf{g}_j)/\tau)}. \quad (3)$$

When $\mathbf{c}_i^a \neq \mathbf{c}_i^b$ Eq. 3 can be seen as the “stochastic-batch” version of Eq. 2, which splits an intact loss into several batches, similar to SGD and GD. When $\mathbf{c}_i^a = \mathbf{c}_i^b$, Eq. 3 can still gather instances, which solves the problem. This formulation looks similar to previous clustering-based methods, but remember that we still conduct group contrasting instead of instance classification that previous methods do, since group features \mathbf{c}_i and \mathbf{g}_j can propagate the gradient, thus has the same status with $f_\theta(\mathbf{x}_i^a)$ in contrasting and directly guide the learning direction.

3.2 Synchronous Momentum Grouping

SMoG’s network structure and training settings are identical to conventional contrastive learning frameworks as Fig. 3 shows. That is generating \mathbf{x}_i^a and

\mathbf{x}_i^b through two different groups of augmentation $t_a, t_b \sim \mathcal{T}$, and gaining their features through a Siamese network. Following MoCo [23], we adopt a momentum network: $\hat{f}_\eta \leftarrow \alpha * \hat{f}_\eta + (1 - \alpha) * f_\theta$, where α is the momentum ratio and here, f_θ denotes layers that also exist in \hat{f}_η . So far, SMOG is consistent with the typical contrastive learning methods. By adding the group assigning function ϕ , we can get the complete SMOG method.

Generating Group Features As mentioned in last section, group-level contrasting optimizes instance features through group features. Thus, group features must represent the latest instance ones to propagate gradients, namely, \mathbf{c}_i needs to be synchronously updated with $f_\theta(\mathbf{x}_i)$ and calculated by a differentiable function. Directly adopting conventional global clustering as ϕ is not feasible to synchronously update \mathbf{c}_i with $f_\theta(\mathbf{x}_i)$ at each iteration due to the computing cost. Thus, we simply modify it to the momentum grouping scheme, generating \mathbf{c}_i through an iterative algorithm synchronized with representation learning.

Before the training start, we initialize the l group features $\{\mathbf{g}_1, \dots, \mathbf{g}_k, \dots, \mathbf{g}_l\}$ randomly or using a clustering method such as k-means. Then along with the training process, we update \mathbf{g}_k and get \mathbf{c}_i in each iteration by:

$$\begin{aligned} \mathbf{c}_i &= \operatorname{argmin}_{\mathbf{g}_k} (\operatorname{sim}(f_\theta(\mathbf{x}_i), \mathbf{g}_k)) \\ \mathbf{g}_k &\leftarrow \beta * \mathbf{g}_k + (1 - \beta) * \operatorname{mean}_{\mathbf{c}_t = \mathbf{g}_k} f_\theta(\mathbf{x}_t), \end{aligned} \quad (4)$$

where β is the momentum ratio, \mathbf{x}_t comes from a mini-batch, and we omit the normalization. Importantly, this mechanism does not introduce extra limitation.

The momentum grouping scheme assigns each instance to the closest group and iteratively updates the group features in a momentum manner. In this way, the group features are always the latest representative of the instance visual features and more importantly, for each iteration, the updated \mathbf{g}_k is adopted to conduct contrasting and the gradient can be back-propagated through \mathbf{g}_k to $f_\theta(\mathbf{x}_i)$. This is the main difference from previous clustering-based methods.

Avoiding Collapse The momentum grouping scheme ensures that the group features \mathbf{g}_k always represent the latest learned instance representations. However, because the algorithm is based on the iterative updating with a local subset of instances (a batch), at the early training stage, it might be unstable. The scale of the groups may be unbalanced or all the instances collapse into few groups. To deal with this, we periodically apply an extra grouping process on a cached relatively large feature set to relocate the groups. There will be long intervals among groupings which are light and quick, thus, the overhead can be ignored.

3.3 Compared with Previous Clustering Methods

Group discrimination is not a first proposed concept. Previous clustering based methods DeepClustering [5] and SwAV [6] also conduct it and in terms of the loss function, the three algorithms have a similar form. The main difference lies on

the method to generate the groups and the way to utilize them. Our SMoG aims at conducting the contrasting among the groups like Eq. 2 shows. This is impossible for previous clustering-based methods since their group features have to be detached out and cannot back-propagate the gradients to the parameters. Thus, they do not contrast the groups but classify the instances into groups. To achieve the group contrasting, we propose the momentum grouping scheme which allows us directly contrasting the groups and propagate gradients. Although we modify the final loss to Eq. 3 for better performance, our SMoG is still a group-level contrastive algorithm, since the group features directly guide the optimization direction. This is the reason we call our method "grouping contrasting" instead of "clustering" to distinguish them. SMoG conducts contrasting among groups instead of classifying instances based on clusters.

3.4 Implementation Details

Augmentation. We adopt the asymmetrical augmentation method used in BYOL [21], where there are two augmentation schemes for the two streams of the Siamese network. The two schemes adopt the same color jittering but one has a stronger Gaussian blur and the other one has stronger solarization. Since the contrasting loss we adopt is also asymmetrical (contrasting among instance and group features), the two streams of the Siamese network are not assigned to a fixed augmentation scheme, instead, they adopt one of them in an alternate manner. In this way, the two streams can generate instance features under the same distribution so that they can be mapped to the same group feature set.

SMoG Setting. Like many previous works [21, 10, 68], both f_θ and \hat{f}_η have a backbone and a projection head. And \hat{f}_η has an extra prediction head stacked on the projection head. The backbone’s output is adopted as the learning representation. The two kinds of head are both a two-layer MLP and their hidden layer is followed by a BatchNorm [28] and an activation function (ReLU for ResNet and GELU [26] for SwinTransformer [40]). The output layer of projection head also has BN. The dimension of the hidden layer is 2048, while for the output layer, it is 128 for CNN and 256 for Transformer. The Siamese’s momentum ratio α is 0.999. In experiments, we group all the instances into 3k groups and the group features \mathbf{g} are initialized with the K-means algorithm. The momentum ratio β of grouping follows a linear schedule from 1.0 to 0.99. To avoid collapse, we reset the group features every 300 iterations with K-means on cached features of the past 300 iterations and \hat{f}_η is synchronously reset with the parameters of f_θ .

Pre-Training Details. We pre-train the models on ImageNet dataset [12]. For CNN backbones, we train with the LARS [69] optimizer with 4096 mini-batchsize on 64 GPUS (when adopting ResNet50). The base learning rate is set to $lr = 0.3 \times \text{batchsize}/256$, following first a 10-epoch warm-up and then the cosine scheduler. The weight decay and temperature τ are set to 10^{-6} and 0.1. For Transformer based backbones, we adopt the Adamw optimizer [42]. The base learning rate and weight decay are $5e^{-4} \times \text{batchsize}/2048$ and 0.05. Other settings

are the same with CNN backbones. For efficient training, we also adopt the multi-crop training scheme [6, 7] with two large views (224×224) and 4 small views (96×96). When adopting multi-crop, the scale of the random crops are [0.2, 1.0] and [0.05, 0.2] respectively for large and small views.

4 Experiments

We evaluate SMOG firstly on the standard benchmark on ImageNet with both CNN and Transformer backbones. We then compare the performance on several downstream tasks and give a detailed ablation study.

4.1 Linear Evaluation

Following the standard benchmarks [23, 8, 21, 6, 7, 10, 68], we first evaluate the representation of ResNet-50 [25] and Swin Transformer Tiny [40] trained with the proposed SMOG by the linear protocol: linear classification on frozen features.

Results on ResNet-50 is shown in Tab. 1 and we can see that previous group-based methods have no advantages over instance contrastive ones (about -2% top-1 accuracy) due to the introduced restrictions. While the proposed SMOG improves the best performance of group-based methods by 2.3% top-1 accuracy and achieves the SOTA performances among all the unsupervised representation methods. More importantly,

after adopting the multi-crop training strategy [6, 7], for the first time, the unsupervised representation of ResNet-50 surpasses the performance of the vanilla supervised one (+ 0.3%) on ImageNet.

We adopt Swin-Transformer to evaluate the performance of SMOG on Transformer backbones. Results are shown in Tab. 2. For the two-view setting, most Transformer-based backbones perform worse than the CNNs with similar parameters. But Transformers benefit more from the multi-view training. We hold

Table 1: Linear protocol results on ImageNet. ResNet50 is adopted. † denotes the model adopts the multi-crop training strategy. “acc” means accuracy.

model	epoch	batchsize	top1 acc	top5 acc
supervised	100ep	256	76.1	92.7
<i>instance contrastive method</i>				
ReSSL [74]	200	256	69.6	-
MoCoV2 [23]	800	256	71.1	90.1
SimSiam [9]	800	256	71.3	-
InfoMin Aug. [60]	800	256	73.0	91.1
MoCoV3 [10]	400	4096	73.1	-
MoCoV3	800	4096	73.8	-
BYOL [21]	400	4096	73.2	-
BYOL	800	4096	74.3	91.6
Barlow Twins [70]	1000	2048	73.2	91.0
RELIC [47]	1000	4096	74.8	92.2
SSL-HSIC [37]	1000	4096	74.8	92.2
<i>group-based method</i>				
DeepCluster	400	256	52.2	-
ODC [71]	400	256	57.6	-
PCL [36]	200	256	67.6	-
DeepClusterV2	400	4096	70.2	-
SwAV	400	4096	70.1	-
CoKe [54]	800	1024	72.2	-
SMoG	400	2048	73.6	91.3
SMoG	800	4096	74.5	91.9
<i>with multi-crop</i>				
SwAV†	400	4096	74.6	-
SwAV†	800	4096	75.3	-
DC-v2†	800	4096	75.2	-
DINO † [7]	800	4096	75.3	-
UniGrad † [59]	800	4096	75.5	-
NNCLR † [15]	1000	4096	75.6	92.4
SMoG†	400	4096	76.4	93.1

the opinion that the self-attention, a global operator, needs more data to train. Thus, stronger augmentation leads to better performances. SMOG, in the two-view setting, achieves SOTA performances compared to DINO (+2%), MoCo (+2%), EsViT (+4%), and MOBY (-0.5%). Note that, for fair comparison, we report the performance of EsViT with only \mathcal{L}_V . In the multi-crops setting, SMOG achieves the supervised performance without Mixup [72] augmentation. Since mixup needs labels, it is not straightforwardly suitable for the unsupervised setting.

This comparison reveals a landmark progress in contrastive learning.

Large Architectures Tab. 3 shows the results on several variants of ResNet-50 with larger widths [31]. The performance increases with a similar trend of supervised method and previous unsupervised framework. And it is worth noting that in previous works [8, 21], with larger architectures, their gap with the supervised learning decreases, but it is a pity that we do not observe an increasing superiority of SMOG over the supervised one. In Fig. 4, we demonstrate that on both CNN and Transformer, SMOG is comparable with the vanilla supervised method.

4.2 Semi-Supervised Fine-Tune Evaluation

Next, we evaluate the proposed SMOG through semi-supervised fine-tuning the unsupervised representation on a subset of training data of ImageNet. We follow the protocol adopted in [8] and the 1% and 10% labeled splits of ImageNet we adopt are the fixed ones provided in [8].

Table 2: Linear protocol results on ImageNet. † denotes adopting the multi-crop training strategy. The throughput (im/s) is calculated on a NVIDIA V100 GPU with 128 samples per forward. We report performances without Mixup and standard ones for supervised methods.

model	backbone	throughput	param	top1 acc
supervised	SwinT	808	28	77.8 / 81.3
Supervised	DeiT-S/16	1007	21	77.5 / 79.8
MoBY	DeiT-S/16	1007	21	72.8
MoBY	SwinT	808	28	75.0
MoCoV3	DeiT-S/16	1007	21	72.5
MoCoV3	ViT-B/16	312	85	76.7
DINO	DeiT-S/16	1007	21	72.5
DINO †	DeiT-S/16	1007	21	77.0
EsViT [35]	SwinT	808	28	70.5
EsViT †	SwinT	808	28	77.0
SMoG	SwinT	808	28	74.5
SMoG †	SwinT	808	28	77.7

Table 3: Linear protocol results on ImageNet with larger backbones. We experiment on wider ResNet. The pre-training details are the same with the ResNet-50 $\times 1$.

model	backbone	param	top1 acc	top5 acc
Supervised	Res50 (x2)	188	77.8	93.8
	Res50 (x4)	375	78.9	94.5
BYOL	Res50 (x2)	188	77.4	93.6
	Res50 (x4)	375	78.6	94.2
SwAV	Res50 (x2)	188	77.3	-
	Res50 (x4)	375	77.9	-
SMoG	Res50 (x2)	188	78.0	93.9
	Res50 (x4)	375	79.0	94.4

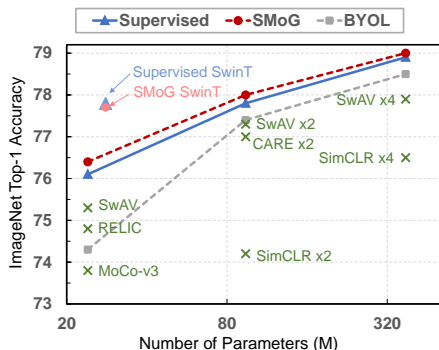


Fig. 4: ResNet linear evaluation on ImageNet compared with supervised training. We can see that for different model sizes, our SMOG achieves comparable performance with supervised method.

As shown in Tab. 4, SMOG consistently outperforms all the previous methods in both 1% and 10% settings on ResNet-50 $\times 1$ and $\times 2$. Besides, we also fine-tune the unsupervised representation on the full ImageNet. With SMOG pre-training, ResNet-50 achieves 78.3% top-1 accuracy under the standard training recipe, surpassing the directly supervised learning by 2.2%. Similarly, the ResNet-50 $\times 2$ also outperforms the directly supervised training by 2.4%.

Table 4: Semi-supervised results on ImageNet.

method	Top-1 Acc (%)			Top-5 Acc (%)		
	1%	10%	100%	1%	10%	100%
<i>resnet-50 $\times 1$</i>						
Supervised	25.4	56.4	76.1	48.4	80.4	92.9
SimCLR	48.3	65.6	76.0	75.5	87.8	93.1
BYOL	53.2	68.8	77.7	78.4	89.0	93.9
SwAV	53.9	70.2	-	78.5	89.9	-
Barlow Twins	55.0	69.7	-	79.2	89.3	-
SSL-HSIC	52.1	67.9	77.2	77.7	88.6	93.6
NNCLR	56.4	69.8	-	80.7	89.3	-
SMoG	58.0	71.2	78.3	81.6	90.5	94.2
<i>resnet-50 $\times 2$</i>						
Supervised	-	-	77.8	-	-	93.8
SimCLR	58.5	71.7	-	83.0	91.2	-
BYOL	62.2	73.5	-	84.1	91.7	-
SMoG	63.6	74.4	80.2	85.6	92.4	95.2

4.3 Transfer to Other Vision Tasks

We further transfer the unsupervised representations of ResNet-50 learned with the proposed SMOG to several downstream tasks, namely semantic segmentation, object detection, and instance segmentation.

We first evaluate SMOG on Cityscapes [11] and VOC-2012 [16] semantic segmentation tasks. For fair comparison, we align all the methods with the same training recipe of FCN [41]. Results

Table 5: Transfer learning results on semantic segmentation task. We fine-tune the representations on VOC2012 and Cityscapes dataset. The segmentation model is FCN with ResNet-50.

model	Cityscapes		VOC-2012	
	mIoU	mAcc	mIoU	mAcc
Supervised	73.83	82.56	73.59	83.74
MoCoV2	74.30	83.37	70.86	80.37
SwAV	74.80	83.01	74.97	84.27
BYOL	74.90	83.73	74.76	84.37
SMoG	76.03	83.97	76.22	85.01

of mean accuracy and mean IoU are provided in Tab. 5, we can see that nearly all the unsupervised representations outperform the conventional supervised one which reveals that in this downstream task, unsupervised representation is already a better choice. And again, SMOG performs better than the original SOTAs on Cityscapes (+1.1 mIoU) and VOC-2012 (+1.4 mIoU) datasets.

Then, we evaluate on object detection and instance segmentation tasks on COCO [38] dataset. Similarly, all the unsupervised representations are transferred with the same fine-tuning recipe of Mask-RCNN [24]. We provide AP results in Tab. 6. The same with semantic segmentation, on these two downstream tasks, our unsupervised pre-training representations are better than the supervised one. And still, SMOG improves the current SOTA results, +0.7 AP on object detection and +1.1 AP on instance segmentation.

4.4 Ablation Study

We provide ablation studies on key components of SMOG. ResNet-50 is adopted. Network is trained for 100 epochs without multi-crop. The representation is evaluated under the linear protocol.

Grouping quality To evaluate the grouping quality, in Fig. 5, we provide the density distribution of group entropy. All the three group-discrimination methods (our SMOG, SwAV, and DeepClusteringv2) have 3k groups. For

each group, its entropy is $\sum_i -p_i * \log(p_i)$ where p_i is the ratio of instances that belong to class i (data annotation) in this group. A lower entropy means a group has more unitary semantics and is more meaningful. Thus, more groups with low entropy indicate higher grouping quality. From Fig. 5, we can see SMOG has twice more low-entropy groups than the two baselines, proving its superiority.

Momentum ratio β of group features From Tab. 7a, we can see that the linearly decreasing schedule for β is much better than the fixed schedule. This is because our momentum grouping scheme updates the group features with only a small part of features (a batch) in each iteration and at the beginning of training, the feature distribution changes drastically, a small β will lead to unstable groups. Thus, we need a large β at the beginning to solve the problem and the linear schedule is a straightforward solution. After adopting the linearly decreasing schedule, our SMOG is not sensitive to the final value of β and we adopt 0.99 as the default value for all the experiments on both CNN and Transformer.

Table 6: Transfer learning results on object detection and instance segmentation tasks. we adopt COCO as the fine-tuning dataset. Mask RCNN with ResNet-50-FPN is the detection and segmentation model. We report the AP metrics.

method	COCO det			COCO instance seg.		
	AP ^{bb}	AP ₅₀ ^{bb}	AP ₇₅ ^{bb}	AP ^{mk}	AP ₅₀ ^{mk}	AP ₇₅ ^{mk}
Rand Init	31.0	49.5	33.2	28.5	46.8	30.4
Supervised	38.9	59.6	42.7	35.4	56.5	38.1
InsDis [67]	37.4	57.6	40.6	34.1	54.6	36.4
PIRL [45]	38.5	57.6	41.2	34.0	54.6	36.2
MoCoV2	39.4	59.9	43.0	35.8	56.9	38.4
SwAV	38.5	60.4	41.4	35.4	57.0	37.7
DC-v2 [6]	38.3	60.3	41.3	35.4	56.7	38.0
BYOL	39.4	59.9	43.0	35.8	56.8	38.5
Barlow Twins	39.2	58.7	42.6	34.3	55.4	36.5
SMoG	40.1	61.6	43.7	36.9	58.7	39.3




Fig. 5: Entropy of each groups. SMOG produces groups with much lower entropy which represents a better grouping quality in terms of high-level semantics.

Avoiding collapse Still, because the group feature updating process cannot access the global feature distribution, there is a possibility of collapse during training. Thus, we adopt the periodical clustering (pd) trick to avoid this. The first two rows of Tab. 7b prove its necessity. We can see that it successfully makes the backbone learn useful representations but the performance is not good enough. We believe this is because the group features are always generated and updated based on the instance features generated by only f_θ , leading to misalignment with features from \hat{f}_η . Thus, we also reset \hat{f}_η with the parameters of f_θ after each clustering. The integrated pd with \hat{f}_η reset leads to ideal performance. And surprisingly, we find that only resetting \hat{f}_η without the clustering can avoid the collapse too. We think the sudden change of \hat{f}_η feature distribution also helps to avoid the network gradually falling into degeneration.

Table 7: Ablation study on SMOG. We adopt ResNet50 as the backbone and train the unsupervised algorithm on ImageNet for 100 epochs without multi-crop training strategy. We report the linear evaluation results (Top-1 accuracy).

(a) Momentum ratio β . Linearly decreasing schedule performs better. (b) Tricks dealing with collapse. The periodical clustering is necessary.

Schedule	Top1 acc	Tricks	Top1 acc
fixed 0.99 β	65.9	None	0.1
1.0 $\beta \rightarrow 0.9$ β	67.0	+ periodical clustering (pd)	53.7
1.0 $\beta \rightarrow 0.99$ β	67.2	+ reset \hat{f}_η periodically	53.4
1.0 $\beta \rightarrow 0.999$ β	67.1	+ pd & reset \hat{f}_η periodically	67.2

(c) Number of groups. 3k is enough for SMOG. (d) Grouping algo. Momentum update better balances grouping and learning.

# Groups	Top1 acc	method	Top1 acc
300	65.2	Randomly select	30.2
1000	67.0	Adopt latest	42.1
3000	67.2	Averaging update	65.8
10000	67.2	Momentum update	67.2

Number of groups In Tab. 7c, we evaluate the influence of the number of groups used in SMOG under the linear protocol. The results show that SMOG is not sensitive to the group number. Even if we tune it in a wide range (1k~30k), the performance is maintained at a stable level (± 0.2), as long as there are enough groups (300 groups are too few to perform well). This is consistent with the conclusion of SwAV [6]. More groups increase the computation time consuming of the momentum grouping algorithm (ϕ) since it needs node communication to synchronize the group features among all the nodes. Thus, we adopt 3K as the default setting for all the experiments.

Method for updating group features In Tab. 7d, we evaluate different operators for updating the group features. We consider 4 operators here: 1) Randomly select (RS): randomly select a group of latest instance features and adopt them as the group features. 2) Adopt latest (AL): $\mathbf{g}_k \leftarrow \text{mean}_{\mathbf{c}_t=\mathbf{g}_k} f_\theta(\mathbf{x}_t)$. 3) Averaging update (AU): $\mathbf{g}_k \leftarrow \mathbf{g}_k + (1/n) * (\text{mean}_{\mathbf{c}_t=\mathbf{g}_k} f_\theta(\mathbf{x}_t) - \mathbf{g}_k)$, where n is the total number of instances belonging to group k , including ones in current iteration. 4) Momentum update (MU): the proposed method described in Eq. 4.

As the random baseline, RS’s poor performance shows us that the grouping algorithm plays an important role in our SMOG method. AL also does not perform well, which implies that the local grouping method will not provide an efficient training signal for representation learning. AU follows the sequential k-means algorithm and achieves a relatively great performance. But compared to our MU, it is not the best choice for conducting grouping together with the representation learning. The specifically designed MU which gives new features more weights is more suitable for the group-level contrastive learning.

5 Conclusion

We extend the recently popular contrastive learning to group level with the proposed SMOG and push the performance of unsupervised representation under linear protocol to the vanilla supervised level. SMOG synchronously conducts the representation learning and grouping process to effectively achieve the group-level contrasting. We hope the newly designed group-level contrastive learning will be useful for the community to further develop visual unsupervised method.

Acknowledgement This work was supported by the National Key R&D Program of China (2021ZD0110700), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), Shanghai Qi Zhi Institute, and SHEITC (018-RGZN-02046).

References

1. Agrawal, P., Carreira, J., Malik, J.: Learning to see by moving. In: ICCV. pp. 37–45 (2015)
2. Asano, Y.M., Patrick, M., Rupprecht, C., Vedaldi, A.: Labelling unlabelled videos from scratch with multi-modal self-supervision. arXiv preprint arXiv:2006.13662 (2020)
3. Asano, Y.M., Rupprecht, C., Vedaldi, A.: Self-labelling via simultaneous clustering and representation learning. arXiv preprint arXiv:1911.05371 (2019)
4. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F., et al.: Analysis of representations for domain adaptation. *NeurIPS* **19**, 137 (2007)
5. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: ECCV. pp. 132–149 (2018)
6. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. arXiv preprint arXiv:2006.09882 (2020)
7. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. arXiv preprint arXiv:2104.14294 (2021)
8. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607. PMLR (2020)
9. Chen, X., He, K.: Exploring simple siamese representation learning. In: CVPR. pp. 15750–15758 (2021)
10. Chen, X., Xie, S., He, K.: An empirical study of training self-supervised visual transformers. arXiv e-prints pp. arXiv–2104 (2021)
11. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR. pp. 3213–3223 (2016)
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255. IEEE (2009)
13. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV. pp. 1422–1430 (2015)
14. Doersch, C., Zisserman, A.: Multi-task self-supervised visual learning. In: ICCV. pp. 2051–2060 (2017)
15. Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. arXiv preprint arXiv:2104.14548 (2021)
16. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *IJCV* **88**(2), 303–338 (2010)
17. Fabbri, M., Lanzi, F., Calderara, S., Palazzi, A., Vezzani, R., Cucchiara, R.: Learning to detect and track visible and occluded body joints in a virtual world. In: ECCV. pp. 430–446 (2018)
18. Fang, H.S., Sun, J., Wang, R., Gou, M., Li, Y.L., Lu, C.: Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In: ICCV. pp. 682–691 (2019)
19. Fernando, B., Bilen, H., Gavves, E., Gould, S.: Self-supervised video representation learning with odd-one-out networks. In: CVPR. pp. 3636–3645 (2017)
20. Ge, C., Liang, Y., Song, Y., Jiao, J., Wang, J., Luo, P.: Revitalizing cnn attentions via transformers in self-supervised visual representation learning. arXiv preprint arXiv:2110.05340 (2021)

21. Grill, J.B., Strub, F., Alché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020)
22. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: International Conference on Artificial Intelligence and Statistics. pp. 297–304. JMLR Workshop and Conference Proceedings (2010)
23. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9729–9738 (2020)
24. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017)
25. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
26. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
27. Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. arXiv preprint arXiv:1808.06670 (2018)
28. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: ICML. pp. 448–456. PMLR (2015)
29. Jayaraman, D., Grauman, K.: Learning image representations tied to ego-motion. In: ICCV. pp. 1413–1421 (2015)
30. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
31. Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: CVPR. pp. 1920–1929 (2019)
32. Larsson, G., Maire, M., Shakhnarovich, G.: Colorization as a proxy task for visual understanding. In: CVPR. pp. 6874–6883 (2017)
33. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436–444 (2015)
34. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: CVPR. pp. 4681–4690 (2017)
35. Li, C., Yang, J., Zhang, P., Gao, M., Xiao, B., Dai, X., Yuan, L., Gao, J.: Efficient self-supervised vision transformers for representation learning. arXiv preprint arXiv:2106.09785 (2021)
36. Li, J., Zhou, P., Xiong, C., Hoi, S.C.: Prototypical contrastive learning of unsupervised representations. arXiv preprint arXiv:2005.04966 (2020)
37. Li, Y., Pogodin, R., Sutherland, D.J., Gretton, A.: Self-supervised learning with kernel dependence maximization. arXiv preprint arXiv:2106.08320 (2021)
38. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755. Springer (2014)
39. Liu, S., Zhong, G., De Mello, S., Gu, J., Jampani, V., Yang, M.H., Kautz, J.: Switchable temporal propagation network. In: ECCV. pp. 87–102 (2018)
40. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
41. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440 (2015)

42. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam (2018)
43. Lotter, W., Kreiman, G., Cox, D.: Deep predictive coding networks for video prediction and unsupervised learning. arXiv preprint arXiv:1605.08104 (2016)
44. Mathieu, M., Couprie, C., LeCun, Y.: Deep multi-scale video prediction beyond mean square error. arXiv preprint arXiv:1511.05440 (2015)
45. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: CVPR. pp. 6707–6717 (2020)
46. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: ECCV. pp. 527–544. Springer (2016)
47. Mitrovic, J., McWilliams, B., Walker, J., Buesing, L., Blundell, C.: Representation learning via invariant causal mechanisms. arXiv preprint arXiv:2010.07922 (2020)
48. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: ECCV. pp. 69–84. Springer (2016)
49. Noroozi, M., Vinjimoor, A., Favaro, P., Pirsiavash, H.: Boosting self-supervised learning via knowledge transfer. In: CVPR. pp. 9359–9367 (2018)
50. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
51. Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q.: Domain adaptation via transfer component analysis. IEEE transactions on neural networks **22**(2), 199–210 (2010)
52. Pang, B., Li, Y., Zhang, Y., Tang, J., Zha, K., Li, J., Lu, C.: Unsupervised representation for semantic segmentation by implicit cycle-attention contrastive learning. In: AAAI
53. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. pp. 2536–2544 (2016)
54. Qian, Q., Xu, Y., Hu, J., Li, H., Jin, R.: Unsupervised visual representation learning by online constrained k-means. arXiv preprint arXiv:2105.11527 (2021)
55. Rezende, D.J., Mohamed, S., Wierstra, D.: Stochastic backpropagation and variational inference in deep latent gaussian models. In: ICML. vol. 2, p. 2. Citeseer (2014)
56. Srivastava, N., Mansimov, E., Salakhudinov, R.: Unsupervised learning of video representations using lstms. In: ICML. pp. 843–852. PMLR (2015)
57. Sun, J., Fang, H.S., Zhu, X., Li, J., Lu, C.: Correlation field for boosting 3d object detection in structured scenes **36**(1) (2022)
58. Sun, J., Li, Y., Chai, L., Fang, H.S., Li, Y.L., Lu, C.: Human trajectory prediction with momentary observation. In: CVPR. pp. 6467–6476 (2022)
59. Tao, C., Wang, H., Zhu, X., Dong, J., Song, S., Huang, G., Dai, J.: Exploring the equivalence of siamese self-supervised learning via a unified gradient framework. arXiv preprint arXiv:2112.05141 (2021)
60. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? arXiv preprint arXiv:2005.10243 (2020)
61. Tung, H.F., Tung, H., Yumer, E., Fragkiadaki, K.: Self-supervised learning of motion capture. CoRR (2017)
62. Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: ICML. pp. 1096–1103 (2008)
63. Vondrick, C., Pirsiavash, H., Torralba, A.: Anticipating visual representations from unlabeled video. In: CVPR. pp. 98–106 (2016)
64. Vondrick, C., Pirsiavash, H., Torralba, A.: Generating videos with scene dynamics. arXiv preprint arXiv:1609.02612 (2016)
65. Wang, X., Liu, Z., Yu, S.X.: Unsupervised feature learning by cross-level instance-group discrimination. In: CVPR. pp. 12586–12595 (2021)

66. Wei, D., Lim, J.J., Zisserman, A., Freeman, W.T.: Learning and using the arrow of time. In: CVPR. pp. 8052–8060 (2018)
67. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR. pp. 3733–3742 (2018)
68. Xie, Z., Lin, Y., Yao, Z., Zhang, Z., Dai, Q., Cao, Y., Hu, H.: Self-supervised learning with swin transformers. arXiv preprint arXiv:2105.04553 (2021)
69. You, Y., Gitman, I., Ginsburg, B.: Scaling sgd batch size to 32k for imagenet training. arXiv preprint arXiv:1708.03888 **6**, 12 (2017)
70. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. arXiv preprint arXiv:2103.03230 (2021)
71. Zhan, X., Xie, J., Liu, Z., Ong, Y.S., Loy, C.C.: Online deep clustering for unsupervised representation learning. In: CVPR. pp. 6688–6697 (2020)
72. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. ICLR (2018)
73. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: ECCV. pp. 649–666. Springer (2016)
74. Zheng, M., You, S., Wang, F., Qian, C., Zhang, C., Wang, X., Xu, C.: Resl: Relational self-supervised learning with weak augmentation. arXiv preprint arXiv:2107.09282 (2021)
75. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR. pp. 2921–2929 (2016)
76. Zhuang, C., Zhai, A.L., Yamins, D.: Local aggregation for unsupervised learning of visual embeddings. In: ICCV. pp. 6002–6012 (2019)