



HAL
open science

Learning Representations by Predicting Bags of Visual Words

Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, Matthieu Cord

► **To cite this version:**

Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, Matthieu Cord. Learning Representations by Predicting Bags of Visual Words. CVPR2020 - Conference on Computer Vision and Pattern Recognition, Jun 2020, Seattle, WA, United States. hal-02505058

HAL Id: hal-02505058

<https://hal.sorbonne-universite.fr/hal-02505058v1>

Submitted on 11 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning Representations by Predicting Bags of Visual Words

Spyros Gidaris¹ Andrei Bursuc¹ Nikos Komodakis² Patrick Pérez¹ Matthieu Cord¹

¹Valeo.ai ²University of Crete

Abstract

Self-supervised representation learning targets to learn convnet-based image representations from unlabeled data. Inspired by the success of NLP methods in this area, in this work we propose a self-supervised approach based on spatially dense image descriptions that encode discrete visual concepts, here called visual words. To build such discrete representations, we quantize the feature maps of a first pre-trained self-supervised convnet, over a k-means based vocabulary. Then, as a self-supervised task, we train another convnet to predict the histogram of visual words of an image (i.e., its Bag-of-Words representation) given as input a perturbed version of that image. The proposed task forces the convnet to learn perturbation-invariant and context-aware image features, useful for downstream image understanding tasks. We extensively evaluate our method and demonstrate very strong empirical results, e.g., our pre-trained self-supervised representations transfer better on detection task and similarly on classification over classes “unseen” during pre-training, when compared to the supervised case.

This also shows that the process of image discretization into visual words can provide the basis for very powerful self-supervised approaches in the image domain, thus allowing further connections to be made to related methods from the NLP domain that have been extremely successful so far.

1. Introduction

The goal of our work is to learn convolutional neural network [38] (convnet) based representations without human supervision. One promising approach towards this goal is the so-called self-supervised representation learning [14, 75, 36, 47, 20, 51], which advocates to train the convnet with an annotation-free pretext task defined using only the information available within an image, e.g., predicting the relative location of two image patches [14]. Pre-training on such a pretext task enables the convnet to learn representations that are useful for other vision tasks of actual interest, such as image classification or object detection. Moreover,

recent work has shown that self-supervision can be beneficial to many other learning problems [61, 18, 74, 28, 10, 10, 29], such as few-shot [61, 18] and semi-supervised [74, 28] learning, or training generative adversarial networks [10].

A question that still remains open is what type of self-supervision we should use. Among the variety of the proposed learning tasks, many follow the general paradigm of first perturbing an image or removing some part/aspect of the image and then training the convnet to reconstruct the original image or the dropped part (e.g., color channel, image region). Popular examples are Denoising AutoEncoders [65], Image Colorization [75, 36], Split-Brain architectures [76], and Image In-painting [51]. However, predicting such low-level image information can be a difficult task to solve, and does not necessarily force the convnet to acquire image understanding “skills”, which is what we ultimately want to achieve. As a result, such reconstruction-based methods have not been very successful so far. In contrast, in Natural Language Processing (NLP), similar self-supervised methods, such as predicting the missing words of a sentence (e.g., BERT [12] and ROBERTA [40]), have proven much more successful at learning strong language representations. The difference of those NLP methods with their computer vision counterparts is that (1) words undoubtedly represent more high-level semantic concepts than raw image pixels. Also, (2) words are defined in a discrete space while images in a continuous one where, without changing the depicted content, small pixel perturbations can significantly alter the target of a reconstruction task.

Spatially dense image quantization into visual words.

Inspired by the above NLP methods, in this work we propose for self-supervised learning in the image domain to use tasks that aim at predicting/reconstructing targets that encode *discrete visual concepts* as opposed, e.g., to (low-level) pixel information. To build such discrete targets, we first take an existing self-supervised method (e.g., rotation prediction [20]) and use it to train an initial convnet, which can learn feature representations that capture mid-to-higher-level image features. Then, for each image, we densely quantize its convnet-based feature map using a k-means-based vocab-




Figure 1: **Learning representations through prediction of Bags of Visual Words.** We first train a feature extractor $\hat{\Phi}(\cdot)$ for a self-supervised task, *e.g.* rotation prediction. Then we compute a visual vocabulary from feature vectors computed from $\hat{\Phi}$ feature maps and compute the corresponding image level BoW vectors. These BoW vectors will serve as ground truth for the next stage. In the second stage we perturb images with $g(\cdot)$ and send them as input to a second network $\Phi(\cdot)$. The BoW prediction module $\Omega(\cdot)$ processes $\Phi(\cdot)$ feature maps to predict BoW vectors corresponding to the original non-perturbed images. Both $\Phi(\cdot)$ and $\Omega(\cdot)$ are trained jointly with cross-entropy loss. The feature extractor $\Phi(\cdot)$ is further used for downstream tasks.

ulary.¹ This results in a spatially dense image description based on discrete codes (*i.e.*, k -means cluster assignments), called visual words hereafter. Such a discrete image representation opens the door to easily adapting self-supervised methods from the NLP community to the image domain. For instance, in this case, one could very well train a BERT-like architecture that, given as input a subset of the patches in an image, predicts the visual words of the missing patches. Although self-supervised methods of this type are definitely something that we plan to explore as future work, in this paper we aim to go one step further and develop (based on the above discrete visual representations) self-supervised tasks that furthermore allow using standard convolutional architectures that are commonly used (and optimized) for the image domain we are interested in. But how should we go about defining such a self-supervised task?

Learning by “reconstructing” bags of visual words. To this end, we take inspiration from the so-called Bag-of-Words [70] (BoW) model in computer vision and propose using as self-supervised task one where we wish (to train a convnet) to predict the histogram of visual words of an image (also known as its BoW representation) when given as input a perturbed version of that image. This type of BoW representations have been very powerful image models, and as such have been extensively used in the past in several computer vision problems (including, *e.g.*, image retrieval, object recognition, and object detection). Interestingly, there is recent empirical evidence that even modern state-of-the-art convnets for image classification exhibit similar behavior to BoW models [7]. By using the above BoW prediction task in

¹Here, by dense quantization, we refer to the fact that each spatial location of the feature map is quantized separately.

the context of self-supervised learning, one important benefit is that it is no longer required to enhance a typical convnet architecture for images (*e.g.*, ResNet-50) with extra network components, such as multiple stacks of attention modules as in [64] or PixelCNN-like autoregressors as in [48], that can make the overall architecture computationally intensive. Furthermore, due to its simplicity, it can be easily incorporated into other types of learning problems (*e.g.*, few-shot learning, semi-supervised learning, or unsupervised domain adaptation), thus allowing to further improve performance for these problems which is an additional advantage.

Concerning the perturbed image (that is used as input to the BoW prediction task), it is generated by applying a set of (commonly used) augmentation techniques such as random cropping, color jittering, or geometric transformations. Therefore, to solve the task of “reconstructing” the BoW histogram of the original image, the convnet must learn to detect visual cues that remain constant (*i.e.*, invariant) to the applied perturbations. Moreover, since the perturbed image can often be only a small part of the original one (due to the cropping transformation), the convnet is also forced to infer the context of the missing input, *i.e.*, the visual words of the missing image regions. This encourages learning of perturbation-invariant and context-aware image features, which, as such, are more likely to encode higher-level semantic visual concepts. Overall, as we show in the experimental results, this has as a result that the proposed self-supervised method learns representations that transfer significantly better to downstream vision tasks than the representations of the initial convnet. As a last point, we note that the above process of defining a convnet-based BoW model and then training another convnet to predict it, can be applied iteratively, which can lead to even better representations.

Contributions. To summarize, the contributions of our work are: (1) We propose the use of discrete visual word representations for self-supervised learning in the image domain. (2) In this context, we propose a novel method for self-supervised representation learning (Fig. 1). Rather than predicting/reconstructing image-pixel-level information, it uses a first self-supervised pre-trained convnet to densely discretize an image to a set of visual words and then trains a second convnet to predict a reduced Bag-of-Words representation of the image given as input perturbed versions of it. (3) We extensively evaluate our method and we demonstrate that it manages to learn high-quality convnet-based image representations, which are significantly superior to those of the first convnet. Furthermore, *our ImageNet-trained self-supervised ResNet-50 representations, when compared to the ImageNet-trained supervised ones, achieve better VOC07+12 detection performance and comparable Places205 classification accuracy, i.e., better generalization on the detection task and similar generalization on the Places205 classes which are “unseen” during self-supervised training.* (4) The simple design of our method allows someone to easily use it on many other learning problems where self-supervision has been shown to be beneficial.

2. Approach

Our goal is to learn in an unsupervised way a feature extractor or convnet model $\Phi(\cdot)$ parameterized by θ that, given an image \mathbf{x} , produces a “good” image representation $\Phi(\mathbf{x})$. By “good” we mean a representation that would be useful for other vision tasks of interest, *e.g.* image classification, object detection. To this end, we assume that we have available a large set of unlabeled images X on which we will train our model. We also assume that we have available an initial self-supervised pre-trained convnet $\hat{\Phi}(\cdot)$. We can easily learn such a model by employing one of the available self-supervised tasks. Here, except otherwise stated, we use RotNet [20] (which is based on the self-supervised task of image rotation prediction) as it is easy to implement and, at the same time, has been shown to achieve strong results in self-supervised representation learning [35].

To achieve our goal, we leverage the initial model $\hat{\Phi}(\cdot)$ to create spatially dense descriptions based on visual words. Then, we aggregate those descriptions into BoW representations and train the model $\Phi(\cdot)$ to “reconstruct” the BoW of an image \mathbf{x} given as input a perturbed version of it. Note that the model $\hat{\Phi}(\cdot)$ remains frozen during the training of the new model $\Phi(\cdot)$. Also, after training $\Phi(\cdot)$, we can set $\hat{\Phi}(\cdot) \leftarrow \Phi(\cdot)$ and repeat the training process.

2.1. Building spatially dense discrete descriptions

Given a training image \mathbf{x} , the first step for our method is to create a spatially dense visual words-based description $q(\mathbf{x})$ using the pre-trained convnet $\hat{\Phi}(\cdot)$. Specifically,

let $\hat{\Phi}(\mathbf{x})$ be a feature map (with \hat{c} channels and $\hat{h} \times \hat{w}$ spatial size) produced by $\hat{\Phi}(\cdot)$ for input \mathbf{x} , and $\hat{\Phi}^u(\mathbf{x})$ the \hat{c} -dimensional feature vector at the location $u \in \{1, \dots, U\}$ of this feature map, where $U = \hat{h} \cdot \hat{w}$. To generate the description $q(\mathbf{x}) = [q^1(\mathbf{x}), \dots, q^U(\mathbf{x})]$, we densely quantize $\hat{\Phi}(\mathbf{x})$ using a predefined vocabulary $V = [\mathbf{v}_1, \dots, \mathbf{v}_K]$ of \hat{c} -dimensional visual word embeddings, where K is the vocabulary size. In detail, for each position u , we assign the corresponding feature vector $\hat{\Phi}^u(\mathbf{x})$ to its closest (in terms of squared Euclidean distance) visual word embedding $q^u(\mathbf{x})$:

$$q^u(\mathbf{x}) = \arg \min_{k=1 \dots K} \|\hat{\Phi}^u(\mathbf{x}) - \mathbf{v}_k\|_2^2. \quad (1)$$

The vocabulary V is learned by applying the k-means algorithm with K clusters to a set of feature maps extracted from the dataset X , *i.e.*, by optimizing the following objective:

$$\min_V \sum_{\mathbf{x} \in X} \sum_u \left[\min_{k=1, \dots, K} \|\hat{\Phi}_p^u(\mathbf{x}) - \mathbf{v}_k\|_2^2 \right], \quad (2)$$

where the visual word embedding \mathbf{v}_k is the centroid of the k -th cluster.

2.2. Generating Bag-of-Words representations

Having generated the discrete description $q(\mathbf{x})$ of image \mathbf{x} , the next step is to create its BoW representation, denoted by $y(\mathbf{x})$. This is a K -dimensional vector whose k -th element $y^k(\mathbf{x})$ either encodes the number of times the k -th visual word appears in image \mathbf{x} ,

$$y^k(\mathbf{x}) = \sum_{u=1, \dots, U} \mathbb{1}[q^u(\mathbf{x}) = k], \quad (3)$$

or indicates if the k -th visual word appears in image \mathbf{x} ,

$$y^k(\mathbf{x}) = \max_{u=1, \dots, U} \mathbb{1}[q^u(\mathbf{x}) = k], \quad (4)$$

where $\mathbb{1}[\cdot]$ is the indicator operator.² Furthermore, to convert $y(\mathbf{x})$ into a probability distribution over visual words, we L_1 -normalize it, *i.e.*, we set $y^k(\mathbf{x}) = \frac{y^k(\mathbf{x})}{\sum_m y^m(\mathbf{x})}$. The resulting $y(\mathbf{x})$ can thus be perceived as a soft categorical label of \mathbf{x} for the K visual words. Note that, although K might be very large, the BoW representation $y(\mathbf{x})$ is actually quite sparse as it has at most U non-zero elements.

2.3. Learning to “reconstruct” BoW

Based on the above BoW representation, we propose the following self-supervised task: given an image \mathbf{x} , we first apply to it a perturbation operator $g(\cdot)$, to get the perturbed image $\tilde{\mathbf{x}} = g(\mathbf{x})$, and then train the model to predict/“reconstruct” the BoW representation $y(\mathbf{x})$ of the original unperturbed image \mathbf{x} from $\tilde{\mathbf{x}}$. This, in turn, means that we

²In our experiments we use the binary version (4) [59, 32] for ImageNet and the histogram version (3) for CIFAR-100 and MiniImageNet.

want to predict the BoW representation $y(\mathbf{x})$ from the feature vector $\Phi(\tilde{\mathbf{x}})$ (where hereafter we assume that $\Phi(\tilde{\mathbf{x}}) \in \mathbb{R}^c$, *i.e.*, the feature representation produced by model $\Phi(\cdot)$ is c -dimensional).³ To this end, we define a prediction layer $\Omega(\cdot)$ that gets $\Phi(\tilde{\mathbf{x}})$ as input and outputs a K -dimensional softmax distribution over the K visual words of the BoW representation. More precisely, the prediction layer is implemented with a linear-plus-softmax layer:

$$\Omega^k(\Phi(\tilde{\mathbf{x}})) = \text{softmax}_k [\gamma \Phi(\tilde{\mathbf{x}})^T \bar{\mathbf{w}}_m]_{m \in [1, \dots, K]}, \quad (5)$$

where $\Omega^k(\Phi(\tilde{\mathbf{x}}))$ is the softmax probability for the k -th visual word, and $W = [\mathbf{w}_1, \dots, \mathbf{w}_K]$ are the K c -dimensional weight vectors (one per visual word) of the linear layer. Notice that, instead of directly applying the weights vectors W to the feature vector $\Phi(\tilde{\mathbf{x}})$, we use their L_2 -normalized versions $\bar{\mathbf{w}}_k = \mathbf{w}_k / \|\mathbf{w}_k\|_2$, and apply a unique learnable magnitude γ for all the weight vectors (γ is a scalar value). The reason for this reparametrization of the linear layer is because the distribution of visual words in the dataset (*i.e.*, how often, or in how many dataset-images, a visual word appears) tends to be unbalanced and, so, without such a reparametrization the network would attempt to make the magnitude of each weight vector proportional to the frequency of its corresponding visual word (thus basically always favoring the most frequently occurring words). In our experiments, the above reparametrization has led to significant improvements in the quality of the learned representations.

Self-supervised training objective. The training loss that we minimize for learning the convnet model $\Phi(\cdot)$ is the expected cross-entropy loss between the predicted softmax distribution $\Omega(\Phi(\tilde{\mathbf{x}}))$ and the BoW distribution $y(\mathbf{x})$:

$$L(\theta, W, \gamma; X) = \mathbb{E}_{\mathbf{x} \sim X} \left[\text{loss}(\Omega(\Phi(\tilde{\mathbf{x}})), y(\mathbf{x})) \right], \quad (6)$$

where $\text{loss}(\alpha, \beta) = -\sum_{k=1}^K \beta^k \log \alpha^k$ is the cross-entropy loss for the discrete distributions $\alpha = (\alpha^k)$ and $\beta = (\beta^k)$, θ are the learnable parameters of $\Phi(\cdot)$, (W, γ) are the learnable parameters of $\Omega(\cdot)$, and $\tilde{\mathbf{x}} = g(\mathbf{x})$.

Image perturbations. The perturbation operator $g(\cdot)$ that we use consists of (a) color jittering (*i.e.*, random changes of the brightness, contrast, saturation, and hue of an image) (b) converting the image to grayscale with probability p , (c) random image cropping, (d) scale or aspect ratio distortions, and (e) horizontal flips. The role served by such an operator is two-fold: to solve the BoW “reconstruction” task after such aggressive perturbations, the convnet must learn image features that (1) are robust w.r.t. the applied perturbations

and at the same time (2) allow predicting the visual words of the original image, even for image regions that are not visible to the convnet due to cropping. To further push towards this direction, we also incorporate the CutMix [72] augmentation technique into our self-supervised method. According to CutMix, given two images $\tilde{\mathbf{x}}_A = g(\mathbf{x}_A)$ and $\tilde{\mathbf{x}}_B = g(\mathbf{x}_B)$, we generate a new synthetic one $\tilde{\mathbf{x}}_S$ by replacing a patch of the first image $\tilde{\mathbf{x}}_A$ with one from the second image $\tilde{\mathbf{x}}_B$. The position and size of the patch is randomly sampled from a uniform distribution. The BoW representation that is used as a reconstruction target for this synthetic image is the convex combination of the BoW targets of the two images, $\lambda y(\mathbf{x}_A) + (1 - \lambda)y(\mathbf{x}_B)$, where $1 - \lambda$ is the patch-over-image area ratio. Hence, with CutMix we force the convnet to infer both (a) the visual words that belong on the patch that was removed from the first image $\tilde{\mathbf{x}}_A$, and (b) the visual words that belong on the image area that surrounds the patch that was copied from second image $\tilde{\mathbf{x}}_B$.

Model initialization and iterated training. We note that the model $\hat{\Phi}(\cdot)$ is used only for building BoW representations and not for initializing the parameters of the $\Phi(\cdot)$ model, *i.e.*, $\Phi(\cdot)$ is randomly initialized before training. Also, as already mentioned, we can apply our self-supervised method iteratively, using each time the previously trained model $\hat{\Phi}(\cdot)$ for creating the BoW representation. We also note, however, that this is not necessary for learning “good” representations; the model learned from the first iteration already achieves very strong results. As a result, only a few more iterations (*e.g.*, one or two) might be applied after that.

3. Related Work

Bag-of-Words. BoW is a popular method for text document representation, which has been adopted and heavily used in computer vision [59, 11]. For visual content, BoW conveniently encapsulates image statistics from hundreds of local features [41] into vector representations. BoW have been studied extensively and leveraged in numerous tasks, while multiple extensions [52, 33] and theoretical interpretations [63] have been proposed. Due to its versatility, BoW has been applied to pre-trained convnets as well to compute image representations from intermediate feature maps [71, 22, 46], however few works have dealt with the integration of BoW in the training pipeline of a convnet. Among them, NetVLAD [2] mimics the BoW-derived VLAD descriptor by learning a visual vocabulary along with the other layers and soft quantizing activations over this vocabulary. Our method differs from previous approaches in training with self-supervision and in predicting directly the BoW vector bypassing quantization and aggregation.

Self-supervision. Self-supervised learning is a recent paradigm aiming to learn representations from data by lever-

³*E.g.*, in the case of ResNet50, $\Phi(\mathbf{x})$ corresponds to the 2048-dimensional feature vector (*i.e.*, $c = 2048$) produced from the global average pooling layer that follows the last block of residual layers.

aging supervision from various intrinsic data signals without any explicit manual annotations and human supervision. The representations learned with self-supervision are then further fine-tuned on a downstream task with limited human annotations available. Numerous creative mechanisms for squeezing out information from data in this manner have been proposed in the past few years: predicting the colors of image [36, 75], the relative position of shuffled image patches [14, 47], the correct order of a set of shuffled video frames [45], the correct association between an image and a sound [3], and many other methods [67, 78, 39, 68].

Learning to reconstruct. Multiple self-supervised methods are formulated as reconstruction problems [65, 43, 34, 78, 21, 76, 51, 75, 1, 54]. The information to be reconstructed can be provided by a different view [21, 78, 54] or sensor [16]. When no such complementary information is available, the current data can be perturbed and the task of the model is now to reconstruct the original input. Denoising an input image back to its original state [65], inpainting an image patch that has been removed from a scene [51], reconstructing images that have been overlaid [1] are some of the many methods of reconstruction from perturbation. While such approaches display impressive results for the task hand, it remains unclear how much structure they can encapsulate beyond the reconstruction of visual patterns [13]. Similar ideas have been initially proposed in NLP where missing words [43] or sentences [34] must be reconstructed. Another line of research employs perturbations for guiding the model towards assigning the same label to both perturbed and original content [15, 5]. Our method also deals with perturbed inputs, however instead of reconstructing the input, we train it to reconstruct the BoW vector of the clean input. This enables learning non-trivial correlations between local visual patterns in the image.

VQ-VAE. These works [49, 56] explore the learning of spatially-dense discrete representations with unsupervised generative models with goal of image generation. Instead, we focus on exploiting discrete image representations in the context of self-supervised image representation learning.

3.1. Discussion

Relation to clustering-based representation learning methods [9, 8, 4]. Our work presents similarities to the Deep(er) Clustering approach [9, 8]. The latter alternates between k-means clustering the images based on their convnet features and using the cluster assignments as image labels for training the convnet. In our case however, we use the k-means clustering for creating BoW representations instead of global image labels. The former leads to richer (more complete) image descriptions compared to the latter as it encodes *multiple local visual concepts* extracted in a *spatially dense* way. For example, a cluster id is not sufficient to describe

an image with multiple objects, like the one in Figure 1, while a BoW is better suited for that. This fundamental difference leads to a profoundly different self-supervised task. Specifically, in our case the convnet is forced to: (1) focus on more localized visual patterns and (2) learn better contextual reasoning (since it must predict the visual words of missing image regions).

Relation to contrastive self-supervised learning methods [15, 5, 25, 44]. Our method bears similarities to recent works exploiting contrastive losses for learning representations that are invariant under strong data augmentation or perturbations [15, 5, 25, 44]. These methods deal with image recognition and the same arguments mentioned above w.r.t. [9, 8] hold. Our contribution departs from this line of approaches allowing our method to be applied to a wider set of visual tasks. For instance, in autonomous driving, most urban images are similar and differ only by few details, e.g. a pedestrian or a car, making image recognition under strong perturbation less feasible. In such cases, leveraging local statistics as done in our method appears as a more appropriate self-supervised task for learning representations.

4. Experiments and results

We evaluate our method (BoWNet) on CIFAR-100, Mini-ImageNet [66], ImageNet [58], Places205 [77], VOC07 [17] classification, and VOC07+12 detection datasets.

4.1. Analysis on CIFAR-100 and MiniImageNet

4.1.1 Implementation details

CIFAR-100. CIFAR-100 consists of $50k$ training images with 32×32 resolution. We train self-supervised WRN-28-10 [73] convnets using those training images. Specifically, we first train a WRN-28-10 based RotNet [20] and build based on that BoW using the feature maps of its last/3rd residual block. Then, we train the BoWNet using those BoW representations. We use a $K = 2048$ vocabulary size. The prediction head of RotNet consists of an extra residual block (instead of just a linear layer); in our experiments this led the feature extractor to learn better representations (we followed this design choice of RotNet for all the experiments in our paper; we provide more implementation details in §C.1).

We train the convnets using stochastic gradient descent (SGD) for 30 epochs of 2000 batch iterations and batch size 128. The learning rate is initialized with 0.1 and multiplied by 0.1 after 12 and 24 epochs. The weight decay is $5e - 4$.

MiniImageNet. Since MiniImageNet is used for evaluating few-shot methods, it has three different splits of classes, train, validation, and test with 64, 16, and 20 classes respectively. Each class has 600 images with 84×84 resolution. We train WRN-28-4 convnets on the 64×600 images that correspond to the training classes following the same training protocol as for CIFAR-100.

4.1.2 Evaluation protocols.

CIFAR-100. To evaluate the learned representations we use two protocols. **(1)** The first is to freeze the learned representations (which in case of WRN-28-10 is a 640-dimensional vector) and train on top of them a 100-way linear classifier for the CIFAR-100 classification task. We use the linear classifier accuracy as an evaluation metric. The linear classifier is trained with SGD for 60 epochs using a learning rate of 0.1 that is multiplied by 0.3 every 10 epochs. The weight decay is 0.001. **(2)** For the second protocol we use a few-shot episodic setting [66] similar to what is proposed on [18]. Specifically, we choose 20 classes from CIFAR-100 and run with them multiple (2000) episodes of 5-way few-shot classification tasks. Essentially, at each episode we randomly sample 5 classes from the 20 ones and then n training examples and $m = 15$ test examples per class (both randomly sampled from the test images of CIFAR-100). For n we use 1, 5, 10, and 50 examples (1-shot, 5-shot, 10-shot, and 50-shot settings respectively). To classify the m examples we use a cosine distance Prototypical-Networks [60] classifier that is applied on top of the frozen representations. We report the mean accuracy over the 2000 episodes. The purpose of this metric is to analyze the ability of the representations to be used for learning with few training examples. More details about this protocol are provided in §C.3.

MiniImageNet. We use the same two protocols as in CIFAR-100. **(1)** The first is to train 64-way linear classifiers on the task of recognizing the 64 training classes of MiniImageNet. Here, we use the same hyper-parameters as for CIFAR-100. **(2)** The second protocol is to use the frozen representations for episodic few-shot classification [66]. The main difference with CIFAR-100 is that here we evaluate using the test classes of MiniImageNet, which were not part of the training set of the self-supervised models. Therefore, with this evaluation we analyze the ability of the representations to be used for learning with few training examples and for “unseen” classes during training. For comparison with this protocol we provide results of the supervised Cosine Classifier (CC) few-shot model [19, 55].

4.1.3 Results

In Tables 1 and 2 we report results for our self-supervised method on the CIFAR-100 and MiniImageNet datasets respectively. By comparing BoWNet with RotNet (that we used for building BoW), we observe that BoWNet improves all the evaluation metrics by at least 10 percentage points, which is a very large performance improvement. Applying BoWNet iteratively (entries BoWNet $\times 2$ and BoWNet $\times 3$) further improves the results (except the 1-shot accuracy). Also, BoWNet outperforms by a large margin the CIFAR-100 linear classification accuracy of the recently proposed AMDIM [5] method (see Table 1), which has been shown to

Method	$n=1$	5	10	50	Linear
Supervised [73]	-	-	-	-	79.5
RotNet	58.3	74.8	78.3	81.9	60.3
Deeper Clustering	65.9	84.6	87.9	90.8	65.4
AMDIM [5]	-	-	-	-	70.2
BoWNet	69.1	86.3	89.2	92.4	71.5
BoWNet $\times 2$	68.5	87.1	90.4	93.8	74.1
BoWNet $\times 3$	68.4	87.2	90.4	93.9	74.5
BoWNet w/o cutmix	68.5	85.8	88.8	92.2	69.7
Sp-BoWNet	67.7	85.8	89.2	92.3	71.3

Table 1: CIFAR-100 linear classifier and few-shot results with WRN-28-10. For few-shot we use $n=1, 5, 10,$ or 50 examples per class. AMDIM uses a higher-capacity custom made architecture.

Method	Classes				Base Linear
	$n=1$	5	10	50	
Supervised CC [19]	56.8	74.1	78.1	82.7	73.7
RotNet	40.8	56.9	61.8	68.1	52.3
RelLoc [14]	40.2	57.1	62.6	68.8	50.4
Deeper Clustering	47.8	66.6	72.1	78.4	60.3
BoWNet	48.7	67.9	74.0	79.9	65.0
BoWNet $\times 2$	49.1	67.6	73.6	79.9	65.6
BoWNet $\times 3$	48.6	68.9	75.3	82.5	66.0

Table 2: MiniImageNet linear classifier and few-shot results with WRN-28-4.

Method	Classes				Base Linear
	$n=1$	5	10	50	
RotNet	40.8	56.9	61.8	68.1	52.3
RelLoc	40.2	57.1	62.6	68.8	50.4
RotNet \rightarrow BoWNet	48.7	67.9	74.0	79.9	65.0
RelLoc \rightarrow BoWNet	51.8	70.7	75.9	81.3	65.2
Random \rightarrow BoWNet	42.4	62.0	68.9	78.1	61.5

Table 3: MiniImageNet linear classifier and few-shot results with WRN-28-4. Impact of base convnet.

achieve very strong results. Finally, the performance of the BoWNet representations on the MiniImageNet novel classes for the 10-shot and especially the 50-shot setting are very close to the that of the supervised CC model (see Table 2).

Impact of CutMix augmentation. In Table 1 we report CIFAR-100 results without CutMix, which confirms that employing CutMix does indeed provide some further improvement on the quality of the learned representations.

Spatial-Pyramid BoW [37]. By reducing the visual words descriptions to BoW histograms, we remove spatial information from the visual word representations. To avoid this, one could divide the image into several spatial grids of

different resolutions, and then extract a BoW representation from each resulting image patch. In Table 1 we provide results for such a case (entry Sp-BoWNet). Specifically, we used 2 levels for the spatial pyramid, one with 2×2 and one with 1×1 resolution, giving in total 5 BoW. Although one would expect otherwise, we observe that adding more spatial information to the BoW via Spatial-Pyramids-BoW, does not improve the quality of the learned representations.

Comparison with Deeper Clustering (DC) [9]. To compare our method against DC we implemented it using the same convnet backbone and the same pre-trained RotNet as for BoWNet. To be sure that we do not disadvantage DC in any way, we optimized the number of clusters (512), the training routine⁴, and applied the same image augmentations (including cutmix) as in BoWNet. We also boosted DC by combining it with rotation prediction and applied our re-parametrization of the linear prediction layer (see equation (5)), which however did not make any difference in the DC case. We observe in Tables 1 and 2 that, although improving upon RotNet, DC has finally a significantly lower performance compared to BoWNet, e.g., several absolute percentage points lower linear classification accuracy, which illustrates the advantage of using BoW as targets for self-supervision instead of the single cluster id of an image.

Impact of base convnet. In Table 3 we provide MiniImageNet results using RelLoc [14] as the initial convnet with which we build BoW (base convnet). RelLoc \rightarrow BoWNet achieves equally strong or better results than in the RotNet \rightarrow BoWNet case. We also conducted (preliminary) experiments with a randomly initialized base convnet (entry Random \rightarrow BoWNet). In this case, to learn good representations, (a) we used in total 4 training rounds, (b) for the 1st round we built BoW from the 1st residual block of the randomly initialized WRN-28-4 and applied PCA analysis before k-means, (c) for the 2nd round we built BoW from 2nd residual block, and (d) for the remaining 2 rounds we built BoW from 3rd/last residual block. We observe that with a random base convnet the performance of BoWNet drops. However, BoWNet still is significantly better than RotNet and RelLoc.

We provide additional experimental results in §B.1.

4.2. Self-supervised training on ImageNet

Here we evaluate BoWNet by training it on the ImageNet dataset that consists of more than $1M$ images coming from 1000 different classes. We use the ResNet-50 (v1) [27] architecture with 224×224 input images for implementing the RotNet and BoWNet models. The BoWNet models are trained with 2 training rounds. For each round we use SGD, 140 training epochs, and a learning rate 0.03 that is multi-

⁴Specifically, we use 30 training epoch (the clusters are updated every 3 epochs), and a constant learning rate of 0.1 (same as in [9]). Each epoch consists of 2000 batch iterations with batch size 512. For simplicity we used one clustering level instead of the two hierarchical levels in [9].

Method	conv4	conv5
ImageNet Supervised [24]	80.4	88.0
RotNet*	64.6	62.8
Jigsaw [24]	64.5	57.2
Colorization [24]	55.6	52.3
BoWNet conv4	73.6	79.3
BoWNet conv5	74.3	78.4

Table 4: VOC07 image classification results for ResNet-50 Linear SVMs. *: our implementation.

Method	ImageNet		Places205	
	conv*	pool5	conv*	pool5
Random [24]	13.7	-	16.6	-
Supervised methods				
ImageNet [24]	75.5	-	51.5	-
ImageNet*	76.0	76.2	52.8	52.0
Places205 [24]	58.2	-	62.3	-
Prior self-supervised methods				
RotNet*	52.5	40.6	45.0	39.4
Jigsaw [24]	45.7	-	41.2	-
Colorization [24]	39.6	-	31.3	-
LA [†] [79]	60.2	-	50.1	-
Concurrent work				
MoCo [25]	-	<u>60.6</u>	-	-
PIRL [44]	63.6	-	49.8	-
CMC [‡] [62]	-	64.1	-	-
BowNet conv4	<u>62.5</u>	62.1	50.9	51.1
BowNet conv5	60.5	60.2	<u>50.1</u>	<u>49.5</u>

Table 5: ResNet-50 top-1 center-crop linear classification accuracy on ImageNet and Places205. pool5 indicates the accuracy for the 2048-dimensional features produced by the global average pooling layer after conv5. conv* indicates the accuracy of the best (w.r.t. accuracy) conv. layer of ResNet-50 (for the full results see §B.3). Before applying classifiers on those conv. layers, we resize their feature maps to around $9k$ dimensions (in the same way as in [24]). [†]: LA [79] uses 10-crops evaluation. [‡]: CMC [62] uses two ResNet-50 feature extractor networks. *: our implementation.

plied by 0.1 after 60, 100, and 130 epochs. The batch size is 256 and the weight decay is $1e - 4$. To build BoW we use a vocabulary of $K = 20000$ visual words created from the 3rd or 4th residual blocks (aka conv4 and conv5 layers respectively) of RotNet (for an experimental analysis of those choices see §B.2). We named those two models BowNet conv4 and BowNet conv5 respectively.

We evaluate the quality of the learned representations on ImageNet classification, Places205 classification, VOC07 classification, and VOC07+12 detection tasks.

Pre-training	AP ⁵⁰	AP ⁷⁵	AP ^{all}
Supervised*	80.8	58.5	53.2
Jigsaw [†] [24]	75.1	52.9	48.9
PIRL [†] [44]	80.7	59.7	54.0
MoCo [25]	81.4	61.2	<u>55.2</u>
BoWNet _{conv4}	80.3	60.4	55.0
BoWNet _{conv5}	<u>81.3</u>	<u>61.1</u>	55.8

Table 6: Object detection with Faster R-CNN fine-tuned on VOC trainval07+12. The detection AP scores (AP⁵⁰, AP⁷⁵, AP^{all}) are computed on test07. All models use ResNet-50 backbone (R50-C4) pre-trained with self-supervision on ImageNet. BoWNet scores are averaged over 3 trials. *: our implementation fine-tuned in the same conditions as BoWNet. †: BatchNorm layers are frozen and used as affine transformation layers.

VOC07 classification results. For this evaluation we use the publicly available code for benchmarking self-supervised methods provided by Goyal *et al.* [24]. [24] implements the guidelines of [50] and trains linear SVMs [6] on top of the frozen learned representations using the VOC07 train+val splits for training and the VOC07 test split for testing. We consider the features of the 3rd (layer conv4) and 4th (layer conv5) residual blocks and provide results in Table 4. Again, BoWNet improves the performance of the already strong RotNet by several points. Furthermore, BoWNet outperforms all prior methods. Interestingly, conv4-based BoW leads to better classification results for the conv5 layer of BoWNet, and conv5-based BoW leads to better classification results for the conv4 layer of BoWNet.

ImageNet and Places205 classification results. Here we evaluate on the 1000-way ImageNet and 205-way Places205 classification tasks using linear classifiers on frozen feature representations. To that end, we follow the guidelines of [24]: for the ImageNet (Places205) dataset we train linear classifiers using Nesterov SGD for 28 (14) training epochs and a learning rate of 0.01 that is multiplied by 0.1 after 10 (5) and 20 (10) epochs. The batch size is 256 and the weight decay is $1e-4$. We report results in Table 5. We observe that BoWNet outperforms all prior self-supervised methods by significant margin. Furthermore, the accuracy gap on Places205 between our ImageNet-trained BoWNet representations and the ImageNet-trained supervised representations is only 0.9 points in p0015. This demonstrates that *our self-supervised representations have almost the same generalization ability to the “unseen” (during training) Places205 classes as the supervised ones.* We also compare against the MoCo [25] and PIRL [44] methods that were recently uploaded on arXiv and essentially are concurrent work. BoWNet outperforms MoCo on ImageNet. When compared to PIRL, BoWNet has around 1 point higher Places205 accuracy but 1 point lower ImageNet accuracy.

VOC detection results. Here we evaluate the utility of our self-supervised method on a more complex downstream task: object detection. We follow the setup considered in prior works [24, 25, 44]: Faster R-CNN [57] with a ResNet50 backbone [26] (R50-C4 in Detectron2 [69]). We fine-tune the pre-trained BoWNet on trainval07+12 and evaluate on test07. We use the same training schedule as [24, 44] adapted for 8 GPUs and freeze the first two convolutional blocks. In detail, we use mini-batches of 2 images per GPU and fine-tune for 25K steps with the learning rate dropped by 0.1 after 17K steps. We set the base learning to 0.02 with a linear warmup [23] of 1,000 steps. We fine-tune BatchNorm layers [30] (synchronizing across GPUs) and use BatchNorm on newly added layers specific to this task.⁵

We compare BoWNet_{conv4} and BoWNet_{conv5} against both classic and recent self-supervised methods and report results in Table 6. Both BoWNet variants exhibit strong performance. Differently from previous benchmarks, the conv5 is clearly better than conv4 on all metrics. This might be due to the fact that here we fine-tune multiple layers and depth plays a more significant role. *Interestingly, BoWNet outperforms the supervised ImageNet pre-trained model, which is fine-tuned in the same conditions as BoWNet. So, our self-supervised representations generalize better to the VOC detection task than the supervised ones.* This result is in line with concurrent works [25, 44] and underpins the utility of such methods in efficiently squeezing out information from data without using labels.

5. Conclusion

In this work we propose BoWNet, a novel method for representation learning that employs spatial dense descriptions based on visual words as targets for self-supervised training. The labels for training BoWNet are provided by a standard self-supervised model. The reconstruction of the BoW vectors from perturbed images along with the discretization of the output space into visual words, enable a more discriminative learning of the local visual patterns in an image. Interestingly, although BoWNet is trained over features learned without label supervision, not only it achieves strong performances, but it also manages to outperform the initial model. This finding along with the discretization of the feature space (into visual words) open additional perspectives and bridges to NLP self-supervised methods that have greatly benefited from this type of approaches in the past few years.

Acknowledgements. We would like to thank Gabriel de Marmiesse for his invaluable support during the experimental implementation and analysis of this work.

⁵He *et al.* [25] point out that features produced by self-supervised training can display different distributions compared to supervised ones and suggest using feature normalization to alleviate this problem.

References

- [1] Jean-Baptiste Alayrac, Joao Carreira, and Andrew Zisserman. The visual centrifuge: Model-free layered video representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2457–2466, 2019. 5
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pfister, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016. 4
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017. 5
- [4] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning, 2019. 5
- [5] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *arXiv preprint arXiv:1906.00910*, 2019. 5, 6
- [6] Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152. ACM, 1992. 8
- [7] Wieland Brendel and Matthias Bethge. Approximating cnns with bag-of-local-features models works surprisingly well on imagenet. In *ICLR*, 2019. 2
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018. 5
- [9] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2959–2968, 2019. 5, 7
- [10] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised generative adversarial networks. *arXiv preprint arXiv:1811.11212*, 2018. 1
- [11] Gabriella Csurka, Christopher Dance, Lixin Fan, Jutta Willamowski, and Cédric Bray. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, volume 1, pages 1–2. Prague, 2004. 4
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1
- [13] Tom van Dijk and Guido de Croon. How do neural networks see depth in single images? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2183–2191, 2019. 5
- [14] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1422–1430, 2015. 1, 5, 6, 7
- [15] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1734–1747, 2015. 5
- [16] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 5
- [17] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. 5
- [18] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. *arXiv preprint arXiv:1906.05186*, 2019. 1, 6
- [19] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018. 6
- [20] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018. 1, 3, 5, 14
- [21] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 5
- [22] Yunchao Gong, Liwei Wang, Ruiqi Guo, and Svetlana Lazebnik. Multi-scale orderless pooling of deep convolutional activation features. In *European conference on computer vision*, pages 392–407. Springer, 2014. 4
- [23] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 8
- [24] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. *arXiv preprint arXiv:1905.01235*, 2019. 7, 8, 15
- [25] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019. 5, 7, 8, 15
- [26] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 8
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [28] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019. 1
- [29] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. *arXiv preprint arXiv:1906.12340*, 2019. 1

- [30] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. [8](#)
- [31] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. On the burstiness of visual elements. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1169–1176. IEEE, 2009. [12](#)
- [32] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. Packing bag-of-features. In *2009 IEEE 12th International Conference on Computer Vision*, pages 2357–2364. IEEE, 2009. [3](#), [14](#)
- [33] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *CVPR 2010-23rd IEEE Conference on Computer Vision & Pattern Recognition*, pages 3304–3311. IEEE Computer Society, 2010. [4](#)
- [34] Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302, 2015. [5](#)
- [35] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. Revisiting self-supervised visual representation learning. *arXiv preprint arXiv:1901.09005*, 2019. [3](#)
- [36] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016. [1](#), [5](#)
- [37] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2169–2178. IEEE, 2006. [6](#)
- [38] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [1](#)
- [39] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017. [5](#)
- [40] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. [1](#)
- [41] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. [4](#)
- [42] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. [12](#)
- [43] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. [5](#)
- [44] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. *arXiv preprint arXiv:1912.01991*, 2019. [5](#), [7](#), [8](#), [15](#)
- [45] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. [5](#)
- [46] Eva Moledano, Kevin McGuinness, Noel E O’Connor, Amaia Salvador, Ferran Marques, and Xavier Giro-i Nieto. Bags of local convolutional features for scalable instance search. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, pages 327–331. ACM, 2016. [4](#)
- [47] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. [1](#), [5](#)
- [48] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#)
- [49] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *arXiv preprint arXiv:1711.00937*, 2017. [5](#)
- [50] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *European conference on computer vision*, pages 801–816. Springer, 2016. [8](#)
- [51] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016. [1](#), [5](#)
- [52] Florent Perronnin and Christopher Dance. Fisher kernels on visual vocabularies for image categorization. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2007. [4](#)
- [53] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. Improving the fisher kernel for large-scale image classification. In *European conference on computer vision*, pages 143–156. Springer, 2010. [12](#)
- [54] Sudeep Pillai, Rareş Ambruş, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9250–9256. IEEE, 2019. [5](#)
- [55] Hang Qi, Matthew Brown, and David G Lowe. Learning with imprinted weights. *arXiv preprint arXiv:1712.07136*, 2017. [6](#)
- [56] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, pages 14837–14847, 2019. [5](#)
- [57] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. [8](#)
- [58] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. [5](#)
- [59] Josef Sivic and Andrew Zisserman. Video google: Efficient visual search of videos. In *Toward category-level object recognition*, pages 127–144. Springer, 2006. [3](#), [4](#), [14](#)
- [60] Jake Snell, Kevin Swersky, and Richard S Zemel. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175*, 2017. [6](#)
- [61] Jong-Chyi Su, Subhansu Maji, and Bharath Hariharan. When does self-supervision improve few-shot learning? *arXiv preprint arXiv:1910.03560*, 2019. [1](#)

- [62] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 7, 15
- [63] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1401–1408, 2013. 4
- [64] Trieu H Trinh, Minh-Thang Luong, and Quoc V Le. Selfie: Self-supervised pretraining for image embedding. *arXiv preprint arXiv:1906.02940*, 2019. 2
- [65] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008. 1, 5
- [66] Oriol Vinyals, Charles Blundell, Tim Lillicrap, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pages 3630–3638, 2016. 5, 6
- [67] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by colorizing videos. In *ECCV*, 2018. 5
- [68] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018. 5
- [69] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2, 2019. 8
- [70] Jun Yang, Yu-Gang Jiang, Alexander G Hauptmann, and Chong-Wah Ngo. Evaluating bag-of-visual-words representations in scene classification. In *Proceedings of the international workshop on Workshop on multimedia information retrieval*, pages 197–206. ACM, 2007. 2
- [71] Joe Yue-Hei Ng, Fan Yang, and Larry S Davis. Exploiting local features from deep networks for image retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 53–61, 2015. 4
- [72] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. *arXiv preprint arXiv:1905.04899*, 2019. 4
- [73] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proc. British Machine Vision Conference*, 2016. 5, 6, 12
- [74] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S⁴I: Self-supervised semi-supervised learning. *arXiv preprint arXiv:1905.03670*, 2019. 1
- [75] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. 1, 5
- [76] Richard Zhang, Phillip Isola, and Alexei A Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1058–1067, 2017. 1, 5
- [77] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems* 27, pages 487–495. Curran Associates, Inc., 2014. 5
- [78] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 5
- [79] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6002–6012, 2019. 7, 15

Appendix




Figure 2: **t-SNE [42] scatter plot of the learnt self-supervised features on CIFAR-100.** Each data point in the t-SNE scatter plot corresponds to the self-supervised feature representation of an image from CIFAR-100 and is colored according to the class that this image belongs to. To reduce clutter, we visualize the features extracted from the images of 20 randomly selected classes of CIFAR-100.

A. Visualizations

A.1. Visualizing the word clusters

In Figure 3 we illustrate visual words used for training our self-supervised method on ImageNet. Since we discover visual words using k-means, to visualize a visual word we depict the 16 patches with the smallest Euclidean distance to the visual word cluster centroid. As can be noticed, visual words encode mid-to-higher-level visual concepts.

A.2. t-SNE scatter plots of the learnt self-supervised features

In Figure 2 we visualize the t-SNE [42] scatter plot of the self-supervised features obtained when applying our method to the CIFAR-100 dataset. For visualizations purposes, we only plot the features corresponding to images that belong to 20 (randomly selected) classes out of the 100 classes of CIFAR-100. As can be clearly observed, the learnt features form class-specific clusters, which indicates that they capture semantic image information.

B. Additional experimental analysis

B.1. CIFAR-100 results

Here we provide an additional ablation analysis of our method on the CIFAR-100 dataset. As in section §4.1, we use the WRN-28-10 [73] architecture.

Models	Linear
BoWNet $K = 512$	69.76
BoWNet $K = 1024$	70.43
BoWNet $K = 2048$	71.01
BoWNet $K = 4096$	70.99

Table 7: CIFAR-100 linear classifier results with WRN-28-10. Impact of vocabulary size. Here we used an initial version of our method implemented with less aggressive augmentation techniques.

Method	$n = 1$	5	10	50	Linear
RotNet	58.3	74.8	78.3	81.9	60.3
BoWNet	69.1	86.3	89.2	92.4	71.5
Additional ablations					
BoWNet - predict $y(\tilde{\mathbf{x}})$	61.7	80.0	83.4	87.4	61.3
BoWNet - linear $\Omega(\cdot)$	70.4	85.6	88.3	90.7	63.3
BoWNet - binary	70.1	86.8	89.5	92.7	71.4

Table 8: CIFAR-100 linear classifier and few-shot results. For these results we use the WRN-28-10.

Impact of vocabulary size. In Table 7 we report linear classification results for different vocabulary sizes K . We see that increasing the vocabulary size from $K = 256$ to $K = 2048$ in CIFAR-100, offers obvious performance improvements. Then, from $K = 2048$ to $K = 4096$, there is no additional improvement.

Reparametrized linear layer for $\Omega(\cdot)$. In section §2.3, we described the linear-plus-softmax prediction layer $\Omega(\cdot)$ implemented with a reparametrized version of the standard linear layer. In this reparametrized version, instead of directly applying the weight vectors $W = [\mathbf{w}_1, \dots, \mathbf{w}_K]$ to a feature vector $\Phi(\tilde{\mathbf{x}})$, we first L_2 -normalize the weight vectors, and then apply a unique learnable magnitude γ for all the weight vectors (see equation (5) of main paper). The goal is to avoid always favoring the most frequently occurring words in the dataset by letting the linear layer of $\Omega(\cdot)$ learn a different magnitude for the weight vector of each word (which is what happens in the case of the standard linear layer). In terms of effect over the weight vector, this reparametrization is similar with the power-law normalization [53] used for mitigating the burstiness effect on BoW-like representations [31]. Here, we examine the impact of the chosen reparametrization by providing in Table 8 results for the case of implementing $\Omega(\cdot)$ with a standard linear layer (entry BoWNet - linear). We see that with the standard linear layer the performance of the BoWNet model deteriorates, especially on the linear classification metric, which validates our design of the $\Omega(\cdot)$ layer.

Predicting $y(\tilde{\mathbf{x}})$ instead of $y(\mathbf{x})$. In our work, given a perturbed image $\tilde{\mathbf{x}}$, we train a convnet to predict the BoW representation $y(\mathbf{x})$ of the original image \mathbf{x} . The purpose of predicting the BoW of the original image \mathbf{x} instead of the perturbed one $\tilde{\mathbf{x}}$, is to force the convnet to learn perturbation-invariant and context-aware features. We examine the impact of this choice by providing in Table 8 results for when the convnet is trained to predict the BoW of




Figure 3: **Examples of visual word cluster members.** The clusters are created by applying k-means on the feature maps produced from the conv4 layer of the ResNet50. For each visual word cluster we depict the 16 patch members with the smallest Euclidean distance to the visual word cluster centroid.

BoW from	linear cls	Vocabulary	linear cls
conv3	42.08	$K = 2048$	45.38
conv4	45.38	$K = 8192$	46.03
conv5	40.37	$K = 20000$	46.45

Vocabulary size $K = 2048$ || BoW from conv4

Table 9: ResNet18 small-scale experiments on ImageNet. Linear classifier results. The accuracy of the RotNet model used for building the BoW representations is 37.61. The left section explores the impact of the layer (of RotNet) that we use for building the BoW representations (with $K = 2048$). The right section explores the impact of the vocabulary size K (with BoW from the conv4).

the perturbed image \tilde{x} (entry BoWNet - predict $y(\tilde{x})$). As expected, in this case there is a significant drop in the BoWNet performance.

Histogram BoW vs Binary BoW In section §2.2 we describe two ways for reducing the visual word description of an image to a BoW representation. Those are, (1) to count the number of times each word appears in the image (see equation (3)), called Histogram BoW, and (2) to just indicate for each word whether it appears in the image (see equation (4)) [59, 32], called Binary BoW. In Table 8 we provide evaluation results with both the histogram version (entry BoWNet) and the binary version (entry BoWNet - binary). We see that they achieve very similar linear classification and few-shot performance.

B.2. Small-scale experiments on ImageNet

Here we provide an additional experimental analysis of our method on the ImageNet dataset. Specifically, we study the impact of the feature block of the base convnet and the vocabulary size that are used for building the BoW representation. Due to the computationally intensive nature of ImageNet, we analyze those aspects of our method by performing “small-scale” ImageNet experiments. By “small-scale” we mean that we use the light-weight ResNet18 architecture and we train using only 20% of ImageNet training images, and for few epochs.

Implementation details. We train the self-supervised models with SGD for 48 epochs. The learning rate is initialized at 0.1 and dropped by a factor of 10 after 15, 30, and 45 epochs. The batch size is 200 and weight decay $5e - 4$.

Evaluation protocols. We evaluate the learned self-supervised representations by freezing them and then training on top of them 1000-way linear classifiers for the ImageNet classification task. The linear classifier is applied on top the feature map of the last residual block of ResNet18, resized to $512 \times 4 \times 4$ with adaptive average pooling. It is trained with SGD for 60 epochs using a learning rate of 0.1 that is dropped by a factor of 10 every 20 epochs. The weight decay is $1e - 3$.

Results. We report results in Table 9. First, we study the impact on the quality of the learned representations of the RotNet feature block that is used for building the BoW representation. In the left section of Table 9 we report results for the cases of (a) conv3 (2nd residual block), (b) conv4 (3rd residual block), and conv5 (4th residual block). We see that the best performance is for the conv4-based BoW. Furthermore, in the right section of Table 9 we examine

the impact of the vocabulary size K on the quality of the learned representations. We see that increasing the vocabulary size from $K = 2048$ to $K = 20000$ leads to significant improvement for the linear classifier. In contrast, in Table 7 with results on CIFAR-100, we saw that increasing the vocabulary size after $K = 2048$ does not improve the quality of the learned representations. Therefore, it seems that the optimal vocabulary size depends on the complexity of the dataset to which we apply the BoW prediction task.

B.3. Full ImageNet and Places205 classification results

In Table 10 we provide the full experimental results of our method on the ImageNet and Places205 classification datasets.

C. Implementation details

C.1. Implementing RotNet

For the implementation of the rotation prediction network, RotNet, we follow the description and settings from Gidaris *et al.* [20]. RotNet is composed of a feature extractor $\hat{\Phi}(\cdot)$ and a rotation prediction module. The rotation prediction module gets as input the output feature maps of $\hat{\Phi}(\cdot)$ and is implemented as a convnet. It consists of a block of residual layers followed by global average pooling and a fully connected classification layer. In the CIFAR-100 experiments where $\hat{\Phi}(\cdot)$ is implemented with a WRN-28-10 architecture, the residual block of the rotation prediction module has 4 residual layers (similar to the last residual block of WRN-28-10), with 640 feature channels as input and output. In the MiniImageNet experiments where $\hat{\Phi}(\cdot)$ is implemented with a WRN-28-4 architecture, the residual block of the rotation prediction module has again 4 residual layers, but with 256 feature channels as input and output. Finally, in ImageNet experiments with ResNet50, the residual block of the rotation prediction module has 1 residual layer with 2048 feature channels as input and output.

During training for each image of a mini-batch, we generate its four rotated copies (0° , 90° , 180° , and 270° 2D rotations) and predict the rotation class of each copy. For supervision we use the cross-entropy loss over the four rotation classes.

After training we discard the rotation prediction module and consider only the feature extractor $\hat{\Phi}(\cdot)$ for the next stages, *i.e.* spatially dense descriptions and visual words.

C.2. Building BoW for self-supervised training

For the ImageNet experiments, given an image, we build its target BoW representation using visual words extracted from both the original and the horizontally flipped version of the image. Also, for faster training we pre-cache the BoW representations. Finally, in all experiments, when computing the target BoW representation we ignore the visual words that correspond to the feature vectors on the edge of the feature maps.

C.3. Few-shot protocol

The typical pipeline in few-shot learning is to first train a model on a set of *base* classes and then to evaluate it on a different set of *novel* classes (each set of classes is split into train and validation subsets). For MiniImageNet experiments we use this protocol, as this dataset has three different splits of classes: 64 training classes,

Method	ImageNet					Places205				
	conv2	conv3	conv4	conv5	pool5	conv2	conv3	conv4	conv5	pool5
Random [24]	13.7	12.0	8.0	5.6	-	16.6	15.5	11.6	9.0	-
Supervised methods										
ImageNet [24]	33.3	48.7	67.9	75.5	-	32.6	42.1	50.8	51.5	-
ImageNet*	32.8	47.0	67.2	76.0	76.2	35.2	42.6	50.9	52.8	52.0
Places205 [24]	31.7	46.0	58.2	51.7	-	32.3	43.2	54.7	62.3	-
Prior self-supervised methods										
RotNet*	30.1	42.0	52.5	46.2	40.6	32.9	40.1	45.0	42.0	39.4
Jigsaw [24]	28.0	39.9	45.7	34.2	-	28.8	36.8	41.2	34.4	-
Colorization [24]	24.1	31.4	39.6	35.2	-	28.4	30.2	31.3	30.4	-
LA [†] [79]	23.3	39.3	49.0	60.2	-	26.4	39.9	47.2	50.1	-
Concurrent work										
MoCo [25]	-	-	-	-	<u>60.6</u>	-	-	-	-	-
PIRL [44]	30.0	40.1	56.6	63.6	-	29.0	35.8	45.3	<u>49.8</u>	-
CMC [‡] [62]	-	-	-	-	64.1	-	-	-	-	-
BowNet conv4	34.4	<u>48.7</u>	<u>60.0</u>	<u>62.5</u>	62.1	36.7	44.7	50.5	50.9	51.1
BowNet conv5	<u>34.2</u>	49.1	60.5	60.4	60.2	36.9	44.7	<u>50.1</u>	49.6	<u>49.5</u>

Table 10: ResNet-50 top-1 center-crop linear classification accuracy on ImageNet and Places205. For the conv2–conv5 layers of ResNet-50, to evaluate linear classifiers, we resize their feature maps to around $9k$ dimensions (in the same way as in [24]). pool5 indicates the accuracy of the linear classifier trained on the 2048-dimensional feature vectors produced by the global average pooling layer after conv5. [†]: LA [79] uses 10 crops for evaluation. [‡]: CMC [62] uses two ResNet-50 feature extractor networks instead of just one. *: our implementation.

16 for validation, and 20 for test. For the few-shot experiments on CIFAR-100 we do not have such splits of classes and we adjust this protocol by selecting a subset of 20 classes and sample from the corresponding test images for evaluation. In this case, the feature extractor $\Phi(\cdot)$ is trained in a self-supervised manner on train images from all 100 classes of CIFAR-100.

Few-shot models are evaluated over a large number of few-shot tasks: we consider here 2,000 tasks. The few-shot evaluation

tasks are formed by first sampling t categories from the set of novel/evaluation classes and then selecting randomly n training samples and m test samples per category. The classification performance is measured on the $t \times m$ test images and is averaged over all sampled few-shot tasks. For few-shot experiments we use $t = 5$, $m = 15$, $n \in \{1, 5, 10, 50\}$.