

OBoW: Online Bag-of-Visual-Words Generation for Self-Supervised Learning

Spyros Gidaris¹, Andrei Bursuc¹, Gilles Puy¹, Nikos Komodakis², Matthieu Cord^{1,3}, Patrick Pérez¹
¹valeo.ai ²University of Crete ³Sorbonne Université

Abstract

Learning image representations without human supervision is an important and active research field. Several recent approaches have successfully leveraged the idea of making such a representation invariant under different types of perturbations, especially via contrastive-based instance discrimination training. Although effective visual representations should indeed exhibit such invariances, there are other important characteristics, such as encoding contextual reasoning skills, for which alternative reconstruction-based approaches might be better suited.

With this in mind, we propose a teacher-student scheme to learn representations by training a convolutional net to reconstruct a bag-of-visual-words (BoW) representation of an image, given as input a perturbed version of that same image. Our strategy performs an online training of both the teacher network (whose role is to generate the BoW targets) and the student network (whose role is to learn representations), along with an online update of the visual-words vocabulary (used for the BoW targets). This idea effectively enables fully online BoW-guided unsupervised learning. Extensive experiments demonstrate the interest of our BoW-based strategy, which surpasses previous state-of-the-art methods (including contrastive-based ones) in several applications. For instance, in downstream tasks such Pascal object detection, Pascal classification and Places205 classification, our method improves over all prior unsupervised approaches, thus establishing new state-of-the-art results that are also significantly better even than those of supervised pre-training. We provide the implementation code at <https://github.com/valeoai/obow>.

1. Introduction

Learning unsupervised image representations based on convolutional neural nets (convnets) has attracted a significant amount of attention recently. Many different types of convnet-based methods have been proposed in this regard, including methods that rely on using annotation-free pretext tasks [15, 27, 43, 48, 52, 76], generative methods that model the image data distribution [16, 17, 20], as well as

clustering-based approaches [4, 6, 7].

Several recent methods opt to learn representations via instance-discrimination training [9, 19, 33, 71], typically implemented in a contrastive learning framework [12, 32]. The primary focus here is to learn low-dimensional image / instance embeddings that are invariant to intra-image variations while being discriminative among different images. Although these methods manage to achieve impressive results, they focus less on other important aspects in representation learning, such as contextual reasoning, for which alternative reconstruction-based approaches [14, 52, 76, 77] might be better suited. For instance, the task of predicting from an image region the contents of the entire image requires to recognize the visual concepts depicted in the provided region and then to infer from them the structure of the entire scene. So, training for such a task has the potential of squeezing out more information from the training images and of learning richer and more powerful representations.

However, reconstructing image pixels is an ambiguous and hard-to-optimize task that forces the convnet to spend a lot of capacity on modeling low-level pixel details. It is all the more unnecessary when the final goal deals with high-level image understanding, such as image classification. To leverage the advantages of the reconstruction principle without focusing on unimportant pixel details, one can focus on reconstruction over high-level visual concepts, referred to as visual words. For instance, BoWNet [25] derived a teacher-student learning scheme following this principle. In this teacher-student setting, given an image, the teacher extracts feature maps that are then quantized in a spatially-dense way over a vocabulary of visual words. Then, the resulting visual words-based image description is exploited for training the student on the self-supervised task of reconstructing the distribution of the visual words of an image, i.e., its bag-of-words (BoW) representation [72], given as input a perturbed version of that same image. By solving this reconstruction task the student is forced to learn perturbation-invariant and context-aware representations while “ignoring” pixel details.

Despite its advantages, the BoWNet approach exhibits some important limitations that do not allow it to fully exploit the potential of the BoW-based reconstruction task. One of them is that it relies on the availability of an already pre-

trained teacher network. More importantly, it assumes that this teacher remains static throughout training. However, due to the fact that during the training process the quality of the student’s representations will surpass the teacher’s ones, a static teacher is prone to offer a suboptimal supervisory signal to the student and to lead to an inefficient usage of the computational budget for training.

In this paper, we propose a BoW-based self-supervised approach that overcomes the aforementioned limitations. To that end, our main technical contributions are three-fold:

1. We design a novel *fully online* teacher-student learning scheme for BoW-based self-supervised training (Fig. 1). This is achieved by online training of the teacher and the student networks.
2. We significantly revisit key elements of the BoW-guided reconstruction task. This includes the proposal of a *dynamic* BoW prediction module used for reconstructing the BoW representation of the student image. This module is carefully combined with adequate online updates of the visual-words vocabulary used for the BoW targets.
3. We enforce the learning of powerful contextual reasoning skills in our image representations. Revisiting data augmentation with aggressive cropping and spatial image perturbations, and exploiting *multi-scale* BoW reconstruction targets, we equip our student network with a powerful feature representation.

Overall, the proposed method leads to a simpler and much more efficient training methodology for the BoW-guided reconstruction task that manages to learn significantly better image representations and therefore achieves or even surpasses the state of the art in several unsupervised learning benchmarks. We call our method OBoW after the *online BoW* generation mechanism that it uses.

2. Related Work

Bags of visual words. *Bag-of-visual-words* representations are powerful image models able to encode image statistics from hundreds of local features. Thanks to that, they were used extensively in the past [13, 39, 53, 59, 64] and continue to be a key ingredient of several recent deep learning approaches [2, 25, 28, 37]. Among them, BoWNet [25] was the first work to use BoWs as reconstruction targets for unsupervised representation learning. Inspired by it, we propose a novel BoW-based self-supervised method with a more simple and effective training methodology that includes generating BoW targets in a fully online manner and further enforcing the learning of context-aware representations.

Self-supervised learning. A prominent paradigm for unsupervised representation learning is to train a convnet on an artificially designed annotation-free pretext task, e.g., [3, 10,

15, 27, 44, 47, 48, 68, 70, 75, 79]. Many works rely on pretext reconstruction tasks [1, 29, 43, 52, 54, 67, 76, 77, 79], where the reconstruction target is defined at image pixel level. This is in stark contrast with our method, which uses a reconstruction task defined over high-level visual concepts (i.e., visual words) that are learnt with a teacher-student scheme in a fully online manner.

Instance discrimination and contrastive objectives. Recently, unsupervised methods based on contrastive learning objectives [8, 9, 22, 23, 33, 35, 38, 46, 49, 62, 69, 71] have shown great results. Among them, contrastive-based instance discrimination training [9, 18, 33, 40, 71] is the most prominent example. In this case, a convnet is trained to learn image representations that are invariant to several perturbations and at the same time discriminative among different images (instances). Our method also learns intra-image invariant representations, since the convnet must predict the same BoW target (computed from the original image) regardless of the applied perturbation. More than that however, our work also places emphasis on learning context-aware representations, which, we believe, is another important characteristic that effective representations should have. In that respect, it is closer to contrastive-based approaches that target to learn context-aware representations by predicting (in a contrastive way) the state of missing image patches [35, 49, 65].

Teacher-student approaches. This paradigm has a long research history [24, 58] and it is frequently used for distilling a single large network or an ensemble, the *teacher*, into a smaller network, the *student* [5, 36, 41, 51]. This setting has been revisited in the context of semi-supervised learning where the teacher is no longer fixed but evolves during training [42, 61, 66]. In self-supervised learning, BoWNet [25] trains a student to match the BoW representations produced by a self-supervised pre-trained teacher. MoCo [33] relies on a slow-moving momentum-updated teacher to generate up-to-date representations to fill a memory bank of negative images. BYOL [31], which is a method concurrent to our work, also uses a momentum-updated teacher and trains the student to predict features generated by the teacher. However, BYOL, similar to contrastive-based instance discrimination methods, uses low-dimensional global image embeddings as targets (produced from the final convnet output) and primarily focuses on making them intra-image invariant. On the contrary, our training targets are produced by converting the intermediate teacher feature maps to high-dimensional BoW vectors that capture multiple *local visual concepts*, thus constituting a richer target representation. Moreover, they are built over an online-updated vocabulary from randomly sampled local features, expressing the current image as statistics over this vocabulary (see § 3.1). Therefore, our BoW targets expose fewer learning “shortcuts” (a critical aspect in self-supervised learning [3, 15, 48, 70]), thus preventing to a larger extent teacher-student collapse and overfitting.

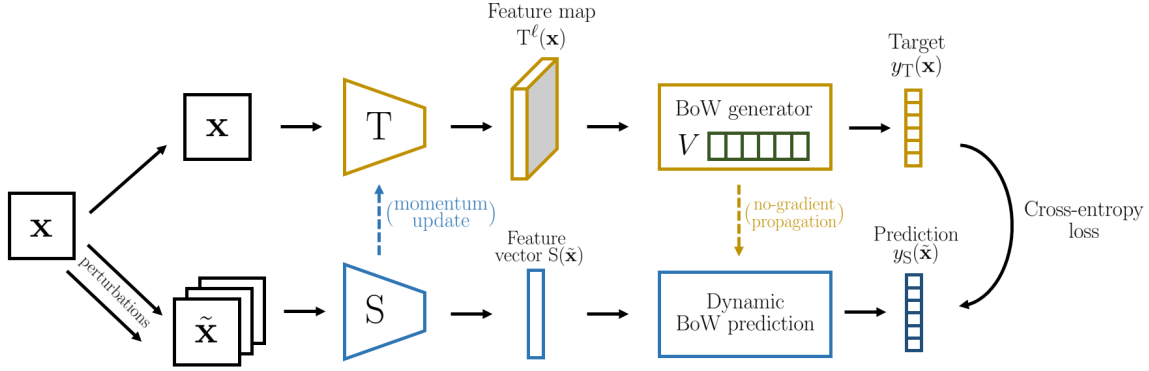


Figure 1: Unsupervised learning with Bag-of-Words guidance. Two encoders T and S learn at different tempos by interacting and learning from each other. An image \mathbf{x} is passed through the encoder T and its output feature maps $T^\ell(\mathbf{x})$ are embedded into a BoW representation $y_T(\mathbf{x})$ over a vocabulary V of features from T. The vocabulary V is updated at each step. The encoder S aims to reconstruct $y_T(\mathbf{x})$ from data-augmented instances $\tilde{\mathbf{x}}$. A dynamic BoW-prediction head learns to leverage the continuously updated vocabulary V to compute the BoW representation from the features $S(\tilde{\mathbf{x}})$. T follows slowly the learning trajectory of S via momentum updates.

Relation to SwAV [8]. OBoW presents some similarity (e.g., using online vocabularies) with SwAV [8]. However, the prediction tasks fundamentally differ: *OBoW exploits a BoW prediction task while SwAV uses an image-cluster prediction task* [4, 6, 7]. BoW targets are much richer representations than image-cluster assignments: a BoW encodes all the local-feature statistics of an image whereas an image-cluster assignment encodes only one global image feature.

3. Our approach

Here we explain our proposed approach for learning image representations by reconstructing bags of visual words. We start with an overview of our method.

Overview. The bag-of-words reconstruction task involves a student convnet $S(\cdot)$ that learns image representations, and a teacher convnet $T(\cdot)$ that generates BoW targets used for training the student network. The student $S(\cdot)$ is parameterized by θ_S and the teacher $T(\cdot)$ by θ_T .

To generate a BoW representation $y_T(\mathbf{x})$ out of an image \mathbf{x} , the teacher first extracts the feature map $T^\ell(\mathbf{x}) \in \mathbb{R}^{c_\ell \times h_\ell \times w_\ell}$, of spatial size $h_\ell \times w_\ell$ with c_ℓ channels, from its ℓ^{th} layer (in our experiments ℓ is either the last L or penultimate $L - 1$ convolutional layer of $T(\cdot)$). It quantizes the c_ℓ -dimensional feature vectors $T^\ell(\mathbf{x})[u]$ at each location $u \in \{1, \dots, h_\ell \times w_\ell\}$ of the feature map over a vocabulary $V = [\mathbf{v}_1, \dots, \mathbf{v}_K]$ of K visual words of dimension c_ℓ . This quantization process produces for each location u a K -dimensional code vector $q(\mathbf{x})[u]$ that encodes the assignment of $T(\mathbf{x})[u]$ to its closest (in terms of squared Euclidean distance) visual word(s). Then, the teacher reduces the quantized feature maps $q(\mathbf{x})$ to a K -dimensional BoW, $\tilde{y}_T(\mathbf{x})$, by channel-wise max-pooling, i.e., $\tilde{y}_T(\mathbf{x})[k] = \max_u q(\mathbf{x})[u][k]$ (alternatively, the reduction can be performed with average pooling), where $q(\mathbf{x})[u][k]$

is the assignment value of the code $q(\mathbf{x})[u]$ for the k^{th} word. Finally, $\tilde{y}_T(\mathbf{x})$ is converted into a probability distribution over the visual words by L_1 -normalization, i.e., $y_T(\mathbf{x})[k] = \frac{\tilde{y}_T(\mathbf{x})[k]}{\sum_{k'} \tilde{y}_T(\mathbf{x})[k']}$.

To learn image representations, the student gets as input a perturbed version of the image \mathbf{x} , denoted as $\tilde{\mathbf{x}}$, and is trained to reconstruct the BoW representation $y_T(\mathbf{x})$, produced by the teacher, of the original unperturbed image \mathbf{x} . To that end, it first extracts a global vector representation $S(\tilde{\mathbf{x}}) \in \mathbb{R}^c$ (with c channels) from the entire image $\tilde{\mathbf{x}}$ and then applies a linear-plus-softmax layer to $S(\tilde{\mathbf{x}})$, as follows:

$$y_S(\tilde{\mathbf{x}})[k] = \frac{\exp(\mathbf{w}_k^\top S(\tilde{\mathbf{x}}))}{\sum_{k'} \exp(\mathbf{w}_{k'}^\top S(\tilde{\mathbf{x}}))}, \quad (1)$$

where $W = [\mathbf{w}_1, \dots, \mathbf{w}_K]$ are the c -dimensional weight vectors (one per word) of the linear layer. The K -dimensional vector $y_S(\tilde{\mathbf{x}})$ is the predicted softmax probability of the target $y_T(\mathbf{x})$. Hence, the training loss that is minimized for a single image \mathbf{x} is the cross-entropy loss

$$\text{CE}(y_S(\tilde{\mathbf{x}}), y_T(\mathbf{x})) = - \sum_{k=1}^K y_T(\mathbf{x})[k] \log(y_S(\tilde{\mathbf{x}})[k]) \quad (2)$$

between the softmax distribution $y_S(\tilde{\mathbf{x}})$ predicted by the student from the perturbed image $\tilde{\mathbf{x}}$, and the BoW distribution $y_T(\mathbf{x})$ of the unperturbed image \mathbf{x} given by the teacher.

Our technical contributions. In the following, we explain (i) in § 3.1, how to construct a fully online training methodology for the teacher, the student and the visual-words vocabulary, (ii) in § 3.2, how to implement a dynamic approach for the BoW prediction that can adapt to continuously-changing vocabularies of visual words, and finally (iii) in § 3.3, how to significantly enhance the learning of contextual reasoning skills by utilizing multi-scale BoW reconstruction targets and by revisiting the image augmentation schemes.

3.1. Fully online BoW-based learning

To make the BoW targets encode more high-level features, BoWNet pre-trains the teacher convnet $T(\cdot)$ with another unsupervised method, such as RotNet [27], and computes the vocabulary V for quantizing the teacher feature maps off-line by applying k-means on a set of teacher feature maps extracted from training images. After the end of the student training, during which the teacher’s parameters remain frozen, the student becomes the new teacher $T(\cdot) \leftarrow S(\cdot)$, a new vocabulary V is learned off-line from the new teacher, and a new student is trained, starting a new training cycle. In this case however, (a) the final success depends on the quality of the first pre-trained teacher, (b) the teacher and the BoW reconstruction targets $y_T(\mathbf{x})$ remain frozen for long periods of time, which, as already explained, results in a suboptimal training signal, and (c) multiple training cycles are required, making the overall training time consuming.

To address these important limitations, in this work we propose a fully online training methodology that allows the teacher to be continuously updated as the student training progresses, with no need for off-line k-means stages. This requires an online updating scheme for the teacher as well as for the vocabulary of visual words used for generating the BoW targets, both of which are detailed below.

Updating the teacher network. Inspired by MoCo [33], the parameters θ_T of the teacher convnet are an exponential moving average of the student parameters. Specifically, at each training iteration the parameters θ_T are updated as

$$\theta_T \leftarrow \alpha \cdot \theta_T + (1 - \alpha) \cdot \theta_S, \quad (3)$$

where $\alpha \in [0, 1]$ is a momentum coefficient. Note that, as a consequence, the teacher has to share exactly the same architecture as the student. With a proper tuning of α , e.g., $\alpha = 0.99$, this update rule allows slow and continuous updates of the teacher, avoiding rapid changes of its parameters, such as with $\alpha = 0$, which would make the training unstable. As in MoCo, for its batch-norm units, the teacher maintains different batch-norm statistics from the student.

Updating the visual-words vocabulary. Since the teacher is continuously updated, off-line learning of V is not a viable option. Instead, we explore two solutions for computing V , *online k-means* and a *queue-based vocabulary*.

Online k-means. One possible choice for updating the vocabulary is to apply online k-means clustering after each training step. Specifically, as proposed in VQ-VAE [50, 56], we use exponential moving average for vocabulary updates. A critical issue that arises in this case is that, as training progresses, the features distribution changes over time. The visual words computed by online k-means do not adapt to this distribution shift leading to extremely unbalanced cluster assignments and even to assignments that collapse to a single cluster. In order to counter this effect, we investigate

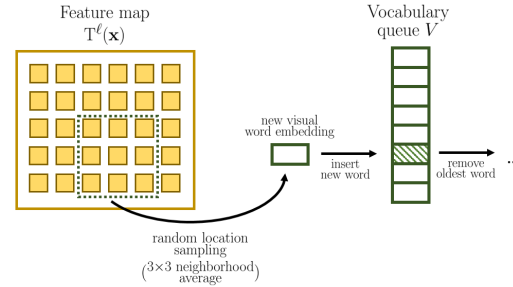


Figure 2: Vocabulary queue from randomly sampled local features. For each input image \mathbf{x} to T , “local” features are pooled from $T^l(\mathbf{x})$ by averaging over 3×3 sliding windows. One of the resulting vectors is selected randomly and added as visual word to the vocabulary queue, replacing the oldest word in the vocabulary.

different strategies: (a) detection of rarely used visual words over several mini-batches and replacement of these words with a randomly sampled feature vector from the current mini-batch; (b) enforcing uniform assignments to each cluster thanks to the Sinkhorn optimization as in, e.g., [4, 8]. For more details see § D in Supplementary Material.

A queue-based vocabulary. In this case, the vocabulary V of visual words is a K -sized queue of random features. At each step, after computing the assignment codes over the current vocabulary V , we update V by selecting one feature vector per image from the current mini-batch, inserting it to the queue, and removing the oldest item in the queue if its size exceeds K . Hence, the visual words in V are always feature vectors from past mini-batches. We explore three different ways to select these local feature vectors: (a) uniform random sampling of one feature vector in $T^l(\mathbf{x})$; (b) global average pooling of $T^l(\mathbf{x})$ (average feature vector of each image); (c) an intermediate approach between (a) and (b) which consists of a local average pooling with a 3×3 kernel (stride 1, padding 0) of the feature map $T^l(\mathbf{x})$ followed by a uniform random sampling of one of the resulting feature vectors (Fig. 2). Our intuition for option (c) is that, assuming that the local features in a 3×3 neighborhood belong to one common visual concept, then local averaging selects a more representative visual-word feature from this neighborhood than simply sampling at random one local feature (option (a)). Likewise, the global averaging option (b) produces a representative feature from an entire image, which however, might result in overly coarse visual word features.

The advantage of the queue-based solution over online k-means is that it is simpler to implement and it does not require any extra mechanism for avoiding unbalanced clusters, since at each step the queue is updated with new randomly sampled features. Indeed, in our experiments, the queue-based vocabulary with option (c) provided the best results.

Generating BoW targets with soft-assignment codes. For generating the BoW targets, we use soft-assignments instead of the hard-assignments used

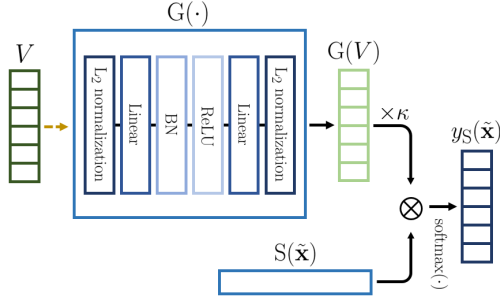


Figure 3: Dynamic BoW-prediction head. $G(\cdot)$ learns to quickly adapt to the visual words in the continuously refreshed vocabulary V . The outputs $G(V)$ are in fact weights that are used for mapping the features $S(\tilde{\mathbf{x}})$ to the corresponding BoW vector $y_S(\tilde{\mathbf{x}})$.

in BoWNet. This is preferable from an optimization perspective due to the fact that the vocabulary of visual words is continuously evolving. We thus compute the assignment codes $q(\mathbf{x})[u]$ as

$$q(\mathbf{x})[u][k] = \frac{\exp(-\frac{1}{\delta} \|\mathbf{T}^\ell(\mathbf{x})[u] - \mathbf{v}_k\|_2^2)}{\sum_{k'} \exp(-\frac{1}{\delta} \|\mathbf{T}^\ell(\mathbf{x})[u] - \mathbf{v}_{k'}\|_2^2)}. \quad (4)$$

The parameter δ is a temperature value that controls the softness of the assignment. We use $\delta = \delta_{\text{base}} \cdot \bar{\mu}_{\text{MSD}}$, where $\delta_{\text{base}} > 0$ and $\bar{\mu}_{\text{MSD}}$ is the exponential moving average (with momentum 0.99) of the mean squared distance of the feature vectors in $\mathbf{T}^\ell(\mathbf{x})$ from their closest visual words. The reason for using an adaptive temperature instead of a constant one is due to the change of magnitude of the feature activations during training, which induces a change of scale of the distances between the feature vectors and the words.

3.2. Dynamic bag-of-visual-word prediction

To learn effective image representations, the student must predict the BoW distribution over V of an image using as input a perturbed version of that same image. However, in our method the vocabulary is constantly updated and the visual words are changing or being replaced from one step to the next. Therefore, predicting the BoW distribution over a continuously updating vocabulary V with a fixed linear layer would make training unstable, if not impossible. To address this issue we propose to use a dynamic BoW-prediction head that can adapt to the evolving nature of the vocabulary. To that end, instead of using fixed weights as in (1), we employ a generation network $G(\cdot)$ that takes as input the current vocabulary of visual words $V = [\mathbf{v}_1, \dots, \mathbf{v}_K]$ and produces prediction weights for them as $G(V) = [G(\mathbf{v}_1), \dots, G(\mathbf{v}_K)]$, where $G(\cdot) : \mathbb{R}^{c_\ell} \rightarrow \mathbb{R}^c$ is parameterized by θ_G and $G(\mathbf{v}_k)$ represents the prediction weight vector for the k^{th} visual word. Therefore, Equation 1 becomes

$$y_S(\tilde{\mathbf{x}})[k] = \frac{\exp(\kappa \cdot G(\mathbf{v}_k)^\top S(\tilde{\mathbf{x}}))}{\sum_{k'} \exp(\kappa \cdot G(\mathbf{v}_{k'})^\top S(\tilde{\mathbf{x}}))}, \quad (5)$$

where κ is a fixed coefficient that equally scales the magnitudes of all the predicted weights $G(V) = [G(\mathbf{v}_1), \dots, G(\mathbf{v}_K)]$, which by design are L_2 -normalized. We implement $G(\cdot)$ with a 2-layer perceptron whose input and output vectors are L_2 -normalized (see Fig. 3). Its hidden layer has size $2 \times c$.

We highlight that dynamic weight-generation modules are extensively used in the context of few-shot learning for producing classification weight vectors of novel classes using as input a limited set of training examples [26, 30, 55]. The advantages of using $G(\cdot)$ instead of fixed weights, which BoWNet uses, are the equivariance to permutations of the visual words, the increased stability to the frequent and abrupt updates of the visual words, a number of parameters $|\theta_G|$ independent from the number of visual words K , hence requiring fewer parameters than a fixed-weights linear layer for large vocabularies.

3.3. Representation learning based on enhanced contextual reasoning

Data augmentation. The key factor for the success of many recent self-supervised representation learning methods [8, 9, 11, 31, 63] is to leverage several image augmentation/perturbation techniques, such as Gaussian blur [9], color jittering and random cropping techniques, as cutmix [73] that substitutes one random-size patch of an image with that of another. In our method, we want to fully exploit the possibility of building strong image representations by hiding local information. As the teacher is randomly initialised, it is important to hide large regions of the original image from the student so as to prevent the student from relying only on low-level image statistics for reconstructing the distributions $y_T(\mathbf{x})$ over the teacher visual words, which capture low-level visual cues at the beginning of the training. Therefore, we carefully design our image perturbations scheme to make sure that the student has access to only a very small portion of the original image. Specifically, similar to [8, 46], we extract from a training image multiple crops with two different mechanisms: one that outputs 160×160 -sized crops that cover less than 60% of the original image and one with 96×96 -sized crops that cover less than 14% of the original image (see Fig. 4). Given those image crops, the student must reconstruct the full bags of visual words from each of them independently. Therefore, our cropping strategy definitively forces the student network to understand and learn spatial dependency between visual parts.

Multi-scale BoW reconstruction targets. We also consider reconstructing BoW from multiple network layers that correspond to different scale levels. In particular, we experiment with using both the last scale level L (i.e., layer `conv5` in ResNet) and the penultimate scale level $L-1$ (i.e., layer `conv4` in ResNet). The reasoning behind this is that the features of level $L-1$ still encode semantically



Figure 4: Reconstructing BoWs from small parts of the original image. Given a training image (left), we extract two types of image crops. The first type (middle) is obtained by randomly sampling an image region whose area covers no more than 60% of the entire image, resizing it to 160×160 and then giving it as input to the student as part of the reconstruction task. The second type (right) is obtained by randomly selecting an area that covers between 60% to 100% of the entire image, resizing it to a 256×256 image, dividing it into 3×3 overlapping patches of size 96×96 , and randomly choosing 5 out of these 9 patches (indicated with red rectangles) that are given as 5 separate inputs to the student. The student must then reconstruct the original BoW target independently for each patch. The blue rectangle on the left image indicates the central 224×224 crop from which the teacher produces the BoW target. Note that, except from horizontal flipping, no other perturbation is applied on the teacher’s inputs.

important concepts but have a smaller receptive field than those in the last level. As a result, the visual words of level $L-1$ that belong to image regions hidden to the student are less likely to be influenced by pixels of image regions the student is given as input. Therefore, by using BoW from this extra feature level, the student is further enforced to learn contextual reasoning skills (and in fact, at a level with higher spatial details due to the higher resolution of level $L-1$), thus learning richer and more powerful representations. When using BoW extracted from two layers, our method includes a separate vocabulary for each layer, denoted by V_L and V_{L-1} for layers L and $L-1$ respectively, and two different weight generators, denoted by $G_L(\cdot)$ and $G_{L-1}(\cdot)$ for layers L and $L-1$, respectively. Regardless of what layer the BoW target comes from, the student uses a single global image representation $S(\tilde{x})$, typically coming from the global average pooling layer after the last convolutional layer (i.e., layer pool_{15} in ResNet), to perform the reconstruction task.

We show empirically in Section 4.1 that the contextual reasoning skills implicitly developed via using the above two schemes are decisive to learn effective image representations with the BoW-reconstruction task.

4. Experiments and results

We evaluate our method (OBoW) on the ImageNet [57], Places205 [78] and VOC07 [21] classification datasets as well as on the VOC07+12 detection dataset.

Implementation details. For our models, the vocabulary size is set to $K = 8192$ words and, as in BoWNet, when computing the BoW targets we ignore the visual words that correspond to the feature vectors on the edge of the teacher

Updating method	Few-shot		Linear
	1-shot	5-shot	
Online k-means			
(a) replacing rare clusters	40.98	60.35	44.45
(b) Sinkhorn-based balancing	37.20	55.22	39.74
Queue-based vocabulary			
(a) local features	40.29	60.81	44.39
(b) globally-averaged features	41.57	62.54	45.79
(c) locally-averaged features	42.11	62.44	45.86
Queue-based vocabulary – multi-scale BoW			
(b) globally-averaged features	41.29	63.09	49.40
(c) locally-averaged features	44.18	64.89	50.89

Table 1: Comparison of online vocabulary-update approaches. The results in the first two sections are with the vanilla version of our method and with the full version in the third section.

feature maps. The momentum coefficient α for the teacher updates is initialized at 0.99 and is annealed to 1.0 during training with a cosine schedule. The hyper-parameters κ and δ_{base} are set to 5 and 1/10 respectively for the results in § 4.1, to 8 and 1/15 respectively for the results in § 4.3. For more implementation details, see § C in Supplementary.

4.1. Analysis

Here we perform a detailed analysis of our method. Due to the computationally intensive nature of pre-training on ImageNet, we use a smaller but still representative version created by keeping only 20% of its images and we implement our model with the light-weight ResNet18 architecture. For training we use SGD for 80 epochs with cosine learning rate initialized at 0.05, batch size 128 and weight decay $5e-4$. We evaluate models trained with two versions of our method, the vanilla version that uses single-scale BoWs and from each training image extracts one 160×160 -sized crop (with which it trains the student), and the full version that uses multi-scale BoWs and extracts from each training image two 160×160 -sized crops plus five 96×96 -sized patches.

Evaluation protocols. After pre-training, we freeze the learned representations and use two evaluation protocols. **(1)** The first one consists in training 1000-way linear classifiers for the ImageNet classification task. **(2)** For the second protocol, our goal is to analyze the ability of the representations to learn with few training examples. To that end, we use 300 ImageNet classes and run with them multiple (200) episodes of 50-way classification tasks with 1 or 5 training examples per class and a Prototypical-Networks [60] classifier.

4.2. Results

Online vocabulary updates. In Tab. 1, we compare the approaches for online vocabulary updates described § 3.1. The

α	lr	Few-shot		Linear
		1-shot	1-shot	
0.99 \rightarrow 1	0.05	42.11	62.44	45.86
0.999	0.05	40.87	61.41	45.76
0.99	0.05	41.19	61.65	46.25
0.9	0.05	40.79	60.92	44.89
0.5	0.05	12.70	23.20	15.41
0.0	0.05	13.19	24.85	17.47
0.5	0.03	39.52	60.18	43.82
0.0	0.01	33.80	55.02	39.90

Table 2: Influence of the momentum coefficient α used for the teacher updates. For these results, we used the vanilla version. In the “0.99 \rightarrow 1” row, α is initialized to 0.99 and annealed to 1.0 with cosine schedule. The other entries use constant α values.

Soft	Dyn	Few-shot		Linear
		1-shot	5-shot	
✓	✓	42.11	62.44	45.86
	✓	38.61	59.98	44.64
✓		2.00	2.00	0.10

Table 3: Ablation of dynamic BoW prediction and soft-quantization. For these results, we used the vanilla version of our method. “Soft”: soft assignment instead of hard assignment. “Dyn”: dynamic weight generation instead of fixed weights.

queue-based solutions achieve in general better results than online k-means. Among the queue-based options, random sampling of locally averaged features, opt. (c), provides the best results. Its advantage over option (b) with global averaging is more evident with multi-scale BoWs where an extra feature level with a higher resolution and more localized features is used, in which case global averaging produces visual words that are too coarse. In all remaining experiments, we use a queue-based vocabulary with option (c).

Momentum for teacher updates. In Table 2, we study the sensitivity of our method w.r.t. the momentum α for the teacher updates (Equation 3). We notice a strong drop in performance when decreasing α from 0.9 to 0.5 (a rapidly-changing teacher), and to 0 (the teacher and student have identical parameters), while keeping the initial learning rate fixed ($lr = 0.05$). However, we noticed that this was not due to any cluster/mode collapse issue. The issue is that the teacher signal is more noisy at low α because of the rapid change of its parameters. This prevents the student to converge when keeping the learning rate as high as 0.05. We notice in Table 2 that a reduction of the learning rate to adapt to the reduction of α reduces the performance gap. This indicates that our method is not as sensitive to the choice of the momentum as MoCo and BYOL were shown to be.

Image crops	Multi-scale	Linear
1×224^2		31.39
$1 \times 224^2 + \text{cutmix}$		39.46
1×160^2		45.86
2×160^2		47.64
5×96^2		44.24
$2 \times 160^2 + 5 \times 96^2$		49.64
2×160^2	✓	49.00
$2 \times 160^2 + 5 \times 96^2$	✓	50.89

Table 4: Evaluation of image crop augmentations and of multi-scale BoWs. See text.

Method	EP	Few-shot		Linear
		$n = 1$	$n = 5$	
BoWNet	200	33.80	55.02	41.30
BoWNet (160 ² crops)	200	29.26	49.68	43.59
OBoW (vanilla)	80	42.11	62.44	45.86
OBoW (full)	80	44.18	64.89	50.89

Table 5: Comparison with BoW-like methods. “EP”: total number of epochs used for pre-training. Note that the BoWNet method consists of 40 epochs for teacher pre-training with the RotNet method followed by two BoWNet training rounds of 80 epochs.

Dynamic BoW prediction and soft quantization. In Table 3, we study the impact of the dynamic BoW prediction and of using soft assignment for the codes instead of hard assignment. We see that (1), as expected, the network is unable to learn useful features without the proposed dynamic BoW prediction, i.e., when using fixed weights; (2) soft assignment indeed provides a performance boost.

Enforcing context-aware representations. In Table 4 we study different types of image crops for the BoW reconstruction tasks, as well as the impact of multi-scale BoW targets. We observe that: (1) as we discussed in § 3.3, smaller crops that hide significant portions of the original image are better suited for our reconstruction task thus leading to dramatic increase in performance (compare entries 1×224^2 with the 1×160^2 and 5×96^2 entries). (2) Randomly sampling two 160×160 -sized crops (entries 2×160^2) and using 96×96 -sized patches leads to another significant increase in performance. (3) Finally, employing multi-scale BoWs improves the performance even further.

BoW-like comparison. In Table 5, we compare our method with the reference BoW-like method BoWNet. For a fair comparison, we implemented BoWNet both with its proposed augmentations, i.e., using one 224×224 -sized crop with cutmix (“BoWNet” row), and with the image augmentation we propose in the vanilla version of our method, i.e., one 160×160 -sized crop (“BoWNet (160² crops)”).

Method	Epochs	Batch	Linear Classification			VOC Detection			Semi-supervised learning	
			ImageNet	Places205	VOC07	AP ⁵⁰	AP ⁷⁵	AP ^{all}	1% Labels	10% Labels
Supervised	100	256	76.5	53.2	87.5	81.3	58.8	53.5	48.4	80.4
BoWNet [25]	325	256	62.1	51.1	79.3	81.3	61.1	55.8	-	-
PCL [45]	200	256	67.6	50.3	85.4	-	-	-	75.3	85.6
MoCo v2 [33]	200	256	67.5	-	-	82.4	63.6	57.0	-	-
SimCLR [9]	200	4096	66.8	-	-	-	-	-	-	-
SwAV [8]	200	256	72.7	56.2 [†]	87.2 [†]	81.8 [†]	60.0 [†]	54.4 [†]	76.7 [†]	88.7 [†]
BYOL [31]	300	4096	72.5	-	-	-	-	-	-	-
OBoW (Ours)	200	256	73.8	56.8	89.3	82.9	64.8	57.9	82.9	90.7
PIRL [46]	800	1024	63.6	49.8	81.1	80.7	59.7	54.0	57.2	83.8
MoCo v2 [33]	800	256	71.1	52.9	87.1	82.5	64.0	57.4	-	-
SimCLR [9]	1000	4096	69.3	53.3	86.4	-	-	-	75.5	87.8
BYOL [31]	1000	4096	74.3	-	-	-	-	-	78.4	89.0
SwAV [8]	800	4096	75.3	56.5	88.9	82.6	62.7	56.1	78.5	89.9

Table 6: Evaluation of ImageNet pre-trained ResNet50 models. The ‘‘Epochs’’ and ‘‘Batch’’ columns provide the number of pre-training epochs and the batch size of each model respectively. The first section includes models pre-trained with a similar number of epochs as our model (second section). We boldfaced the best results among all sections as well as of only the top two. For the linear classification tasks, we provide the top-1 accuracy. For object detection, we fine-tuned Faster R-CNN (R50-C4) on VOC `trainval07+12` and report detection AP scores by testing on `test07`. For semi-supervised learning, we fine-tune the pre-trained models on 1% and 10% of ImageNet and report the top-5 accuracy. Note that, in this case the ‘‘Supervised’’ entry results come from [74] and are obtained by supervised training using only 1% or 10% of the labelled data. All the classification results are computed with single-crop testing. [†]: results computed by us.

We see that our method, even in its vanilla version, achieves significantly better results, while using at least two times fewer training epochs, which validates the efficiency of our proposed fully-online training methodology.

4.3. Self-supervised training on ImageNet

Here we evaluate our method by pre-training with it convnet-based representations on the full ImageNet dataset. We implement the full solution of our method (as described in § 4.1) using the ResNet50 (v1) [34] architecture. We evaluate the learned representations on ImageNet, Places205, and VOC07 classification tasks as well as on VOC07+12 detection task and provide results in Table 6. On the ImageNet classification we evaluate on two settings: (1) training linear classifiers with 100% of the data, and (2) fine-tuning the model using 1% or 10% of the data, which is referred to as *semi-supervised* learning.

Results. Pre-training on the full ImageNet and then transferring to downstream tasks is the most popular benchmark for unsupervised representations and thus many methods have configurations specifically tuned on it. In our case, due to the computational intensive nature of pre-training on ImageNet, no full tuning of OBoW took place. Nevertheless, it achieves very strong empirical results across the board. Its classification performance on ImageNet is 73.8%, which is substantially better than instance discrimination methods

MoCo v2 and SimCLR, and even improves over the recently proposed BYOL and SwAV methods when considering a similar amount of pre-training epochs. Moreover, in VOC07 classification and Places205 classification, it achieves a new state of the art despite using significantly fewer pre-training epochs than related methods. On the semi-supervised ImageNet ResNet50 setting, it significantly surpasses the state of the art for 1% labels, and is also better for 10% labels using again much fewer epochs. On VOC detection, it outperforms previous state-of-the-art methods while demonstrating strong performance improvements over supervised pre-training.

5. Conclusion

In this work, we introduce OBoW, a novel unsupervised teacher-student scheme that learns convnet-based representations with a BoW-guided reconstruction task. By employing an efficient fully-online training strategy and promoting the learning of context-aware representations, it delivers strong results that surpass prior state-of-the-art approaches on most evaluation protocols. For instance, when evaluating the derived unsupervised representations on the Places205 classification, Pascal classification or Pascal object detection tasks, OBoW attains a new state of the art, surpassing prior methods while demonstrating significant improvements over supervised representations.

References

- [1] Jean-Baptiste Alayrac, Joao Carreira, and Andrew Zisserman. The visual centrifuge: Model-free layered video representations. In *CVPR*, 2019. 2
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*, 2016. 2
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *ICCV*, 2017. 2
- [4] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. 1, 3, 4
- [5] Cristian Bucilă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *KDD*, 2006. 2
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 1, 3
- [7] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *ICCV*, 2019. 1, 3
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2, 3, 4, 5, 8
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2, 5, 8
- [10] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised GANs via auxiliary rotation loss. In *CVPR*, 2019. 2
- [11] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv*, 2020. 5
- [12] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 1
- [13] Gabriella Csurka, Christopher Dance, Lixin Fan, Jutta Willamowski, and Cédric Bray. Visual categorization with bags of keypoints. In *ECCVW*, 2004. 2
- [14] Arnaud Dapogny, Matthieu Cord, and Patrick Pérez. The missing data encoder: Cross-channel image completion with hide-and-seek adversarial network. In *AAAI*, 2020. 1
- [15] Carl Doersch, Abhinav Gupta, and Alexei Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015. 1, 2
- [16] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *ICLR*, 2017. 1
- [17] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019. 1
- [18] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Trans. PAMI*, 38(9), 2015. 2
- [19] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NeurIPS*, 2014. 1
- [20] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Massropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. In *ICLR*, 2017. 1
- [21] Mark Everingham, Luc Van Gool, Chris Williams, John Winn, and Andrew Zisserman. The Pascal visual object classes (VOC) challenge. *IJCV*, 88(2), 2010. 6
- [22] William Falcon and Kyunghyun Cho. A framework for contrastive self-supervised learning and designing a new approach. *arXiv*, 2020. 2
- [23] Jonathan Frankle, David J Schwab, and Ari Morcos. Are all negatives created equal in contrastive instance discrimination? *arXiv*, 2020. 2
- [24] Elizabeth Gardner and Bernard Derrida. Three unfinished works on the optimal storage capacity of networks. *Journal of Physics A: Mathematical and General*, 22(12), 1989. 2
- [25] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Learning representations by predicting bags of visual words. In *CVPR*, 2020. 1, 2, 8
- [26] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. 5
- [27] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 1, 2, 4
- [28] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. ActionVLAD: Learning spatio-temporal aggregation for action classification. In *CVPR*, 2017. 2
- [29] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017. 2
- [30] Faustino Gomez and Jürgen Schmidhuber. Evolving modular fast-weight networks for control. In *ICANN*, 2005. 5
- [31] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, 2020. 2, 5, 8
- [32] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 1
- [33] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 4, 8
- [34] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8
- [35] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *ICML*, 2020. 2
- [36] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS*, 2014. 2
- [37] Himalaya Jain, Spyros Gidaris, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. QuEST: Quantized embedding space for transferring knowledge. In *ECCV*, 2020. 2
- [38] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *arXiv*, 2020. 2

- [39] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *CVPR*, 2010. 2
- [40] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. In *NeurIPS*, 2020. 2
- [41] Anoop Korattikara Balan, Vivek Rathod, Kevin P Murphy, and Max Welling. Bayesian dark knowledge. In *NeurIPS*, 2015. 2
- [42] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2
- [43] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Learning representations for automatic colorization. In *ECCV*, 2016. 1, 2
- [44] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *ICCV*, 2017. 2
- [45] Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 8
- [46] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *CVPR*, 2020. 2, 5, 8
- [47] Ishan Misra, Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. 2
- [48] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 1, 2
- [49] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, 2018. 2
- [50] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *NeurIPS*, 2017. 4
- [51] George Papamakarios. Distilling model knowledge. *arXiv*, 2015. 2
- [52] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 1, 2
- [53] Florent Perronnin and Christopher Dance. Fisher kernels on visual vocabularies for image categorization. In *CVPR*, 2007. 2
- [54] Sudeep Pillai, Rareş Ambruş, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *ICRA*, 2019. 2
- [55] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018. 5
- [56] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In *NeurIPS*, 2019. 4
- [57] 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. *IJCV*, 115(3), 2015. 6
- [58] David Saad and Sara A Solla. Dynamics of on-line gradient descent learning for multilayer neural networks. In *NeurIPS*, 1996. 2
- [59] Josef Sivic and Andrew Zisserman. Video google: Efficient visual search of videos. In *Toward category-level object recognition*. Springer, 2006. 2
- [60] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. 6
- [61] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2
- [62] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. 2
- [63] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning. In *NeurIPS*, 2020. 5
- [64] Giorgos Tolias, Yannis Avrithis, and Hervé Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *ICCV*, 2013. 2
- [65] Trieu H Trinh, Minh-Thang Luong, and Quoc V Le. Selfie: Self-supervised pretraining for image embedding. *arXiv*, 2019. 2
- [66] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. In *IJCAI*, 2019. 2
- [67] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008. 2
- [68] Carl Vondrick, Abhinav Shrivastava, Alireza Fathi, Sergio Guadarrama, and Kevin Murphy. Tracking emerges by coloring videos. In *ECCV*, 2018. 2
- [69] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, 2020. 2
- [70] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *CVPR*, 2018. 2
- [71] Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. In *CVPR*, 2018. 1, 2
- [72] Jun Yang, Yu-Gang Jiang, Alexander G Hauptmann, and Chong-Wah Ngo. Evaluating bag-of-visual-words representations in scene classification. In *MIR*, 2007. 1
- [73] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. CutMix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 5
- [74] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S⁴L: Self-supervised semi-supervised learning. In *ICCV*, 2019. 8
- [75] Liheng Zhang, Guo-Jun Qi, Liqiang Wang, and Jiebo Luo. AET vs. AED: Unsupervised representation learning by auto-encoding transformations rather than data. In *CVPR*, 2019. 2
- [76] Richard Zhang, Phillip Isola, and Alexei Efros. Colorful image colorization. In *ECCV*, 2016. 1, 2
- [77] Richard Zhang, Phillip Isola, and Alexei Efros. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In *CVPR*, 2017. 1, 2

[78] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In *NeurIPS*, 2014. 6

[79] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017. 2