

Granularity-aware Adaptation for Image Retrieval over Multiple Tasks

Jon Almazán¹ Byungsoo Ko² Geonmo Gu²
Diane Larlus¹ Yannis Kalantidis¹

¹NAVER LABS Europe ²NAVER Corp.

Abstract. Strong image search models can be learned for a specific domain, *i.e.* set of labels, provided that some labeled images of that domain are available. A practical visual search model, however, should be versatile enough to solve multiple retrieval tasks simultaneously, even if those cover very different specialized domains. Additionally, it should be able to benefit from even unlabeled images from these various retrieval tasks. This is the more practical scenario that we consider in this paper. We address it with the proposed **Grappa**, an approach that starts from a strong pretrained model, and adapts it to tackle multiple retrieval tasks concurrently, using only unlabeled images from the different task domains. We extend the pretrained model with multiple independently trained sets of adaptors that use pseudo-label sets of different sizes, effectively mimicking different pseudo-granularities. We reconcile all adaptor sets into a single unified model suited for all retrieval tasks by learning fusion layers that we guide by propagating pseudo-granularity attentions across neighbors in the feature space. Results on a benchmark composed of six heterogeneous retrieval tasks show that the unsupervised **Grappa** model improves the zero-shot performance of a state-of-the-art self-supervised learning model, and in some places reaches or improves over a task label-aware oracle that selects the most fitting pseudo-granularity per task.

1 Introduction

The last few years have witnessed progress on image retrieval: successful models can be trained, provided that a set of labeled images from the domain of interest (not necessary from the same categories) is available for training, as in the common deep metric learning scenario. Those models are as powerful as they are specialized: it has been shown, and we confirm in our experiments, that one model carefully tailored for one domain (*e.g.* bird species) tend to perform poorly to a neighboring yet different domain (*e.g.* dog breeds).

Here, we argue that a practical visual search system should be able to solve multiple retrieval tasks simultaneously, without needing to explicitly specialize for each task. Consider for example a visual search system specialized to fauna and flora. In such a system, the image database covers a broad range of fine-grained domains, *e.g.* from searching among different insect species to different kinds of mushrooms. For the system to also handle coral species, it should be as simple as providing a set of unlabeled coral images.

In parallel, the field has worked towards pretraining large and generic models for visual representations that can be used, often as a black box, to extract features for new tasks. Among those, models trained in a self-supervised way have shown to be versatile to various target tasks, including image retrieval [22,9].

In this work, we assume access to such a large pretrained model that already provides good zero-shot performance. We also assume access to an unlabeled set of images possibly from multiple tasks. We propose to adapt the initial model so it performs even better on multiple image retrieval tasks *simultaneously*, *i.e.* when this same adapted model is used to extract features for all tasks.

This raises two questions. First, *how should we perform adaptation?* Fine-tuning is prohibitively costly especially for large pretrained models, and does not always transfer well. As an alternative to fine-tuning, and inspired by an early work on multi-task training [47] and a recent trend in natural language processing [28,42], we propose to use adaptor layers. Adaptors are embedded in between architecture blocks, and are the only weights learned, the ones from the original pretrained model remaining fixed. Our experiments show that this composite architecture allows for a versatile adaptation of a strong initial model by adjusting only a small percentage of the model parameters.

Second, *how should we reconcile various retrieval tasks in a single model?* A retrieval task focuses on a given set of visual concepts, often associated to a particular granularity. Yet, unlike in classification for which the granularity is known beforehand, the granularity of a retrieval task is context dependent, and depends on the gallery of images where visual search is performed. We therefore propose learning different sets of adaptors, each set tailored to one specific granularity. As we assume that training images are unlabeled, not even to indicate the retrieval task they correspond to, we propose to automatically define levels of granularity by partitioning the training set into more and more clusters. As a result, each partition corresponds to a different set of pseudo-labels. We then independently train one set of adaptors for each pseudo-granularity.

Next, we need to reconcile these different sets of adaptors into a single multi-purpose retrieval model. One option is to combine them with a naive fusion mechanism. The resulting model improves results on all retrieval tasks, showing the clear benefit of a multi-granularity understanding of the data. Another option is to go one step further and to achieve adaptor fusion via attention propagation. In this case, we require consistency between the adaptor attention of




Fig. 1: **Grappa** is an unsupervised method that trains a *single model* with higher zero-shot performance (measured with \mathcal{R} -Precision or \mathcal{RP}) than the pretrained DINO [9] model, over several retrieval tasks.

nearest neighbors in the feature space. We observe this fusion mechanism further improves the model.

To summarize, our contribution is threefold. First, we palliate the absence of image and task labels by creating sets of pseudo-labels, with the goal of approximating any possible granularities in a given set of retrieval tasks. Second, we propose a way to extend transformer-based architectures with adaptors, and a training framework that tailors individual sets of adaptors to different pseudo-granularities. Third, we propose a number of ways for fusing the adapter features, *e.g.* via augmentation invariance or via propagating attention from neighbors in the image features space. We validate our approach on a collection of datasets for deep metric learning and we show that **Grappa** improves over the successful DINO pretrained model, a model known to already obtain strong zero-shot performance on all these retrieval tasks (see Fig. 1).

2 Related Work

The task we tackle in this paper strongly relates to deep metric learning. It requires specific architectural changes of neural networks to extend them with adaptors. Note that our task can be seen as solving a zero-shot problem, *i.e.* it requires no labeled data from the downstream datasets and learns a single model for all tasks, something fairly uncommon in transfer learning.

Deep metric learning (DML). DML aims to learn a metric between data points that reflects the semantic similarity between them. It plays an important role in a wide range of tasks such as image clustering [27,7], unsupervised learning [10,25,8], and visual search [6,20,48]. Recent DML approaches typically learn visual similarity using either a pair-based loss [12,54,41,23,31] which considers pair-wise similarities, a proxy-based loss [59,58,15,24,32], which considers the similarity between samples and class representative proxies, or a contextual classification loss [66,5,17,51]. In most cases, DML approaches finetune an ImageNet pretrained model for *each* target retrieval task, and each of those finetuned models fall short when applied to other retrieval tasks. We aim at a more versatile visual search system that handles multiple retrieval tasks with a *single* model.

Neural architectures with adaptation layers. Adaptation layers (or adaptors) have emerged [47,28,42,60] as a way to avoid common problems rising in sequential finetuning or multi-task learning when trying to finetune large pretrained models to solve multiple tasks, namely the issues of catastrophic forgetting [37] and task imbalance. Rebuffi *et al.* [47] were the first to introduce adaptors to visual recognition tasks, adapting a convolutional model to many classification tasks. Adaptors have also been used with transformer architectures for natural language processing [28]; bottleneck layers are added to all the blocks of a pretrained model and finetuned, keeping the underlying model fixed.

Recently, Pfeiffer *et al.* [42] introduced a way to share knowledge between adaptors using an adaptor fusion layer within a two-stage learning framework: adaptors are trained independently in the first stage, they are kept fixed while

only the fusion layer is trained in the second stage. All the methods mentioned above still result in models that specialize to a single task; *e.g.* [42] learns a separate fusion layer per downstream task, whereas we would like to learn a single model for all tasks.

Zero-shot problems. The field has recently taken an interest in pretraining large models, sometimes called zero-shot models, using large quantities of data. Those have been shown to be versatile and applicable to many target tasks. Among them, self-supervised models [11,8,25,65,9] are trained using self-defined pseudo-labels as supervision and typically millions of images (*e.g.* from ImageNet [14]). Recent works [21,55] exploit even larger yet uncured, sets of unlabeled images to enhance the quality of the learned representations. Others [46,29,67] have leveraged multiple modalities, *e.g.* training visual representations so they are similar to the textual representations of their associated text. Those self-supervised or multimodal methods offer excellent initialization to be finetuned for a wide range of downstream tasks. Sometimes they are used in a zero-shot setting: a single model is used as a feature extractor, typically to solve multiple tasks. This is the regime we study here, but we further assume that a small amount of unlabeled data from the downstream tasks exists.

Relation to other transfer tasks. The idea of transferring a model trained for a given task to a related one has become central to computer vision [52,50], and appears in many research fields such as task transfer [64], domain adaptation [13] or self-supervised learning [19,40,9]. Yet, in all those, the initial model is only a starting point and it is typically not only extended, but also retrained for each task of interest, leading to a multitude of specialized models. In our work, we need a *single* model to perform well across retrieval tasks. In that regard, this work is closer to zero-shot transfer of the large pretrained models discussed above. Also related are Mixtures of Experts (MoE) [63,53,45,49], an ensembling technique that decomposes a predictive problem into subtasks, training one expert for each. Although MoE architectures may look similar to ours at first glance, they typically rely on gating and pooling mechanisms that learn to predict, in a supervised way, which experts to trust, and how to combine them. Similar to typical transfer approaches, they build one specialized model for each target task. Here, we focus on a purely unsupervised task: no labels are provided to indicate image semantic content nor the retrieval task images belong to.

3 A granularity-aware multi-purpose retrieval model

In this section we present **Grappa**, a method for adapting a pretrained model to multiple retrieval tasks simultaneously, in an unsupervised way. We first formalize our task, *i.e.* visual search over several retrieval tasks using a single model (Sec. 3.1). We then present an overview of the approach (Sec. 3.2). Next, we detail each step, *i.e.* building multiple granularities (Sec. 3.3), learning adaptors using granularity-aware pseudo-labels (Sec. 3.4), and learning to fuse them by propagating adaptor attention across feature space neighbors (Sec. 3.5).

3.1 Background

Our task of interest, **visual search on multiple retrieval tasks**, can be seen as a variant of the standard deep metric learning (DML) task. The most common protocol in DML is to a) split the *classes*¹ into disjoint train and test sets of labels; b) learn a separate model for each retrieval task on the corresponding train split; c) perform retrieval on all images of the (unseen) test split of classes.

Our setting has several key differences. First, we solve multiple retrieval tasks *simultaneously*. This means that we do not learn one model for each but a *single* model that will be used for all tasks. Second, we only assume access to a set of *unlabeled* images from each retrieval task, and do not have access to labeled training sets, unlike standard DML methods. Even more challenging, unlabeled training images are provided jointly without knowing which target retrieval task they correspond to nor the total number of retrieval tasks.

More formally, let \mathcal{T} be the set of m retrieval tasks that we want to simultaneously tackle. Each task \mathcal{T}^t is associated with a training and a test set. At train time, we are provided with a fused training set \mathcal{D} composed of the union of all training sets of the m datasets in \mathcal{T} . As mentioned earlier, images are not associated to any class or task label.

With so many unknowns on the target retrieval tasks, an obvious choice is to **start with a large pretrained model**. Self-[11,25,9] or weakly-[46,29] supervised learning have been shown to lead to strong models that generalize well and exhibit high zero-shot transfer performance. We assume that we are given such a model. Here, we base our work on the recently proposed Visual Transformer (ViT) [16], a popular, efficient, and highly performing architecture, pretrained in a self-supervised way with DINO [9].

We set our pretrained model \mathcal{M} to be a ViT with L transformer layers and an input patch size $P \times P$ pixels. Input image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ is reshaped into a sequence of T flattened 2D patches where $T = HW/P^2$. The transformer uses constant latent vector size D through all of its layers, so flattened patches are first mapped to D dimensions with a trainable linear projection and concatenated in \mathbf{h}^0 , together with a prepended learnable [class] token and added position embeddings. The transformer encoder [56] consists of alternating blocks of multi-headed self-attention (MSA) and MLP (which contain two layers with a GELU non-linearity). LayerNorm (LN) is applied before every block, and residual connections after every block. Formally, each layer of \mathcal{M} (shown with a gray background on Fig. 3, left) is given by:

$$\mathbf{h}^l = \text{MLP}(\text{LN}(\tilde{\mathbf{h}}^l)) + \tilde{\mathbf{h}}^l, \quad \tilde{\mathbf{h}}^l = \text{MSA}(\text{LN}(\mathbf{h}^{l-1})) + \mathbf{h}^{l-1}, \quad (1)$$

for $l = \{1 \dots L\}$. The image representation \mathbf{z} is the output of the [class] token after the last layer \mathbf{h}^L , *i.e.* $\mathbf{z} = \text{LN}(\mathbf{h}^L)[\text{class}]$. We refer the reader to [16] for more details about the ViT architecture.

¹ We will use the term *classes* to refer to sets of images with the same label, whether the latter represents object instances or fine-grained classes.

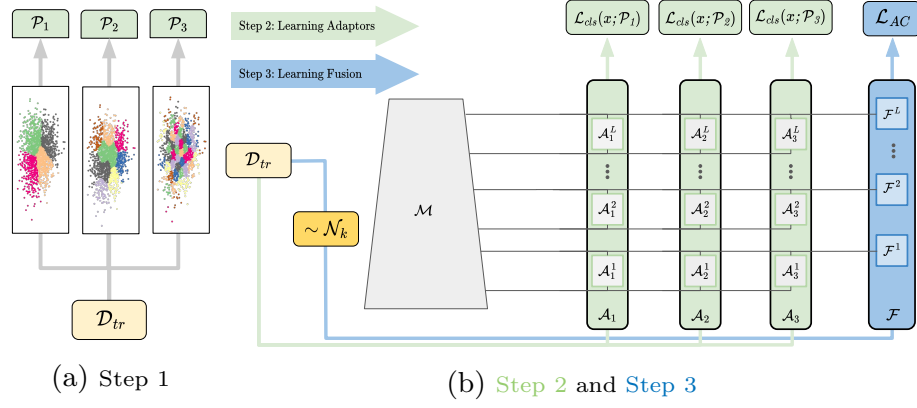


Fig. 2: **Training of the proposed Grappa.** Left: granularities correspond to pseudo-labels \mathcal{P}_i obtained by multiple clusterings of the feature space (Step 1). Right: we learn the granularity-aware adaptors (Step 2, in green), and then learn how to fuse them (Step 3, in blue). For this example, $N=3$.

3.2 Method overview

Our method builds on the ViT [16] model \mathcal{M} pretrained with DINO [9], that we treat as an architectural backbone. We extend and train it in an unsupervised way using \mathcal{D} . The training process consists of three steps (summarized in Fig. 2):

- **Step 1: Learning pseudo-labels.** We build multiple sets of pseudo-labels. Each set partitions the feature space using clustering and corresponds to a different pseudo-granularity. This process is illustrated in Fig. 2a.
- **Step 2: Learning adaptors for each pseudo-label set independently.** We learn a set of adaptors specific to each pseudo-granularity using a classification loss. This process is depicted by the green arrows in Fig. 2b.
- **Step 3: Learning to fuse adaptors.** We learn a set of fusion layers to merge the outputs of multiple adaptors using a transformation invariance or an attention propagation loss, *i.e.* neighboring images should have similar attentions over adaptors. This process is depicted by the blue arrows in Fig. 2b.

These three stages lead to a single model we denote as \mathcal{M}^* , that unifies the multiple granularities, and consists of: the pretrained model \mathcal{M} used as a *frozen* backbone (its parameters are kept fixed during the entirety of the process), embedded adaptors \mathcal{A}_i , and fusion layers \mathcal{F} . This single model is used as a unique feature extractor for *all* retrieval tasks considered in our benchmark. We denote our method **Grappa**, that stands for learning **G**ranularity-aware **A**daptors by **A**ttention **P**ropagation. The following subsections detail the learning stages.




Fig. 3: Architecture of the l^{th} layer of the Grappa model, for $N=3$ adaptors.

3.3 Step 1: Learning pseudo-labels

We would like to build multiple sets of pseudo-labels such that they partition the feature space at different ‘granularities’. We can approximate this partitioning by estimating multiple sets of clusters while varying the number of centers.

In practice, we extract features using the pretrained model \mathcal{M} . Let $\mathbf{z} = f(\mathbf{x}; \mathcal{M})$ be the feature of an image $\mathbf{x} \in \mathcal{D}$. Let the set of all features for training set \mathcal{D} be $\mathcal{Z} = \{f(\mathbf{x}; \mathcal{M}), \forall \mathbf{x} \in \mathcal{D}\}$. To get multiple sets of pseudo-labels, we cluster the full set of features \mathcal{Z} into sets of centroids $\mathcal{C}_i, i = 1..N$, of respectively k_i clusters, where k_i gets monotonically larger as i approaches N . This produces N sets of pseudo-labels $\mathcal{P}_1, \dots, \mathcal{P}_N$. For each pseudo-label set \mathcal{P}_i , an image $\mathbf{x} \in \mathcal{D}$, is associated to a pseudo-label given by $\mathcal{P}_i(\mathbf{x}) = \arg \min_{\mathbf{c} \in \mathcal{C}_i} \|\mathbf{z} - \mathbf{c}\|$, for $\mathbf{z} = f(\mathbf{x}; \mathcal{M})$. We rely on the vanilla k -means clustering algorithm [35] with k -means++ [1] initialization, a common choice for the size of our benchmark. For even larger datasets, more scalable variants could be used, like hierarchical [43], approximate [2], or quantized [3] k -means. Note that other works have used k -means to define pseudolabels [7, 61]. Yet, our work is the first to learn multiple sets and subsequently use all of them.

3.4 Step 2: Learning adaptors for each pseudo-label set

Given the N sets of pseudo-labels computed in the previous step, we now would like to learn adaptors tailored to each pseudo-label set, *i.e.* to each pseudo-granularity. We use the pretrained model \mathcal{M} as a backbone and extend it by embedding an adaptor at every layer. We then learn the adaptor parameters while keeping the backbone ones frozen. We learn a set of L adaptors for *each* pseudo-granularity in an independent way.

Adaptor architecture. Recent works in natural language processing [28, 44, 60] have embedded adaptor layers in transformer architectures. We follow a similar design and embed L adaptors, one at the end of each transformer layer of \mathcal{M} .

Formally, we learn a separate set of adaptors \mathcal{A}_i for each pseudo-label set $\mathcal{P}_i, i = \{1, \dots, N\}$. Each set \mathcal{A}_i consists of L adaptors, denoted as $\mathcal{A}_i^1, \dots, \mathcal{A}_i^L$. These adaptors are bottleneck layers with an intermediate dimensionality of D' (where $D' < D$), a GELU layer [26] in between, and a residual connection at

the end. Since we are modifying the architecture of \mathcal{M} by interleaving it with blocks, we need to revisit notations. The output of layer l in \mathcal{M} (after the basic ViT block) is now defined as $\bar{\mathbf{h}}^l = \text{MLP}(\text{LN}(\tilde{\mathbf{h}}^l)) + \tilde{\mathbf{h}}^l$. The output of the new layer l (the original ViT block combined with an adaptor) can still be denoted as \mathbf{h}^l . Details of the overall architecture are shown in Fig. 3.

Learning the adaptors. Given a set of pseudo-labels \mathcal{P}_i , we can learn the parameters of the set of adaptors \mathcal{A}_i via a supervised cross entropy loss. Specifically, we use the norm-softmax loss [58,59,34,66] that, for image \mathbf{x} , is given by:

$$\mathcal{L}_{cls}(\mathbf{x}; y) = \log \frac{\exp(\gamma \cos \theta_y)}{\sum_{y'=1}^{k_i} \exp(\gamma \cos \theta_{y'})}, \quad (2)$$

where γ is a scale factor, $\cos \theta_y$ is the cosine similarity to the classifier of class y , and the loss is guided by the pseudo-labels, *i.e.* $y = \mathcal{P}_i(\mathbf{x})$. After learning the parameters for each set \mathcal{P}_i , we keep the adaptors and discard the classifiers.

3.5 Step 3: Learning to fuse adaptors

The process described in Sec. 3.4 leads to N separate sets of adaptors, each tailored to a different pseudo-granularity. The next step is to unify all adaptor sets into a single architecture. To that end, we append (*i.e.* stack) the N adaptors for each layer *in parallel*, as shown in Fig. 3. We then concatenate adaptor outputs in a tensor $\mathbf{U}^l \in \mathbb{R}^{N \times T \times D}$ for each layer $l = \{1, \dots, L\}$, where each row corresponds to the output of one adaptor for this layer. Here, another residual connection is added, giving the model the opportunity to bypass the adapter if needed. Tensor \mathbf{U}^l is therefore given by $\mathbf{U}^l = \{\mathcal{A}_i^l(\bar{\mathbf{h}}^l) + \text{MLP}(\text{LN}(\bar{\mathbf{h}}^l)), i = 1..N\}$ and is then fed, together with $\bar{\mathbf{h}}^l$, to a fusion layer, as detailed below.

First Option: Fusion by average pooling. A straightforward way of fusing the outputs of the N adaptors is to treat them as equally important and average them. The fusion layer therefore is simply an average pooling layer that takes tensor $\mathbf{U}^l \in \mathbb{R}^{N \times T \times D}$ as input and computes the mean over its first dimension. We refer to this simpler version of our approach as **Grappa-avg**.

Second option: Learning to fuse. Treating all adaptors as equally important for any input image goes against our intuition that different retrieval tasks are more related to certain granularities, and hence more suited for the corresponding adaptors. We therefore design a fusion layer with trainable parameters, that can learn to weigh the different adaptor outputs. We use a simple dot-product self-attention architecture over the sequence of N adaptor outputs. Yet, we make two crucial modifications to the vanilla *query-key-value* self-attention: a) To learn an *image-level* attention, we average over the T spatial tokens; b) Given that we want to fuse the adaptors but do not want to alter the adaptor representations, we omit the linear projection of the value branch, and only learn projections for the query and key branches that affect the re-weighting of adaptor features.

Specifically, the fusion layer learns an attention vector of size N over the adaptors, given inputs $\bar{\mathbf{h}}^l$ and \mathbf{U}^l , by $\mathcal{F}^l(\bar{\mathbf{h}}^l, \mathbf{U}^l) = \alpha^l(\bar{\mathbf{h}}^l, \mathbf{U}^l)\mathbf{U}^l$, where vector $\alpha^l(\bar{\mathbf{h}}^l, \mathbf{U}^l) \in \mathbb{R}^N$ is given by:

$$\alpha^l(\bar{\mathbf{h}}^l, \mathbf{U}^l) = \text{softmax} \left(\frac{(\mathbf{Q} \sum_T \bar{\mathbf{h}}^l) (\mathbf{K} \sum_T \mathbf{U}^l)^T}{\sqrt{D}} \right) \quad (3)$$

where $l = \{1, \dots, L\}$, \mathbf{Q} and \mathbf{K} are linear projections of size $D \times D$. A final residual connection is added after the fusion layer. The architecture details of a complete layer are shown in Fig. 3. The latter comprises the ViT block, the adaptors, and the fusion layer, all appended in a residual fashion.

Given pretrained model \mathcal{M} and multiple sets of adaptors, one way to build a single model is to select one set of adaptors per image. This amounts to guessing which pseudo-granularity best fits each image. We argue that, in a generic visual search system, “picking a granularity” for a query image depends less on the image content than on the retrieval task, *i.e.* the gallery used at test time. Given a dog image query, for example, the only way to know if we are looking for any dog image or only images of the same dog breed, is by looking at the local structure of the gallery around that image. Both scenarios might favor different representations; our system reconciles them by learning a combination of adaptors. Obviously, we do *not* have access to the test images during training. Yet, we assume access to unlabeled set of images \mathcal{D} , representative of the target retrieval tasks, or at least of their granularity. Again, these images are provided without *task labels*, we do not know which retrieval task they correspond to.

Without any other supervisory signal, we argue that the local neighborhood in the feature space of the training set \mathcal{D} can be used to approximate the “granularity” of a query. In other words, we assume that visually similar images from \mathcal{D} should yield similar attention vectors over the sets of adaptors. We therefore propose to learn to fuse adaptors using a loss on neighboring image pairs in the feature space. In this step, the backbone model \mathcal{M} and the adaptors remain frozen. The fusion layer only learns \mathbf{K} and \mathbf{Q} , two linear projections that are multiplied to give the attention vectors α^l , for each ViT encoder l . This means that any loss applied to this fusion step *only re-weights* adaptor features. We denote the final model, composed of the backbone with all embedded adaptors and their fusion, as \mathcal{M}^* , and the corresponding feature extractor as $f^*(\mathbf{x}, \mathcal{M}^*)$.

Attention propagation loss. As mentioned, we propose to train the fusion layer leveraging the assumption that neighboring image pairs in the feature space should use similar attentions over adaptors. Let $\mathcal{N}_k(\mathbf{x}; \mathcal{D})$ denote the k nearest neighbors of \mathbf{x} from dataset \mathcal{D} . We define neighbors $(\mathbf{x}_i, \mathbf{x}_j)$ as a pair of inputs such that $\mathbf{x}_j \in \mathcal{N}_k(\mathbf{x}_i; \mathcal{D})$. Although neighbors could be built using the pretrained model $\mathbf{z} = f(\mathbf{x}, \mathcal{M})$ (static k -NN), the representations $\tilde{\mathbf{z}} = f^*(\mathbf{x}, \mathcal{M}^*)$ from the learned model \mathcal{M}^* provide a better estimation. This requires to periodically update neighbors during training (in practice we do it at every epoch). Given a pair of neighboring features, we bring their adaptor attentions close to each other and strive for *attention consistency*.

Attention consistency is enforced using the pairwise Barlow Twins loss [65]. Given a batch of image pairs, the loss is defined over the output representations $\tilde{\mathbf{z}}_i = f^*(\mathbf{x}_i; \mathcal{M}^*)$, $\tilde{\mathbf{z}}_j = f^*(\mathbf{x}_j; \mathcal{M}^*)$ of our model, computed over the $D \times D$

cross-correlation matrix C and averaged over the batch, *i.e.*:

$$\mathcal{L}_{BT} = \sum_n (1 - C^{nn})^2 + \beta \sum_n \sum_{m \neq n} (C^{nm})^2, C_{nm} = \frac{\sum_b g(\hat{\mathbf{z}}_i)^{b,n} g(\hat{\mathbf{z}}_j)^{b,m}}{\sqrt{\sum_b (g(\hat{\mathbf{z}}_i)^{b,n})^2} \sqrt{\sum_b (g(\hat{\mathbf{z}}_j)^{b,m})^2}}, \quad (4)$$

where b iterates over pairs in the batch, n and m iterate over feature dimensions, β is a hyperparameter and $g(\cdot)$ is a MLP projector appended to the model and discarded after training. We refer the reader to [65] for more details.

Originally, *i.e.* in [65], this loss was defined over two transformed versions of the same image ($\mathbf{x}_i = t(\mathbf{x}), \mathbf{x}_j = t(\mathbf{x})$). When image pairs are created using image transformations, Eq.(4) defines a transformation consistency (TC) loss or \mathcal{L}_{TC} . This is a variant that we consider in our benchmark, referred to as **Grappa-T**. However, we are interested in applying this loss on neighboring pairs in the feature space, $(\mathbf{x}_i, \mathbf{x}_j)$ such as $\mathbf{x}_j \in \mathcal{N}_k(\mathbf{x}_i; \mathcal{D})$, and using it for attention propagation. In this case, we depart from [65] and follow the recent TLDR method [30] which uses the Barlow Twins loss over neighbor pairs for learning a feature encoder for dimensionality reduction. Similarly, we use the Barlow Twins loss on image pairs defined using the k-NN graph. We denote the loss in Eq.(4) as an *attention consistency* (AC) loss, or \mathcal{L}_{AC} , and refer to this variant as **Grappa-N**.

4 Experiments

In this section we validate the proposed **Grappa** on several retrieval tasks. These tasks are collected in a new benchmark that we introduce, called MRT. It unifies 6 fine-grained classification datasets under a retrieval setting. We show statistics and present the evaluation protocol of this benchmark in Sec. 4.1, we present the methods we compare in Sec. 4.2 and we report all results in Sec. 4.3.

Implementation details. We use ViT-Small [16] as a backbone architecture, with a patch size of 16 pixels and DINO [9] pre-trained weights. We generate $N=8$ sets of pseudo-labels on the training set of MRT, respectively composed of 256, 1024, 4096, 8,192, 16,384, 32,768, 65,536, and 131,072 clusters. We learn a set of adaptors for each pseudo-label set, using the norm-softmax loss from Eq.(2) and an Adam optimizer with a learning rate and weight decay of 0.001. We train the fusion layer over these adaptors using the Barlow Twins loss [65] and the LARS optimizer [62]. We used the same hyper-parameters for the scaling and β as suggested in [65], and a learning rate and weight decay of 0.5 and 0.001.

Evaluation metrics. Recent works [38,18] in DML have questioned standard evaluation metrics (*i.e.* Recall@1) and argue they are not fair. Thus, we report the R-Precision (\mathcal{RP}) and MAP@R metrics recently introduced in [38].

4.1 Multiple Retrieval Tasks (MRT) Benchmark

Data. The Multiple Retrieval Tasks (MRT) benchmark combines the 6 following fine-grained datasets under a retrieval setting: Aircraft [36], Cars [33], CUB [57],

Table 1: **Statistics of the Multiple Retrieval Tasks (MRT) benchmark.** It is composed of 6 datasets. Classes in train and test are *disjoint*. We provide the number of classes as a reference, but labels are *never* used during training.

Dataset	Training		Testing	
	# Classes	# Images	# Classes	# Images
Aircraft [36]	50	5,000	50	5,000
Cars [33]	98	8,054	98	8,131
CUB [57]	100	5,864	100	5,924
Flowers [39]	51	3,870	51	4,319
Food-101 [4]	51	51,000	50	50,000
Products [41]	11,318	59,551	11,316	60,502
MRT	11,668	133,339	11,665	133,876

Flowers [39], Food-101 [4], and Stanford online products (Products) [41]. We follow standard practice in the DML community and, for each dataset, assign the first half of the classes (ordered alphabetically) for training and the second half for testing. We then combine images from all the training splits into a single training set \mathcal{D} of 133,339 images and discard their class and dataset labels. This is the training set we use to learn the pseudo-labels, as well as the adaptor and the fusion parameters. We show statistics for all datasets in Table 1.

Evaluation protocol. Models are trained on the combined training set \mathcal{D} without task nor class labels. Evaluation is performed on the test split of each task independently, following a leave-one-out protocol: each image is used as a query once to rank all the other images in the test set. For evaluation, relevance is defined according to class labels and we report mean average precision (MAP@R or mAP) and \mathcal{R} -Precision ($\mathcal{R}P$) over all queries.

4.2 Compared methods

Baselines. First and foremost, we compare with the DINO pretrained visual transformer of [9], a self-supervised model trained on ImageNet1K. It obtains impressive zero-shot performance on retrieval and constitutes a very strong baseline.² We denote this baseline as **DINO** or simply \mathcal{M} in Table 2 and Fig. 4.

The **Grappa** architecture adds an extra set of parameters in the form of adaptors and fusion layers. To verify that improvements do not simply come from these extra parameters, we report results with a second baseline that has the same number of parameters as the **Grappa** models, but uses no pseudo-granularity-based adaptors nor the proposed attention consistency loss. Instead of following Step 2, we randomly initialize adaptors and finetune them when learning the fusion. For the latter, we use the Barlow Twins loss from Eq.(4), with transformation consistency, and train it on the training set of MRT, similar to **Grappa**. We denote this baseline: \mathcal{M}^* (**random**).

² We chose DINO over the CLIP [46] model as the training set of CLIP is not public. Therefore the data in MRT might be part of its 400M image-text training pairs.

Table 2: **Results on the Multiple Retrieval Task (MRT) benchmark.** We report MAP@R (mAP) and \mathcal{RP} on the six datasets of MRT, obtained by a single model from those listed in Sec. 4.2, all unsupervised. The oracle (in gray) is not comparable as it selects the set of adaptors that performs best for each task.

	Aircraft		Cars		CUB		Flowers		Food-101		Products	
	RP	mAP	RP	mAP	RP	mAP	RP	mAP	RP	mAP	RP	mAP
DINO (M)	17.5	9.2	9.0	3.5	33.6	22.9	62.5	57.0	27.0	16.1	34.1	31.5
Results without pseudo-granularity adaptors												
M + A ₁	15.4	7.8	7.9	2.7	16.2	7.8	62.8	56.9	22.1	11.6	32.0	29.5
M* (random)	11.8	5.0	6.2	1.7	11.0	4.5	62.7	56.2	21.9	11.2	28.9	26.5
Results using only a single set of pseudo-granularity adaptors												
M + P ₁	18.0	9.5	8.9	3.4	33.3	22.7	63.3	57.6	29.0	17.9	32.2	29.6
M + P ₂	18.0	9.5	9.1	3.6	34.9	24.1	66.7	61.3	29.3	18.1	32.7	30.1
M + P ₃	18.3	9.7	9.6	3.8	35.3	24.5	67.9	62.7	29.5	18.4	34.5	31.9
M + P ₄	18.1	9.5	9.7	3.9	35.2	24.3	67.9	62.6	29.6	18.4	35.8	33.1
M + P ₅	17.6	9.1	9.8	3.9	33.5	22.6	67.5	62.2	29.5	18.4	37.7	34.9
M + P ₆	17.3	8.7	9.8	3.9	31.4	20.3	66.9	61.5	29.4	18.3	39.7	36.8
M + P ₇	16.3	8.0	9.4	3.7	26.5	15.8	64.9	59.1	28.8	17.8	42.4	39.6
M + P ₈	12.6	5.4	7.5	2.7	14.8	6.6	56.7	49.8	22.1	11.9	46.5	43.7
Oracle (O)	18.3	9.7	9.8	3.9	35.3	24.5	67.9	62.7	29.6	18.4	46.5	43.7
i.e., only	P ₃		P ₅ –P ₆		P ₃		P ₃ –P ₄		P ₅		P ₈	
Results using 8 pseudo-granularity adaptors and fusion (M*)												
Grappa-avg	17.9	9.3	9.8	3.9	34.1	23.1	66.5	61.2	30.4	19.2	37.1	34.4
Grappa-T	18.1	9.5	9.8	3.9	34.4	23.5	66.5	61.2	30.4	19.2	35.9	33.2
Grappa-N	18.1	9.5	9.9	4.0	35.1	24.1	67.2	61.9	30.5	19.3	36.3	33.6
vs. DINO	(↑ 0.6)	(↑ 0.3)	(↑ 0.9)	(↑ 0.5)	(↑ 1.5)	(↑ 1.2)	(↑ 4.7)	(↑ 4.9)	(↑ 3.5)	(↑ 3.2)	(↑ 2.2)	(↑ 2.1)

Proposed. We report results for models with adaptor fusion described in Sec 3 together with results for individual adaptors. More precisely, we build N pseudo-label sets on the training set of MRT and train N sets of adaptors on these pseudo-labels independently. We report their individual performance as $\mathcal{M} + \mathcal{P}_i$. As mentioned above, we use DINO as a backbone and keep it frozen.

Then, using MRT’s training set again, we combine these N adaptors into a single model and train the fusion layer using i) a Barlow Twins loss on the final representation when creating pairs from two augmented views, reported as **Grappa-T**, and ii) our proposed Attention Consistency framework which relies on the local neighborhood, reported as **Grappa-N**. We also report results for the case where the fusion is an average pooling layer, denoted as **Grappa-avg**.

An adaptor selector oracle. What if we could choose the best performing pseudo-granularity for each retrieval task? Obviously, this requires access to the test set labels. This also results in a different representation per task, which departs from the universal representation we seek to learn. For these reasons, we only consider this variant as an oracle, and its results should not be compared with others. We still provide it as a reference, showing how much could be achieved if we set the attention as a one-hot vector that only enables the best possible pseudo-granularity adaptor. We denote the oracle as \mathcal{O} .

4.3 Results

We present our results in Table 2 and Fig. 4. Again, note that, unlike the common DML experimental setting, we use *a unique model for all retrieval datasets and no class nor task labels during training*. We make the following observations.

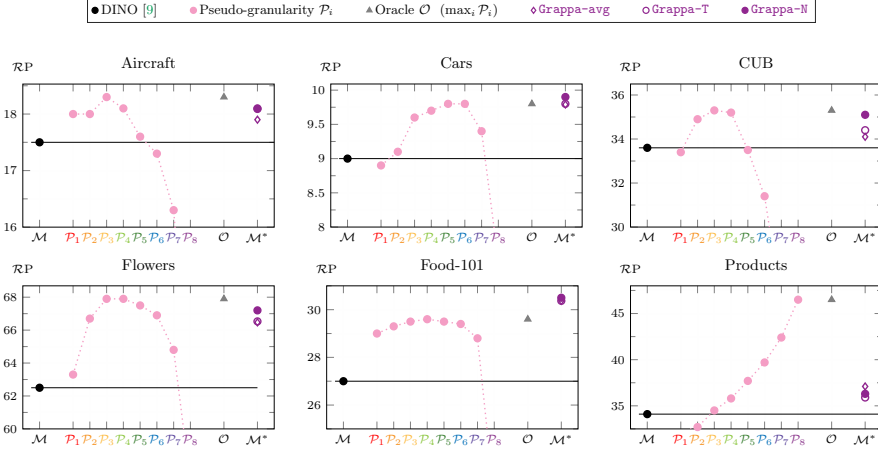


Fig. 4: **Results per dataset in MRT.** All 6 datasets use the same model.

Baselines. We confirm the initial observation [9] that, for common DML datasets, *DINO* is a very strong baseline. It achieves good performance even on the more challenging metrics \mathcal{RP} and MAP@R . Also, it turned out to be very challenging to improve over DINO by keeping its backbone frozen and embedding extra modules. We did our best to learn the additional modules from scratch, but were unsuccessful. Rows 2-3 of Table 2 report the best results after hyper-parameter tuning for a single set (no fusion) and for 8 sets of adaptors with fusion. Embedding randomly initialized modules typically deteriorates the performance. This makes separately trained pseudo-granularity adaptors all the more important.

Using a single set of adaptors. In rows 4-12 of Table 2 we report results when only using adaptors from a single pseudo-granularity; these results are also visualized as lavender-colored points in Fig. 4. We observe that, for each dataset, there exists at least one pseudo-label set that improves over DINO, and that the best one (reported as Oracle O) is different for different retrieval tasks. The oracle results use separate models, and selecting the best one for each task requires access to labels. Considering each set of adaptors as a separate model, some improve over DINO on several retrieval tasks, showing that even individual pseudo-granularities-specific sets of adaptors can be useful. Yet, as we will see next, higher gains can be achieved by fusing multiple adaptors into a single model.

Fusion and attention consistency. The bottom part of Table 2 shows the performance of the final model M^* that combines adaptors from all eight pseudo-granularities ($P_1 \dots P_8$). We see that using the proposed attention consistency (AC) loss over neighborhood pairs results in an even stronger model that improves over DINO and over the Grappa-T variant that uses pairs from different augmentations of the same image, in all datasets. In fact, we see that in most cases (*i.e.* apart from Products), the proposed AC loss is able to give results




Fig. 5: **Qualitative results.** Queries (first column) from Flowers and Food-101 and their top 5 retrieved results (columns 2-6) by DINO [9] and **Grappa-N**.

on-par with the oracle \mathcal{O} that would select the right pseudo-granularity for each dataset. Moreover, for Food-101 and Cars, **Grappa** outperforms the best adaptor for the dataset, showing that this oracle is not necessarily an upper bound: combining pseudo-granularities is beneficial. It is also worth noting that the simpler and parameter-free fusion of the **Grappa-avg** model is still improving over DINO in all cases, while it is also the best performing **Grappa** variant on the instance-level Products dataset.

Qualitative results. Fig. 5 presents qualitative results for two queries from the MRT benchmark. We present two cases where **Grappa-N** achieves significantly higher recall than DINO, *i.e.* cases where the local feature space is adapted in a way that images from the correct class are closer together.

Failure cases. As we see from Table 2 and Fig. 4, our fusion mechanism is not able to learn that \mathcal{P}_8 is the best pseudo-granularity for Products. We attribute this to the different topology of that dataset and the fact that hyperparameters were chosen to optimize performance for the union of the six benchmarks.

5 Conclusions

We present **Grappa**, an unsupervised approach for adapting a large pretrained backbone to simultaneously tackle multiple retrieval tasks, given only an unlabeled set of training images associated to these retrieval tasks. We show that one can adapt a large pretrained visual transformer using a set of pseudo-granularity adaptors and simple fusion layers. Our **Grappa** models bring consistent gains over the strong DINO [9] baseline on all six retrieval tasks we adapt to. We envision this work as a first step towards models that dynamically adapt.

Acknowledgements. MIAI@Grenoble Alpes (ANR-19-P3IA-0003)

References

1. Arthur, D., Vassilvitskii, S.: k-means++: The advantages of careful seeding. Tech. rep., Stanford (2006)
2. Avrithis, Y., Kalantidis, Y.: Approximate gaussian mixtures for large scale vocabularies. In: Proc. ECCV (2012)
3. Avrithis, Y., Kalantidis, Y., Anagnostopoulos, E., Emiris, I.Z.: Web-scale image clustering revisited. In: Proc. ICCV (2015)
4. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101—mining discriminative components with random forests. In: Proc. ECCV (2014)
5. Boudiaf, M., Rony, J., Ziko, I.M., Granger, E., Pedersoli, M., Piantanida, P., Ayed, I.B.: A unifying mutual information view of metric learning: cross-entropy vs. pairwise losses. In: Proc. ECCV (2020)
6. Cao, B., Araujo, A., Sim, J.: Unifying deep local and global features for image search. In: Proc. ECCV (2020)
7. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: Proc. ECCV (2018)
8. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: Proc. NeurIPS (2020)
9. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. Proc. ICCV (2021)
10. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proc. ICML (2020)
11. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: Proc. ICML (2020)
12. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: Proc. CVPR (2005)
13. Csurka, G. (ed.): Domain Adaptation in Computer Vision Applications. Advances in Computer Vision and Pattern Recognition, Springer (2017)
14. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Proc. CVPR (2009)
15. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proc. CVPR (2019)
16. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: Proc. ICLR (2021)
17. Elezi, I., Vascon, S., Torcinovich, A., Pelillo, M., Leal-Taixé, L.: The group loss for deep metric learning. In: Proc. ECCV (2020)
18. Fehervari, I., Ravichandran, A., Appalaraju, S.: Unbiased evaluation of deep metric learning algorithms. arXiv preprint arXiv:1911.12528 (2019)
19. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: Proc. ICLR (2018)
20. Gordo, A., Almazán, J., Revaud, J., Larlus, D.: Deep image retrieval: Learning global representations for image search. In: Proc. ECCV (2016)
21. Goyal, P., Caron, M., Lefaudeaux, B., Xu, M., Wang, P., Pai, V., Singh, M., Liptchinsky, V., Misra, I., Joulin, A., et al.: Self-supervised pretraining of visual features in the wild. arXiv preprint arXiv:2103.01988 (2021)

22. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning. In: *Proc. NeurIPS* (2020)
23. Gu, G., Ko, B.: Symmetrical synthesis for deep metric learning. In: *Proc. AAAI* (2020)
24. Gu, G., Ko, B., Kim, H.G.: Proxy synthesis: Learning with synthetic classes for deep metric learning. In: *Proc. AAAI* (2021)
25. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proc. CVPR* (2020)
26. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). *arXiv:1606.08415* (2016)
27. Hershey, J.R., Chen, Z., Le Roux, J., Watanabe, S.: Deep clustering: Discriminative embeddings for segmentation and separation. In: *Proc. ICASSP* (2016)
28. Houlshby, N., Stanislaw Jastrzebski, A.G., Morrone, B., de Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for NLP. In: *Proc. ICML* (2019)
29. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *Proc. ICML* (2021)
30. Kalantidis, Y., Lassance, C., Almazán, J., Larlus, D.: TLDR: Twin learning for dimensionality reduction. *TMLR* (2022)
31. Ko, B., Gu, G.: Embedding expansion: Augmentation in embedding space for deep metric learning. In: *Proc. CVPR* (2020)
32. Ko, B., Gu, G., Kim, H.G.: Learning with memory-based virtual classes for deep metric learning. In: *Proc. ICCV* (2021)
33. Krause, J., Deng, J., Stark, M., Li, F.F.: Collecting a large-scale dataset of fine-grained cars. In: *Proc. ICCV-W* (2013)
34. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphreface: Deep hypersphere embedding for face recognition. In: *Proc. CVPR* (2017)
35. Lloyd, S.: Least squares quantization in pcm. *TIT* **28**(2), 129–137 (1982)
36. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. *arXiv:1306.5151* (2013)
37. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation* **24**, 104–169 (1989)
38. Musgrave, K., Belongie, S., Lim, S.N.: A metric learning reality check. In: *Proc. ECCV* (2020)
39. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Proc. ICCVGIP* (2008)
40. Noroozi, M., Favaro, P.: Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In: *Proc. ECCV* (2016)
41. Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: *Proc. CVPR* (2016)
42. Pfeiffer, J., Kamath, A., Rücklé, A., Cho, K., Gurevych, I.: AdapterFusion: Non-destructive task composition for transfer learning. In: *Proc. EACL* (2021)
43. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: *Proc. CVPR* (2007)
44. Philip, J., Berard, A., Gallé, M., Besacier, L.: Monolingual adapters for zero-shot neural machine translation. In: *Proc. EMNLP* (2020)

45. Puigcerver, J., Riquelme, C., Mustafa, B., Renggli, C., Pinto, A.S., Gelly, S., Keysers, D., Houlsby, N.: Scalable transfer learning with expert models. In: Proc. ICLR (2021)
46. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Proc. ICML (2021)
47. Rebuffi, S.A., Bilen, H., Vedaldi, A.: Learning multiple visual domains with residual adapters. In: Proc. NeurIPS (2017)
48. Revaud, J., Almazán, J., Rezende, R., de Souza, C.: Learning with average precision: Training image retrieval with a listwise loss. In: Proc. ICCV (2019)
49. Riquelme, C., Puigcerver, J., Mustafa, B., Neumann, M., Jenatton, R., Pinto, A.S., Keysers, D., Houlsby, N.: Scaling vision with sparse mixture of experts. In: Proc. NeurIPS (2021)
50. Sariyildiz, M.B., Kalantidis, Y., Larlus, D., Alahari, K.: Concept generalization in visual representation learning. In: Proc. ICCV (2021)
51. Seidenschwarz, J., Elezi, I., Leal-Taixé, L.: Learning intra-batch connections for deep metric learning. In: Proc. ICML (2021)
52. Sharif Razavian, A., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: An astounding baseline for recognition. In: Proc. CVPR-W (2014)
53. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In: Proc. ICLR (2017)
54. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. Proc. NeurIPS (2016)
55. Tian, Y., Henaff, O.J., van den Oord, A.: Divide and contrast: Self-supervised learning from uncurated data. In: Proc. ICCV (2021)
56. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Proc. NeurIPS (2017)
57. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. rep., California Institute of Technology (2011)
58. Wang, F., Cheng, J., Liu, W., Liu, H.: Additive margin softmax for face verification. IEEE Signal Processing Letters (2018)
59. Wang, F., Xiang, X., Cheng, J., Yuille, A.L.: Normface: L2 hypersphere embedding for face verification. In: Proc. ACM Multimedia (2017)
60. Wang, R., Tang, D., Duan, N., Wei, Z., Huang, X., Ji, J., Cao, G., Jiang, D., Zhou, M.: K-adapter: Infusing knowledge into pre-trained models with adapters. In: ACL/IJCNLP (Findings) (2021)
61. Yan, X., Misra, I., Gupta, A., Ghadiyaram, D., Mahajan, D.: Clusterfit: Improving generalization of visual representations. In: Proc. CVPR (2020)
62. You, Y., Gitman, I., Ginsburg, B.: Large batch training of convolutional networks. In: Proc. ICLR (2018)
63. Yuksel, S.E., Wilson, J.N., Gader, D., P.: Twenty years of mixture of experts. Transactions on Neural Networks and Learning Systems (2012)
64. Zamir, A., Sax, A., Shen, W., Guibas, L., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: Proc. CVPR (2018)
65. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. In: Proc. ICML (2021)
66. Zhai, A., Wu, H.Y.: Classification is a strong baseline for deep metric learning. In: Proc. BMVC (2019)
67. Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., Beyer, L.: LiT: Zero-shot transfer with locked-image text tuning. In: Proc. CVPR (2022)