

Unsupervised Deep Metric Learning with Transformed Attention Consistency and Contrastive Clustering Loss

Yang Li¹[0000-0002-8372-1481], Shichao Kan²[0000-0003-0097-6196], and Zhihai He¹[0000-0002-2647-8286]

¹ University of Missouri, Columbia, MO, USA
yltb5@mail.missouri.edu, hezhi@missouri.edu

² Beijing Jiaotong University, Beijing, China
16112062@bjtu.edu.cn

Abstract. Existing approaches for unsupervised metric learning focus on exploring self-supervision information within the input image itself. We observe that, when analyzing images, human eyes often compare images against each other instead of examining images individually. In addition, they often pay attention to certain keypoints, image regions, or objects which are discriminative between image classes but highly consistent within classes. Even if the image is being transformed, the attention pattern will be consistent. Motivated by this observation, we develop a new approach to unsupervised deep metric learning where the network is learned based on self-supervision information across images instead of within one single image. To characterize the consistent pattern of human attention during image comparisons, we introduce the idea of transformed attention consistency. It assumes that visually similar images, even undergoing different image transforms, should share the same consistent visual attention map. This consistency leads to a pairwise self-supervision loss, allowing us to learn a Siamese deep neural network to encode and compare images against their transformed or matched pairs. To further enhance the inter-class discriminative power of the feature generated by this network, we adapt the concept of triplet loss from supervised metric learning to our unsupervised case and introduce the contrastive clustering loss. Our extensive experimental results on benchmark datasets demonstrate that our proposed method outperforms current state-of-the-art methods for unsupervised metric learning by a large margin.

Keywords: Unsupervised Metric Learning, Attention Map, Consistency Loss, Contrastive Loss.

1 Introduction

Deep metric learning aims to learn discriminative features that can aggregate visually similar images into compact clusters in the high-dimensional feature

space while separating images of different classes from each other. It has many important applications, including image retrieval [38, 14, 11], face recognition [37], visual tracking [33] and person re-identification [43, 15]. In supervised deep metric learning, we assume that the labels for training data are available. In this paper, we consider unsupervised deep metric learning where the image labels are not available. Learning directly and automatically from images in an unsupervised manner without human supervision represents a very important yet challenging task in computer vision and machine learning.

Clustering is one of the earliest methods developed for unsupervised learning. Recently, motivated by the remarkable success of deep learning, researchers have started to develop unsupervised learning methods using deep neural networks [2]. Auto-encoder trains an encoder deep neural network to output feature representations with sufficient information to reconstruct input images by a paired decoder [44]. As we know, during deep neural network training, the network model is updated and learned in an iterative and progressive manner so that the network output can match the target. In other works, deep neural networks need human supervision to provide ground-truth labels. However, in unsupervised learning, there are no labels available. To address this issue, researchers have exploited the unique characteristics of images and videos to create various self-supervised labels, objective functions, or loss functions, which essentially convert the unsupervised learning into a supervised one so that the deep neural networks can be successfully trained. For example, in DeepCluster [2], clustering is used to generate pseudo labels for images. Various supervised learning methods have been developed to train networks to predict the relative position of two randomly sampled patches [6], solve Jigsaw image puzzles [26], predict pixel values of missing image patches [5], classify image rotations of four discrete angles [10], reconstruct image transforms [44], etc. Once successfully trained by these pretext tasks, the baseline network should be able to generate discriminative features for subsequent tasks, such as image retrieval, classification, matching, etc [13].

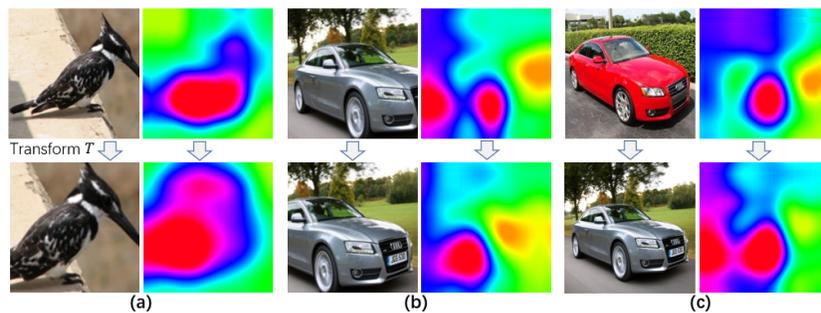


Fig. 1. Consistency of visual attention across images under transforms.

In this work, we propose to explore a new approach to unsupervised deep metric learning. We observe that existing methods for unsupervised metric learning focus on learning a network to analyze the input image itself. As we know, when examining and classifying images, human eyes compare images back and forth in order to identify discriminative features [9]. In other words, comparison plays an important role in human visual learning. When comparing images, they often pay attention to certain keypoints, image regions, or objects which are discriminative between image classes but highly consistent across image within classes. Even when the image is being transformed, the attention areas will be consistent. To further illustrate this, we provide three examples in Fig. 1. In (a), human eyes can easily tell the top image A of the first column and the bottom image B are the same bird since they have the same visual characteristics. The attention will be on the feather texture and head shape. In the pixel domain, A and B are up to a spatial transform, specifically, cropping plus resizing. When the human eyes moves from image A to its transformed version B , the attention will be also transformed so that it can be still focused on the head and feather. If we represent this attention using the attention map in deep neural networks, the attention map $\mathcal{M}(A)$ for image A and the attention map $\mathcal{M}(B)$ for image B should also follow the same transform, as shown in the second column of Fig. 1(a). We can also see this consistency of attention across image under different transforms in other examples in Figs. 1(b) and (c).

This lead to our idea of transformed attention consistency. Based on this idea, we develop a new approach to unsupervised deep metric learning based on image comparison. Specifically, using this consistency, we can define a pairwise self-supervision loss, allowing us to learn a Siamese deep neural network to encode and compare images against their transformed or matched pairs. To further enhance the inter-class discriminative power of the feature generated by this network, we adapt the concept of triplet loss from supervised metric learning to our unsupervised case and introduce the contrastive clustering loss. Our extensive experimental results on benchmark datasets demonstrate that our proposed method outperforms current state-of-the-art methods by a large margin.

2 Related Work and Major Contributions

This work is related to deep metric learning, self-supervised representation learning, unsupervised metric learning, and attention mechanisms.

(1) Deep metric learning. The main objective of deep metric learning is to learn a non-linear transformation of an input image by deep neural networks. In a common practice [35, 42], the backbone in deep metric learning can be pre-trained on 1000 classes ImageNet [29] classification, and is then jointly trained on the metric learning task with an additional linear embedding layer. Many recent deep metric learning methods are built on pair-based [4, 12, 27] and triplet relationships [22, 30, 40]. Triplet loss [30] defines a positive pair and a negative pair based on the same anchor point. It encourages the embedding distance of positive pair to be smaller than the distance of negative pair by a given mar-

gin. Multi-similarity loss [35] considers multiple similarities and provides a more powerful approach for mining and weighting informative pairs by considering multiple similarities. The ability of mining informative pairs in existing methods is limited by the size of mini-batch. Cross-batch memory (XBM) [36] provides a memory bank for the feature embeddings of past iterations. In this way, the informative pairs can be identified across the dataset instead of a mini-batch.

(2) Self-supervised representation learning. Self-supervised representation learning directly derives information from unlabeled data itself by formulating predictive tasks to learn informative feature representations. DeepCluster [2] uses k -means clustering to assign pseudo-labels to the features generated by the deep neural network and introduces a discriminative loss to train the network. Gidaris *et al.* [10] explore the geometric transformation and propose to predict the angle (0° , 90° , 180° , and 270°) of image rotation as a four-way classification. Zhang *et al.* [44] propose to predict the randomly sampled transformation from the encoded features by Auto-encoding transformation (AET). The encoder is forced to extract the features with visual structure information, which are informative enough for the decoder to decode the transformation. Self-supervision has been widely used to initialize and pre-train backbone on unlabeled data, and is then fine-tuned on a labeled training data for evaluating different tasks.

(3) Unsupervised metric learning. Unsupervised metric learning is a relatively new research topic. It is a more challenging task since the training classes have no labels and it does not overlap with the testing classes. Iscen *et al.* [17] propose an unsupervised method to mine hard positive and negative samples based on manifold-aware sampling. The feature embedding can be trained with standard contrastive and triplet loss. Ye *et al.* [42] propose to utilize the instance-wise relationship instead of class information in the learning process. It optimizes the instance feature embedding directly based on the positive augmentation invariant and negative separated properties.

(4) Attention mechanism. The goal of the attention mechanism is to capture the most informative feature in the image. It explores important parts of features and suppress unnecessary parts [1, 23, 18]. Convolutional block attention module (CBAM) [39] is an effective attention method with channel and spatial attention module which can be integrated into existing convolutional neural network architectures. Fu *et al.* [8] propose to produce the attention proposals and train the attention module and embedding module in an iterative two-stage manner. Chen *et al.* [3] propose the hybrid-attention system by random walk graph propagation for object attention and the adversary constraint for channel attention.

Compared to existing methods, the *unique contributions* of this paper can be summarized as follows. (1) Unlike existing methods which focus on information analysis of the input image only, we explore a new approach for unsupervised deep metric learning based on image comparison and cross-image consistency. (2) Motivated by the human visual experience, we introduce the new approach of transformed attention consistency to effectively learn a deep neural network which can focus on discriminative features. (3) We extend the existing triplet

loss developed for supervised metric learning to unsupervised learning using k -mean clustering to assign pseudo labels and memory bank to allow its access to all training samples, instead of samples in the current mini-batch. (4) Our experimental results demonstrate that our proposed approach has improved the state-of-the-art performance by a large margin.

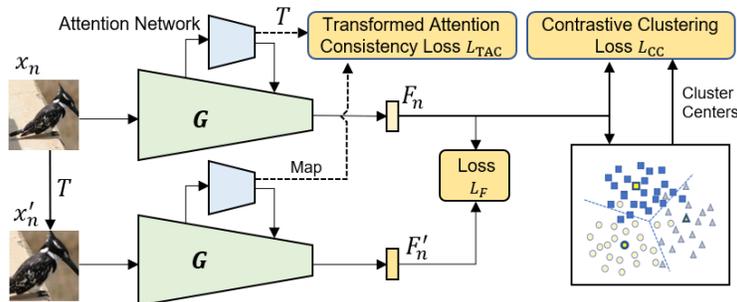


Fig. 2. Overview of the proposed approach for unsupervised deep metric learning with transformed attention consistency and contrastive clustering loss.

3 Method

3.1 Overview

Suppose that we have a set of unlabeled images $\mathcal{X} = \{x_1, x_2, \dots, x_N\}$. Our goal is to learn a deep neural network to extract their features $\mathbf{G}(x_n) \in \mathbb{R}^d$, where d is the feature dimension. Fig. 2 shows the overall design of our proposed method for unsupervised deep metric learning based on transformed attention consistency and contrastive clustering loss (TAC-CCL). Given an input image x_n , we apply a transform T , which is randomly sampled from a set of image transforms \mathcal{T} , to x_n , to obtain its transformed version $x'_n = T(x_n)$. In our experiments, we mainly consider spatial transforms, including cropping (sub-image), rotation, zooming, and perspective transform. Each transform is controlled by a set of transform parameters. For example, the cropping is controlled by its bounding box. The perspective transform is controlled by its 6 parameters. Image pairs (x_n, x'_n) are inputs to the Siamese deep neural network. These two identical networks will be trained to extract features F_n and F'_n for these two images. As illustrated in Fig. 2, each network is equipped with an attention network to learn the attention map which will modulate the output feature map. The attention map can enforce the network to focus on discriminative local features for the specific learning tasks or datasets. Let M_n and M'_n be the attention maps for images x_n and x'_n , respectively. According to the transformed attention consistency, we

shall have

$$M'_n = T(M_n). \quad (1)$$

Based on this constraint, we introduce the transformed attention consistency loss L_{TAC} to train the feature embedding network \mathbf{G} , which will be further explained in Section 3.3. Besides this attention consistency, we also require that the output features F_n and F'_n should be similar to each other since the corresponding input images x_n and x'_n are visually the same. To enforce this constraint, we introduce the feature similarity loss $L_F = \|F_n - F'_n\|_2$ which is the L_2 -normal between these two features. To ensure that image features from the same class aggregate into compact clusters while image features from different classes are pushed away from each other in the high-dimensional feature space, we introduce the contrastive clustering loss L_{CC} , which will be further explained in the Section 3.3.

3.2 Baseline System

In this work, we first design a baseline system. Recently, a method called multi-similarity (MS) loss [35] has been developed for supervised deep metric learning. In this work, we adapt this method from supervised metric learning to unsupervised metric learning using k -means clustering to assign pseudo labels. Also, the original MS method computes the similarity scores between image samples in the current mini-batch. In this work, we extend this similarity analysis to the whole training set using the approach of memory bank [36]. The features of all training samples generated by the network are stored in the memory bank by the enqueue method. When the memory bank is full, the features and corresponding labels of the oldest mini-batch are removed by the dequeue method. Using this approach, the current mini-batch has access to the whole training set. We can then compute the similarity scores between all samples in the mini-batch and all samples in the training set. Our experiments demonstrate that this enhanced similarity matrix results in significantly improved performance in unsupervised metric learning. In this work, we use this network as the baseline system, denoted by TAC-CLL (baseline).

3.3 Loss Functions

To further improve the performance of the baseline system, we introduce the ideas of transformed attention consistency and contrastive clustering loss, which are explained in the following.

The transformed attention consistency aims to enforce the feature embedding network \mathbf{G} to focus visually important features instead of other background noise. Let $M_n(u, v)$ and $M'_n(u, v)$ be the attention maps for input image pair x_n and x'_n , where (u, v) represents a point location in the attention map. Under the transform T , this point is mapped to a new location denoted by $(T_u(u, v), T_v(u, v))$. According to (1), if we transform the attention map $M_n(u, v)$ for the original image x_n by T , it should match the attention map $M'_n(u, v)$ for

the transformed image $x'_n = T(x_n)$. Based on this, the proposed transformed attention consistency loss L_{TAC} is defined as follows

$$L_{TAC} = \sum_{(u,v)} |M_n(u,v) - M'_n(T_u(u,v), T_v(u,v))|^2, \quad (2)$$

where $u' = T_u(u,v)$ and $v' = T_v(u,v)$ are the mapped location of (u,v) in image x'_n .

The contrastive clustering loss extends the triplet loss [30] developed in supervised deep metric learning, where an anchor sample x is associated with a positive sample x_+ and a negative sample x_- . The triplet loss aims to maximize the ratio $S(x, x_+)/S(x, x_-)$, where $S(\cdot, \cdot)$ represents the cosine similarity between two features. It should be noted that this triplet loss requires the knowledge of image labels, which however are not available in our unsupervised case. To extend this triplet loss to unsupervised metric learning, we propose to cluster the image features into K clusters. In the high-dimensional feature space, we wish these clusters are compact and are well separated from each other by large margins. Let $\{C_k\}$, $1 \leq k \leq K$, be the cluster centers. Let $C_+(F_n)$ be the nearest center which has the minimum distance to the input image feature F_n and the corresponding distance is denoted by $d_+(F_n) = \|F_n - C_+(F_n)\|_2$. Let $C_-(F_n)$ be the cluster center which has the second minimum distance to F_n and the corresponding distance is denoted by $d_-(F_n) = \|F_n - C_-(F_n)\|_2$. If the contrastive ratio of $d_+(F_n)/d_-(F_n)$ is small, then this feature has more discriminative power. We define the following contrastive clustering loss

$$L_{CC} = \mathcal{E}_{F_n} \left\{ \frac{\|F_n - C_+(F_n)\|_2}{\|F_n - C_-(F_n)\|_2} \right\}, \quad (3)$$

which is the average contrastive ratio of all input image features. During the training process, the network \mathbf{G} , as well as the feature for each input, is progressively updated. For example, the clustering is performed and the cluster centers are updated for every 20 epochs.

3.4 Transformed Attention Consistency with Cross-Images Supervision

Note that, in our proposed approach, we transform or augment the input image x_n to create its pair x'_n . These two are from the same image source. We also notice that most of existing self-supervision methods, such as predicting locations of image patches and classifying the rotation of an image [10], and reconstructing the transform of the image [44], all focus on self-supervision information within the image itself. The reason behind this is that image patches from the same image will automatically have the same class label. This provides an important self-supervision constraint to train the network. However, this one-image approach will limit the learning capability of the network since the network is not able to compare multiple images. As we know, when human eyes are examining images to determine which features are discriminative, they need to compare

multiple images to determine which set of features are consistent across images and which set of features are background noise [9]. Therefore, in unsupervised learning, it is highly desirable to utilize the information across images.

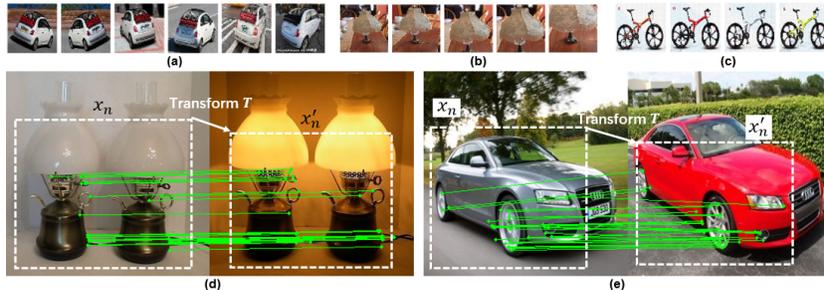


Fig. 3. Sub-image matching for cross-image supervision.

Figs. 3(a)-(c) show image samples from the Cars and SOP benchmark datasets. We can see that images from the same class exhibit strong similarity between images, especially in the object regions. The question is how to utilize these unlabeled images to create reliable self-supervision information for unsupervised learning? In this work, we propose to perform keypoint or sub-image matching across images. Specifically, as illustrated in Fig. 3(d) and (e), for a given image sample I_n , in the pre-processing stage, we perform affine-SIFT [24] keypoint matching between I_n and other images in the dataset and find the top matches with confidence scores about a very high threshold. We then crop out the sub-images containing high-confidence keypoints as x_n and x'_n which are related by a transform T . This high-confidence constraint aims to ensure that x_n and x'_n are having the same object class or semantic label. In this way, for each image in the k -means cluster, we can find multiple high-confidence matched sub-images. For example, on the CUB dataset, for top-2 matching in each cluster, the label accuracy is 77.0%, which is much higher than the true positive pair rate obtained by k -means clustering (39.1%). This will significantly augment the training set, establish cross-image self-supervision, and provide significantly enhanced visual diversity for the network to learn more robust and discriminative features. In this work, we combine this cross-image supervision with the transformed attention consistency. Let $\{(u_i, v_i)\}$ and $\{(u'_i, v'_i)\}$, $1 \leq i \leq N$, be the set of matched keypoints in x_n and x'_n . We wish that, within the small neighborhoods of these matched keypoints, the attention maps M_n and M'_n are consistent. To define a small neighborhood around a keypoint (u_i, v_i) , we use the following 2-D Gaussian kernel,

$$\phi(u - u_i, v - v_i) = e^{-\frac{(u - u_i)^2}{2\sigma_u^2} - \frac{(v - v_i)^2}{2\sigma_v^2}}. \quad (4)$$

Let

$$\Gamma(u, v) = \sum_{i=1}^M \phi(u - u_i, v - v_i), \quad \Gamma'(u, v) = \sum_{i=1}^M \phi(u - u'_i, v - v'_i), \quad (5)$$

which define two masks to indicate the neighborhood areas around these matched keypoints in these two attention maps. The extended transformed attention consistency becomes

$$L_{TAC} = \sum_{(u,v)} |M_n(u, v) \cdot \Gamma(u, v) - M'_n(u, v) \cdot \Gamma'(u, v)|^2, \quad (6)$$

which compares the difference between these two attention maps around these matched keypoints. Compared to the label propagation method developed for semi-supervised learning [45, 46], our cross-image supervision method is unique in the following aspects: (1) it discovers sub-images of the same label (with very high probability) from unlabeled images. (2) It establishes the transform between these two sub-images and combines with the transformed attention consistency to achieve efficient unsupervised deep metric learning.

4 Experimental Results

In this section, we conduct extensive experiments on benchmark datasets in image retrieval settings to evaluate the proposed TAC-CCL method for unsupervised deep metric learning.

4.1 Datasets

We follow existing papers on unsupervised deep metric learning [42] to evaluate our proposed methods on the following three benchmark datasets. **(1) CUB-200-2011 (CUB)** [34] is composed of 11,788 images of birds from 200 classes. The first 100 classes (5864 images) are used for training, with the rest 100 classes (5924 images) for testing. **(2) Cars-196 (Cars)** [21] contains 16,185 images of 196 classes of car models. We use the first 98 classes with 8054 images for training, and remaining 98 classes (8131 images) for testing. **(3) Stanford Online Products (SOP)** [27] has 22,634 classes (120,053 images) of online products. We use the first 11,318 products (59,551 images) for training and the remaining 11,316 products (60,502 images) for testing. The training classes are separated from the test classes. We use the standard image retrieval performance metric (Recall@K), for performance evaluations and comparisons.

4.2 Implementation Details

We implement our proposed method by PyTorch and follow the standard experimental settings in existing papers [27, 35, 42] for performance comparison. We

use the same GoogLeNet [32] pre-trained on ImageNet as the backbone network [25, 27, 31] and a CBAM [39] attention module is placed after the *inception_5b* layer. A fully connected layer is then added on the top of the network as the embedding layer. The default dimension of embedding is set as 512. For the clustering, we set the number of clusters K to be 100 for the CUB and Cars datasets, and $K = 10000$ for the SOP dataset. For each batch, we follow the data sampling strategy in multi-similarity loss [35] to sample 5 images per class. For data augmentation, images in the training set are randomly cropped at size 227×227 with random horizontal flipping, while the images in testing set is center cropped. Adam optimizer [20] is used in all experiments and the weigh decay is set as $5e^{-4}$.

Table 1. Recall@ K (%) performance on CUB and Cars datasets in comparison with other methods.

Methods	Backbone	CUB				Cars			
		R@1	R@2	R@4	R@8	R@1	R@2	R@4	R@8
<i>Supervised Methods</i>									
ABIER [28]	GoogLeNet	57.5	68.7	78.3	86.2	82.0	89.0	93.2	96.1
ABE [19]	GoogLeNet	60.6	71.5	79.8	87.4	85.2	90.5	94.0	96.1
Multi-Similarity [35]	BN-Inception	65.7	77.0	86.3	91.2	84.1	90.4	94.0	96.5
<i>Unsupervised Methods</i>									
Exemplar [7]	GoogLeNet	38.2	50.3	62.8	75.0	36.5	48.1	59.2	71.0
NCE [41]	GoogLeNet	39.2	51.4	63.7	75.8	37.5	48.7	59.8	71.5
DeepCluster [2]	GoogLeNet	42.9	54.1	65.6	76.2	32.6	43.8	57.0	69.5
MOM [17]	GoogLeNet	45.3	57.8	68.6	78.4	35.5	48.2	60.6	72.4
Instance [42]	GoogLeNet	46.2	59.0	70.1	80.2	41.3	52.3	63.6	74.9
TAC-CCL (baseline)	GoogLeNet	53.9	66.2	76.9	85.8	43.0	53.8	65.3	76.0
TAC-CCL	GoogLeNet	57.5	68.8	78.8	87.2	46.1	56.9	67.5	76.7
Gain		+11.3	+9.8	+8.7	+7.0	+4.8	+4.6	+3.9	+1.8

4.3 Performance Comparisons with State-of-the-Art Methods

We compare the performance of our proposed methods with the state-of-the-art unsupervised methods on image retrieval tasks. The mining on manifolds (MOM) [17] and the invariant and spreading instance feature method (denoted by Instance) [42] are current state-of-the-art methods for unsupervised metric learning. They both use the GoogLeNet [32] as the backbone encoder. In the Instance paper [42], the authors have also implemented three other state-of-the-art methods originally developed for feature learning and adapted them to unsupervised metric learning tasks: Exemplar [7], NCE (Noise-Contrastive Estimation) [41], and DeepCluster [2]. We include the results of these methods for comparisons. We have also included the performance of recent supervised deep metric learning methods for comparison so that we can see the performance difference

between unsupervised metric learning and supervised one. These methods include: ABIER [28], and ABE [19], and MS (Multi-Similarity) [35]. Both ABIER and ABE methods are using the GoogLeNet as the backbone encoder. The MS method is using the BN-Inception network [16] as the backbone encoder.

The results for the CUB, Cars, and SOP datasets are summarized in Tables 1 and 2, respectively. We can see that our proposed TAC-CCL method achieves new state-of-the-art performance in unsupervised metric learning on both fine-grained CUB and Cars datasets and the large-scale SOP dataset. On the CUB dataset, our TAC-CCL improves the Recall@1 by 11.3% and is even competitive to some supervised metric learning methods, e.g., ABIER [28]. On the Cars dataset, our TAC-CCL outperforms the current state-of-the-art Instance method [42] by 4.8%. On SOP, our method achieves 63.9% and outperforms existing methods by a large margin of 15%. For other Recall@ K rates with large values of k , the amount of improvement is also very significant. Note that our baseline system achieves a large improvement over existing methods. The proposed TAC-CCL approach further improves upon this baseline system by another 1.4-3.6%.

Fig. 4 shows examples of retrieval results from the CUB, Cars, and SOP datasets. In each row, the first image highlighted with a blue box is the query image. The rest images are the top 15 retrieval results. Images highlighted with red boxes are from different classes. It should be noted that some classes have very small number of samples. We can see that our TAC-CCL can learn discriminative features to achieve satisfying retrieval results, even for these challenging tasks. For example, at the first row of the SOP dataset, our model is able to learn the glass decoration feature under the lampshade, which is a unique feature of the query images. In addition, the negative retrieved results are also visually closer to the query images.

Table 2. Recall@ K (%) performance on SOP dataset in comparison with other methods.

Methods	Backbone	SOP		
		R@1	R@10	R@100
<i>Supervised Methods</i>				
ABIER [28]	GoogLeNet	74.2	86.9	94.0
ABE [19]	GoogLeNet	76.3	88.4	94.8
Multi-Similarity [35]	BN-Inception	78.2	90.5	96.0
<i>Unsupervised Methods</i>				
Exemplar [7]	GoogLeNet	45.0	60.3	75.2
NCE [41]	GoogLeNet	46.6	62.3	76.8
DeepCluster [2]	GoogLeNet	34.6	52.6	66.8
MOM [17]	GoogLeNet	43.3	57.2	73.2
Instance [42]	GoogLeNet	48.9	64.0	78.0
TAC-CCL (baseline)	GoogLeNet	62.5	76.5	87.2
TAC-CCL	GoogLeNet	63.9	77.6	87.8
	Gain	+15.0	+13.6	+9.8

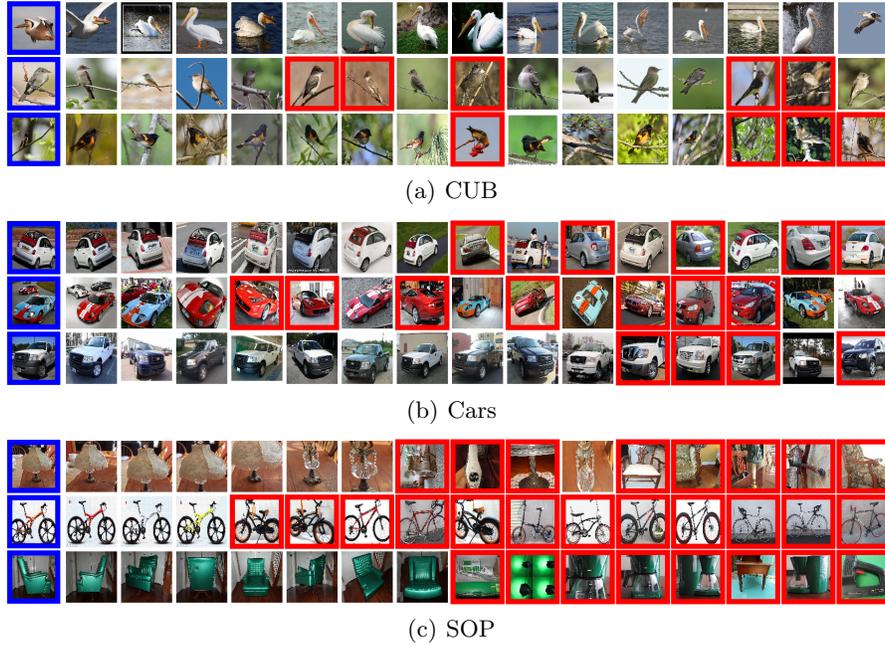


Fig. 4. Retrieval results of some example queries on CUB, Cars, and SOP datasets. The query images and the negative retrieved images are highlighted with blue and red.

4.4 Ablation Studies

In this section, we conduct ablation studies to perform in-depth analysis of our proposed method and its different components.

(1) Impact of the number of clusters. The proposed contrastive clustering loss is based on clustering in the feature space. The number of clusters K is a critical parameter for the proposed method since it determines the number of pseudo labels. We conduct the following ablation study experiment on the CUB data to study the impact of K . The first plot in Fig. 5(a) shows the Recall@1 results with different values of K : 50, 100, 200, 500, and 1000. The other three plots show the results for Recall@2, 4, and 8. We can see that, on this dataset, the best value of K is 100, which is the number of test classes in the CUB dataset. The performance drops when K increases. This study suggests that the best value of K is close to the truth number test classes of the dataset.

(2) Impact of different embedding sizes. In this ablation study, we follow existing supervised metric learning methods [35, 28] to study the impact of different embedding sizes, or the size of the embedded feature. For example, the feature size ranges from 64, 128, 256, 512, to 1024. The first plot of Fig. 5(b) shows the Recall@1 results for different embedding size. The results for

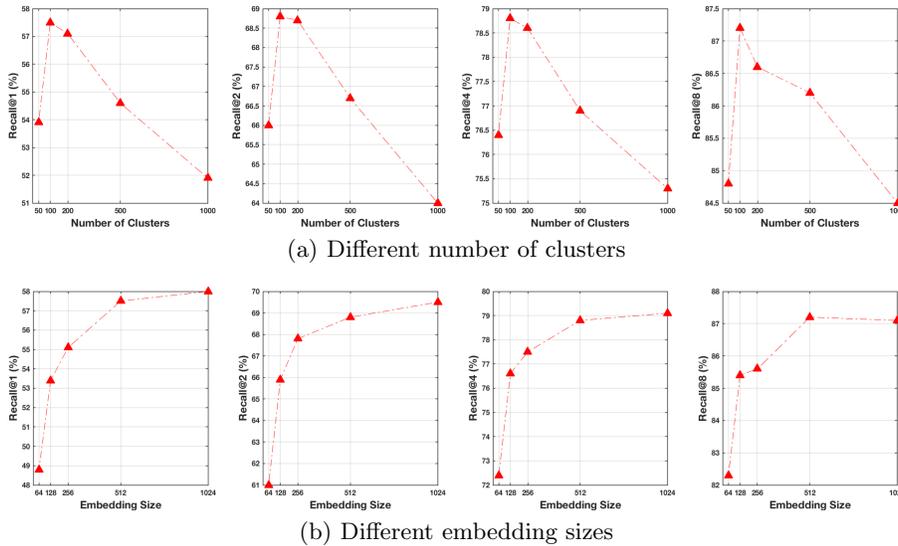


Fig. 5. Recall@ K (%) performance on CUB dataset in comparison with different number of clusters and different embedding size.

Recall@2, 4, 8 are shown in the other three plots. We can see that unsupervised metric learning performance increases with the embedding size since it contains more feature information with enhanced discriminative power.

Table 3. Recall@ K (%) performance on SOP dataset using Resnet-18 network without pre-trained parameters.

Methods	SOP		
	R@1	R@10	R@100
Random	18.4	29.4	46.0
Exemplar [7]	31.5	46.7	64.2
NCE [41]	34.4	49.0	65.2
MOM [17]	16.3	27.6	44.5
Instance [42]	39.7	54.9	71.0
Ours	47.0	62.6	77.5

(3) Impact of the pre-trained model. We follow the recent state-of-the-art unsupervised metric learning Instance method [42] and evaluate the performance of our proposed method on the large-scale SOP dataset by using the Resnet-18 network without pre-trained parameters. From Table 3, we can see our proposed method outperforms Instance method [42] by more than 7%.

(4) Performance contributions of different algorithm components.

Our proposed system has three major components: the baseline system for unsupervised deep metric learning, transformed attention consistency (TAC), and contrastive clustering loss (CCL). In this ablation study, we aim to identify the contribution of each algorithm component on different datasets. Table 4 summarizes the performance results on the CUB, Cars, and SOP datasets using three different method configurations: (1) the baseline system, (2) baseline with CCL, and (3) baseline with CCL + TAC. We can see that both the CCL and TAC approaches significantly improve the performance.

In our Supplementary Material, we will provide additional method implementation details, experimental results, and ablation studies.

Table 4. The performance of different components from our TAC-CCL method on CUB, Cars, and SOP datasets.

	CUB				Cars				SOP		
	R@1	R@2	R@4	R@8	R@1	R@2	R@4	R@8	R@1	R@10	R@100
Baseline	53.9	66.2	76.9	85.8	43.0	53.8	65.3	76.0	62.5	76.5	87.2
+CCL	55.7	67.8	77.5	86.2	44.7	55.6	65.9	75.7	63.0	76.8	87.2
+TAC	57.5	68.8	78.8	87.2	46.1	56.9	67.5	76.7	63.9	77.6	87.8

5 Conclusions

In this work, we have developed a new approach to unsupervised deep metric learning based on image comparisons, transformed attention consistency, and contrastive clustering loss. This transformed attention consistency leads to a pairwise self-supervision loss, allowing us to learn a Siamese deep neural network to encode and compare images against their transformed or matched pairs. To further enhance the inter-class discriminative power of the feature generated by this network, we have adapted the concept of triplet loss from supervised metric learning to our unsupervised case and introduce the contrastive clustering loss. Our extensive experimental results on benchmark datasets demonstrate that our proposed method outperforms current state-of-the-art methods by a large margin.

6 Acknowledgement

This work was supported in part by National Science Foundation under grants 1647213 and 1646065. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Ba, J., Mnih, V., Kavukcuoglu, K.: Multiple object recognition with visual attention. arXiv preprint arXiv:1412.7755 (2014)
2. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 132–149 (2018)
3. Chen, B., Deng, W.: Hybrid-attention based decoupled metric learning for zero-shot image retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2750–2759 (2019)
4. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 1, pp. 539–546. IEEE (2005)
5. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)
6. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1422–1430 (2015)
7. Dosovitskiy, A., Fischer, P., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with exemplar convolutional neural networks. IEEE transactions on pattern analysis and machine intelligence **38**(9), 1734–1747 (2015)
8. Fu, J., Zheng, H., Mei, T.: Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4438–4446 (2017)
9. Gazzaniga, M.S.: The cognitive neurosciences. MIT Press (2009)
10. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. arXiv preprint arXiv:1803.07728 (2018)
11. Grabner, A., Roth, P.M., Lepetit, V.: 3d pose estimation and 3d model retrieval for objects in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3022–3031 (2018)
12. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1735–1742. IEEE (2006)
13. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. arXiv preprint arXiv:1911.05722 (2019)
14. He, X., Zhou, Y., Zhou, Z., Bai, S., Bai, X.: Triplet-center loss for multi-view 3d object retrieval. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1945–1954 (2018)
15. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. arXiv preprint arXiv:1703.07737 (2017)
16. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)
17. Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Mining on manifolds: Metric learning without labels. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7642–7651 (2018)
18. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in neural information processing systems. pp. 2017–2025 (2015)

19. Kim, W., Goyal, B., Chawla, K., Lee, J., Kwon, K.: Attention-based ensemble for deep metric learning. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 736–751 (2018)
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
21. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13). Sydney, Australia (2013)
22. Kumar, V., Harwood, B., Carneiro, G., Reid, I., Drummond, T.: Smart mining for deep metric learning. arXiv preprint arXiv:1704.01285 **2** (2017)
23. Mnih, V., Heess, N., Graves, A., et al.: Recurrent models of visual attention. In: Advances in neural information processing systems. pp. 2204–2212 (2014)
24. Morel, J.M., Yu, G.: Asift: A new framework for fully affine invariant image comparison. SIAM journal on imaging sciences **2**(2), 438–469 (2009)
25. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 360–368 (2017)
26. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision. pp. 69–84. Springer (2016)
27. Oh Song, H., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4004–4012 (2016)
28. Opitz, M., Waltner, G., Possegger, H., Bischof, H.: Deep metric learning with bier: Boosting independent embeddings robustly. IEEE transactions on pattern analysis and machine intelligence (2018)
29. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision **115**(3), 211–252 (2015)
30. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
31. Sohn, K.: Improved deep metric learning with multi-class n-pair loss objective. In: Advances in neural information processing systems. pp. 1857–1865 (2016)
32. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1–9 (2015)
33. Tao, R., Gavves, E., Smeulders, A.W.: Siamese instance search for tracking. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1420–1429 (2016)
34. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
35. Wang, X., Han, X., Huang, W., Dong, D., Scott, M.R.: Multi-similarity loss with general pair weighting for deep metric learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5022–5030 (2019)
36. Wang, X., Zhang, H., Huang, W., Scott, M.R.: Cross-batch memory for embedding learning. arXiv preprint arXiv:1912.06798 (2019)
37. Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: European conference on computer vision. pp. 499–515. Springer (2016)

38. Wohlhart, P., Lepetit, V.: Learning descriptors for object recognition and 3d pose estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3109–3118 (2015)
39. Woo, S., Park, J., Lee, J.Y., So Kweon, I.: Cbam: Convolutional block attention module. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 3–19 (2018)
40. Wu, C.Y., Manmatha, R., Smola, A.J., Krahenbuhl, P.: Sampling matters in deep embedding learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2840–2848 (2017)
41. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3733–3742 (2018)
42. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6210–6219 (2019)
43. Yu, R., Dou, Z., Bai, S., Zhang, Z., Xu, Y., Bai, X.: Hard-aware point-to-set deep metric for person re-identification. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 188–204 (2018)
44. Zhang, L., Qi, G.J., Wang, L., Luo, J.: Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2547–2555 (2019)
45. Zhu, X., Ghahramani, Z.: Learning from labeled and unlabeled data with label propagation. Tech. rep., Citeseer (2002)
46. Zhu, X.J.: Semi-supervised learning literature survey. Tech. rep., University of Wisconsin-Madison Department of Computer Sciences (2005)