

Offline-Online Associated Camera-Aware Proxies for Unsupervised Person Re-Identification

Menglin Wang¹, Jiachen Li, Baisheng Lai, Xiaojin Gong², *Member, IEEE*, and Xian-Sheng Hua³, *Fellow, IEEE*

Abstract—Recently, unsupervised person re-identification (Re-ID) has received increasing research attention due to its potential for label-free applications. A promising way to address unsupervised Re-ID is clustering-based, which generates pseudo labels by clustering and uses the pseudo labels to train a Re-ID model iteratively. However, most clustering-based methods take each cluster as a pseudo identity class, neglecting the intra-cluster variance mainly caused by the change of cameras. To address this issue, we propose to split each single cluster into multiple proxies according to camera views. The camera-aware proxies explicitly capture local structures within clusters, by which the intra-ID variance and inter-ID similarity can be better tackled. Assisted with the camera-aware proxies, we design two proxy-level contrastive learning losses that are, respectively, based on offline and online association results. The offline association directly associates proxies according to the clustering and splitting results, while the online strategy dynamically associates proxies in terms of up-to-date features to reduce the noise caused by the delayed update of pseudo labels. The combination of two losses enables us to train a desirable Re-ID model. Extensive experiments on three person Re-ID datasets and one vehicle Re-ID dataset show that our proposed approach demonstrates competitive performance with state-of-the-art methods. Code will be available at: <https://github.com/Terminator8758/O2CAP>.

Index Terms—Unsupervised person re-identification, contrastive learning, proxies.

I. INTRODUCTION

PERSON re-identification (Re-ID) is the task of identifying the same person in non-overlapping cameras. Due to its significance in video surveillance and public security, this task has been extensively studied for decades. State-of-the-art performance is achieved mostly by supervised methods [1], [2], requiring full labels that are expensive and time-consuming to annotate. Recently, semi-supervised [3], [4] and unsupervised [5], [6] Re-ID have been attracting more and

more research interest, in a hope to reduce annotation cost and make the techniques more practical to real-world deployments. Although considerable progress has been achieved in these tasks, there is still a big gap in performance compared to the supervised counterparts.

This work focuses on the purely unsupervised Re-ID task that requires no labels and thus fully releases the annotation burden. A promising research line for unsupervised Re-ID is clustering-based, which generates pseudo labels from the clustering results and uses the pseudo labels to learn a Re-ID model iteratively. With the popularity of contrastive learning [7], [8], [9], the integration of clustering and contrastive learning demonstrates great potential in recent unsupervised studies [10], [11], [12]. However, the pseudo labels generated by clustering is far from perfect, impeding Re-ID models from getting further improved.

The noise of clustering-based pseudo labels mainly stems from two aspects. On one hand, pedestrian images are of high intra-ID variance and inter-ID similarity, which makes it extremely challenging for clustering algorithms such as DBSCAN [13] or K-Means to achieve accurate clusters. On the other hand, most methods perform clustering in an offline manner based on the features extracted at the beginning of each epoch, while instance features used for matching are extracted via a model updated on the fly. It implies that the generated pseudo labels might be out of date. To deal with the inevitable label noise, efforts have been made in label refinement [14], [15], [16], hybrid contrastive learning [10], [12], [17] that combines cluster- and instance-level contrasts together, online pseudo label generation [18], and other techniques [19], [20].

In this work, we propose camera-aware proxies to better deal with the intra-ID variance and inter-ID similarity. Our approach is inspired by the following observation. Since severe intra-ID variance is mainly caused by the change of camera views, an ID's intra-camera images tend to gather more tightly in a feature space than its inter-camera images. As a result, the clusters obtained by unsupervised clustering often present multiple sub-clusters, roughly corresponding to different camera views. Therefore, as shown in Figure 1, we propose to split each cluster, which is obtained by a camera-agnostic clustering method, into multiple camera-aware proxies according to camera views. Based on the pseudo labels that are generated from these camera-aware proxies, we design a proxy-level contrastive learning (CL) method that is more effective than cluster-level CL and requires less memory footprints than instance-level CL methods [10], [12], [17].

Manuscript received 13 January 2022; revised 26 June 2022 and 12 August 2022; accepted 28 September 2022. Date of publication 14 October 2022; date of current version 20 October 2022. This work was supported by the Major Scientific Research Project of Zhejiang Laboratory, China, under Grant 2019DB0ZX01. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Tao Mei. (*Corresponding author: Xiaojin Gong.*)

Menglin Wang is with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China (e-mail: lynnwang6875@gmail.com).

Jiachen Li and Xiaojin Gong are with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310027, China, and also with the Zhejiang Provincial Key Laboratory of Information Processing, Communication and Networking (IPCAN), Hangzhou 310027, China (e-mail: lijachen_isee@zju.edu.cn; gongxj@zju.edu.cn).

Baisheng Lai and Xian-Sheng Hua are with the Alibaba Group, Hangzhou 311121, China (e-mail: baisheng.lbs@alibaba-inc.com; huaxiansheng@gmail.com).

Digital Object Identifier 10.1109/TIP.2022.3213193

1941-0042 © 2022 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

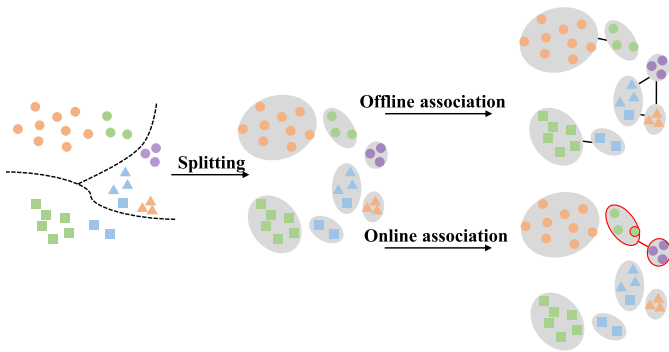


Fig. 1. An illustration of the proposed camera-aware proxies, together with offline and online associations. Each cluster obtained by a camera-agnostic clustering step is split into multiple camera-aware proxies according to camera views. Proxies are further associated via an offline strategy and an online strategy to perform proxy-level contrastive learning. The offline strategy straightforwardly associates proxies split from the same cluster together. The online strategy dynamically associates positive proxies for an anchor instance (marked by red circles) in terms of up-to-date features. (Instances belonging to the same ID are represented by the same shape and those coming from the same camera are in the same color).

To deal with the noise arisen from the offline clustering criterion as well as the delayed update of pseudo labels, we further propose a strategy that employs offline and online associations to complementarily mine positive and hard negative proxies. Offline association is based on the offline clustering results. It straightforwardly associates the proxies split from a positive cluster as the positive ones. Online association aims to utilize up-to-date features of instances and proxies to associate positive proxies. To this end, we design an instance-proxy balanced similarity and a camera-aware nearest neighbor criterion to produce reliable associations, which are vital to make online association effective. Finally, two proxy-level contrastive learning losses are defined, respectively, according to the results of offline association and online association.

The main contributions are summarized as follows:

- Instead of using camera-agnostic clusters, we propose camera-aware proxies that explicitly capture local structures within clusters. They enable us to perform proxy-level contrastive learning, which can tackle the intra-ID variance and inter-ID similarity better.
- With the assistance of the camera-aware proxies, we design a strategy that combines offline and online association based proxy-level contrastive learning together. This strategy can, to some extent, conquer the noise caused by the offline clustering criterion and the delayed update of pseudo labels.
- Extensive experiments on three person Re-ID datasets and one vehicle Re-ID dataset show that the proposed method is competitive to state-of-the-art purely unsupervised and UDA-based Re-ID methods.

Note that this work is an extended version of our preliminary work CAP [11]. In contrast to CAP, we make the following improvements:

- This work proposes an online association scheme to utilize up-to-date features for proxy association, based on which an additional proxy-level contrastive learning loss is designed.

- We discard the intra-camera contrastive learning loss that is used in CAP and analyze why this loss is not helpful to the new model in this work.
- More experiments on ablation studies and sensitivity analysis, as well as on a vehicle Re-ID dataset, are conducted to make a thorough validation.
- Experiments show that this work improves CAP by a considerable margin on all datasets. Especially, on the most challenging dataset MSMT17, 4.6% Rank-1 and 5.5% mAP improvements are gained.

II. RELATED WORK

A. Unsupervised Person Re-ID

According to whether external labeled datasets are used or not, previous unsupervised methods can be classified into purely unsupervised and UDA-based groups.

Purely unsupervised person Re-ID requires no annotations and therefore is more attractive. Existing methods commonly resort to pseudo labels for learning. Clustering [5], [20], k-NN [21], [22], graph [23], [24], or hypergraph [15] based techniques have been developed to generate pseudo labels. Clustering-based methods such as BUC [5] and HCT [20] conduct learning in a camera-agnostic way, which can capture the similarity within IDs but neglect the intra-ID variance mostly caused by the change of camera views. Alternatively, TAUDL [21], DAL [22], UGA [24], and IICS [25] divide the Re-ID task into intra- and inter-camera learning stages, by which the discriminative ability learned within cameras can further facilitate ID association across cameras. Our preliminary CAP [11] proposes camera-aware proxies to deal with the intra-ID variance and conducts the unsupervised learning also from both intra- and inter-camera perspectives. But, in this extension work we find out that the intra-camera learning is not necessary, or even harmful, to perform effective learning when intra-camera pseudo labels are noisy. We therefore merely focus on the inter-camera learning while propose a combination of offline and online association to boost performance.

Unsupervised domain adaptation (UDA) based person Re-ID demands a source dataset that is fully labeled, but leaves the target dataset unlabeled. To address this task, existing methods either transfer image styles [26], [27], [28], [29] or reduce distribution discrepancy [30], [31], [32] between different domains. These methods pay much attention to transfer knowledge from source domain to target domain. In addition, to sufficiently exploit unlabeled data in the target domain, clustering [10], [32], [33], [34], [35] or k-NN [6], [29] based methods have also been adopted, analogous to those introduced in the purely unsupervised task. Differently, these methods either take both original and transferred data [6], [10], [29], [33] into account, or integrate a clustering procedure together with an adversarial learning step [34]. Although external labeled datasets are used, UDA-based methods do not gain noticeable advantage against recent purely unsupervised counterparts.

In both purely unsupervised and UDA-based person Re-ID, early methods [21], [26], [27], [33] often use ID classification loss or triplet loss [36] for learning. Recently, contrastive

learning (CL) [7], [8], [9] has attracted a surge of research interest. Various CL-based methods [6], [10], [11], [12], [18], [32], [35], [37], [38] have been developed for unsupervised Re-ID. Moreover, the issue of pseudo label noise occurred in unsupervised Re-ID has also been noticed recently and various methods [14], [16], [18], [19], [20] have been proposed to address it. We will introduce these related work and state the difference of our work in Section II-C and II-E.

B. Utilization of Camera Information

As one type of valuable meta information, camera index has been extensively utilized in previous unsupervised Re-ID methods. For instance, CAMEL [39] and UCDA-CCE [30] aim to align distributions under different cameras by learning camera-specific projections or performing camera-aware domain adaptation. TAUDL [21], UTAL [40], UGA [24], IICS [25], and MGH [15] combine intra-camera and inter-camera learning together to boost unsupervised Re-ID performance. Besides, we notice that camera information is also implicitly used in SpCL [10]. Although SpCL [10] does not explicitly use camera information in its model, it designs a MultiGallerySampler to ensure in-batch images of each ID to be sampled from diverse cameras, which enhances its performance further. These methods leverage camera information in various ways, but they do not pay attention to the intra-cluster variance caused by camera discrepancy. In this work, we propose camera-aware proxies to capture the camera-specific variance within each cluster, based on which offline and online association strategies are designed.

C. Pseudo Label Refinement

Clustering-based methods take a dominant role in recent unsupervised Re-ID. However, the pseudo labels generated by an offline clustering technique such as DBSCAN [13] inevitably contain noise. In order to reduce the label noise, various methods [14], [16], [18], [19], [20] have been developed recently. For instance, RLCC [16] refines pseudo labels via clustering consensus over consecutive training generations. MGH [15] adopts a hypergraph to propagate and refine pseudo labels generated by global DBSCAN, intra-camera DBSCAN, and global KNN. HCT [20] utilizes hierarchical clustering and Zheng et al. [18] propose an online clustering strategy to generate high quality pseudo labels. MMT [14] refines pseudo labels by the combined use of offline hard pseudo labels and online soft pseudo labels. In contrast to these methods, we propose camera-aware proxies (CAP) to conquer the clustering noise resulted from high complex data structures, and further design offline and online association strategies specific to CAP to deal with the noise stemmed from the clustering criterion as well as the delayed update of pseudo labels.

D. Metric Learning With Proxies

Metric learning plays an important role in person Re-ID and many other vision tasks. A loss extensively utilized in metric learning is the triplet loss [36]. It measures the distances of an anchor to a positive instance and a negative instance that are

usually sampled within a batch. Along with the increment of the batch size, the number of triplets increases dramatically, resulting in a slow convergence and an inferior performance. To overcome this issue, Proxy-NCA [41] and Center loss [42] propose to use proxies or centers, which represent sets of data instances, for the measurement of similarity and dissimilarity. The use of proxies instead of instances is able to capture more contextual information and greatly reduces the total number of triplets, enabling metric learning to achieve better performance. Further, with the awareness of intra-class variances, Magnet [43], MaPML [44], SoftTriple [45], GEORGE [46], and ProxyGML [47] adopt multiple proxies to represent a single cluster, by which local structures can be represented better. Our work is inspired by these studies. However, in contrast to set a fixed number of proxies for each class or design a complex strategy to find an appropriate proxy number, we split a cluster into a variant number of proxies simply according to the involved camera views, making our proxies more suitable for the Re-ID task.

E. Contrastive Learning

Contrastive learning (CL) performs learning also via comparing similarities of samples and therefore belongs to a special metric learning technique. In recent years, CL has attracted great interest due to its success in unsupervised representation learning tasks [7], [8], [9]. Some losses such as Instance loss [48] adopt a parametric form to conduct instance-level discrimination, while most typical contrastive losses such as InfoNCE [7], [8], [9], [49] perform in a non-parametric way, aiming to pull positive samples together while push negative samples apart. The losses are primitively focused on the contrast of instances and later extended to prototypes or proxies [50], [51] as well. A crucial problem in CL is how to select positive and negative instances/proxies for effective comparison. Nowadays, this problem still remains open although various hard negative mining strategies [52], [53] for contrastive learning have been proposed.

Recently, contrastive learning has also been widely applied to UDA-based [10], [18], [32], [35] and purely unsupervised [11], [12], [37], [38], [54] person Re-ID tasks. For instance, SpCL [10] constructs a hybrid memory and executes contrastive learning at source-domain class-level, target-domain cluster-level, and un-clustered instance-level. Our CAP [11] constructs a proxy-level memory bank to perform intra- and inter-camera contrastive learning at proxy-level. Later on, ICE [12] and Liu et al. [37] boost SpCL and CAP via augmenting the models with instance-level contrastive learning, while MGH [15] and Isobe et al. [32] integrate contrastive learning with hypergraph and Fourier augmentation. In contrast to them [12], [15], [17], [32], [37], [38], this work extends CAP via sticking on the proxy-level contrastive learning alone to keep the approach simple yet effective.

III. A CLUSTERING-BASED RE-ID BASELINE

We first set up a baseline model for the purely unsupervised Re-ID task. As the common practice in clustering-based methods [5], [20], [33], the baseline learns a Re-ID

model iteratively and, at each epoch, it alternates between a clustering step and a model updating step. In contrast to previous methods that utilize an ID classification loss [33] or a triplet loss [20], we adopt a non-parametric Softmax loss [7] for the model updating. This non-parametric loss, also termed as an InfoNCE loss [9], [55], plays an important role in recent contrastive learning techniques that have been successfully applied to various unsupervised learning tasks [5], [6], [7], [8], [9], [10]. It also makes our baseline model effective and extensible. In the followings, we briefly introduce the details.

Given an unlabeled dataset $\mathcal{D} = \{x_i\}_{i=1}^N$, where x_i is the i -th image and N is the total number of images. We build the baseline model on a convolutional neural network (CNN) f_θ that is parameterized by θ . The parameters are initialized from an ImageNet-pretrained [56] network. When image x is input, the network extracts a d -dimensional feature $f_\theta(x)$. Then, at each epoch, we adopt DBSCAN [13] to cluster the features of all images, and further select reliable clusters by simply discarding isolated outliers. All images within each cluster are assigned with a same pseudo ID label. By this means, we get a labeled dataset $\mathcal{D}' = \{(x_i, \tilde{y}_i)\}_{i=1}^{N'}$, in which $\tilde{y}_i \in \{1, \dots, Y\}$ is a generated pseudo label. N' is the number of images remained in the selected clusters and Y is the number of clusters.

Once pseudo labels are generated, we adopt the contrastive learning technique for model updating. It is implemented via an external memory bank and an InfoNCE loss. Specifically, we construct a cluster-level memory bank $\mathcal{K} \in R^{d \times Y}$. During back-propagation, when image x_i is input, we update the memory entry of its target ID class \tilde{y}_i via a moving average scheme. That is,

$$\mathcal{K}[\tilde{y}_i] \leftarrow \mu \mathcal{K}[\tilde{y}_i] + (1 - \mu) f_\theta(x_i), \quad (1)$$

where $\mathcal{K}[\tilde{y}_i]$ is the \tilde{y}_i -th entry of the memory bank, storing the updated feature centroid of class \tilde{y}_i , and $\mu \in [0, 1]$ is an updating rate.

Then, the loss of the baseline model is defined by

$$\mathcal{L}_{Base} = - \sum_{i=1}^B \log \frac{\exp(\mathcal{K}[\tilde{y}_i]^T f_\theta(x_i)/\tau)}{\sum_{j=1}^Y \exp(\mathcal{K}[j]^T f_\theta(x_i)/\tau)}, \quad (2)$$

where τ is a temperature factor and B is the batch size. This loss is a cluster-level contrastive learning loss, which aims to pull an instance close to the centroid of its class while push it away from the centroids of all other classes.

Considering that hard negative mining strategies may boost the performance of contrastive learning [52], [53], we additionally define an alternative loss of the baseline model as follows:

$$\mathcal{L}_{Base2} = - \sum_{i=1}^B \log \frac{\exp(\mathcal{K}[\tilde{y}_i]^T f_\theta(x_i)/\tau)}{\sum_{j \in \{\tilde{y}_i\} \cup Q_0} \exp(\mathcal{K}[j]^T f_\theta(x_i)/\tau)}, \quad (3)$$

where Q_0 contains the memory indexes of the hard negatives that are sampled via selecting the K_1 -nearest negative clusters.

IV. THE CAMERA-AWARE PROXY ASSISTED METHOD

Like previous clustering-based methods [5], [20], [33], [34], the above-mentioned baseline model conducts the clustering

and model updating steps in a camera-agnostic way. This way is able to maintain the similarity within each cluster, but may neglect the intra-cluster variance. Considering that severe intra-cluster variance is mainly caused by the change of camera views, we split each single cluster into multiple camera-specific proxies. Each proxy represents the instances coming from the same camera. The obtained camera-aware proxies provide us with an explicit way to deal with the variance within clusters. Besides, the proxies enable the contrastive learning to pay more attention to the hardest negative instances, which helps to reduce the inter-ID similarity. Therefore, with the assistance of the proxies, we design two inter-camera contrastive learning losses, which respectively take advantage of offline and online associated proxies, for the model updating. The entire framework is illustrated in Figure 2, in which the modified clustering step and the improved model updating step are alternatively iterated.

More specifically, at each epoch, after the camera-agnostic clustering we split the clusters into camera-aware proxies, and generate a new set of pseudo labels that are assigned in a per-camera manner. That is, the proxies within each camera view are independently labeled. It also means that two proxies split from the same cluster may be assigned with two different labels. We denote the newly labeled dataset of the c -th camera by $\mathcal{D}_c = \{(x_i, \tilde{y}_i, \tilde{z}_i, c_i)\}_{i=1}^{N_c}$. Here, image x_i , which previously is annotated with a global pseudo label \tilde{y}_i , is additionally annotated with an intra-camera pseudo label $\tilde{z}_i \in \{1, \dots, Z_c\}$ and a camera label $c_i = c \in \{1, \dots, C\}$. N_c and Z_c are, respectively, the number of images and proxies in camera c , and C is the camera number. Then, the entire labeled dataset is $\mathcal{D}'' = \bigcup_{c=1}^C \mathcal{D}_c$.

Consequently, we construct a proxy-level memory bank $\mathcal{K}' \in R^{d \times Z}$, where $Z = \sum_{c=1}^C Z_c$ is the total number of proxies in all cameras. Each entry of the memory stores a proxy, which is updated by the same scheme as introduced in Eq. (1) but in proxy-wise. With the proxies stored in the memory bank, we design offline and online association strategies to match the per-camera labeled proxies over all cameras, based on which two contrastive learning losses are proposed.

A. The Contrastive Learning on Offline Associated Proxies

We first design a contrastive learning loss according to the camera-agnostic clustering and camera-aware splitting that are conducted offline. As pointed out in [52], it is crucial to select appropriate positive and negative samples in order to perform effective contrastive learning. Fortunately, our camera-aware splitting strategy provides us with a straightforward way to find positive and negative proxies.

Specifically, given image x_i , we retrieve its positive proxies from all cameras, which share the same global pseudo label \tilde{y}_i . That is, all proxies split from the cluster \tilde{y}_i are associated as positive ones. We refer to this association way as *offline association* because the proxies are associated according to the results of offline clustering and splitting. The memory index set of these retrieved positive proxies is denoted by \mathcal{P}_1 . Besides, we retrieve the K_1 -nearest negative proxies from all

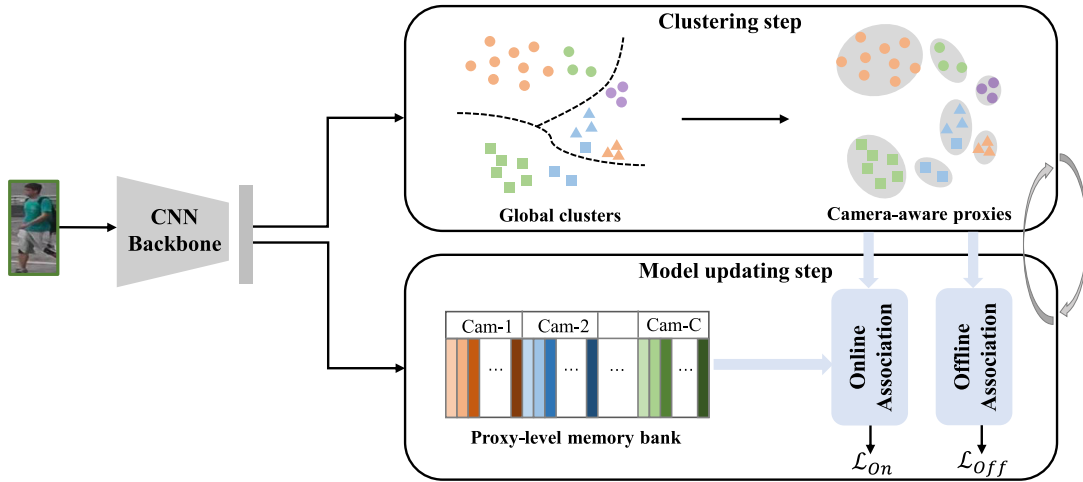


Fig. 2. An overview of our proposed method. It iteratively alternates between a clustering step and a model updating step. In the clustering step, a global clustering is performed and then each cluster is split into multiple camera-aware proxies. In the model updating step, assisted with an external memory bank, two proxy-level contrastive learning losses \mathcal{L}_{off} and \mathcal{L}_{on} , which are respectively based on the results of offline and online associations, are optimized.

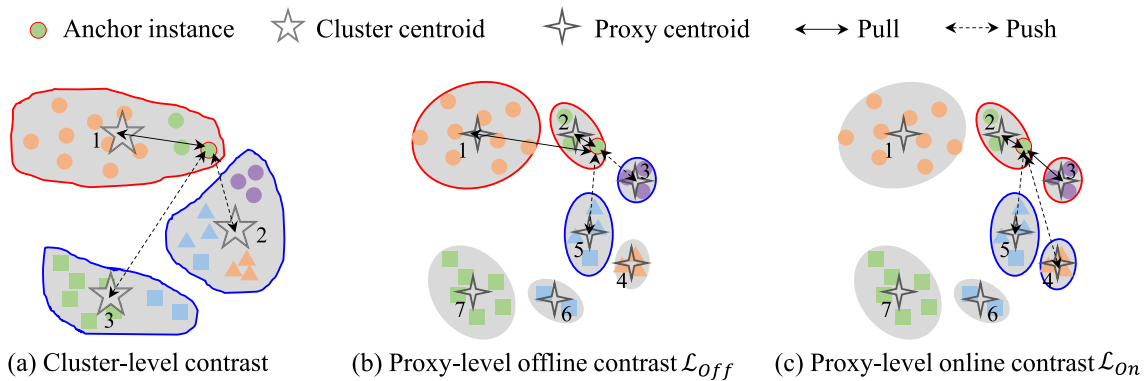


Fig. 3. An illustration of different contrastive learning losses. An anchor instance and its positives are marked with red boundary, and its hard negatives are marked with blue boundaries. (a) The cluster-level contrast loss pulls the anchor instance close to the centroid of the cluster it belongs to, and pushes the anchor away from other cluster centroids. (b) The proxy-level offline contrast loss \mathcal{L}_{off} pulls the anchor instance close to its positive proxies (e.g. Proxy 1 and 2) which are split from the positive cluster, and pushes the anchor away from the sampled hard negative proxies (e.g. Proxy 3 and 5). (c) The online contrast loss \mathcal{L}_{on} pulls the anchor close to the positive proxies (e.g. Proxy 2 and 3) that are associated online, while repels it away from its hard negative proxies (e.g. Proxy 4 and 5). In this toy example, Proxy 3 is a false negative recalled by offline association but it is correctly associated as a positive proxy via online association.

remaining proxies as the hard negative ones, whose memory indexes are recorded by a set \mathcal{Q}_1 . By this means, we define the first contrastive learning loss as follows.

$$\mathcal{L}_{off} = - \sum_{i=1}^B \left(\frac{1}{|\mathcal{P}_1|} \sum_{u \in \mathcal{P}_1} \log \frac{S(u, x_i)}{\sum_{p \in \mathcal{P}_1} S(p, x_i) + \sum_{q \in \mathcal{Q}_1} S(q, x_i)} \right), \quad (4)$$

in which $S(u, x_i) = \exp(\mathcal{K}'[u]^T f_\theta(x_i)/\tau)$, and $|\mathcal{P}_1|$ is the cardinality of \mathcal{P}_1 .

Note that this loss is an inter-camera contrastive learning term as both positive and negative proxies are retrieved across cameras. It in essence maximizes the multiplication of probabilities of x_i being recognized as each positive proxy class. Thus, this loss pulls an instance close to all positive proxy centroids, which encourages a balanced learning for

instance-rich and instance-deficient proxies within each cluster, leading to a high intra-ID compactness. Meanwhile, this loss also pushes the instance away from its hard negative proxies. In contrast to the cluster-level loss defined in Eq. (3), the proxy-level hard negatives can capture local structures at a finer granularity and pay more attention to those hardest negative instances, as shown in Figure 3 (b). Therefore, the proxy-level learning can conquer the inter-ID similarity better.

B. The Contrastive Learning on Online Associated Proxies

The contrastive learning loss defined above still suffers from noise. As shown in Figure 3 (b), this loss may push false negatives away or pull false positives together due to the inaccurate clustering results. Considering that the inaccuracy arises from the density based clustering criterion (DBSCAN) as well as the out-of-date pseudo labels generated by offline clustering the features extracted at the beginning of each epoch, we propose an *online association* strategy that utilizes the nearest neighbor

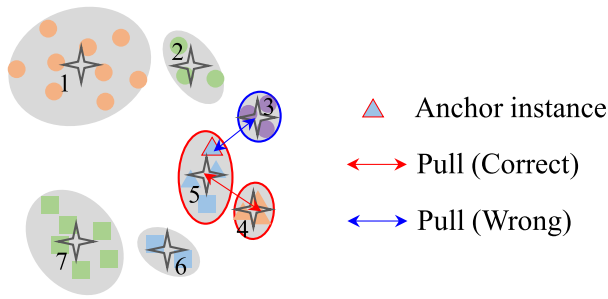


Fig. 4. An example of the instance-proxy balanced similarity. For the anchor instance, the instance-to-proxy similarity tends to pull a false positive proxy (Proxy 3) close, while the proxy-to-proxy similarity pulls a true positive proxy (Proxy 4) together.

criterion and up-to-date information for rectification. More specifically, it takes advantage of instance features extracted via the up-to-date Re-ID model together with the updated proxy entries to dynamically associate positive proxies for each anchor instance. In order to retrieve positive proxies more accurately, we design an *instance-proxy balanced similarity* and a *camera-aware nearest neighbor* criterion for association. Further, we define a contrastive learning loss based on the online association results.

The *instance-proxy balanced similarity* measures the similarity between an instance and a proxy based on the combination of an instance-to-proxy similarity and a proxy-to-proxy similarity. It is defined by

$$\text{sim}(f_{\theta}(x_i), \mathcal{K}'[j]) = wf_{\theta}^T(x_i)\mathcal{K}'[j] + (1-w)\mathcal{K}'[\tilde{z}_i]^T\mathcal{K}'[j]. \quad (5)$$

Here, \tilde{z}_i is the pseudo label of the proxy that x_i belongs to, termed as the self-proxy, and j is the index of any proxy. $w \in [0, 1]$ is a weight to balance the instance-to-proxy and proxy-to-proxy similarities. In contrast to a single instance-to-proxy similarity that is commonly used elsewhere, this balanced similarity is less sensitive to noise. A toy example is illustrated in Figure 4. For the anchor instance marked with a red triangle, the instance-to-proxy similarity tends to associate a false proxy (Proxy 3) as a positive one, while the proxy-to-proxy similarity can associate a positive proxy (Proxy 4) correctly. The balanced similarity is helpful for such scenarios.

The *camera-aware nearest neighbor criterion* is designed to associate at most one positive proxy within each camera. We propose this criterion based on the following observation: an ID's instances coming from the same camera are very likely to be grouped into the same cluster and thus into the same proxy. It implies that, for most anchor instances, there is at most one positive proxy existing in each camera. Therefore, directly using the global KNN, which tends to associate multiple proxies within one camera, may result in false positives. To reduce false positives, we propose the camera-aware nearest neighbor criterion. Specifically, given an instance image x_i , it first selects the 1-nearest neighbor proxy in each camera according to the instance-proxy balanced similarity, and then chooses Top- K_2 proxies from the selected ones as the positive proxies (K_2 is less than the number of cameras). The memory indexes of the associated positives are

Algorithm 1 The O2CAP Algorithm

Input: Unlabeled training set \mathcal{D} , network f_{θ} , number of training epochs $maxEpoch$, number of batches $numBatch$ in one epoch, memory updating rate μ , temperature τ , and balancing weight w ;

Output: Trained network f_{θ} ;

```

1: for epoch = 1 to  $maxEpoch$  do
2:   Perform a global clustering;
3:   Split clusters into camera-aware proxies;
4:   Construct proxy-level memory bank  $\mathcal{K}'$ ;
5:   Generate the pseudo labeled dataset  $\mathcal{D}''$ ;
6:   for b = 1 to  $numBatch$  do
7:     Sample a mini-batch  $\mathcal{B}' = \{(x_i, \tilde{y}_i, \tilde{z}_i, c_i)\}_{i=1}^B$ ;
8:     Obtain  $\mathcal{P}_1$  and  $\mathcal{Q}_1$  by offline association;
9:     Obtain  $\mathcal{P}_2$  and  $\mathcal{Q}_2$  by online association;
10:    Compute the loss  $\mathcal{L}$  defined in Eq. (7);
11:    Backward to update network parameter  $\theta$ ;
12:    Update proxy entries in the memory  $\mathcal{K}'$ ;

```

recorded in a set \mathcal{P}_2 . In addition, we select K_1 -nearest proxies from the remaining in terms of the instance-to-proxy similarity as the hard negatives and denote them by \mathcal{Q}_2 .

Based on the online associated proxies, we define the second contrastive learning loss as

$$\begin{aligned} \mathcal{L}_{On} &= - \sum_{i=1}^B \left(\frac{1}{|\mathcal{P}_2|} \sum_{u \in \mathcal{P}_2} \log \frac{S(u, x_i)}{\sum_{p \in \mathcal{P}_2} S(p, x_i) + \sum_{q \in \mathcal{Q}_2} S(q, x_i)} \right). \end{aligned} \quad (6)$$

C. A Summary of the Proposed Method

The proposed method iteratively alternates between the camera-aware proxy clustering step and the contrastive learning based model updating step. The entire loss used for model updating is

$$\mathcal{L} = \mathcal{L}_{Off} + \mathcal{L}_{On}, \quad (7)$$

To enable better understanding of our offline-online associated camera-aware proxies (O2CAP) based method, we summarize the overall procedure in Algorithm 1.

A *proxy-balanced sampling strategy*. A mini-batch in Algorithm 1 involves an update to the Re-ID model using a small set of samples. Apart from the loss optimization, the strategy of choosing appropriate samples in each batch is also important for model updating, especially when the data distribution is imbalanced [57], [58]. In this work, we propose a proxy-balanced sampling strategy that randomly chooses P proxies and K samples per proxy in each mini-batch. In contrast to the commonly used instance- or cluster-balanced sampling [36], the proposed sampling strategy makes sure each proxy gets equal chance to be sampled. It thus facilitates the learning of image-deficient proxies, as well as the learning of the proxy-level contrastive losses.

TABLE I

THE STATISTICS OF EACH DATASET. #CAMERAS, #IDS, AND #IMAGES ARE THE NUMBER OF CAMERAS, IDS, AND IMAGES, RESPECTIVELY. CID IS THE AVERAGED CAMERA-PER-ID VALUE AND IID IS THE AVERAGED IMAGE-PER-ID VALUE

Dataset	Training Set					Gallery Set		Query Set	
	#cameras	# IDs	#images	CID	IID	#IDs	#images	#IDs	#images
Market-1501	6	751	12,936	4.34	17.23	751	15,913	750	3,368
DukeMTMC-reID	8	702	16,522	3.13	23.54	1,110	17,661	702	2,228
MSMT17	15	1,041	32,621	4.63	31.34	3,060	82,161	3,060	11,659
VeRi-776	20	576	37,778	8.93	65.59	200	11,579	200	1,678

D. An Alternative of O2CAP

In the O2CAP model, we adopt two contrastive learning losses to, respectively, utilize the proxies obtained via offline and online association. Intuitively, there is another alternative way to utilize these offline and online associated proxies. That is, we get a positive proxy set via merging the offline and online associated positive proxies, i.e. $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$. And then we select K_1 -nearest proxies from the remaining as the hard negatives and denote them by \mathcal{Q}_3 . By this means, instead of \mathcal{L} defined in Eq. (7), we use the following alternative loss \mathcal{L}_{Merge} for the model training,

$$\mathcal{L}_{Merge} = - \sum_{i=1}^B \left(\frac{1}{|\mathcal{P}_3|} \sum_{u \in \mathcal{P}_3} \log \frac{S(u, x_i)}{\sum_{p \in \mathcal{P}_3} S(p, x_i) + \sum_{q \in \mathcal{Q}_3} S(q, x_i)} \right). \quad (8)$$

However, in experiments we will show that \mathcal{L}_{Merge} is inferior to the loss \mathcal{L} that considers offline and online associated proxies separately.

E. Discussion on the Intra-Camera Contrastive Learning

All losses introduced above belong to inter-camera contrastive learning. In our preliminary CAP [11], an intra-camera contrastive learning loss is defined to learn the discriminative ability within cameras. We here make a brief introduction and discussion about it. Given image x_i , together with its per-camera pseudo label \tilde{z}_i and camera label c_i , we set $A = \sum_{c=1}^{c_i-1} Z_c$ to be the total proxy number accumulated from the first to the $c_i - 1$ -th camera, and $j = A + \tilde{z}_i$ to be the index of the corresponding entry in the proxy-level memory. Then, the intra-camera contrastive learning loss is defined by

$$\mathcal{L}_{Intra} = - \sum_{c=1}^C \left(\frac{1}{N_c} \sum_{x_i \in \mathcal{D}_c} \log \frac{\exp(\mathcal{K}'[j]^T f_\theta(x_i)/\tau)}{\sum_{k=A+1}^{A+N_{c_i}} \exp(\mathcal{K}'[k]^T f_\theta(x_i)/\tau)} \right). \quad (9)$$

This loss performs contrastive learning within each camera. It pulls an instance close to the proxy to which it belongs, while pushes away from all other proxies in the same camera and ignores the proxies in other cameras.

The combination of this loss with \mathcal{L}_{Off} (denoted as \mathcal{L}_{Inter} in CAP [11]), together with the proxy-balanced sampling strategy, enable CAP to gain the best performance among its model variants. However, during experiments we observe

that adding this intra-camera loss to the O2CAP model is not necessary any more. Considering that the per-camera pseudo labels generated by camera-aware proxies inevitably contain noise, we conjecture that this loss is more effective when a Re-ID model is relatively weak. When O2CAP has already achieved great discriminative ability by the synergy of offline and online association, the intra-camera loss learned from noisy labels may bring confusion to the model and lead to a performance degeneration. Therefore, we discard this loss in our O2CAP model.

V. EXPERIMENTS

A. Datasets and Evaluation Metrics

We evaluate the proposed method on three person Re-ID datasets: Market-1501 [59], DukeMTMC-reID [60], [61], and MSMT17 [26]. The Market-1501 and DukeMTMC-reID datasets are collected on university campus, containing outdoor scenarios only. The MSMT17 dataset is a larger and more challenging dataset, which contains both indoor and outdoor scenarios under different weather conditions. In order to validate the generalization ability of our method, we additionally evaluate it on a vehicle Re-ID dataset, VeRi-776 [62]. The statistics of these datasets, including the number of cameras, IDs, and images contained in training, gallery, and query sets, are summarized in Table I. Meanwhile, the averaged Camera-per-ID (CID) value and the averaged Image-per-ID (IID) value are also provided for reference.

For performance evaluation, we adopt the commonly used mean Average Precision (mAP) and Cumulative Matching Characteristic (CMC) as the metrics. The CMC metric is reported via Rank-1, Rank-5, and Rank-10. To make a fair comparison, we do not use any post-processing techniques (e.g. Re-ranking [63]) during evaluation.

B. Implementation Details

We adopt a slightly modified ResNet-50 [64] as the network backbone. The modifications are as follows: we discard the classification layer in ResNet-50 but additionally add a Batch Normalization (BN) layer right after the Global Average Pooling (GAP) layer, following [36], [65], [66]. The scale and shift parameters of the additional BN layer are initialized as 1 and 0 respectively. The remaining parameters of the backbone are initialized from the original ResNet-50 that is trained on ImageNet. The BN layer outputs a 2048-dimensional feature for each image. Each feature is further normalized by L_2 norm and then used for the computation of the losses as well as the update of memory entries during training. The L_2 normalized

TABLE II

COMPARISON OF THE PROPOSED METHOD AND ITS VARIANTS. *PBsampling* IS THE PROXY-BALANCED SAMPLING STRATEGY. WHEN *PBsampling* IS NOT SELECTED, THE MODELS USE THE CONVENTIONAL CLUSTER-BALANCED SAMPLING STRATEGY. NOTE THAT \mathcal{L}_{Off} IS THE SAME AS \mathcal{L}_{Inter} IN CAP [11]. AND DUE TO RE-IMPLEMENTATION WITH SOME HYPER-PARAMETERS ALTERED (E.G. ITERATION NUMBER PER EPOCH), THE PERFORMANCE VALUES OF CAP AND ITS VARIANTS ARE NOT EXACTLY THE SAME WITH THOSE REPORTED IN THE PRELIMINARY WORK [11]

Models	Components				Market-1501		DukeMTMC-ReID		MSMT17		VeRi-776	
	\mathcal{L}_{Intra}	\mathcal{L}_{Off}	\mathcal{L}_{On}	<i>PBsampling</i>	R1	mAP	R1	mAP	R1	mAP	R1	mAP
Baseline1					84.4	67.7	65.8	48.1	29.4	13.0	46.2	18.0
Baseline2					84.9	68.9	72.2	55.3	34.3	15.3	52.0	19.4
					\mathcal{L}_{Base}							
					\mathcal{L}_{Base2}							
CAP1	✓				76.4	57.5	74.2	57.4	43.7	20.3	76.4	35.1
CAP2	✓			✓	83.6	67.1	76.5	60.8	42.1	18.3	79.1	36.5
CAP3		✓			89.3	76.1	77.1	61.1	63.7	33.2	78.7	33.6
CAP4		✓		✓	92.5	82.8	79.6	64.0	66.6	36.9	80.6	35.5
CAP5	✓	✓			90.1	78.7	77.3	61.2	61.6	32.2	82.7	39.5
CAP	✓	✓		✓	92.2	81.6	81.2	68.3	65.9	35.5	83.8	38.7
O2CAP1	✓	✓	✓	✓	92.0	82.0	82.0	69.0	71.3	41.7	87.1	41.3
O2CAP		✓	✓	✓	92.5	82.7	83.9	71.2	72.0	42.4	87.5	41.9
O2CAP2			✓	✓	91.4	80.2	82.0	68.6	57.2	28.2	80.9	35.2
O2CAP3				✓	92.1	81.8	83.7	71.3	67.5	38.3	85.2	38.9
					\mathcal{L}_{Merge}							

features are also used for the distance computation during test time.

The hyper-parameters involved in our model are empirically set as follows. The memory updating rate μ is 0.2 and the temperature factor τ in all contrastive losses is set to 0.07. The number of hard negatives, i.e. K_1 , is fixed to 50. The balancing weight w in Eq. (5) is 0.15. The number of positive proxies associated online, i.e. K_2 , is set to 3 for both Market-1501 and MSMT17, 2 for DukeMTMC-reID, and 8 for VeRi-776, roughly close to but less than their CID values listed in Table I. At the beginning of each epoch, we compute Jaccard distance [63] of all features and then use DBSCAN [13] with a threshold of 0.5 and neighborhood $eps = 4$ to conduct the camera-agnostic clustering.

In addition, we adopt the commonly used random flipping, cropping, and erasing to augment data for training. The model is trained by ADAM [67] optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and weight decay of 0.0005. The learning rate is initially set to 0.00035 with a warmup in the first 10 epochs, and is divided by 10 after each 20 epochs. The number of total epochs is 50. The number of iterations in each epoch is set to 400, as the common practice [10], [12], [14]. The batch size is 32. Following the proposed proxy-balanced sampling strategy, we randomly sample 8 camera-aware proxies and 4 images per proxy within each batch. Our model is implemented with the Pytorch [68] framework. All experiments are run on a single GTX 1080Ti GPU. The training phase (50 epochs) takes about 2.5 hours for Market-1501 and DukeMTMC-reID, 3.5 hours for MSMT17, and 4 hours for VeRi-776 dataset.

C. Ablation Studies

In this subsection, we conduct a series of experiments to validate the effectiveness of each proposed component. The performance of our full model and its variants are presented in Table II.

1) *Effectiveness of the Proxy-Level Contrastive Learning:* Let us first compare CAP3 with Baseline2. The CAP3 model (using \mathcal{L}_{Off} only) and the Baseline2 model (using \mathcal{L}_{Base2})

perform the contrastive learning, respectively, at the proxy-level and cluster-level while keep other settings the same. They pull an instance, respectively, close to its positive proxies or its positive cluster. Since offline association straightforwardly associates the proxies split from one positive cluster as the positive proxies, both models in essence encourage high compactness within the same clusters. However, as shown in Figure 3, the proxy-level learning achieves a more balanced learning of instance-rich and instance-deficient proxies, leading to a higher intra-cluster compactness. The other difference of these two models lies in that CAP3 pushes an instance away from its *hard negative proxies* while Baseline2 pushes the instance away from its *hard negative clusters*. Therefore, CAP3 pays more attention to the hardest negative instances so that the inter-ID similarity is better conquered. Last but not least, the pseudo labels generated from proxies are less noisy than those obtained from clusters, which also benefit the model learning. Due to these reasons, CAP3 outperforms Baseline2 by a great margin, especially on the complex datasets such as MSMT17 and VeRi-776, as shown in Table II.

2) *Effectiveness of the Combination of Offline and Online Association:* From O2CAP1 vs. CAP and O2CAP vs. CAP4, we observe that the models additionally integrated with the online association based loss (\mathcal{L}_{On}) can boost the performance by a considerable margin on the datasets including DukeMTMC-ReID, MSMT17, and VeRi-776. Besides, the combination of offline and online association in O2CAP is implemented via the sum of two contrastive losses \mathcal{L}_{Off} and \mathcal{L}_{On} . To validate the effectiveness of this combination way, we compare it with an alternative way that directly merges the positive proxy sets obtained offline and online into one set and adopts one loss \mathcal{L}_{Merge} to perform contrastive learning. As shown in Table II, the O2CAP model (using \mathcal{L}_{Off} and \mathcal{L}_{On}) performs comparable with the O2CAP3 model (using \mathcal{L}_{Merge}) on Market-1501 and DukeMTMC-ReID, but significantly outperforms O2CAP3 on MSMT17 and VeRi-776, demonstrating the superiority of the two-loss combination way in complex scenarios.

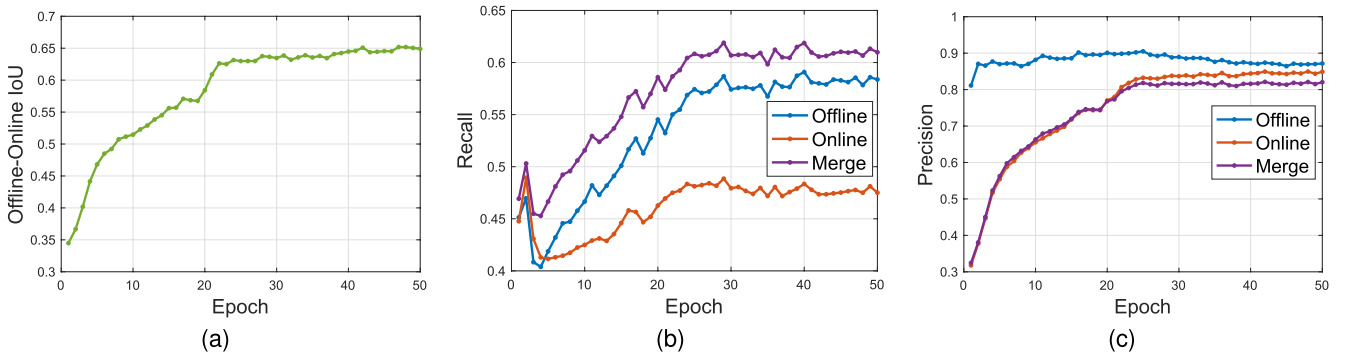


Fig. 5. The statistics of associated positive proxies during training on MSMT17. (a) is the averaged IoU of offline and online associated proxies for each instance. (b) includes the averaged recall of offline, online associated positive proxies and their union. (c) includes the averaged precision of offline, online associated positive proxies and their union.

To further analyze the complementarity of offline and online association, we investigate some typical association statistics. Specifically, the offline and online proxy sets associated to each anchor instance x_i are denoted as \mathcal{P}_1^i and \mathcal{P}_2^i , respectively, and the union set is $\mathcal{P}_3^i = \mathcal{P}_1^i \cup \mathcal{P}_2^i$. We further assign each proxy with a ground-truth ID label via taking the ID of its majority instances, by which the ground-truth positive proxy set \mathcal{P}_{gt}^i of x_i can be roughly obtained. Then, the averaged Intersection over Union (IoU) of offline and online associated proxies, together with the averaged precision and recall of the proxies associated in different ways, are computed as followings:

$$IoU = \frac{1}{N'} \sum_{i=1}^{N'} \frac{|\mathcal{P}_1^i \cap \mathcal{P}_2^i|}{|\mathcal{P}_1^i \cup \mathcal{P}_2^i|}, \quad (10)$$

$$precision = \frac{1}{N'} \sum_{i=1}^{N'} \frac{|\mathcal{P}_r^i \cap \mathcal{P}_{gt}^i|}{|\mathcal{P}_r^i|}, \quad (11)$$

$$recall = \frac{1}{N'} \sum_{i=1}^{N'} \frac{|\mathcal{P}_r^i \cap \mathcal{P}_{gt}^i|}{|\mathcal{P}_{gt}^i|} \quad (12)$$

in which N' is the total training images and $r = 1, 2$, or 3 denotes the offline, online, or union association set.

Figure 5 plots the detailed association statistics varying during training on MSMT17. As shown in Figure 5 (a), the IoU between offline and online associated proxies increases from an initial value of 0.34 to a converged value of 0.65. It indicates that the two associations are getting more consistent as training goes on, but a part of them are still unique, leading to their complementarity and synergy. Figure 5 (b) and (c) present the recall and precision of two association schemes. An intuitive observation is that the precision and recall of two associations are improving during training, and offline association achieves higher performance than online association. Nevertheless, the union of offline and online associations generates higher recall than either alone, showing that offline and online association indeed complement each other, and together they retrieve more positive proxies to benefit model learning. On the other hand, the precision of the union association is lower than offline or online association, which may explain the compromised performance of O2CAP3.

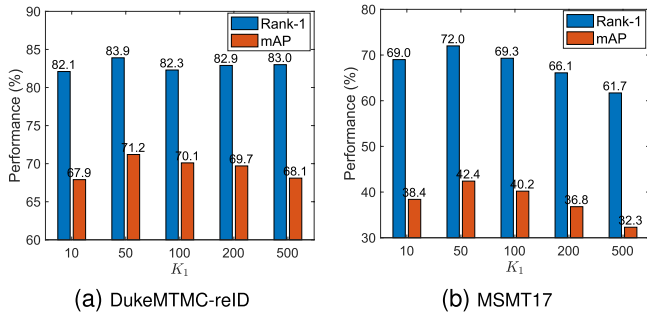
TABLE III

COMPARISON OF THE ONLINE ASSOCIATION STRATEGIES AND THEIR COUNTERPARTS. EXPERIMENTS ARE CONDUCTED ON THE O2CAP MODEL. GLOBAL KNN CHOOSES K-NEAREST NEIGHBORS GLOBALLY AND 'CA' IS OUR CAMERA-AWARE CRITERION. FOR THE SIMILARITY MEASUREMENTS, 'BALANCED' REFERS TO THE BALANCED SIMILARITY AND 'ORIGINAL' IS THE SINGLE INSTANCE-TO-PROXY SIMILARITY

Settings			DukeMTMC-reID		MSMT17	
KNN	Positives	Negatives	R1	mAP	R1	mAP
Global	Original	Original	80.7	67.2	69.1	39.7
Global	Balanced	Balanced	81.3	66.9	60.0	32.2
Global	Balanced	Original	80.6	66.3	50.0	24.2
CA	Original	Original	82.8	70.7	68.9	39.2
CA	Balanced	Balanced	83.4	69.2	71.1	41.1
CA	Balanced	Original	83.9	71.2	72.0	42.4

3) *Effectiveness of the Strategies in Online Association:* In the design of our online association, we propose an instance-proxy balanced similarity and a camera-aware nearest neighbor criterion to select positive proxies. In order to validate their effectiveness, we compare them, respectively, with the original instance-to-proxy similarity and the global KNN that are ordinarily used [6], [15]. Table III presents the comparison results. Compared to the global KNN, the camera-aware criterion can avoid to associate multiple proxies within one camera and therefore reduce false associations. It plays a vital role to make online association effective. From Table III and Table II we see that, when the global KNN is adopted, the online association is sensitive to the similarity measurement and is not always leading to performance enhancement. On the contrary, the models adopting the camera-aware criterion consistently improve performance no matter which similarity is used. When the camera-aware criterion is adopted, the balanced similarity can boost the performance further. The best performance is achieved when the balanced similarity is used for the association of positive proxies while the instance-to-proxy similarity is used for negative mining.

4) *Effectiveness of the Proxy-Balanced Sampling Strategy:* In order to validate the effectiveness of the proposed proxy-balanced sampling strategy, we compare it to the commonly used cluster-balanced strategy that takes no consideration of intra-cluster distributions. From CAP2 vs. CAP1, CAP4 vs.

Fig. 6. The performance evaluation with different values of K_1 .

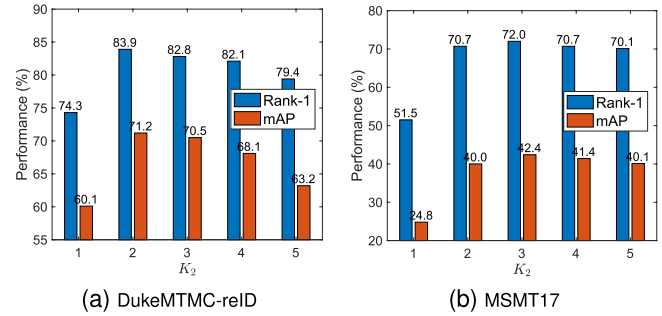
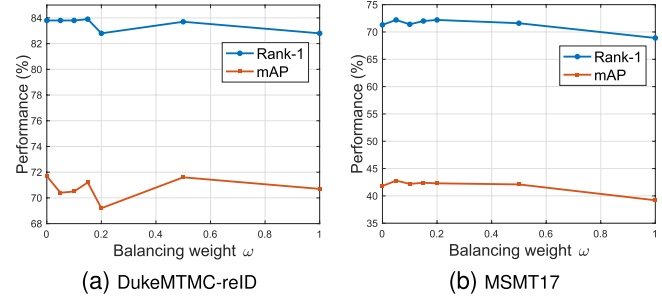
CAP3, and CAP vs. CAP5, we observe that the models using the proxy-balanced sampling strategy consistently outperform the counterparts using the cluster-balanced strategy on almost all datasets. The results show that the proxy-balanced sampling strategy facilitates the learning of all camera-specific proxies and naturally fits to the proxy-level contrastive learning.

5) *Analysis on the Intra-Camera Contrastive Learning*: The intra-camera contrastive learning loss (\mathcal{L}_{Intra}) was proposed in CAP [11] to take advantage of the per-camera pseudo labels to learn discriminative ability within cameras and further boost the learning of global discrimination. As shown by CAP1 and CAP2, this loss enables the models to gain considerable discriminative ability. And from the full model CAP, we see that the integration of \mathcal{L}_{Intra} and \mathcal{L}_{Off} does improve the performance over the model CAP4 that utilizes \mathcal{L}_{Off} only. However, when comparing O2CAP with O2CAP1, we find out that the integration of \mathcal{L}_{Intra} , on the contrary, degrades the performance on all datasets. Our conjecture is as follows. The per-camera pseudo labels generated from camera-aware proxies are, although more reliable than the cluster-level pseudo labels, still noisy. Thus, the noisy label based intra-camera learning may bring confusion to the O2CAP model that has already achieved great discriminative ability.

D. Parameter Analysis

In this subsection, we analyze the sensitivity of the hyper-parameters involved in O2CAP. Among all hyper-parameters, the memory updating rate μ and the temperature factor τ have been investigated in many other works so that we simply follow [4], [11] to set them. Here, we conduct experiments on DukeMTMC-reID and MSMT17 to investigate the sensitivity of the remaining hyper-parameters, which include the number of hard negative proxies (K_1), the number of positive proxies associated online (K_2), and the weight w in the instance-proxy balanced similarity. In addition, the sensitivity of the training epochs is also investigated as a reference.

1) *The Sensitivity of K_1* : K_1 is the number of negative proxies mined in offline or online association. Figure 6 presents the performance of O2CAP when K_1 varies from 10 to 500. We observe that the performance goes up when the number of hard negatives increases from 10 to 50. But the performance gradually drops when more negative proxies are taken into consideration. It indicates that easy negatives may hamper the contrastive learning. Focusing on a small number of

Fig. 7. The performance evaluation with different values of K_2 .Fig. 8. The performance evaluation with different values of w .

most informative negative proxies helps our model to better discriminate confusing instances.

2) *The Sensitivity of K_2* : K_2 is the number of positive proxies associated online for each instance. Figure 7 presents the performance of O2CAP when K_2 varies from 1 to 5. We see that the performance increases first and then degenerates. We also notice that the best performed K_2 value is closely related to the CID value listed in Table I. If K_2 is larger than the CID value, false positives will be inevitably included and if K_2 is too small, no enough positive proxies will be recalled. Therefore, we set K_2 to be the value that is one less than CID for each dataset.

3) *The Sensitivity of w* : The weight w is to balance the instance-to-proxy and proxy-to-proxy parts in the balanced similarity. When $w = 1$, the balanced similarity degenerates to the instance-to-proxy similarity and when $w = 0$, it becomes the proxy-to-proxy similarity. Figure 8 presents the performance varying along with the change of w . As shown in the figure, the performance on DukeMTMC-reID is less sensitive to w when it increases from 0 to 1. On the contrary, setting w to be a small value is more beneficial on MSMT17. Therefore, we set $w = 0.15$ throughout all other experiments to achieve a trade-off on all datasets.

4) *The Sensitivity of Training Epochs*: Finally, we investigate the performance sensitivity with respect to the number of training epochs. Figure 9 presents the performance of O2CAP when the training epoch varies from 1 to 100. We observe that both Rank-1 and mAP go up quickly in the first 30 epochs, and then slowly converge at the 50-th epoch. After 50 epochs, the performance fluctuates very slightly and no discernible improvement is gained. Therefore, we set the total training epoch to 50 in order to achieve a trade-off between training time and model accuracy.

TABLE IV

COMPARISON WITH STATE-OF-THE-ART METHODS. BOTH PURELY UNSUPERVISED AND UDA-BASED METHODS ARE INCLUDED. WE ALSO PROVIDE SEVERAL FULLY SUPERVISED METHODS FOR REFERENCE. O2CAP: OUR PROPOSED METHOD IN THIS WORK. O2CAP(IBN): O2CAP WITH IBN-RESNET50 AS BACKBONE. O2CAP(IBN+GEMPOOL): O2CAP WITH IBN-RESNET50 AND GEM POOLING. O2CAP(w/GT): OUR PROPOSED METHOD WITH GROUND TRUTH, INDICATING THE UPPER BOUND

Methods	Reference	Market-1501				DukeMTMC-ReID				MSMT17			
		R1	R5	R10	mAP	R1	R5	R10	mAP	R1	R5	R10	mAP
Purely Unsupervised													
BUC [5]	AAAI19	66.2	79.6	84.5	38.3	47.4	62.6	68.4	27.5	-	-	-	-
UGA [24]	ICCV19	87.2	-	-	70.3	75.0	-	-	53.3	49.5	-	-	21.7
SSL [69]	CVPR20	71.7	83.8	87.4	37.8	52.5	63.5	68.9	28.6	-	-	-	-
HCT [20]	CVPR20	80.0	91.6	95.2	56.4	69.6	83.4	87.4	50.7	-	-	-	-
CycAs [70]	ECCV20	84.8	-	-	64.8	77.9	-	-	60.1	50.1	-	-	26.7
IICS [25]	CVPR21	89.5	95.2	97.0	72.9	80.0	89.0	91.6	64.4	56.4	68.8	73.4	26.9
RLCC [16]	CVPR21	90.8	96.3	97.5	77.7	83.2	91.6	93.8	69.2	56.5	68.4	73.1	27.9
CAP [11]	AAAI21	91.4	96.3	97.7	79.2	81.1	89.3	91.8	67.3	67.4	78.0	81.4	36.9
GroupSampling [54]	arXiv21	92.3	96.6	97.8	79.2	82.7	91.1	93.5	69.1	56.2	67.3	71.5	24.6
ICE [12]	ICCV21	93.8	97.6	98.4	82.3	83.3	91.5	94.1	69.9	70.2	80.5	84.4	38.9
ICE(IBN-ResNet) [12]	ICCV21	94.2	97.6	98.5	82.5	83.6	91.9	93.9	70.7	70.7	81.0	84.6	40.6
MGH [15]	MM21	93.2	96.8	98.1	81.7	83.7	92.1	93.7	70.2	70.2	81.2	84.5	40.6
MGCE-HCL [17]	EC21	92.1	-	-	79.6	82.5	-	-	67.5	-	-	-	-
Liu et al. [37]	arXiv21	93.0	97.5	-	82.4	84.9	92.3	-	72.2	68.6	79.4	-	38.4
HHCL [38]	NIDC21	93.4	97.7	98.5	84.2	85.1	92.4	94.6	73.3	58.9	71.3	75.8	31.8
O2CAP	This work	92.5	96.9	98.0	82.7	83.9	91.3	93.4	71.2	72.0	81.9	85.4	42.4
O2CAP(IBN)	This work	93.1	97.4	98.1	83.7	85.2	91.9	93.5	72.8	75.5	84.8	87.7	46.9
O2CAP(IBN+GeMPool)	This work	93.3	96.9	97.7	85.0	85.3	91.4	93.3	73.2	77.3	85.6	88.2	48.3
Unsupervised Domain Adaptation													
PUL [33]	TOMM18	45.5	60.7	66.7	20.5	30.0	43.4	48.5	16.4	-	-	-	-
SPGAN [27]	CVPR18	51.5	70.1	76.8	22.8	41.1	56.6	63.0	22.3	-	-	-	-
ECN [6]	CVPR19	75.1	87.6	91.6	43.0	63.3	75.8	80.4	40.4	30.2	41.5	46.8	10.2
SSG [71]	ICCV19	80.0	90.0	92.4	58.3	73.0	80.6	83.2	53.4	32.2	-	51.2	13.3
AD-Cluster [34]	CVPR20	86.7	94.4	96.5	68.3	72.6	82.5	85.5	54.1	-	-	-	-
MMCL [72]	CVPR20	84.4	92.8	95.0	60.4	72.4	82.9	85.0	51.4	43.6	54.3	58.9	16.2
Zhong et al. [29]	TPAMI21	84.1	92.8	95.4	63.8	74.0	83.7	87.4	54.4	42.5	55.9	61.5	16.0
MMT [14]	ICLR20	87.7	94.9	96.9	71.2	78.0	88.8	92.5	65.1	50.1	63.9	69.8	23.3
MEB-Net [73]	ECCV20	89.9	96.0	97.5	76.0	79.6	88.3	92.2	66.1	-	-	-	-
SpCL [10]	NeurIPS20	90.3	96.2	97.7	76.7	82.9	90.1	92.5	68.8	53.1	65.8	70.5	26.5
SpCL(IBN-ResNet) [10]	NeurIPS20	91.5	96.9	98.0	79.2	83.4	91.0	93.1	69.9	58.9	70.4	75.2	31.8
Isobe et al. [32]	ICCV21	94.2	-	-	83.4	83.5	-	-	70.8	66.6	-	-	36.3
Zheng et al. [18]	ICCV21	91.5	-	-	80.0	82.2	-	-	70.1	56.1	-	-	29.3
MCRN [35]	AAAI22	93.8	97.5	98.5	83.8	84.5	91.7	93.8	71.5	67.5	77.9	81.6	35.7
Fully Supervised													
PCB [74]	ECCV18	93.8	-	-	81.6	83.3	-	-	69.2	68.2	-	-	40.4
ABD-Net [1]	ICCV19	95.6	-	-	88.3	89.0	-	-	78.6	82.3	90.6	-	60.8
FlipReID [75]	EUVIP21	95.3	-	-	88.5	89.4	-	-	79.8	83.3	-	-	64.3
TransReID [2]	ICCV21	95.2	-	-	89.5	90.7	-	-	82.6	86.2	-	-	69.4
O2CAP(w/ GT)	This work	94.4	98.0	98.8	86.1	88.2	94.1	96.0	76.3	79.0	89.4	92.0	53.7
O2CAP(IBN,w/ GT)	This work	95.0	98.0	98.7	86.9	89.5	94.7	96.1	78.3	81.0	90.3	92.7	56.7
O2CAP(IBN+GeMPool,w/ GT)	This work	95.0	98.0	98.6	87.3	89.6	95.6	96.6	79.9	82.8	91.1	93.0	58.8

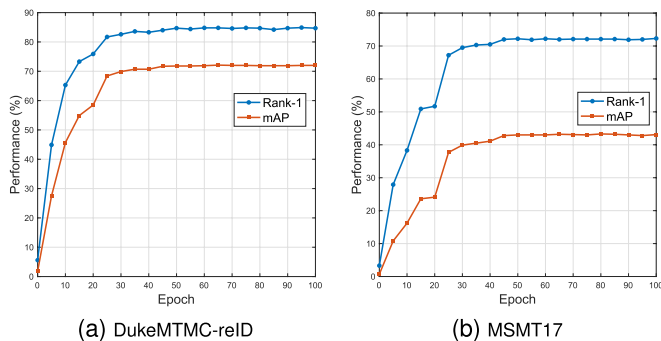


Fig. 9. The performance evaluation with different training epochs.

E. Comparison to State-of-The-Arts

In this subsection, we compare the proposed method (named as O2CAP) with state-of-the-art methods on both person and vehicle Re-ID datasets. The comparison results are summarized in Table IV and Table V.

1) *Comparison With Purely Unsupervised Methods on Person Re-ID*: 14 representative or recent purely unsupervised methods are included for comparison. Most of these methods are clustering-based, and CAP [11] (our preliminary work), ICE [12], MGH [15], MGCE-HCL [17], Liu et al. [37], and HHCL [38] also adopt the contrastive learning technique same as this work. Both ICE and MGH are built upon CAP. The former extends CAP via introducing the instance-level contrastive learning and the latter additionally introduces hypergraph for label refinement. In contrast, our O2CAP sticks to the proxy-level contrastive learning and improves CAP via the combination of offline and online associations, keeping the model simple yet effective. From the results we see that O2CAP outperforms CAP and earlier methods by a significant margin. When compared to other contrastive learning based methods, O2CAP achieves competitive results on Market-1501 and DukeMTMC-ReID, while demonstrates a considerable superiority on MSMT17.

TABLE V
COMPARISON WITH STATE-OF-THE-ART METHODS ON VeRi-776.
† INDICATES AN UDA-BASED METHOD WORKING UNDER
THE PURELY UNSUPERVISED SETTING

Methods	Reference	VeRi-776			
		R1	R5	R10	mAP
Purely Unsupervised					
SSML [77]	IROS21	74.5	80.3	-	26.7
SpCL [†] [10]	NeurIPS20	79.9	86.8	89.9	36.9
RLCC [16]	CVPR21	83.4	88.8	90.9	39.6
O2CAP	This work	87.5	92.7	94.4	41.9
O2CAP(IBN)	This work	89.6	93.5	94.7	42.4
O2CAP(IBN+GeMPool)	This work	89.7	93.8	95.1	43.0
Fully Supervised					
VSCR [78]	IJCV21	94.1	97.9	98.6	75.5
VehicleNet [79]	TMM20	96.8	-	-	83.4
TransReID [2]	ICCV21	97.1	-	-	82.0
HRCN [80]	ICCV21	97.3	98.9	-	83.1
O2CAP(w/ GT)	This work	93.2	97.9	98.9	73.6
O2CAP(IBN,w/ GT)	This work	94.6	98.0	98.9	75.1
O2CAP(IBN+GeMPool,w/ GT)	This work	94.4	98.1	99.0	75.0

2) *Comparison With UDA-Based Methods on Person Re-ID*: Table IV also presents 13 unsupervised domain adaptation based methods for comparison. Among them, ECN [6], SpCL [10], Isobe et al. [32], Zheng et al. [18] and MCRN [35] utilize contrastive learning as well. Although all UDA-based methods exploit external labeled data to boost the Re-ID performance on target datasets, they do not gain noticeable advantage when compared to recent purely unsupervised counterparts [11], [12], [15]. For instance, our O2CAP performs on par or better than all UDA-based methods on DukeMTMC-ReID and MSMT17. Especially on the most complex dataset MSMT17, O2CAP outperforms them by a great margin due to the exploit of camera-aware proxies.

3) *Comparison With Fully Supervised Methods on Person Re-ID*: We additionally provide four representative fully supervised methods for reference, including PCB [74], ABD-Net [1], FlipReID [75] and TransReID [2]. Besides, we also report the performance of our network backbone trained with ground-truth labels, which indicates the upper bound performance of our method. The results show that our unsupervised O2CAP has already outperforms the well-known PCB on all datasets with respect to all metrics, except Rank-1 on Market-1501. The performance gap between the unsupervised O2CAP and its supervised counterpart has also been greatly mitigated. Note that recent supervised methods such as FlipReID [75] and TransReID [2] have set up new state-of-the-art performance on all Re-ID datasets. It implies the potential of our unsupervised method to further improve performance if a more advanced backbone network could be adopted.

4) *O2CAP With a More Advanced Backbone*: The backbone used in O2CAP is ResNet50, which is a relatively plain network. In order to investigate the generalization ability of our method to different backbones, we additionally conduct an experiment using the IBN-ResNet50 backbone. It replaces batch normalization in ResNet50 via instance batch normalization (IBN) [76], which has been proved effective to boost the Re-ID performance [10], [12]. As shown in Table IV, O2CAP with IBN-ResNet50 is able to improve the performance further. And using GeM (Generalized Mean) pooling together with IBN-ResNet50 gives additional improvement. Especially on MSMT17, 5.3% Rank-1 and 5.9% mAP improvements have been gained. The results show that our

method is orthogonal to network design and could further benefit from better backbones.

5) *Comparison on Vehicle Re-ID*: Table V presents the comparison results on a vehicle Re-ID dataset VeRi-776 [62]. Three state-of-the-art unsupervised methods including SSML [77], SpCL[†] [10] and RLCC [16] are taken for comparison. Four fully supervised method VSCR [78], VehicleNet [79], TransReID [2] and HRCN [80], together with our network trained with ground truth, are also provided for reference. Compared to RLCC [16], our method achieves 4.1% Rank-1 and 2.3% mAP improvements, validating its effectiveness for the unsupervised vehicle Re-ID task.

VI. CONCLUSION


In this paper, we have presented a camera-aware proxy assisted method for the purely unsupervised person Re-ID. The proposed camera-aware proxies are able to deal with the large intra-ID variance via explicitly considering the variance within clusters resulted from the change of camera views. They can also better tackle the inter-ID similarity by paying more attention to the hardest negative instances when compared to the cluster-level counterparts. With the assistance of camera-aware proxies, two proxy-level contrastive learning losses based on offline and online associations are designed to optimize the Re-ID model. Extensive experiments on both person and vehicle Re-ID datasets, especially on the most challenging ones, have demonstrated the superiority of our method.

REFERENCES

- [1] T. Chen et al., "ABD-NET: Attentive but diverse person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8351–8361.
- [2] S. He, H. Luo, P. Wang, F. Wang, H. Li, and W. Jiang, "TransReID: Transformer-based object re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 15013–15022, doi: 10.1109/ICCV48922.2021.01474.
- [3] Q. Yang, A. Wu, and W.-S. Zheng, "Deep semi-supervised person re-identification with external memory," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2019, pp. 1096–1101, doi: 10.1109/ICME.2019.00192.
- [4] M. Wang, B. Lai, J. Huang, X. Gong, and X.-S. Hua, "Graph-induced contrastive learning for intra-camera supervised person re-identification," *IEEE ACCESS*, vol. 9, pp. 20850–20860, 2021, doi: 10.1109/ACCESS.2021.3055266.
- [5] Y. Lin, X. Dong, L. Zheng, Y. Yan, and Y. Yang, "A bottom-up clustering approach to unsupervised person re-identification," in *Proc. AAAI*, 2019, pp. 8738–8745, doi: 10.1609/aaai.v33i01.33018738.
- [6] Z. Zhong, L. Zheng, Z. Luo, S. Li, and Y. Yang, "Invariance matters: Exemplar memory for domain adaptive person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 598–607, doi: 10.1109/CVPR.2019.00069.
- [7] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin, "Unsupervised feature learning via non-parametric instance discrimination," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3733–3742, doi: 10.1109/CVPR.2018.00393.
- [8] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. ICML*, 2020, pp. 1597–1607.
- [9] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9729–9738, doi: 10.1109/CVPR42600.2020.00975.
- [10] Y. Ge, D. Chen, F. Zhu, R. Zhao, and H. Li, "Self-paced contrastive learning with hybrid memory for domain adaptive object re-id," in *Proc. NeurIPS*, 2020, pp. 11309–11321.

- [11] M. Wang, B. Lai, J. Huang, X. Gong, and X.-S. Hua, "Camera-aware proxies for unsupervised person re-identification," in *Proc. AAAI*, 2021, pp. 2764–2772.
- [12] H. Chen, B. Lagadec, and F. Bremond, "ICE: Inter-instance contrastive encoding for unsupervised person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 14960–14969, doi: [10.1109/ICCV48922.2021.01469](https://doi.org/10.1109/ICCV48922.2021.01469).
- [13] M. Ester et al., "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proc. KDD*, 1996, pp. 226–231.
- [14] Y. Ge, D. Chen, and H. Li, "Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification," in *Proc. ICLR*, 2020, pp. 1–15.
- [15] Y. Wu, X. Wu, X. Li, and J. Tian, "MGH: Metadata guided hypergraph modeling for unsupervised person re-identification," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 1571–1580, doi: [10.1145/3474085.3475296](https://doi.org/10.1145/3474085.3475296).
- [16] X. Zhang, Y. Ge, Y. Qiao, and H. Li, "Refining pseudo labels with clustering consensus over generations for unsupervised object re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 3436–3445, doi: [10.1109/CVPR46437.2021.00344](https://doi.org/10.1109/CVPR46437.2021.00344).
- [17] H. Sun, M. Li, and C.-G. Li, "Hybrid contrastive learning with cluster ensemble for unsupervised person re-identification," 2021, *arXiv:2201.11995*.
- [18] Y. Zheng et al., "Online pseudo label generation by hierarchical cluster dynamics for adaptive person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 8371–8381, doi: [10.1109/ICCV48922.2021.00826](https://doi.org/10.1109/ICCV48922.2021.00826).
- [19] H. Feng, M. Chen, J. Hu, D. Shen, H. Liu, and D. Cai, "Complementary pseudo labels for unsupervised domain adaptation on person re-identification," *IEEE Trans. Image Process.*, vol. 30, pp. 2898–2907, 2021, doi: [10.1109/TIP.2021.3056212](https://doi.org/10.1109/TIP.2021.3056212).
- [20] K. Zeng, M. Ning, Y. Wang, and Y. Guo, "Hierarchical clustering with hard-batch triplet loss for person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 13657–13665, doi: [10.1109/CVPR42600.2020.01367](https://doi.org/10.1109/CVPR42600.2020.01367).
- [21] M. Li, X. Zhu, and S. Gong, "Unsupervised person re-identification by deep learning tracklet association," in *Proc. ECCV*, 2018, pp. 737–753, doi: [10.1007/978-3-030-01225-0_45](https://doi.org/10.1007/978-3-030-01225-0_45).
- [22] Y. Chen, X. Zhu, and S. Gong, "Deep association learning for unsupervised video person re-identification," in *Proc. BMVC*, 2018, pp. 737–753.
- [23] M. Ye, A. J. Ma, L. Zheng, J. Li, and P. C. Yuen, "Dynamic label graph matching for unsupervised video re-identification," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5142–5150, doi: [10.1109/ICCV.2017.550](https://doi.org/10.1109/ICCV.2017.550).
- [24] J. Wu, H. Liu, Y. Yang, Z. Lei, S. Liao, and S. Li, "Unsupervised graph association for person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8321–8330, doi: [10.1109/ICCV.2019.00841](https://doi.org/10.1109/ICCV.2019.00841).
- [25] S. Xuan and S. Zhang, "Intra-inter camera similarity for unsupervised person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 11926–11935.
- [26] L. Wei, S. Zhang, W. Gao, and Q. Tian, "Person transfer GAN to bridge domain gap for person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 79–88, doi: [10.1109/CVPR.2018.00016](https://doi.org/10.1109/CVPR.2018.00016).
- [27] W. Deng, L. Zheng, Q. Ye, G. Kang, Y. Yang, and J. Jiao, "Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 994–1003, doi: [10.1109/CVPR.2018.00110](https://doi.org/10.1109/CVPR.2018.00110).
- [28] J. Liu, Z.-J. Zha, D. Chen, R. Hong, and M. Wang, "Adaptive transfer network for cross-domain person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 7202–7211, doi: [10.1109/CVPR.2019.00737](https://doi.org/10.1109/CVPR.2019.00737).
- [29] Z. Zhong, L. Zheng, Z. Luo, S. Li, and Y. Yang, "Learning to adapt invariance in memory for person re-identification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 8, pp. 2723–2738, Aug. 2020, doi: [10.1109/TPAMI.2020.2976933](https://doi.org/10.1109/TPAMI.2020.2976933).
- [30] L. Qi, L. Wang, J. Huo, L. Zhou, Y. Shi, and Y. Gao, "A novel unsupervised camera-aware domain adaptation framework for person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8080–8089, doi: [10.1109/ICCV.2019.00817](https://doi.org/10.1109/ICCV.2019.00817).
- [31] A. Wu, W.-S. Zheng, and J.-H. Lai, "Unsupervised person re-identification by camera-aware similarity consistency learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6922–6931, doi: [10.1109/ICCV.2019.00702](https://doi.org/10.1109/ICCV.2019.00702).
- [32] T. Isobe, D. Li, L. Tian, W. Chen, Y. Shan, and S. Wang, "Towards discriminative representation learning for unsupervised person re-identification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 8526–8536, doi: [10.1109/ICCV48922.2021.00841](https://doi.org/10.1109/ICCV48922.2021.00841).
- [33] H. Fan, L. Zheng, C. Yan, and Y. Yang, "Unsupervised person re-identification: Clustering and fine-tuning," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 14, no. 4, pp. 1–18, Nov. 2018, doi: [10.1145/3243316](https://doi.org/10.1145/3243316).
- [34] Y. Zhai et al., "AD-Cluster: Augmented discriminative clustering for domain adaptive person re-identification," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9021–9030, doi: [10.1109/CVPR42600.2020.00904](https://doi.org/10.1109/CVPR42600.2020.00904).
- [35] Y. Wu et al., "Multi-centroid representation network for domain adaptive person re-ID," in *Proc. AAAI*, 2022, pp. 2750–2758.
- [36] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," 2017, *arXiv:1703.07737*.
- [37] T. Liu, Y. Lin, and B. Du, "Unsupervised person re-identification with stochastic training strategy," 2021, *arXiv:2108.06938*.
- [38] Z. Hu, C. Zhu, and G. He, "Hard-sample guided hybrid contrast learning for unsupervised person re-identification," in *Proc. 7th IEEE Int. Conf. Neww. Intell. Digit. Content (IC-NIDC)*, Nov. 2021, pp. 91–95, doi: [10.1109/IC-NIDC54101.2021.9660560](https://doi.org/10.1109/IC-NIDC54101.2021.9660560).
- [39] H.-X. Yu, A. Wu, and W.-S. Zheng, "Cross-view asymmetric metric learning for unsupervised person re-identification," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 994–1002, doi: [10.1109/ICCV.2017.113](https://doi.org/10.1109/ICCV.2017.113).
- [40] M. Li, X. Zhu, and S. Gong, "Unsupervised tracklet person re-identification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 7, pp. 1770–1782, Jul. 2020, doi: [10.1109/TPAMI.2019.2903058](https://doi.org/10.1109/TPAMI.2019.2903058).
- [41] Y. Movshovitz-Attias, A. Toshev, T. K. Leung, S. Ioffe, and S. Singh, "No fuss distance metric learning using proxies," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 360–368, doi: [10.1109/ICCV.2017.47](https://doi.org/10.1109/ICCV.2017.47).
- [42] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *Proc. ECCV*, 2016, pp. 499–515, doi: [10.1007/978-3-319-46478-7_31](https://doi.org/10.1007/978-3-319-46478-7_31).
- [43] O. Rippel, M. Paluri, P. Dollar, and L. Bourdev, "Metric learning with adaptive density discrimination," in *Proc. ICLR*, 2016, pp. 1–15.
- [44] Q. Qian, J. Tang, H. Li, S. Zhu, and R. Jin, "Large-scale distance metric learning with uncertainty," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8542–8550, doi: [10.1109/CVPR.2018.00891](https://doi.org/10.1109/CVPR.2018.00891).
- [45] Q. Qian et al., "SoftTriple loss: Deep metric learning without triplet sampling," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6449–6457, doi: [10.1109/ICCV.2019.00655](https://doi.org/10.1109/ICCV.2019.00655).
- [46] N. Sohoni, J. A. Dunnmon, G. Angus, A. Gu, and C. Ré, "No subclass left behind: Fine-grained robustness in coarse-grained classification problems," in *Proc. NeurIPS*, 2020, pp. 19339–19352, doi: [10.5555/3495724.3497346](https://doi.org/10.5555/3495724.3497346).
- [47] Y. Zhu, M. Yang, C. Deng, and W. Liu, "Fewer is more: A deep graph metric learning perspective using fewer proxies," in *Proc. NeurIPS*, 2020, pp. 17792–17803, doi: [10.5555/3495724.3497217](https://doi.org/10.5555/3495724.3497217).
- [48] Z. Zheng, L. Zheng, M. Garrett, Y. Yang, M. Xu, and Y.-D. Shen, "Dual-path convolutional image-text embeddings with instance loss," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 16, no. 2, pp. 1–23, May 2020, doi: [10.1145/3383184](https://doi.org/10.1145/3383184).
- [49] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proc. AISTATS*, 2010, pp. 297–304, doi: [10.1.1.161.7404](https://doi.org/10.1.1.161.7404).
- [50] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," in *Proc. NeurIPS*, 2020, pp. 9912–9924, doi: [10.5555/3495724.3496555](https://doi.org/10.5555/3495724.3496555).
- [51] J. Li, P. Zhou, C. Xiong, and S. Hoi, "Prototypical contrastive learning of unsupervised representations," in *Proc. ICLR*, 2021, pp. 1–16.
- [52] J. D. Robinson, C.-Y. Chuang, S. Sra, and S. Jegelka, "Contrastive learning with hard negative samples," in *Proc. ICLR*, 2021, pp. 1–13.
- [53] Y. Kalantidis, M. B. Sariyildiz, N. Pion, P. Weinzaepfel, and D. Larlus, "Hard negative mixing for contrastive learning," in *Proc. NeurIPS*, 2020, pp. 21798–21809, doi: [10.5555/3495724.3497553](https://doi.org/10.5555/3495724.3497553).
- [54] X. Han et al., "Rethinking sampling strategies for unsupervised person re-identification," 2021, *arXiv:2107.03024*.
- [55] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," 2018, *arXiv:1807.03748*.
- [56] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NeurIPS*, 2012, pp. 84–90, doi: [10.1145/3065386](https://doi.org/10.1145/3065386).


- [57] D. Mahajan et al., “Exploring the limits of weakly supervised pre-training,” in *Proc. ECCV*, 2018, pp. 185–201, doi: [10.1007/978-3-030-01216-8_12](https://doi.org/10.1007/978-3-030-01216-8_12).
- [58] Y. Zhang, X.-S. Wei, B. Zhou, and J. Wu, “Bag of tricks for long-tailed visual recognition with deep convolutional neural networks,” in *Proc. AAAI*, 2021, pp. 3447–3455, doi: [10.1609/aaai.v35i4.16458](https://doi.org/10.1609/aaai.v35i4.16458).
- [59] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, “Scalable person re-identification: A benchmark,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1116–1124, doi: [10.1109/ICCV.2015.133](https://doi.org/10.1109/ICCV.2015.133).
- [60] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, “Performance measures and a data set for multi-target, multi-camera tracking,” in *Proc. ECCV*, 2016, pp. 17–35, doi: [10.1007/978-3-319-48881-3_2](https://doi.org/10.1007/978-3-319-48881-3_2).
- [61] Z. Zheng, L. Zheng, and Y. Yang, “Unlabeled samples generated by GAN improve the person re-identification baseline in vitro,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3774–3782, doi: [10.1109/ICCV.2017.405](https://doi.org/10.1109/ICCV.2017.405).
- [62] X. Liu, W. Liu, T. Mei, and H. Ma, “PROVID: Progressive and multimodal vehicle reidentification for large-scale urban surveillance,” *IEEE Trans. Multimedia*, vol. 20, no. 3, pp. 645–658, Mar. 2018, doi: [10.1109/TMM.2017.2751966](https://doi.org/10.1109/TMM.2017.2751966).
- [63] Z. Zhong, L. Zheng, D. Cao, and S. Li, “Re-ranking person re-identification with k-reciprocal encoding,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3652–3661, doi: [10.1109/CVPR.2017.389](https://doi.org/10.1109/CVPR.2017.389).
- [64] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778, doi: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90).
- [65] Y. Sun, L. Zheng, W. Deng, and S. Wang, “SVDNet for pedestrian retrieval,” in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3820–3828, doi: [10.1109/ICCV.2017.410](https://doi.org/10.1109/ICCV.2017.410).
- [66] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, “Bag of tricks and a strong baseline for deep person re-identification,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2019, pp. 1487–1495, doi: [10.1109/CVPRW.2019.00190](https://doi.org/10.1109/CVPRW.2019.00190).
- [67] D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” 2014, *arXiv:1412.6980*.
- [68] A. Paszke et al., “Pytorch: An imperative style, high-performance deep learning library,” in *Proc. NeurIPS*, 2019, pp. 8026–8037, doi: [10.5555/3454287.3455008](https://doi.org/10.5555/3454287.3455008).
- [69] Y. Lin, L. Xie, Y. Wu, C. Yan, and Q. Tian, “Unsupervised person re-identification via softened similarity learning,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 3387–3396, doi: [10.1109/CVPR42600.2020.00345](https://doi.org/10.1109/CVPR42600.2020.00345).
- [70] Z. Wang et al., “CycAs: Self-supervised cycle association for learning re-identifiable descriptions,” in *Proc. ECCV*, 2020, pp. 72–88, doi: [10.1007/978-3-030-58621-8_5](https://doi.org/10.1007/978-3-030-58621-8_5).
- [71] Y. Fu et al., “Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6111–6120, doi: [10.1109/ICCV.2019.00621](https://doi.org/10.1109/ICCV.2019.00621).
- [72] D. Wang and S. Zhang, “Unsupervised person re-identification via multi-label classification,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10978–10987, doi: [10.1109/CVPR42600.2020.01099](https://doi.org/10.1109/CVPR42600.2020.01099).
- [73] Y. Zhai, Q. Ye, S. Lu, M. Jia, R. Ji, and Y. Tian, “Multiple expert brainstorming for domain adaptive person re-identification,” in *Proc. ECCV*, 2020, pp. 594–611, doi: [10.1007/978-3-030-58571-6_35](https://doi.org/10.1007/978-3-030-58571-6_35).
- [74] Y. Sun, L. Zheng, Y. Yang, Q. Tian, and S. Wang, “Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline),” in *Proc. ECCV*, 2018, pp. 501–518, doi: [10.1007/978-3-030-01225-0_30](https://doi.org/10.1007/978-3-030-01225-0_30).
- [75] X. Ni and E. Rahtu, “FlipReID: Closing the gap between training and inference in person re-identification,” in *Proc. 9th Eur. Workshop Vis. Inf. Process. (EUVIP)*, Jun. 2021, pp. 1–6.
- [76] X. Pan, P. Luo, J. Shi, and X. Tang, “Two at once: Enhancing learning and generalization capacities via ibn-net,” in *Proc. ECCV*, 2018, pp. 484–500, doi: [10.1007/978-3-030-01225-0_29](https://doi.org/10.1007/978-3-030-01225-0_29).
- [77] J. Yu and H. Oh, “Unsupervised vehicle re-identification via self-supervised metric learning using feature dictionary,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 3806–3813, doi: [10.1109/IROS51168.2021.9636545](https://doi.org/10.1109/IROS51168.2021.9636545).
- [78] S. Teng, S. Zhang, Q. Huang, and N. Sebe, “Viewpoint and scale consistency reinforcement for UAV vehicle re-identification,” *Int. J. Comput. Vis.*, vol. 129, no. 2385, pp. 719–735, 2020, doi: [10.1007/s11263-020-01402-2](https://doi.org/10.1007/s11263-020-01402-2).
- [79] Z. Zheng, T. Ruan, Y. Wei, Y. Yang, and T. Mei, “VehicleNet: Learning robust visual representation for vehicle re-identification,” *IEEE Trans. Multimedia*, vol. 23, pp. 2683–2693, 2021, doi: [10.1109/TMM.2020.3014488](https://doi.org/10.1109/TMM.2020.3014488).
- [80] J. Zhao, Y. Zhao, J. Li, K. Yan, and Y. Tian, “Heterogeneous relational complement for vehicle re-identification,” in *Proc. ICCV*, 2021, pp. 205–214, doi: [10.1109/ICCV48922.2021.00027](https://doi.org/10.1109/ICCV48922.2021.00027).



Menglin Wang received the B.S. degree from Xidian University, China, in 2016, and the Ph.D. degree from the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China, in 2021. Her research interests include computer vision, unsupervised/semi-supervised learning, and person re-identification.



Jiachen Li received the B.S. degree from the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China, in 2020, where he is currently pursuing the Ph.D. degree. His research interests include person and vehicle re-identification.



Baisheng Lai received the B.S. and Ph.D. degrees from the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China, in 2013 and 2018, respectively. He is currently a Staff Algorithm Engineer at Damo Academy, Alibaba Group, Hangzhou. His research interests include semantic segmentation and image retrieval.



Xiaojin Gong (Member, IEEE) received the B.A. and M.A. degrees from Tsinghua University, Beijing, China, and the Ph.D. degree in electrical and computer engineering from Virginia Tech, Blacksburg, VA, USA, in 2009. She is currently an Associate Professor with the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China. Her research interests include computer vision, image processing, and artificial intelligence.



Xian-Sheng Hua (Fellow, IEEE) received the B.S. and Ph.D. degrees in applied mathematics from Peking University, Beijing, in 1996 and 2001, respectively. In 2001, he joined Microsoft Research Asia as a Researcher. Since 2013, he has been a Senior Researcher at Microsoft Research Redmond. In 2015, he became a Researcher and the Senior Director of the Alibaba Group. He has authored or coauthored over 250 research articles and filed over 90 patents. His research interests include multimedia search, advertising, understanding, and mining, pattern recognition, and machine learning. He is an ACM Distinguished Scientist. He was honored as one of the recipient of MIT35. He served as the Program Co-Chair for the IEEE ICME 2013, the ACM Multimedia 2012, and the IEEE ICME 2012. He served on the Technical Directions Board for the IEEE Signal Processing Society.