

Group-aware Label Transfer for Domain Adaptive Person Re-identification

Kecheng Zheng^{1*}, Wu Liu^{2*}, Lingxiao He², Tao Mei², Jiebo Luo³, Zheng-Jun Zha^{1†}

¹University of Science and Technology of China, ²AI Research of JD, ³University of Rochester
zkcys001@mail.ustc.edu.cn, {liuwu1,helinxiao3,tmei}@jd.com, jluo@cs.rochester.edu, zhazj@ustc.edu.cn

Abstract

Unsupervised Domain Adaptive (UDA) person re-identification (ReID) aims at adapting the model trained on a labeled source-domain dataset to a target-domain dataset without any further annotations. Most successful UDA-ReID approaches combine clustering-based pseudo-label prediction with representation learning and perform the two steps in an alternating fashion. However, offline interaction between these two steps may allow noisy pseudo labels to substantially hinder the capability of the model. In this paper, we propose a Group-aware Label Transfer (GLT) algorithm, which enables the online interaction and mutual promotion of pseudo-label prediction and representation learning. Specifically, a label transfer algorithm simultaneously uses pseudo labels to train the data while refining the pseudo labels as an online clustering algorithm. It treats the online label refinery problem as an optimal transport problem, which explores the minimum cost for assigning M samples to N pseudo labels. More importantly, we introduce a group-aware strategy to assign implicit attribute group IDs to samples. The combination of the online label refining algorithm and the group-aware strategy can better correct the noisy pseudo label in an online fashion and narrow down the search space of the target identity. The effectiveness of the proposed GLT is demonstrated by the experimental results (Rank-1 accuracy) for Market1501→DukeMTMC (82.0%) and DukeMTMC→Market1501 (92.2%), remarkably closing the gap between unsupervised and supervised performance on person re-identification.¹

1. Introduction

Person re-identification (ReID) [27, 50, 19] is the important task of matching person images captured from non-overlapping camera networks, which is widely used in prac-

*Equal contribution.

†Corresponding author.

¹Full codes are available in <https://github.com/JDAI-CV/fast-reid> and <https://github.com/zkcys001/UDAStrongBaseline>.

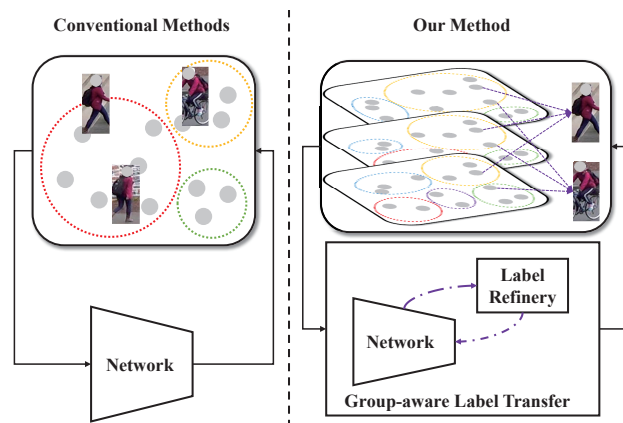


Figure 1. Illustration of conventional methods and our group-aware label transfer method. In our method, each instance is assigned to multiple prototypes with different granularity for generating multi-group pseudo labels, and then the noisy multi-group pseudo labels are online refined. With guidance of refined multi-group pseudo labels, our method can learn an embedding space that encodes the semantic multi-granularity structure of data.

tical applications such as automatic surveillance, content-based retrieval, and behavior analysis [7, 23, 52, 48, 44, 24, 30, 12, 28, 17, 29]. It has been proved that existing approaches can achieve remarkable performance when the training and testing data are collected from the same application scenario but often fail to generalize well to other scenarios due to the inevitable domain gaps. Therefore, it is necessary for both academia and industry to study the Unsupervised Domain Adaptive (UDA) person re-identification problem.

Existing UDA ReID approaches [14, 26, 35, 45, 55] typically include three steps: feature pre-training with labeled source domain data, clustering-based pseudo-label prediction for the target domain data, and feature representation learning/fine-tuning with the pseudo-labels. The last two steps are usually iteratively conducted to strengthen or promote each other. However, the first problem is that the pseudo-labels assigned through clustering usually contain

incorrect labels due to the divergence/domain gap between the source and target data, and the imperfect nature of the clustering algorithm. Such noisy labels may mislead the feature learning and harm the domain adaptation performance. Although the label refining objective in clustering is tractable, it is an offline time-consuming scheme as it requires a pass over the entire dataset. Therefore, online refining those incorrect samples when training can help model learn more robust and accurate representation.

Another problem is that the target domain lacks ID information, so it is difficult to cluster the person images according to human identity. However, in the real world, each person has their own characteristics in his or her appearance. There may be common appearances shared by a group of people but they are not the same identity (e.g. two men with the same red coats and similar black pants as shown in Fig. 1). Therefore, group-based description [21] that involves common characteristics in a pseudo group, can be useful to narrow down the set of candidates and beneficial to identify the exact persons [47, 3, 2]. This group-aware strategy can cluster a person into multi-group clustering prototypes, and is able to efficiently embed a significant number of people and briefly describe an unknown person. Inspired by this, as shown in Fig. 1, combining the online label refining algorithm with the group-aware strategy may be beneficial to the success of domain adaptation.

In this paper, we propose a Group-aware Label Transfer (GLT) algorithm that facilitates the online interaction and mutual promotion of the pseudo labels prediction and feature learning. Specifically, our method simultaneously uses pseudo labels to train the data while online refining the pseudo labels via the label transfer algorithm. This label transfer algorithm regards the resulting label refining problem as optimal transport, which explores the minimum cost for assigning M samples to N pseudo labels. This problem can therefore be solved by the Sinkhorn-Knopp algorithm [6] of linear programming in polynomial time. Meanwhile, the label transfer algorithm is scalable, which can be trained with several batches or the whole dataset and can scale to unlimited amounts of data. More importantly, we introduce a group-aware strategy to assign implicit attribute group IDs to samples. Explicit grouping requires manual annotation. Therefore, we adopt a group-aware clustering algorithm to generate multi-group pseudo labels. By adopting the concept of grouping, the ReID network can reduce the search space and flexibly embed a significant number of identities into an embedding feature. As shown in Fig. 1, we combine the online label refining algorithm with the multi-group pseudo labels. The combination of the online label refining algorithm and the group-aware strategy can better correct the noisy pseudo label online and narrow down the search space of the target identity and predict more accurate pseudo labels. In addition, we design a target instance mem-

ory bank combined with weighted contrastive loss [36] to make the model more powerful for features representation. The proposed GLT framework achieves state-of-the-art performance on Market-to-Duke, Duke-to-Market, Market-to-MSMT, and Duke-to-MSMT with unsupervised domain adaptation, and in particular the *i.e.* Duke-to-Market performances (92.2 Rank-1 and 79.5 mAP) are almost comparable with the supervised learning performances (94.1 Rank-1 and 85.7 mAP). Our method significantly closes gap between unsupervised and supervised performance on person re-identification, *i.e.*, 92.2 vs. 94.1 top-1 accuracy in DukeMTMC to Market1501 transfer.

The main contributions of this paper can be summarized in four aspects:

- We make the first attempt towards integrating clustering and feature learning in a unified framework through the label transfer method for UDA-ReID. It can online refine the predicted pseudo labels to improve the feature representation ability of the model on the target domain.
- We propose a group-aware feature learning strategy based on label transfer to refine multi-group pseudo labels, which provides good latent pseudo-label groups for improving the quality of representation learning.
- The GLT framework achieves significant performance improvements compared to state-of-the-art approaches on Market→Duke, Duke→Market, Market→MSMT, Duke→MSMT ReID tasks. Even for the supervised learning methods, our algorithm is remarkably closing the gap.

2. Related Work

Clustering-based Methods in general generate hard or soft pseudo labels based on clustering results and then fine-tune/train the models based on the pseudo labels [11, 51, 42, 14, 43, 55, 13, 20, 4, 8, 49], which are widely used due to their superior performance. PUL [11] iteratively obtains the hard clustering labels and trains the model in a self-training manner. SSG [42] exploits the potential similarity of the global and local body parts, respectively, to build multiple independent clusters. PAST [51] develop a self-training method with progressive augmentation framework to promote the model performance progressively.

Pseudo label noises caused by unsupervised clustering are always an obstacle to UDA Re-ID. Such noisy labels would mislead the feature learning and impede the achievement of high performance. Recently, some methods introduce mutual learning among two or three collaborative networks to mutually exploit the refined soft pseudo labels of the peer networks as supervision [14, 46]. To suppress the noises in the pseudo labels, NRMT [53] maintains two

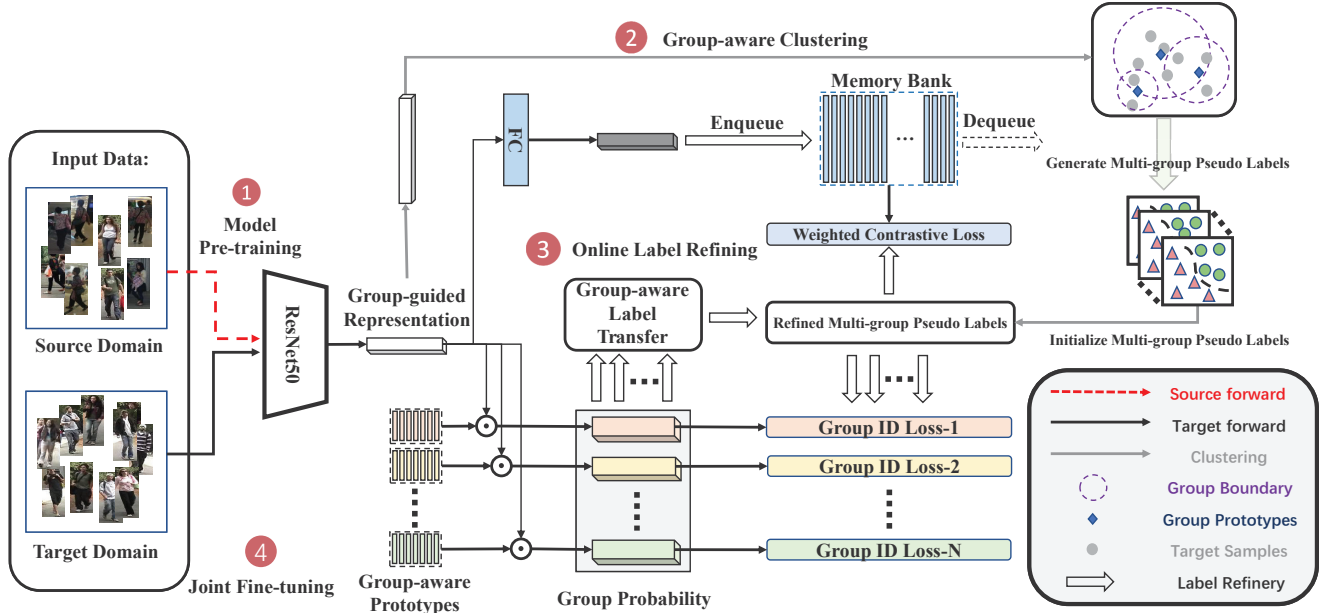


Figure 2. Illustration of our group-aware label transfer framework (GLT), where the group-aware label transfer module and the ReID model are alternately optimized to improve each other towards the final objective of accurate person re-ID. In the model pre-training stage, we pre-train the network on source domain labeled data. Then, in the group-aware clustering stage, we perform clustering algorithm on the unlabeled target data using the multi-group strategy and assign multi-group pseudo labels based on the clustering results. In the label online refining stage, the label transfer algorithm treats the resulting label refining problem as an optimal transport problem, and refine the multi-group pseudo labels by Sinkhorn-Knopp algorithm [6] of linear programming. In the joint fine-tuning stage, we use the refined multi-group pseudo labels to train the networks. Stage 3 and Stage 4 are performed alternately in the online scheme.

networks during training to perform collaborative clustering and mutual instance selection, which reduces the fitting to noisy instances by using mutual supervision and the reliable instance selection. These approaches need mutual learning of two or more networks and are somewhat complicated. Besides, the selection of reliable instances in NRMT [53] is a hard selection, which requires a careful determination of threshold parameters and may lose the chance of exploiting useful information of the abandoned samples.

Domain translation based methods [9, 40] consider to translate target images by the source images. Then they use these translated source-to-target images and their corresponding ground-truth identities to fine-tune the target-domain model. PTGAN [40] utilizes the color consistency to impose pixel-level constraints during the domain translation. SPGAN [9] further minimizes the feature-level similarities between translated images and the original ones. This kind of approach does not apply to real-world scenarios, because the GAN-based method is hardly used to train and may introduce extra computing cost.

Memory bank based methods have been widely used for unsupervised representation learning which facilitates the introduction of contrastive loss [15] for the general tasks. He *et al.* [15] leverage the memory bank to better train the model to exploit the similarity between a sample and the

instances in the global memory bank. Wang *et al.* [39] propose using the memory bank to facilitate hard negative instance mining across batches. ECN [55] leverages memory bank to enforce exemplar-invariance, camera-invariance, and neighborhood-invariance over global training data (instead of the local batch) for UDA person ReID. We introduce contrastive loss to the target instance memory bank to enable the joint optimization of positive pairs and negative pairs for a query/anchor sample over all the samples in the memory bank, which serves as our strong baseline.

3. Group-aware Label Transfer

We first introduce the overall architecture, then derive our group-aware label transfer framework and finally introduce the memory bank with weighted contrastive loss.

3.1. Clustering-based Method Revisit

Formulation. Unsupervised domain adaptive person re-identification aims at adapting the model trained on a source domain dataset with annotations to another unlabeled target domain dataset. Formally, given the source domain data $\mathbb{D}_s = \{(x_i^s, y_i^s) |_{i=1}^{N_s}\}$ with the manual annotation, we can access target domain images $\mathbb{D}_t = \{x_i |_{i=1}^{N_t}\}$ without any

ground-truth labels. N_s and N_t denote the number of the labeled source data and the unlabeled target data, respectively.

Overview. Deep-cluster framework is the general pipeline due to the conciseness and effectiveness of this framework. To be specific, a ReID model $F(\cdot|\theta)$ firstly is pre-trained via a supervised task on source domain images, where θ denotes the parameters of network. Then we adopt this well-pretrained model to extract the features of target domain images which denote as $\{\mathbf{f}_i = F(\mathbf{x}_i|\theta)\}_{i=1}^{N_t} \in \mathbb{R}^D$. After that, the target domain images are grouped into K classes by clustering algorithm. Let $\{\tilde{y}_i\}_{i=1}^{N_t} \in \{1, \dots, K\}$ denote the pseudo labels generated by target domain images $\{\mathbf{x}_i\}_{i=1}^{N_t}$, which contain the noisy label due to the imperfect nature of clustering algorithm.

To calculate a identification loss, a classification head $FC_k : \mathbb{R}^D \rightarrow \mathbb{R}^K$ is adopted to convert the embedding feature to a score vector, in the form of $p(\tilde{y}_i|\mathbf{x}_i) = \text{softmax}(FC(F(\mathbf{x}_i|\theta)))$. The network parameters θ and a learnable target-domain classifier FC_k are then optimized with respect to a cross-entropy loss $\mathcal{L}_{id}(\theta)$ and a triplet loss $\mathcal{L}_{tri}(\theta)$ in the form of,

$$\begin{aligned} \mathcal{L}_{id}(\theta) &= -\frac{1}{m_t} \sum_{i=1}^{m_t} \log p(\tilde{y}_i|\mathbf{x}_i), \\ \mathcal{L}_{tri}(\theta) &= \frac{1}{m_t} \sum_{i=1}^{m_t} \max(0, s_i^n - s_i^p + m), \end{aligned} \quad (1)$$

where m_t is the mini-batch size of target data, s_i^p denotes the hardest positive cosine similarity, s_i^n denotes the hardest negative cosine similarity, scripts i, p and j, n indicate the indices of hardest positive and negative feature in each mini-batch for the sample x_i , and $m = 0$ denotes the triplet distance margin.

Such two steps, pseudo label generation by clustering and feature learning with pseudo labels, are alternated until the training converges. However, because of the domain gap, the pseudo-labels are not always reliable and there are noisy labels. This would mislead the feature representation learning towards learning noisy information. So, we need to adopt an online clustering strategy to refine those samples with wrong pseudo labels. As illustrated in Fig. 2, the framework is alternated to gradually train the ReID model on the target domain with the guidance of pseudo labels.

3.2. Online Clustering by Group-aware Label Transfer

Each image \mathbf{x}_i is transformed into \mathbf{f}_i by a ReID model. We then compute a group probability \mathbf{p}_i from this feature by mapping \mathbf{f}_i to a set of K trainable prototype vectors,

$\mathbf{C} = \{c_1, c_2, \dots, c_K\}$. We denote the matrix by \mathbf{C} whose columns are the $\{c_1, c_2, \dots, c_K\}$. Then, the pseudo labels $\mathbf{Q} \in [0, 1]^{K \times N}$ are adopted to supervise the group probability \mathbf{P} , and we use Label Transfer to online update the pseudo labels as training goes on. We now first introduce how to update the prototypes online.

Prototype Prediction Problem. Our prototypes can be implemented by a non-parametric prototypical classifier or a parametric linear classifier. The optimization can be viewed to optimize the cluster-assignment probability through the cross-entropy loss, where the prototypes \mathbf{C} represent the trainable weights for a linear classifier. With K-means clustering, the linear classifier has a fixed set of weights as the mean vectors for the representations in each cluster.

The prototypes are learned by the non-parametric or a parametric way noted above. Here, we will introduce the prototype prediction problem. Firstly, we rewrite the cross-entropy loss of eq. (1) by encoding the labels as posterior distributions $q(\tilde{y}_i|\mathbf{x}_i) \in \mathbf{Q}$, which is generated by clustering algorithm:

$$\begin{aligned} \mathcal{L}_g(p, q; \theta) &= -\frac{1}{N} \sum_N q(\tilde{y}_i|\mathbf{x}_i) \log p(\tilde{y}_i|\mathbf{x}_i), \\ \text{where } p(\tilde{y}_i|\mathbf{x}_i) &= \frac{\exp(\frac{1}{\tau} \mathbf{f}_i \mathbf{c}_k)}{\sum_{k'} \exp(\frac{1}{\tau} \mathbf{f}_i \mathbf{c}_{k'})}, \end{aligned} \quad (2)$$

where τ is a temperature parameter. If we set the posterior distributions $q(\tilde{y}_i|\mathbf{x}_i) = \delta(\tilde{y}_i - y_i)$ to be deterministic, this eq. (2) is similar to the general cross-entropy loss. We can use a non-parametric prototypical classifier performs or a parametric linear classifier to update the prototypes. This loss function is jointly minimized respect to the prototypes \mathbf{C} and the image encoder parameters θ .

Online Refining Pseudo Labels via Label Transfer. On the Clustering-based UDA methods, inevitable noisy labels caused by the clustering procedure bring a negative effect on network training. To address this issue, we propose a label transfer method to correct the noisy pseudo labels in an online training scheme by combining the representation learning and clustering algorithm. The general pipeline alternately tackles the optimization of \mathbf{q} and representational learning \mathbf{p} . However, separate optimizations cannot well generate accurate pseudo labels in representational learning because of restraining the interaction of these two steps. In order to make our method online, we try to utilize the prototypes and group probability to refine noisy pseudo labels. We compute refinery pseudo labels using the prototypes \mathbf{C} such that all the examples in a batch are equally partitioned by the prototypes. This equipartition constraint ensures that the refinery pseudo labels for different images in several

batches or a whole dataset are distinct, which preventing the trivial solution where every image has the same pseudo label:

$$\min_{p,q} \mathcal{L}_g(p, q; \theta) \forall \tilde{y}_i : q(\tilde{y}_i | \mathbf{x}_i) \in \{0, 1\}$$

$$\sum_{i=1}^N q(\tilde{y}_i = k | \mathbf{x}_i) = \frac{N}{k}, \quad (3)$$

where each sample x_i is assigned to exactly one label and the N data points are split among the K classes for alleviating a trivial solution. Inspired by [1], they enforce an equal partition by constraining the matrix \mathbf{Q} to belong to the transportable prototype. We also propose to adapt their solution to work on the whole dataset by restricting the transportable prototype.

The objective in eq. (3) is combinatorial in q and thus is very difficult to optimize by an end-to-end manner. However, the optimal transport algorithm [6] is able to solve this issue. Let $P_{\tilde{y}_i, i} = p(\tilde{y}_i | \mathbf{x}_i) \frac{1}{N}$ be the (i, \tilde{y}_i) element of $K \times N$ matrix as the joint probabilities estimated by the model, which represents the group probability that the i -th sample is predicted to become the \tilde{y}_i -th pseudo-label. Similarly, $Q_{y_i, i} = q(\tilde{y}_i | \mathbf{x}_i) \frac{1}{N}$ is $K \times N$ matrix as the pseudo label matrix. Using the notation of [6], we relax matrix \mathbf{Q} to be an element of the transportable prototype:

$$H(\mathbf{w}, \mathbf{c}) := \{ \mathbf{Q} \in \mathbb{R}_+^{K \times N} | \mathbf{Q} \cdot \mathbf{1} = \mathbf{w}, \mathbf{Q}^\top \cdot \mathbf{1} = \mathbf{c} \},$$

$$\mathbf{c} = \frac{1}{N} \cdot \mathbf{1}, \mathbf{w} = \frac{1}{N} \cdot \mathbf{1}, \quad (4)$$

where $\mathbf{1}$ are vectors of all ones of the appropriate dimensions, so that \mathbf{c} and \mathbf{w} are the marginal projections of matrix \mathbf{Q} onto its rows and columns, respectively. In our case, we require \mathbf{Q} to be a matrix of conditional probability distributions that is similar to the number distribution of source domain.

In practice, the computational complexity is enormous, including millions of data points and thousands of classes. Traditional algorithms for solving this problem are difficult to be applied. Thus, we adopt the Sinkhorn-Knopp algorithm [6] to address this issue. This amounts to introducing a regularization term:

$$\min_{\mathbf{Q} \in H(\mathbf{w}, \mathbf{c})} \langle \mathbf{Q}, -\log \mathbf{P} \rangle + \frac{1}{\lambda} H(\mathbf{Q}), \quad (5)$$

where $\langle \cdot \rangle$ is the Frobenius dot-product between two matrices, \mathbf{R} denotes to the number distribution of source domain, $\langle \mathbf{Q}, -\log \mathbf{P} \rangle = -\frac{1}{N} \sum_{i=1}^N \sum_{\tilde{y}_i=1}^K q(\tilde{y}_i | \mathbf{x}_i) \log p(\tilde{y}_i | \mathbf{x}_i)$, \log is applied element-wise, H is the entropy function, $H(\mathbf{Q}) = -\sum_{q \in \mathbf{Q}} q \log q$ and λ is a parameter that controls the smoothness of the mapping. We find that using a low λ can have a strong entropy regularization, which generally leads to a trivial solution where all samples collapse into an

Algorithm 1 Group-aware Label Transfer

Require: Source domain data $\mathbb{D}_s = \{(\mathbf{x}_i^s, \mathbf{y}_i^s) |_{i=1}^{N_s}\}$, target domain images $\mathbb{D}_t = \{\mathbf{x}_i |_{i=1}^{N_t}\}$, multi-group k values $\{k_1, \dots, k_M\}$;

- 1: Pre-training source-domain encoder $F(\cdot | \theta)$ by minimizing Eq. (1) on \mathbb{D}_s ;
- 2: Generating multi-group pseudo labels $\{\mathbf{Q}_{k_m}\}_{m=1}^M$ by clustering the target domain features $\{F(\mathbf{x}_i | \theta)\}_{i=1}^{N_t}$;
- 3: **for** i in $[1, iterations]$ **do**
- 4: Updating parameters θ by the gradient descent of objective function equation (11) under the guidance of multi-group pseudo labels $\{\mathbf{Q}_{k_m}\}_{m=1}^M$;
- 5: **for** k_m in $\{k_1, \dots, k_M\}$ **do**
- 6: Extracting the multi-group probability matrix \mathbf{P}_{k_m} of target images \mathbb{D}_t by the current model $F(\cdot | \theta)$;
- 7: **while** $\|\alpha_{j-1} - \alpha_j\|_1 < 0.1$ **do**
- 8: $\alpha_j \leftarrow [\mathbf{P}_{k_m}^\lambda \beta_{j-1}]^{-1}$, $\beta_j \leftarrow [\alpha_{j-1}^\top \mathbf{P}_{k_m}^\lambda]^{-1}$;
- 9: **end while**
- 10: $\mathbf{Q}'_{k_m} = \text{diag}(\alpha) \mathbf{P}_{k_m}^\lambda \text{diag}(\beta)$ # Refining noisy multi-group pseudo labels;
- 11: **end for**
Refining the multi-group pseudo labels $\mathbf{Q}_{k_m} = \mathbf{Q}'_{k_m}$;
- 12: **end for**

unique representation and are all assigned uniformly to all prototypes. Hence, in practice we keep a high λ . According to the Sinkhorn-Knopp algorithm [1], this minimizer of regularization term of eq. (5) can be written as:

$$\mathbf{Q} = \text{diag}(\alpha) \mathbf{P}^\lambda \text{diag}(\beta) = \text{diag}(\alpha) (\mathbf{C}^\top \mathbf{F})^\lambda \text{diag}(\beta), \quad (6)$$

where $\mathbf{F} = [\mathbf{f}_1, \dots, \mathbf{f}_N]$ denotes to the feature vectors of images, $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_K]$ denotes to the prototypes of pseudo labels, α and β are two renormalization vectors of scaling coefficients, the resulting matrix \mathbf{Q} is also a probability matrix that represent the refining pseudo labels (see [6] for a derivation). By the iterative Sinkhorn-Knopp algorithm [1], the vectors α and β can be obtained via a simple matrix scaling iteration:

$$\forall y : \alpha_y \leftarrow [\mathbf{P}^\lambda \beta]_y^{-1} \quad \forall i : \beta_i \leftarrow [\alpha^\top \mathbf{P}^\lambda]_i^{-1}, \quad (7)$$

where $\alpha_0 = \mathbf{c}$ and $\beta_0 = \mathbf{w}$. When λ is very large, optimizing eq. (7) is mostly equivalent to optimizing eq. (6). However, even for intermediate values of λ , the results still tend to approximate the optimal solution [6]. In the original optimal transport problem, λ adjust the balance between the convergence speed and optimum approximation. In our case, we only regard the λ as the hype-parameter, because we are more interested in obtaining a better performance in the final clustering and representation learning results, rather than solving the optimal transport problem.

After using a group probability matrix \mathbf{P} to update the pseudo label matrix \mathbf{Q} , we use the current \mathbf{Q} as the pseudo label to train the ReID model for optimizing the group probability \mathbf{P} . Such two steps, pseudo label prediction and feature learning with pseudo labels are alternated until the training converges.

Group-aware Label Transfer. The actual identification number k of the target domain is unknown, and it is hard to use the one value of k to generate accurate pseudo labels. Therefore, group-based description [21] that involves common characteristics in the pseudo groups, can be useful to narrow down the set of candidates, which is beneficial to identify the exact unlabeled persons. This group-aware strategy can cluster a person into multi-group clustering prototypes, which is able to efficiently embed a significant number of people and briefly describe an unknown person. Then, we introduce the group-aware strategy to assign implicit attribute group IDs to samples.

In our formulation, we firstly adopt the multi-k setting $\{k_1, \dots, k_M\}$ to generate the multi-group pseudo label via multi-k K-means. Then, we use them to train the ReID model:

$$\mathcal{L}_G = -\frac{1}{N} \sum_{m=1}^M \sum_N q_{k_m}(\tilde{y}_i | \mathbf{x}_i) \log p_{k_m}(\tilde{y}_i | \mathbf{x}_i), \quad (8)$$

Based on the multi-group pseudo labels, the group-aware label transfer module adopts the multi-group optimization strategy to refine these pseudo labels. The different groups of pseudo labels have different classification heads, while sharing the parameters of feature extractor θ among them. In formulation, we optimize the multi-group objective functions of the eq. (9).

$$\sum_{m=1}^M \min_{\mathbf{Q}_{k_m} \in H_{k_m}} (\langle \mathbf{Q}_{k_m}, -\log \mathbf{P}_{k_m} \rangle + \frac{1}{\lambda} H(\mathbf{Q}_{k_m})),$$

$$H_{k_m} := \left\{ \mathbf{Q}_{k_m} \in \mathbb{R}_+^{K_m \times N} \mid \mathbf{Q}_{k_m} \cdot \mathbf{1} = \mathbf{r}, \mathbf{Q}_{k_m}^\top \cdot \mathbf{1} = \mathbf{c} \right\}, \quad (9)$$

$$\mathbf{c} = \frac{1}{N} \cdot \mathbf{1}, \mathbf{r} = \frac{1}{N} \cdot \mathbf{1},$$

where $\{\mathbf{P}_{k_m} = \mathbf{C}_{k_m}^\top \cdot \mathbf{F}\}_{m=1}^M$ denotes to the multi-group probability and $\{\mathbf{C}_{k_m}\}_{m=1}^M$ denotes to the multi-group prototypes. Each update involves M matrix-vector multiplication with complexity $\mathcal{O}(M N K)$, so it is relatively quick.

3.3. Target Instance Memory Bank

We first describe our target instance memory bank module, with an updating mechanism. Then we show that a weighting scheme focusing on informative pairs is beneficial for contrastive loss with multiple positive samples.

Updating mechanism. We adopt the memory bank [16] to remember the previous feature embeddings. We maintain the memory bank as a queue of data samples while mini-batches of samples arrive. This allows us to reuse feature embeddings from the immediately preceding mini-batch. The size of our memory banks can be much larger than the typical mini-batch size and can be set flexibly and independently to hyperparameters. The samples in the memory bank are progressively replaced. In addition, the current

mini-batch are enqueued into the dictionary and the oldest mini-batch in the queue are removed. The dictionary always represents a sampled subset of all data, and the extra computation to maintain this dictionary is manageable. In conjunction with this, it is beneficial to remove the oldest mini-batch as it has the most outdated features and is therefore the most inconsistent with the latest key.

Weighted contrastive loss. We use the cosine distance to calculate the similarity s of features between mini-batch and memory bank. Moreover, given a query/anchor sample, we split the samples to L negative pairs and K positive pairs according to the pseudo labels, where weighted contrastive loss [36] is suitable for this issue. Thus, weighted contrastive loss is proposed to better calculate the loss in terms of multiple positive samples, which is beneficial for the optimization of model:

$$\mathcal{L}_{WCL} = \log \left[1 + \sum_{j=1}^L \exp(\gamma \alpha_j^n (s_j^n - m)) \right. \\ \left. \sum_{i=1}^P \exp(-\gamma \alpha_k^p (s_k^p - 1 + m)) \right], \quad (10)$$

where $\alpha_j^n = [m + s_j^n]_+$ and $\alpha_k^p = [1 + m - s_k^p]_+$ are non-negative weighting factors respectively, $[\cdot]_+$ is the ‘‘cut-off at zero’’ operation, m refers to a margin in order to better separate similarity, γ is a scale factor, and L and P denote the number of negative pairs and positive pairs in the memory bank for the sample x_i of mini-batch.

We integrate the losses mentioned above. The total loss could be formulated as:

$$\mathcal{L}_{target} = \lambda_{tri} \mathcal{L}_{TRI} + \lambda_g \mathcal{L}_G + \lambda_{wcl} \mathcal{L}_{WCL}. \quad (11)$$

where λ_{tri} , λ_g , and λ_{wcl} are weighting factors.

4. Experiments

4.1. Datasets

We mainly evaluate our method on three person ReID datasets, including DukeMTMC-reID [34], Market-1501 [54] and MSMT17 [40]. DukeMTMC-reID [34] consists of 36,411 images which include 702 identities for training and 702 identities for testing, where all images are captured from 8 cameras. Market-1501 [54] contains 12,936 images of 751 identities for training and 19,281 images of 750 identities for testing, which are captured by 6 cameras. MSMT [40] contains 126,441 images of 4101 identities, of which 1041 identities are used for training, and these images are captured by 15 cameras. We adopt Mean average precision (mAP) and CMC top-1/5/10 accuracy as evaluation.

Table 1. Performance comparison to the unsupervised domain adaptative person Re-ID state-of-the-art methods on DukeMTMC-reID [34], Market-1501 [54] and MSMT17 [40] datasets. “BOT” denotes to “bag of tricks” method [31], which is a strong baseline in the ReID task.

Methods	DukeMTMC→Market1501				Market1501→DukeMTMC			
	mAP	top-1	top-5	top-10	mAP	top-1	top-5	top-10
UMDL [32] (CVPR’16)	12.4	34.5	52.6	59.6	7.3	18.5	31.4	37.6
PTGAN [40] (CVPR’18)	-	38.6	-	66.1	-	27.4	-	50.7
PUL [11] (TOMM’18)	20.5	45.5	60.7	66.7	16.4	30.0	43.4	48.5
TJ-AIDL [38] (CVPR’18)	26.5	58.2	74.8	81.1	23.0	44.3	59.6	65.0
SPGAN [9] (CVPR’18)	22.8	51.5	70.1	76.8	22.3	41.1	56.6	63.0
ATNet [26](CVPR’19)	25.6	55.7	73.2	79.4	24.9	45.1	59.5	64.2
SPGAN+LMP [10](CVPR’18)	26.7	57.7	75.8	82.4	26.2	46.4	62.3	68.0
CFSM [5] (AAAI’19)	28.3	61.2	-	-	27.3	49.8	-	-
BUC [25] (AAAI’19)	38.3	66.2	79.6	84.5	27.5	47.4	62.6	68.4
ECN [55] (CVPR’19)	43.0	75.1	87.6	91.6	40.4	63.3	75.8	80.4
UCDA [33] (ICCV’19)	30.9	60.4	-	-	31.0	47.7	-	-
PDA-Net [22] (ICCV’19)	47.6	75.2	86.3	90.2	45.1	63.2	77.0	82.5
PCB-PAST [51] (ICCV’19)	54.6	78.4	-	-	54.3	72.4	-	-
SSG [42] (ICCV’19)	58.3	80.0	90.0	92.4	53.4	73.0	80.6	83.2
ACT [41] (AAAI’20)	60.6	80.5	-	-	54.5	72.4	-	-
MPLP+MMCL [37] (CVPR’20)	60.4	84.4	92.8	95.0	51.4	72.4	82.9	85.0
DAAM [18] (AAAI’20)	48.8	71.3	82.4	86.3	53.1	77.8	89.9	93.7
AD-Cluster [45] (CVPR’20)	68.3	86.7	94.4	96.5	54.1	72.6	82.5	85.5
MMT [14] (ICLR’20)	71.2	87.7	94.9	96.9	65.1	78.0	88.8	92.5
NRMT [53](ECCV’20)	71.7	87.8	94.6	96.5	62.2	77.8	86.9	89.5
B-SNR+GDS-H [20](ECCV’20)	72.5	89.3	-	-	59.7	76.7	-	-
MEB-Net [46](ECCV’20)	76.0	89.9	96.0	97.5	66.1	79.6	88.3	92.2
Source Pretrain	33.9	65.3	79.2	84.0	35.2	53.3	67.7	73.1
Our GLT	79.5	92.2	96.5	97.8	69.2	82.0	90.2	92.8
Supervised learning (BOT [31])	85.7	94.1	-	-	75.8	86.2	-	-

Methods	Marke1501→MSMT17				DukeMTMC→MSMT17			
	mAP	top-1	top-5	top-10	mAP	top-1	top-5	top-10
PTGAN [40] (CVPR’18)	2.9	10.2	-	24.4	3.3	11.8	-	27.4
ENC [55] (CVPR’19)	8.5	25.3	36.3	42.1	10.2	30.2	41.5	46.8
SSG [42] (ICCV’19)	13.2	31.6	-	49.6	13.3	32.2	-	51.2
DAAM [18] (AAAI’20)	20.8	44.5	-	-	21.6	46.7	-	-
MMT [14] (ICLR’20)	22.9	49.2	63.1	68.8	23.3	50.1	63.9	69.8
Source Pretrain	7.3	19.8	29.4	34.5	11.2	31.5	42.3	47.7
Our GLT	26.5	56.6	67.5	72.0	27.7	59.5	70.1	74.2
Supervised learning (BOT [31])	48.3	72.3	-	-	48.3	72.3	-	-

4.2. Implementation Details

Training data organization. For our joint training strategy, each mini-batch contains 64 source-domain images of 4 ground-truth identities (16 for each identity) and 64 target-domain images of 4 pseudo identities. The pseudo identities are assigned by the clustering algorithm. All images are resized to 256×128 , and random perturbations are applied to each image, *e.g.* randomly erasing, cropping and flipping.

Hyper-parameters. We tune the hyper-parameters of the proposed framework on the task of Market→Duke, and the chosen hyper-parameters are directly applied to all the other tasks. ADAM optimizer is adopted to optimize the networks with weighting factors $\lambda_{tri} = 1$, $\lambda_g = 1$, $\lambda_{mcl} = 0.05$ and the triplet margin $m = 0.3$. The initial learning rates (lr) are set to 0.00035. The source-domain pre-training iterates for 30 epochs and the learning rate decreases to 1/10 of its previous value every 10 epochs. We use the K-means to initialize the pseudo label, and we utilize the group-aware

label transfer to refine the pseudo label per epoch. We set $k = \{500, 1000, 1500, 2000\}$ for all datasets in the K-means setting.

4.3. Comparison with the State-of-the-art Methods

We compare our proposed GLT with the state-of-the-art methods on four domain adaptation settings in Tab. 1, *i.e.* Duke→Market, Market→Duke, Market→MSMT, and Duke→MSMT. Specifically, we surpass the clustering-based SSG [42] by considerable margins of 21.2%, 15.8%, 13.3%, and 14.4% improvements in terms of mAP on these four tasks with simpler network architectures and no extra computation of local features. DAAM [18] also adopts the clustering-based pipeline with the guidance of attention mechanism and our method shows a noticeable 30.7%, 15.1%, 5.7% and 6.1% improvements in terms of mAP on these four tasks without the help of attention mechanism. Meanwhile, our GLT adopts the one model to

significantly surpass the two-model methods MMT [14] that uses the same backbone, showing a noticeable 8.3%, 4.1%, 3.6% and 4.4% improvements in terms of mAP. Our GLT significantly outperforms the second best UDA method MEB-Net by 3.5% and 3.1% in mAP accuracy, for Duke→Market, Market→Duke, respectively. Moreover, without any manual annotations, our unsupervised domain adaptive method significantly bridges the gap between unsupervised and supervised performance on person re-identification, i.e., 92.2% vs. 94.1% top-1 accuracy in Duke→Market task [31].

4.4. Ablation Studies

In this section, we evaluate each components of our proposed framework by conducting ablation studies on Duke→Market and Market→Duke tasks. The experimental results are shown in Tab. 2.

Effectiveness of proposed components. We conduct extensive ablation studies to investigate the effectiveness of the proposed components in Tab. 2, i.e., label transfer, group-aware label transfer, and target instance memory bank. When introducing the label transfer algorithm, the performance is improved by 7.8% rank-1 accuracy and 6.9% mAP on Duke→Market task. By sensibly adopting memory bank with weighted contrastive loss, better representation learning enables the model to increase the rank-1 accuracy from 88.1% to 89.0% and mAP from 66.4% to 68.2%, respectively. By taking account of the group-aware label transfer, the performance is improved by 11.3% rank-1 accuracy and 3.0% mAP, which indicates that the group-aware label transfer can provides good latent pseudo-label groups for improving the quality of representation learning. In addition, group-aware strategy without the guidance of label transfer bring a little improvement due to the lack of label transfer for pseudo label refinery. This illustrates the importance of group-aware label transfer module, not just group-aware strategy. These experimental results prove the necessity and effectiveness of our group-aware label transfer for correcting noise pseudo labels.

Effectiveness of loss functions. We validate the effectiveness of multi-group ID loss, weighted contrastive loss and triplet loss in our proposed GLT framework. Triplet loss brings 3.2% and 2.1% improvements in terms of mAP, and weighted contrastive loss brings 2.6% and 2.8% improvements in terms of mAP in the Market→Duke and Duke→Market tasks, respectively. Triplet loss utilizes the hardest information within a mini-batch of instance pairs in each iteration. Meanwhile, Weighted contrastive loss in the memory bank mines the global negative and positive instances among the samples of whole dataset. The both losses can help each other.

Table 2. Ablation studies for our proposed framework on individual components. **TMB:** target instance memory bank. **G:** baseline with the multi-group pseudo labels. **LT:** label transfer. **GLT:** group-aware label transfer.

Methods	Duke→Market			
	mAP	top-1	top-5	top-10
Supervised learning	85.7	94.1	-	-
Source Pretrain	38.5	65.7	79.2	84.0
Baseline	59.5	80.3	90.4	93.0
Baseline + G	63.0	84.7	91.4	93.7
Self-labeling [1]	40.9	66.6	80.6	86.3
Baseline + TMB	62.4	82.1	91.8	94.9
Baseline + LT	66.4	88.1	94.4	96.2
Baseline + TMB + LT	68.2	89.0	94.5	96.1
Baseline + G + TMB	65.2	84.3	94.8	96.7
Baseline + G + GLT	78.7	91.4	96.7	97.4
Baseline + G + TMB + GLT	79.5	92.2	96.5	97.8

Table 3. Evaluation of the effectiveness of triplet loss and weighted contrastive loss, “w/o WCL” denotes to using the weighted contrastive loss [16].

Loss Functions	Duke→Market		Market→Duke	
	mAP	top-1	mAP	top-1
GLT w/o TRI loss	77.3	90.9	67.1	81.2
GLT w/o WCL loss	77.9	90.8	66.4	80.7
GLT	79.5	92.2	69.2	82.0

5. Conclusion

To solve the problem of UDA person ReID, we propose a novel framework, the group-aware label transfer method (GLT), to combine the accurate pseudo-label prediction and effective ReID representation learning in one unified optimization objective. The holistic and immediate interaction between these two steps in the training process can greatly help the UDA person ReID task. To achieve this goal, we consider the online label refinement problem as an optimal transport problem, which explores the minimal cost of assigning M samples to N pseudo labels. More importantly, we introduce a group-aware strategy to assign implicit attribute group IDs to samples. The combination of the online label refining algorithm and the group-aware strategy can better correct the noisy pseudo label in an online fashion and narrow down the search space of the target identity. Our method not only achieves state-of-the-art performance but also significantly closes the gap between supervised and unsupervised performance on person re-identification.

6. Acknowledgement

This work was supported by the National Key R&D Program of China under Grand No. 2020AAA0103800 and National Natural Science Foundation of China (NSFC) under Grants U19B2038.

References

- [1] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. 5, 8
- [2] Junjie Cai, Zheng-Jun Zha, Meng Wang, Shiliang Zhang, and Qi Tian. An attribute-assisted reranking model for web image search. *TIP*, 24(1):261–272, 2014. 2
- [3] Junjie Cai, Zheng-Jun Zha, Wengang Zhou, and Qi Tian. Attribute-assisted reranking for web image retrieval. In *ACM MM*, pages 873–876, 2012. 2
- [4] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020. 2
- [5] Xiaobin Chang, Yongxin Yang, Tao Xiang, and Timothy M Hospedales. Disjoint label space transfer learning with common factorised space. In *AAAI*, 2019. 7
- [6] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, pages 2292–2300, 2013. 2, 3, 5
- [7] Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. Generalizable person re-identification with relevance-aware mixture of experts. In *CVPR*, 2021. 1
- [8] Yongxing Dai, Jun Liu, Yan Bai, Zekun Tong, and Ling-Yu Duan. Dual-refinement: Joint label and feature refinement for unsupervised domain adaptive person re-identification. *arXiv preprint arXiv:2012.13689*, 2020. 2
- [9] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *CVPR*, pages 994–1003, 2018. 3, 7
- [10] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *CVPR*, 2018. 7
- [11] Hehe Fan, Liang Zheng, Chenggang Yan, and Yi Yang. Unsupervised person re-identification: Clustering and fine-tuning. *TOMM*, 14(4):1–18, 2018. 2, 7
- [12] Chuang Gan, Naiyan Wang, Yi Yang, Dit-Yan Yeung, and Alex G Hauptmann. Devnet: A deep event network for multimedia event detection and evidence recounting. In *CVPR*, pages 2568–2577, 2015. 1
- [13] Chuang Gan, Ting Yao, Kuiyuan Yang, Yi Yang, and Tao Mei. You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images. In *CVPR*, pages 923–932, 2016. 2
- [14] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020. 1, 2, 7, 8
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 3
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 6, 8
- [17] Lingxiao He, Yinggang Wang, Wu Liu, He Zhao, Zhenan Sun, and Jiashi Feng. Foreground-aware pyramid reconstruction for alignment-free occluded person re-identification. In *ICCV*, pages 8450–8459, 2019. 1
- [18] Yangru Huang, Peixi Peng, Yidong Li Yi Jin, Junliang Xing, and Shiming Ge. Domain adaptive attention model for unsupervised cross-domain person re-identification. In *AAAI*, 2020. 7
- [19] Yukun Huang, Zheng-Jun Zha, Xueyang Fu, and Wei Zhang. Illumination-invariant person re-identification. In *ACM MM*, pages 365–373, 2019. 1
- [20] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Global distance-distributions separation for unsupervised person re-identification. In *ECCV*, 2020. 2, 7
- [21] Yonghyun Kim, Wonpyo Park, Myung-Cheol Roh, and Jongju Shin. Groupface: Learning latent groups and constructing group-based representations for face recognition. In *CVPR*, pages 5621–5630, 2020. 2, 6
- [22] Yu-Jhe Li, Ci-Siang Lin, Yan-Bo Lin, and Yu-Chiang Frank Wang. Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation. In *ICCV*, 2019. 7
- [23] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, pages 2197–2206, 2015. 1
- [24] Ji Lin, Chuang Gan, Kuan Wang, and Song Han. Tsm: Temporal shift module for efficient and scalable video understanding on edge devices. *TPAMI*, 2020. 1
- [25] Yutian Lin, Xuanyi Dong, Liang Zheng, Yan Yan, and Yi Yang. A bottom-up clustering approach to unsupervised person re-identification. In *AAAI*, 2019. 7
- [26] Jiawei Liu, Zheng-Jun Zha, Di Chen, Richang Hong, and Meng Wang. Adaptive transfer network for cross-domain person re-identification. In *CVPR*, pages 7202–7211, 2019. 1, 7
- [27] Jiawei Liu, Zheng-Jun Zha, Xuejin Chen, Zilei Wang, and Yongdong Zhang. Dense 3d-convolutional neural network for person re-identification in videos. *TOMM*, 15(1s):1–19, 2019. 1
- [28] Wu Liu, Cheng Zhang, Huadong Ma, and Shuangqun Li. Learning efficient spatial-temporal gait features with deep learning for human identification. *Neuroinformatics*, 16(3):457–471, 2018. 1
- [29] Xinchun Liu, Wu Liu, Meng Zhang, Jingwen Chen, Lianli Gao, Chenggang Yan, and Tao Mei. Social relation recognition from videos via multi-scale spatial-temporal reasoning. In *CVPR*, pages 3566–3574, 2019. 1
- [30] Xiang Long, Gerard De Melo, Dongliang He, Fu Li, Zhizhen Chi, Shilei Wen, and Chuang Gan. Purely attention based local feature integration for video classification. *TPAMI*, 2020. 1

- [31] H. Luo, W. Jiang, Y. Gu, F. Liu, X. Liao, S. Lai, and J. Gu. A strong baseline and batch normalization neck for deep person re-identification. *TMM*, pages 1–1, 2019. 7, 8
- [32] Peixi Peng, Tao Xiang, Yaowei Wang, Massimiliano Pontil, and Yonghong Tian. Unsupervised cross-dataset transfer learning for person re-identification. In *CVPR*, 2016. 7
- [33] Lei Qi, Lei Wang, Jing Huo, Luping Zhou, Yinghuan Shi, and Yang Gao. A novel unsupervised camera-aware domain adaptation framework for person re-identification. In *ICCV*, 2019. 7
- [34] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCVW*, pages 17–35, 2016. 6, 7
- [35] Liangchen Song, Cheng Wang, Lefei Zhang, Bo Du, Qian Zhang, Chang Huang, and Xinggong Wang. Unsupervised domain adaptive re-identification: Theory and practice. *arXiv preprint arXiv:1807.11334*, 2018. 1
- [36] Yifan Sun, Changmao Cheng, Yuhang Zhang, Chi Zhang, Liang Zheng, Zhongdao Wang, and Yichen Wei. Circle loss: A unified perspective of pair similarity optimization. In *CVPR*, 2020. 2, 6
- [37] Dongkai Wang and Shiliang Zhang. Unsupervised person re-identification via multi-label classification. In *CVPR*, 2020. 7
- [38] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*, pages 2275–2284, 2018. 7
- [39] Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R Scott. Cross-batch memory for embedding learning. In *CVPR*, pages 6388–6397, 2020. 3
- [40] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, pages 79–88, 2018. 3, 6, 7
- [41] Fengxiang Yang, Ke Li, Zhun Zhong, Zhiming Luo, Xing Sun, Hao Cheng, Xiaowei Guo, Feiyue Huang, Rongrong Ji, and Shaozi Li. Asymmetric co-teaching for unsupervised cross domain person re-identification. In *AAAI*, 2019. 7
- [42] Fu Yang, Wei Yunchao, Wang Guanshuo, Zhou Yuqian, Shi Honghui, and Huang Thomas. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *ICCV*, 2019. 2, 7
- [43] Hong-Xing Yu, Wei-Shi Zheng, Ancong Wu, Xiaowei Guo, Shaogang Gong, and Jian-Huang Lai. Unsupervised person re-identification by soft multilabel learning. In *CVPR*, 2019. 2
- [44] Zheng-Jun Zha, Jiawei Liu, Di Chen, and Feng Wu. Adversarial attribute-text embedding for person search with natural language query. *TMM*, 22(7):1836–1846, 2020. 1
- [45] Yunpeng Zhai, Shijian Lu, Qixiang Ye, Xuebo Shan, Jie Chen, Rongrong Ji, and Yonghong Tian. Ad-cluster: Augmented discriminative clustering for domain adaptive person re-identification. In *CVPR*, 2020. 1, 7
- [46] Yunpeng Zhai, Qixiang Ye, Shijian Lu, Mengxi Jia, Rongrong Ji, and Yonghong Tian. Multiple expert brainstorming for domain adaptive person re-identification. In *ECCV*, 2020. 2, 7
- [47] Hanwang Zhang, Zheng-Jun Zha, Shuicheng Yan, Jingwen Bian, and Tat-Seng Chua. Attribute feedback. In *ACM MM*, pages 79–88, 2012. 2
- [48] Jing Zhang and Dacheng Tao. Empowering things with intelligence: A survey of the progress, challenges, and opportunities in artificial intelligence of things. *IEEE Internet of Things Journal*, 2020. 1
- [49] Qiming Zhang, Jing Zhang, Wei Liu, and Dacheng Tao. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *NeurIPS*, pages 433–443, 2019. 2
- [50] Wei Zhang, Shengnan Hu, Kan Liu, and Zhengjun Zha. Learning compact appearance representation for video-based person re-identification. *TCSVT*, 29(8):2442–2452, 2018. 1
- [51] Xinyu Zhang, Jiewei Cao, Chunhua Shen, and Mingyu You. Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In *ICCV*, 2019. 2, 7
- [52] Yiheng Zhang, Dong Liu, and Zheng-Jun Zha. Improving triplet-wise training of convolutional neural network for vehicle re-identification. In *ICME*, pages 1386–1391, 2017. 1
- [53] Fang Zhao, Shengcai Liao, Guo-Sen Xie, Jian Zhao, Kaihao Zhang, and Ling Shao. Unsupervised domain adaptation with noise resistible mutual-training for person re-identification. In *ECCV*, 2020. 2, 3, 7
- [54] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pages 1116–1124, 2015. 6, 7
- [55] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *CVPR*, 2019. 1, 2, 3, 7