

# Global Distance-distributions Separation for Unsupervised Person Re-identification

Xin Jin<sup>1,2\*</sup>    Cuiling Lan<sup>2\*\*</sup>    Wenjun Zeng<sup>2</sup>    Zhibo Chen<sup>1\*\*</sup>

<sup>1</sup> University of Science and Technology of China

<sup>2</sup> Microsoft Research Asia, Beijing, China

jinxustc@mail.ustc.edu.cn    {culan, wezeng}@microsoft.com  
chenzhibo@ustc.edu.cn

**Abstract.** Supervised person re-identification (ReID) often has poor scalability and usability in real-world deployments due to domain gaps and the lack of annotations for the target domain data. Unsupervised person ReID through domain adaptation is attractive yet challenging. Existing unsupervised ReID approaches often fail in correctly identifying the positive samples and negative samples through the distance-based matching/ranking. The two distributions of distances for positive sample pairs (Pos-distr) and negative sample pairs (Neg-distr) are often not well separated, having large overlap. To address this problem, we introduce a global distance-distributions separation (GDS) constraint over the two distributions to encourage the clear separation of positive and negative samples from a global view. We model the two global distance distributions as Gaussian distributions and push apart the two distributions while encouraging their sharpness in the unsupervised training process. Particularly, to model the distributions from a global view and facilitate the timely updating of the distributions and the GDS related losses, we leverage a momentum update mechanism for building and maintaining the distribution parameters (mean and variance) and calculate the loss on the fly during the training. Distribution-based hard mining is proposed to further promote the separation of the two distributions. We validate the effectiveness of the GDS constraint in unsupervised ReID networks. Extensive experiments on multiple ReID benchmark datasets show our method leads to significant improvement over the baselines and achieves the state-of-the-art performance.

**Keywords:** Unsupervised learning, person re-identification, global distance-distributions separation, momentum update, hard mining

## 1 Introduction

Person re-identification aims to identify the same person across images. In recent years, significant progress has been made on fully supervised person re-identification (ReID) [30, 47, 11, 26, 31, 46, 1, 15, 18, 16], where groundtruth labels

---

\* This work was done when Xin Jin was an intern at MSRA.

\*\* Corresponding Author.

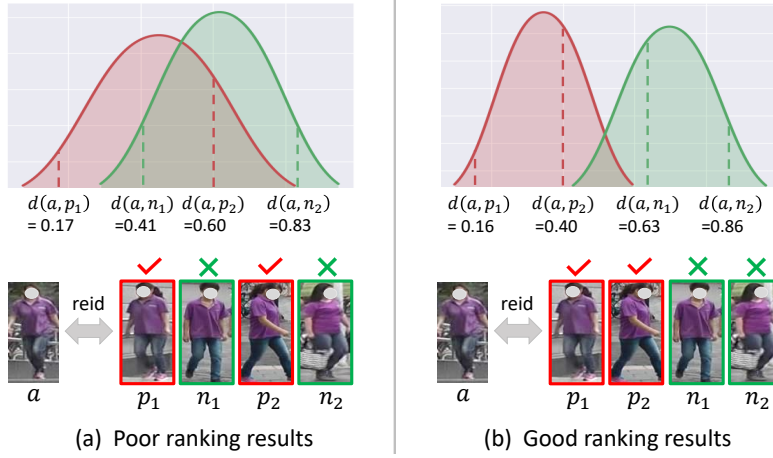


Fig. 1: Illustration of our motivation. The red curves denote the distribution of the distances of positive sample pairs (shorted as Pos-distr) while the green curves denote that of negative sample pairs (shorted as Neg-distr). (a) The two distributions are often not well separated for the current unsupervised ReID models, resulting in poor retrieval/ranking results. (b) Once the two distance distributions are well separated, the positive and negative samples can be well identified based on distance and superior ranking results can be obtained.

are accessible for training. However, they often have poor scalability and usability in real-world deployments. First, they typically perform well on the trained dataset but suffer from significant performance degradation when testing on a previously unseen dataset due to the domain gaps. There are usually large style discrepancies across domains/datasets, due to the discrepancy of imaging devices and environments (*e.g.*, lighting conditions, background, viewing angles) [23]. Second, it is costly to annotate images for each newly deployed environment. One popular solution is unsupervised domain adaptation, which transfers the knowledge from the labeled source domain to the unlabeled target domain.

Many efforts have been made to develop unsupervised domain adaptation for person re-identification (UDA-ReID) [20, 34, 32, 23, 25, 7, 42, 38, 43, 44]. Pseudo label based approaches usually pre-train the network with source domain data and predict pseudo labels for the unlabeled target images, *e.g.*, by clustering, followed by fine-tuning with the pseudo labels [40, 41, 7, 29, 45, 9, 37]. Transfer-based approaches often transfer the labeled source images to have the style of the target domain [35, 5, 23]. These approaches suffer from either noisy labels or noisy images due to incorrect pseudo labels or unrealisticness of the transferred images [37]. Such distractions lead to serious inseparability of positive and negative samples for ReID.

Person ReID inference/test can be considered as a retrieval task, which aims to identify the images of the same person from a gallery image set by comparing

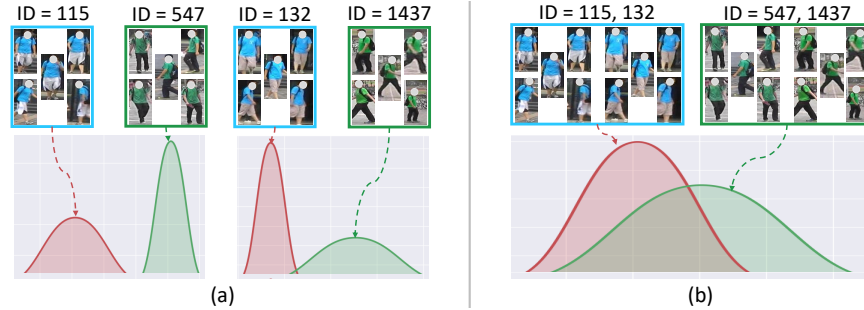


Fig. 2: Local separability is not enough. (a) Within each batch (or subset), the Pos-distr and Neg-distr are well separated. (b) The distributions are inseparable when mixing the two batches together.

feature distances between the query image and the gallery images [49, 39]. Fig. 1(a) illustrates the two distributions of distances of positive sample pairs (red curve) and negative sample pairs (green curve). The inseparability of the distance distributions leads to poor retrieval/ranking results. Given an anchor image  $a$ , as shown in Fig. 1(a), its distance ( $d(a, n_1) = 0.41$ ) to the negative sample  $n_1$  is even smaller than its distance ( $d(a, p_2) = 0.60$ ) to the positive sample  $p_2$ . In comparison, when the two distance distributions are better separated as illustrated in Fig. 1(b), the positive samples can be correctly ranked.

Such poor separability is often observed for unsupervised ReID but is unfortunately under-explored. There is a mismatch between the optimization objective design and the ReID purpose. For ReID inference, it is a *global* ranking/retrieval problem (*e.g.*, based on the feature distances to the query image). The separability of the positive and negative global distance distributions is important for this *global* ranking problem. However, most of the current ReID optimizations/losses are designed from a local perspective. Triplet-based losses [13] optimize the embedding space to encourage the distance of the negative pair to be larger than the positive pair. Such constraint is enforced in three instances, *i.e.*, (an anchor sample, a positive sample, and a negative sample). However, as shown in the examples in Fig. 1(a), both the triplets of  $(a, p_1, n_1)$ , and  $(a, p_2, n_2)$  meet the marginal constraints but the identification of positive/negative samples fails when ranking them all together. In practice, it is impossible to optimize over all possible triplets of a dataset within an acceptable duration. We argue that dataset-wise (*i.e.*, global) constraint is imperative to address this problem. Ustinova *et al.* go a step further and propose a Histogram loss for deep embedding to encourage the separation of distributions within each batch [33]. Unfortunately, as illustrated in Fig. 2, being still a local solution, the separation is easily broken across batches and still cannot guarantee the superiority of the global retrieval performance. Kumar *et al.* use a global loss to optimize the separation of dataset-wise distributions for learning local image descriptors [19]. However,

their design is less efficient which prohibits the timely update of distributions and leads to inaccurate loss calculation.

In this paper, we propose to optimize unsupervised person ReID from a global distance-distribution perspective by introducing a global distance-distributions separation (GDS) constraint for effective ReID feature learning. Different from the local constraints (*e.g.*, triplet loss) for learning embedding features, we model the global (dataset-wise) distributions of the distances for positive pairs (Pos-distr) and negative pairs (Neg-distr) and promote their clear separability. Particularly, we model the two global distance distributions as Gaussian distributions with *updatable mean and variance*. To model the distributions from a global view and ensure timely updates of the distributions and GDS-related losses, we leverage the momentum update mechanism for building and maintaining the distribution variables and calculate the GDS loss on the fly during the training. The GDS loss helps push away the two distributions while encouraging the sharpness of individual distribution. Moreover, to better separate the Pos-distr and Neg-distr, we introduce a distribution-based hard mining by pushing the right-tail of the Pos-distr and the left-tail of the Neg-distr to have a soft margin. We summarize our main contributions as follows:

- We are the first to propose to optimize unsupervised person ReID from a global distance-distribution perspective by encouraging the global separation of positive and negative samples. We address the problem of inseparability of distance distributions in existing unsupervised ReID models by introducing a global distance-distributions separation (GDS) constraint.
- We maintain and update the distribution variables through a momentum update mechanism, enabling the timely update of the distribution variables and accurate estimation of the loss for each batch.
- To further promote the separation of the Pos-distr and Neg-distr, we introduce a distribution-based hard mining mechanism in GDS.

GDS is simple yet effective and can be potentially used as a plug-and-play tool in many unsupervised ReID frameworks for performance enhancement. We validate the effectiveness of our GDS constraint in two representative UDA-ReID approaches, the clustering-based approach with pseudo labelling, and a style transfer-based approach. Extensive experiments are conducted on multiple ReID benchmark datasets and ours achieves the best performance.

## 2 Related Work

### 2.1 Unsupervised Person Re-identification

Unsupervised domain adaptive person ReID aims to learn a ReID model from a labeled source domain and an unlabeled target domain. It is attractive for real-world deployments as it does not require the expensive annotation efforts while exploiting the source domain knowledge. The domain gap between datasets results in poor generalization of a source domain trained model to another domain. Many domain adaptation techniques have been designed for person ReID. They

can mainly be grouped into three categories: *style transfer*, *attribute recognition*, and *clustering-based pseudo label estimation*. Style transfer based approaches translate the source labeled images to ones with the style of the target domain by using image style translation techniques (*e.g.*, Cycle-GAN [56]) for adaptation [35, 5, 23]. Their performance is much inferior to recent clustering-based approaches, since there is still a gap between the generated images and the target domain images [37]. Attribute-based approaches [34, 22] aim to share the source domain knowledge with the target domain through learning of some cues, such as person body attributes, to regularize the feature learning. Such external cues rely on manual annotation and thus limit their applicability. Clustering-based approaches are popular with superior performance [29, 7, 38, 45, 9, 37]. The basic idea is to exploit the similarity of unlabeled samples by feature clustering for predicting pseudo labels and then use them for fine-tuning. Such clustering and training process are usually alternatively performed until the model is stable.

These unsupervised ReID models usually suffer from the interference from noisy pseudo labels or noisy generated images. The inseparability of the distance distributions of positive pairs and negative pairs is serious, there exists undesired small inter-class distances and large intra-class distances (see Fig. 1(a)).

## 2.2 Metric Learning for Person Re-identification

In person ReID tasks, metric learning aims to make the features of the same identity closer while pushing the features of different identities further apart [14, 55, 49, 36, 39]. Loss function designs play the role of metric learning to guide the feature representation learning. For person ReID, there are several widely studied loss functions with their variants. *Identity loss* treats the training process of person ReID as an image classification problem in guiding the feature learning [50]. This enforces a distance/correlation constraint on samples with respect to the “class center” but ignores the direct comparisons between samples. ReID inference is actually a comparison based on the feature similarity/distance between a query image and each image in the gallery set. *Triplet-based loss* and its variants [13] consider the relative ordering of the positive pair distance and negative pair distance by constructing a triplet (an anchor point  $a$ , a positive point  $p$ , a negative point  $n$ ). They optimize the features to encourage the distance between the negative pair to be larger than the distance between the positive pair by a hard or soft margin. The relative ordering of positive pairs and negative pairs matches the retrieval purpose better. In recent years, triplet-based loss is popular and has become one of the default losses for most of the ReID networks [45, 9, 37–39]. However, it only jointly considers the separability among three samples and is far from satisfactory for the global comparison/ranking problem of ReID. Histogram loss [33] goes a step further and encourages the separability of the positive and negative distributions within a batch. As explained in Fig. 2, it is still a local constraint and results in poor separability.

Besides capturing the local structure of data with triplet-based loss, Zhang *et al.* [45] use global information by appending a changeable classification layer

to the model with the number of classes being the number of clusters. This encourages the separability of clustered centers but is less effective in encouraging the separation of distance distributions of positive pairs and negative pairs. Besides, the classification layer needs to be re-trained together with the change of clustering results which may result in a slow and unstable convergence.

To address the poor separability between the distance distributions of the positive pairs and negative pairs observed in existing unsupervised ReID approaches, we introduce a simple yet effective global distributions separation constraint with a momentum updating mechanism. The separation of global distributions has been investigated in learning local image descriptors, with applications in several problems involving the matching of local image patches, such as wide baseline stereo, structure from motion, image classification [19]. Ours is superior to theirs. First, our momentum updating mechanism enables the timely and more accurate updating of losses and networks, leading to optimized performance. In contrast, theirs calculates the distribution statistics from all the samples to obtain the global loss before the next epoch (an epoch means that the entire dataset is passed). This loss is then used for all the mini-batches even though the actual statistics and losses vary with the optimization of each batch. Second, we propose a distribution-based hard mining for effectively promoting the separation. More importantly, we identify one key problem in unsupervised person ReID and address the challenge through a simple yet effective loss design and updating mechanism.

### 3 Unsupervised ReID with GDS Constraint

We aim at enhancing the performance of unsupervised person ReID, by addressing the inseparability of global distance-distributions of positive pairs and negative pairs. As illustrated in Fig. 1, the inseparability could seriously degrade the ReID performance. We propose a global distance-distributions separation (GDS) constraint powered by a momentum update mechanism, which we also refer to as the GDS algorithm. Conceptually, the GDS algorithm is not limited to any specific unsupervised ReID network and it is general. To facilitate the better understanding of the GDS algorithm, we describe it under the representative and popular clustering-based framework for unsupervised person ReID.

Clustering-based approaches in general consist of three stages with the last two stages executed alternatively for many iterations [29, 7, 38, 45, 9, 37]. As shown in Fig. 3, the first stage, *model pre-training*, is to pretrain the ReID network with the source domain labeled data for feature learning. The second stage, *clustering stage*, aims at assigning pseudo labels through the clustering results to the unlabeled target domain data. In the third stage of *adaptation*, the pseudo labels are used for the fine-tuning of the network for domain adaptation, where triplet-based losses are often used for optimization. Triplet-based constraint is enforced on a triplet of instances and cannot assure the correct relative order of the distances for extensive pairs of samples (see Fig. 1(a) for the illustration of the inseparability). We apply our GDS algorithm in the third




Fig. 3: Flowchart of the clustering-based UDA-ReID framework with our Global Distance-distributions Separation (GDS) constraint. In the third stage for adaptation, our GDS constraint with momentum updating optimizes the feature learning by encouraging the separation of the global Pos-distr (red curve) and Neg-distr (green curve). Particularly, we (a) enlarge the distance of the means of the two distributions; (b) encourage the sharpness of each distribution; and (c) introduce distribution-based hard mining to enforce a soft margin between the right-tail of the Pos-distr and the left-tail of the Neg-distr.

stage for more effective feature learning. GDS encourages the separation of the global distance-distributions of positive sample pairs and negative sample pairs.

### 3.1 Global Distance-distribution Modeling with Momentum Update

We have observed the dataset-wise distance distributions of positive sample pairs and negative sample pairs from ReID models (including unsupervised and supervised methods) and found they exhibit Gaussian-like distributions. Therefore, we model the distance distribution of sample pairs by Gaussian distribution with mean  $\mu$  and variance  $\sigma^2$ . We denote the distance distribution of positive sample pairs as  $\mathcal{N}(\mu^+, \sigma^{+2})$  and that of negative sample pairs as  $\mathcal{N}(\mu^-, \sigma^{-2})$ . We aim at designing optimization metrics and strategies to encourage the separation of the two dataset-wise (global) distributions in guiding the feature learning, which would benefit the distance-based ranking across all the images in the gallery set for ReID inference.

Instead of simultaneously exploring all samples, a convolutional neural network is optimized at the mini-batch level, where the loss calculated from the batch is back-propagated to optimize the network to facilitate timely update and avoid the requirement of large memory. Motivated by this, we maintain the

dataset-wise distribution parameters and update them at every batch. This enables timely update of the estimated distribution variables and the corresponding GDS loss, where the distributions change with the updating of the feature extraction network. We formulate the update process here. Particularly, we maintain the global distance distribution of positive sample pairs with variables  $\mu_g^+$  and  $\sigma_g^{+2}$  and that of negative sample pairs with variables  $\mu_g^-$  and  $\sigma_g^{-2}$ . We denote the local mean and variance corresponding to the distances of positive sample pairs within a local batch as  $\mu_l^+$  and  $\sigma_l^{+2}$  (and  $\mu_l^-$  and  $\sigma_l^{-2}$  for negative pairs). We formulate the momentum update of the global distance-distribution for the distance of positive pairs as

$$\left\{ \begin{array}{l} \mu_g^+ \leftarrow \beta \mu_g^+ + (1 - \beta) \mu_l^+ \\ \sigma_g^{+2} \leftarrow \beta \sigma_g^{+2} + (1 - \beta) \sigma_l^{+2} \end{array} \right\}, \text{ with } \left\{ \begin{array}{l} \mu_l^+ = \frac{1}{N^+} \sum_{i=1}^{N^+} d_i^+ \\ \sigma_l^{+2} = \frac{1}{N^+} \sum_{i=1}^{N^+} (d_i^+ - \mu_g^+)^2 \end{array} \right\} \quad (1)$$

where the hyper-parameter  $\beta \in [0, 1]$  is a momentum coefficient,  $N^+$  denotes the number of sampled positive pairs.  $d_i^+$  denotes the Euclidean distance of the  $i^{th}$  positive pair, where  $d(\mathbf{x}_{i_1}, \mathbf{x}_{i_2}) = \frac{1}{2} \|\mathbf{x}_{i_1} - \mathbf{x}_{i_2}\| \in [0, 1]$  for the two normalized feature vectors  $\mathbf{x}_{i_1}$  and  $\mathbf{x}_{i_2}$  of two samples. The update for negative sample pairs is similar. Please see **Supplementary** for the mathematical analysis of the rationality of our momentum update design.

### 3.2 Global Distance-distributions Separation (GDS) Constraint

To encourage the separation of the global distance-distributions of positive sample pairs and negative sample pairs, we design a GDS loss with distribution-based hard mining, *i.e.*, GDS-H loss, which optimizes the network at each mini-batch. The GDS-H loss which consists of the basic GDS loss  $\mathcal{L}_{GDS}$  and distribution-based hard mining loss  $\mathcal{L}_H$  is defined as

$$\mathcal{L}_{GDS-H} = \mathcal{L}_{GDS} + \lambda_h \mathcal{L}_H, \quad (2)$$

where  $\lambda_h$  is a hyper-parameter that balances their importance.

**GDS Loss.**  $\mathcal{L}_{GDS}$  is defined as

$$\mathcal{L}_{GDS} = \text{Softplus}(\mu_g^+ - \mu_g^-) + \lambda_\sigma (\sigma_g^{+2} + \sigma_g^{-2}), \quad (3)$$

where  $\lambda_\sigma$  is another hyper-parameter that balances the importance of mean and variance (please see the study on  $\lambda_h$  and  $\lambda_\sigma$  in the experiment).  $\text{Softplus}(\cdot) = \ln(1 + \exp(\cdot))$ , as a soft-margin formulation, has similar behavior to the hinge function but it decays exponentially instead of having a hard cut-off [13]. As illustrated in Fig. 3, the first item  $\text{Softplus}(\mu_g^+ - \mu_g^-)$  encourages the separation of the centers of the two distributions while the second item  $(\sigma_g^{+2} + \sigma_g^{-2})$  encourages the sharpness of the two distributions by minimizing their variances.

**Distribution-based Hard Mining.** To better promote the separation of the two distributions, as illustrated in Fig. 3(c), we expect the hard samples lying in the potential overlap regions can also be separated. We achieve this by introducing a distribution-based hard mining loss  $\mathcal{L}_H$  as

$$\mathcal{L}_H = \text{Softplus}((\mu_g^+ + \kappa \sigma_g^+) - (\mu_g^- - \kappa \sigma_g^-)), \quad (4)$$



where  $\kappa$  is a hyper-parameter which controls the strength of hard mining. The larger of  $\kappa$ , the more areas of the distribution are covered. Motivated by the *three-sigma rule of thumb* which expresses a conventional heuristic that nearly all values (99.7%) are taken to lie within three times of standard deviations of the mean [12], we set  $\kappa$  to 3 and also experimentally found that this results in superior performance. As illustrated in Fig. 3(c), the physical meaning of Eq. (4) is to encourage the right-tail of the Pos-distr to be smaller than the left-tail of the Neg-distr by a soft margin. This enables the relative ordering of the two *distributions* (instead of only ordering of the centers of the distributions).

**Differentiability of GDS-H Loss for Optimization.** Importantly, the GDS-H loss in Eq. (2) is differentiable w.r.t. the batch-level statistics  $\mu_l$  and  $\sigma_l^2$  and thus the batch samples. To facilitate the description, we denote the means of the maintained global distance-distribution of positive sample pairs as  $\mu_g^+$  and  $\mu_g^{+'}$  before and after the update using the current batch, respectively, where  $\mu_g^{+'} = \beta\mu_g^+ + (1 - \beta)\mu_l^+$ . Similarly, we denote  $\sigma_g^{+2'} = \beta\sigma_g^{+2} + (1 - \beta)\sigma_l^{+2}$ . For ease of presentation, we rewrite the variables  $\sigma_g^{+2'}$  and  $\sigma_g^{+2}$  as  $\rho_g^{+'}$  and  $\rho_g^+$ , respectively. Then  $\rho_g^{+'} = \beta\rho_g^+ + (1 - \beta)\sigma_l^{+2}$ . Similarly, we can define these for the global distance-distribution of negative sample pairs. The gradient of GDS-H loss ( $\mathcal{L}_{GDS-H} = \mathcal{L}$ ) w.r.t. the sample feature  $\mathbf{x}_{i_1}$  is as

$$\begin{aligned} \frac{\partial \mathcal{L}_{GDS-H}}{\partial \mathbf{x}_{i_1}} &= \frac{\partial \mathcal{L}}{\partial \mu_g^{+'}} \frac{\partial \mu_g^{+'}}{\partial \mu_l^+} \frac{\partial \mu_l^+}{\partial d_i^+} \frac{\partial d_i^+}{\partial \mathbf{x}_{i_1}} + \frac{\partial \mathcal{L}}{\partial \rho_g^{+'}} \frac{\partial \rho_g^{+'}}{\partial \sigma_l^{+2}} \frac{\partial \sigma_l^{+2}}{\partial d_i^+} \frac{\partial d_i^+}{\partial \mathbf{x}_{i_1}} \\ &\quad + \frac{\partial \mathcal{L}}{\partial \mu_g^-} \frac{\partial \mu_g^-}{\partial \mu_l^-} \frac{\partial \mu_l^-}{\partial d_i^-} \frac{\partial d_i^-}{\partial \mathbf{x}_{i_1}} + \frac{\partial \mathcal{L}}{\partial \rho_g^-} \frac{\partial \rho_g^-}{\partial \sigma_l^{+2}} \frac{\partial \sigma_l^{+2}}{\partial d_i^-} \frac{\partial d_i^-}{\partial \mathbf{x}_{i_1}}, \end{aligned} \quad (5)$$

this reveals that our method could enable the batch-level gradient back-propagation for timely update of the network for feature learning.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

To evaluate the effectiveness of our GDS constraint for unsupervised person ReID and to be consistent with what were done in prior works for performance comparisons, we conduct extensive experiments on the public ReID datasets, including Market1501 [48], DukeMTMC-reID [51], CUHK03 [21], and the large-scale MSMT17 [35]. We denote Market1501 by M, DukeMTMC-reID by Duke or D, and CUHK03 by C for short. Given a labeled source dataset A and an unlabeled target dataset B, we denote this unsupervised ReID setting as A  $\rightarrow$  B.

We follow common practices and use the cumulative matching characteristics (CMC) at Rank-1, and mean average precision (mAP) for evaluation.

### 4.2 Implementation Details

We build our *Baseline* following the clustering-based pseudo label approach [29], which uses ResNet-50 [2, 46, 24, 37, 29] as the backbone network for feature ex-

Table 1: Effectiveness of the proposed GDS loss and the distribution-based hard mining loss (H).

| (a) Study on the clustering-based method. |             |             |             |             | (b) Study on style transfer based method. |             |             |             |             |
|---|-------------|-------------|-------------|-------------|---|-------------|-------------|-------------|-------------|
| Clustering-based                          | M→D         |             | D→M         |             | Style Transfer based                      | M→D         |             | D→M         |             |
|   | mAP         | Rank-1      | mAP         | Rank-1      |   | mAP         | Rank-1      | mAP         | Rank-1      |
| Direct transfer                           | 19.8        | 35.3        | 21.8        | 48.3        | SPGAN [5]                                 | 22.3        | 41.1        | 22.8        | 51.5        |
| Baseline                                  | 48.4        | 67.1        | 52.1        | 74.3        | SPGAN+ <b>GDS</b>                         | 25.3        | 45.9        | 25.9        | 56.4        |
| Baseline+ <b>GDS</b>                      | 52.9        | 71.4        | 57.1        | 78.5        | SPGAN+ <b>GDS-H</b>                       | <b>27.5</b> | <b>47.7</b> | <b>29.4</b> | <b>60.8</b> |
| Baseline+ <b>GDS-H</b>                    | <b>55.1</b> | <b>73.1</b> | <b>61.2</b> | <b>81.1</b> |   |             |             |             |             |

traction. In the first stage, we pretrain the network using the source dataset, supervised by classification loss [31, 10] and triplet loss with batch hard mining [13]. In the clustering stage, we perform clustering using DBSCAN [6] based on the extracted features of the target dataset, for the purpose of pseudo label assignment. In the adaptation stage, with the pseudo labels, triplet loss with batch hard mining [13] is used to finetune the network. The second stage and the third stage are executed alternatively for 30 iterations. In our schemes, on top of *Baseline*, we add the proposed GDS losses in the adaptation stage. The input image resolution is  $256 \times 128$ . See **Supplementary** for more training details.

### 4.3 Ablation Study

We perform most of the ablation studies based on the powerful clustering-based method, *i.e.*, *Baseline*. We follow the popular settings [9, 37, 54] for ablation study, *i.e.*, using Market1501→Duke (M→D), and Duke→Market1501 (D→M). **Effectiveness of Our GDS Constraint.** Table 1(a) shows the comparisons. *Direct transfer* denotes the case that the network is only trained on source dataset and is directly tested on the target dataset. *Baseline+GDS* denotes our basic scheme where our proposed GDS loss is added on top of *Baseline*. *Baseline+GDS-H* denotes our final scheme where our GDS-H loss (GDS loss + distribution-based hard-mining loss, with momentum update of distributions) is added on top of *Baseline*. We have the following observations.

- 1) Thanks to the encouragement of the separation of the global distance distributions, our final scheme *Baseline+GDS-H* significantly outperforms *Baseline* by **6.7%** and **9.1%** in mAP for M→D and D→M, respectively.
- 2) Our distribution-based hard mining promotes the separation of the two distributions by enabling the relative ordering of the two distributions (instead of only their centers). Such hard mining further brings **2.2%** and **4.1%** improvements in mAP for M→D and D→M (*Baseline+GDS-H* vs. *Baseline+GDS*).
- 3) *Baseline* which uses a clustering-based approach for adaptation outperforms *Direct transfer* by a large margin, indicating the necessity of the adaptation with target domain data.

**Effectiveness Evaluation in terms of ROC and PR Curves.** For some applications of ReID such as people tracking, the determination of whether two

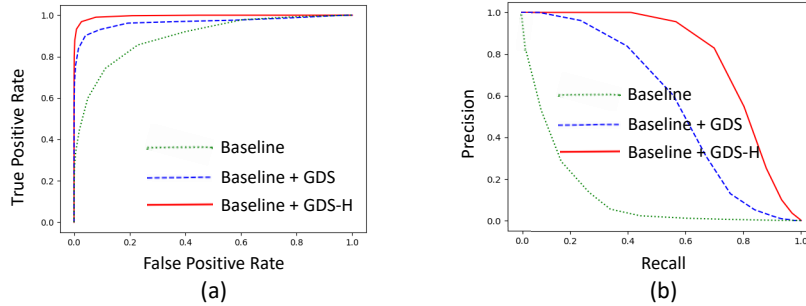


Fig. 4: (a) ROC curve and (b) PR curve on the test set of the target dataset Duke (Market1501→Duke) for *Baseline*, *Baseline+GDS* and *Baseline+GDS-H*.

Table 2: Performance comparison of different loss designs.

| Different losses                    | M→D         |             | D→M         |             |
|-------------------------------------|-------------|-------------|-------------|-------------|
|                                     | mAP         | Rank-1      | mAP         | Rank-1      |
| Baseline                            | 48.4        | 67.1        | 52.1        | 74.3        |
| Baseline + Histogram loss [33]      | 50.2        | 69.4        | 53.7        | 76.3        |
| Baseline + F-Statistic loss [27]    | 51.1        | 70.3        | 55.5        | 76.9        |
| Baseline + Classification loss [45] | 49.7        | 68.9        | 53.8        | 75.9        |
| Baseline + Global loss [19]         | 49.6        | 68.8        | 54.3        | 76.5        |
| Baseline + <b>GDS-H (Ours)</b>      | <b>55.1</b> | <b>73.1</b> | <b>61.2</b> | <b>81.1</b> |

person images match or not degrades to a comparison of their feature distance with a pre-defined threshold  $\theta$  [28, 3]. If the distance is smaller than  $\theta$ , they are judged as the same person. By varying the threshold, we obtain the Receiver Operator Characteristic (ROC) curve [8] and Precision-Recall (PR) curve [4] as shown in Fig. 4. Obviously, our final scheme *Baseline+GDS-H* also significantly outperforms *Baseline* under the threshold-based evaluation metrics.

**Generality of our GDS Constraint.** The GDS constraint is not limited to clustering-based approaches and we also validate its effectiveness under a representative style transfer based ReID approach SPGAN [5]. Table 1(b) shows the results. We observe that **1)** our final scheme *SPGAN+GDS-H* significantly outperforms *SPGAN* by **5.2%** and **6.6%** in mAP for M→D and D→M, respectively; **2)** our distribution-based hard mining brings **2.2%** and **3.5%** improvements in mAP for M→D and D→M, respectively (*SPGAN+GDS-H* vs. *SPGAN+GDS*); **3)** this style transfer based approach *SPGAN* is less effective than the clustering based approach *Baseline* due to the unsatisfactory quality of the transferred images.

**Comparison with Other Losses.** We compare our GDS constraint with several other losses by implementing them on the same network *Baseline* at its third stage (adaptation stage). Table 2 shows the results. Both Histogram loss [33] and F-Statistic loss [27] explore the local (batch-level) statistics for optimization. Histogram loss [33] estimates the similarity (or distance) distributions




Fig. 5: Influence on ReID performance in terms of mAP and Rank-1 accuracy of (a) momentum coefficient  $\beta$  (see Eq. (1) in Subsection 3.1), and (b) strength  $\kappa$  in the distribution-based hard mining for M→D and D→M settings.

of positive sample pairs and negative sample pairs by accumulating the similarity (or distance) values to the bins of two histograms and minimizes their overlap. F-Statistic loss [27] borrows a particular statistic from analysis of variance (ANOVA) hypothesis testing for equality of means. Adding a changeable fully connected layer followed by classification loss (with the number of classes equal to the number of clusters) [45] plays a role of global constraint which encourages the separability of cluster centers. Global loss [19] also encourages the global separation of distance distributions but there is a lack of timely update of distributions, resulting in inaccurate loss calculation and poor optimization.

We observe that: **1)** The performance of our GDS-H loss significantly outperforms the other losses. **2)** With the local separation constraints, Histogram loss and F-Statistic loss both improve the performance over *Baseline* but are inferior to our GDS-H loss. **3)** Adding a global classification loss brings about 1.3%~1.7% improvement in mAP over *Baseline* but are not as effective as ours. **4)** Thanks to the effective momentum update mechanism and distribution-based hard mining, our GDS constraint significantly outperforms [19].

#### 4.4 Design Choices of GDS

**Influence of the Momentum Coefficient  $\beta$ .** For the distribution update,  $\beta$  controls the contribution ratio of the batch-level statistics to the maintained distributions (Eq. (1)). Fig. 5(a) shows its influence. We observe that a relatively larger value (*e.g.*,  $\beta = 0.99$ ) could maintain more stable global distributions and works much better than a smaller value. When  $\beta$  is 0, our GDS loss degrades to a batch-wise (local) distance distribution separation constraint. Such degradation results in a large performance drop in comparison to our GDS ( $\beta = 0.99$ ), from the best 55.1%/61.2% to 51.4%/56.2% in mAP for M→D/D→M, demonstrating the importance of modeling *global* distance distributions.

**Influence of the strength  $\kappa$  in the Distribution-based Hard Mining.** We compare the cases of adding different width offsets ( $\kappa = 1 \sim 4$ ) over the mean as the “hard” region definition (Eq. (4)). In Fig. 5(b), we observe that  $\kappa = 3$  leads to the best performance. Based on *three-sigma rule of thumb*, 99.7% values lie within three times of standard deviations of the mean. This could well cover the entire distribution while excluding the side effects of some extreme outliers.




Fig. 6: Histograms of the distances of the positive sample pairs (red) and negative sample pairs (green) on the test set of the target dataset Duke (Market1501→Duke) for schemes of (a) *Direct transfer*, (b) *Baseline*, (c) *Baseline+GDS*, and (d) *Baseline+GDS-H*. Here we use all the 48018 positive sample pairs and 48018 randomly sampled negative sample pairs for visualization.

#### 4.5 Visualization of Dataset-wise (Global) Distance Distributions

To better understand GDS constraint, we visualize the dataset-wise Pos-distr and Neg-distr in Fig. 6. Thanks to the adaptation on the unlabeled target dataset, the distance distributions of *Baseline* present a much better separability than that of *Direct transfer*. But, there is still a large overlap between the two distributions. By introducing our GDS constraint, our final scheme *Baseline+GDS-H* greatly reduces the overlap of distributions. Besides, our distribution-based hard mining loss is very helpful in promoting the separation ((c) vs. (d)).

#### 4.6 Comparison with State-of-the-Arts

Thanks to the capability of encouraging the separation of the two global distance distributions, our proposed GDS constraint effectively addresses the distance distributions inseparability problem observed in existing unsupervised ReID models. We evaluate the effectiveness of our GDS constraint by comparing with the state-of-the-art approaches on three datasets of Market-1501 (M), DukeMTMC-reID (D) and CUHK03 (C) with six settings in Table 3. More results about the largest dataset MSMT17 can be found in **Supplementary**.

With the effective loss design of GDS, our final scheme *Baseline+GDS-H* achieves the best performance for five dataset settings without any increase in computation complexity in inference. Ours is inferior to *PCB-R-PAST\** [45] only for the C→D setting. But we should not look too much into this comparison as it is not a fair one. First, *PCB-R-PAST\** applied re-ranking [52] as post-operation but we do not. Second, *PCB-R-PAST\** is built on top of a more powerful ReID specific model structure of PCB [31]. Ours uses ResNet-50. Third, the input resolution of *PCB-R-PAST\** is  $384 \times 128$  while ours is  $256 \times 128$ . Our *Baseline+GDS-H* outperforms the second best approach *ACT* [37] on all the settings, achieving 4.6% and 9.9% gain in mAP for D→C and C→D, respectively. Besides, our loss design is simple in implementation and does not require the complicated co-training process used in [37]. Conceptually, our GDS loss is complementary to many approaches like *ACT* [37], *PCB-R-PAST\** [45] and could be used to further improve their performance.

Table 3: Performance (%) comparisons with the state-of-the-art approaches for unsupervised person ReID. \* means applying a re-ranking method of k-reciprocal encoding [52]. Note that *Baseline* is built following [29] with ResNet-50 backbone and thus has nearly the same performance as *Theory* [29].

| Unsupervised ReID | Venue     | M→D  |        | D→M  |        | M→C  |        | D→C  |        | C→M  |        | C→D  |        |
|-------------------|-----------|------|--------|------|--------|------|--------|------|--------|------|--------|------|--------|
|                   |           | mAP  | Rank-1 | mAP  | Rank-1 | mAP  | Rank-1 | mAP  | Rank-1 | mAP  | Rank-1 | mAP  | Rank-1 |
| PTGAN [35]        | CVPR'18   | —    | 27.4   | —    | 38.6   | —    | —      | —    | —      | —    | 31.5   | —    | 17.6   |
| SPGAN [5]         | CVPR'18   | 22.3 | 41.1   | 22.8 | 51.5   | —    | —      | —    | —      | 19.0 | 42.8   | —    | —      |
| TJ-AIDL [34]      | CVPR'18   | 23.0 | 44.3   | 26.5 | 58.2   | —    | —      | —    | —      | —    | —      | —    | —      |
| HHL [53]          | ECCV'18   | 27.2 | 46.9   | 31.4 | 62.2   | —    | —      | —    | —      | 29.8 | 56.8   | 23.4 | 42.7   |
| MAR [42]          | CVPR'19   | 48.0 | 67.1   | 40.0 | 67.7   | —    | —      | —    | —      | —    | —      | —    | —      |
| ECN [54]          | CVPR'19   | 40.4 | 63.3   | 43.0 | 75.1   | —    | —      | —    | —      | —    | —      | —    | —      |
| PAUL [38]         | CVPR'19   | 53.2 | 72.0   | 40.1 | 68.5   | —    | —      | —    | —      | —    | —      | —    | —      |
| SSG [9]           | ICCV'19   | 53.4 | 73.0   | 58.3 | 80.0   | —    | —      | —    | —      | —    | —      | —    | —      |
| PCB-R-PAST* [45]  | ICCV'19   | 54.3 | 72.4   | 54.6 | 78.4   | —    | —      | —    | —      | 57.3 | 79.5   | 51.8 | 69.9   |
| Theory [29]       | PR'2020   | 48.4 | 67.0   | 52.0 | 74.1   | 46.4 | 47.0   | 28.8 | 28.5   | 51.2 | 71.4   | 32.2 | 49.4   |
| ACT [37]          | AAAI'20   | 54.5 | 72.4   | 60.6 | 80.5   | 48.9 | 49.5   | 30.0 | 30.6   | 64.1 | 81.2   | 35.4 | 52.8   |
| Baseline          | This work | 48.4 | 67.1   | 52.1 | 74.3   | 46.2 | 47.0   | 28.8 | 28.4   | 51.2 | 71.4   | 32.0 | 49.4   |
| Baseline+GDS-H    | This work | 55.1 | 73.1   | 61.2 | 81.1   | 49.7 | 50.2   | 34.6 | 36.0   | 66.1 | 84.2   | 45.3 | 64.9   |
| B-SNR[17]         | CVPR'20   | 54.3 | 72.4   | 66.1 | 82.2   | 47.6 | 47.5   | 31.5 | 33.5   | 62.4 | 80.6   | 45.7 | 66.7   |
| B-SNR[17]+GDS-H   | This work | 59.7 | 76.7   | 72.5 | 89.3   | 50.7 | 51.4   | 38.9 | 41.0   | 68.3 | 86.7   | 51.0 | 71.5   |

To further demonstrate the effectiveness of our GDS constraint, we replace the ResNet-50 backbone network of *Baseline* by a more powerful Style Normalization and Restitution (SNR) network [17] (which inserts four light-weight SNR modules to ResNet-50 to tackle the style variation problem for generalizable person ReID) and denote the new baseline scheme as *B-SNR* for simplicity. In Table 3, the scheme *B-SNR+GDS-H* which adopts our GDS constraint also significantly improves the performance of the strong baseline scheme *B-SNR* and achieves the best performance on all settings.

## 5 Conclusions

We propose the use of a Global Distance-distributions Separation (GDS) constraint to enhance the unsupervised person ReID performance. We model the global distance-distributions by Gaussian distributions and encourage their separation. Particularly, we exploit a momentum update mechanism to maintain the variables of the global distributions and enable the timely update of the distributions and the GDS-related loss, facilitating the optimization of the network for each batch. Moreover, we propose distribution-based hard mining to better promote the separation of the distributions. Extensive ablation studies demonstrate the effectiveness of our GDS constraint. We achieve the state-of-the-art performance on the benchmark datasets. The GDS design is simple yet effective. It is conceptually complementary to many of the available approaches and can be taken as a plug-and-play tool for ReID performance enhancement.

## Acknowledgments

This work was supported in part by NSFC under Grant U1908209, 61632001 and the National Key Research and Development Program of China 2018AAA0101400.

## References

1. Ahmed, E., Jones, M., Marks, T.K.: An improved deep learning architecture for person re-identification. In: CVPR (2015)
2. Almazan, J., Gajic, B., Murray, N., Larlus, D.: Re-id done right: towards good practices for person re-identification. arXiv preprint arXiv:1801.05339 (2018)
3. Chen, L., Ai, H., Zhuang, Z., Shang, C.: Real-time multiple people tracking with deeply learned candidate selection and person re-identification. In: ICME. pp. 1–6 (2018)
4. Davis, J., Goadrich, M.: The relationship between precision-recall and roc curves. In: ICML. pp. 233–240 (2006)
5. Deng, W., Zheng, L., Ye, Q., Kang, G., Yang, Y., Jiao, J.: Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In: CVPR (2018)
6. Ester, M., Kriegel, H.P., Sander, J., Xu, X., et al.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: Kdd. vol. 96, pp. 226–231 (1996)
7. Fan, H., Zheng, L., Yan, C., Yang, Y.: Unsupervised person re-identification: Clustering and fine-tuning. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM) (2018)
8. Fawcett, T.: An introduction to roc analysis. Pattern recognition letters **27**(8), 861–874 (2006)
9. Fu, Y., Wei, Y., Wang, G., Zhou, X., Shi, H., Huang, T.S.: Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. ICCV (2019)
10. Fu, Y., Wei, Y., Zhou, Y., et al.: Horizontal pyramid matching for person re-identification. In: AAAI (2019)
11. Ge, Y., Li, Z., Zhao, H., et al.: Fd-gan: Pose-guided feature distilling gan for robust person re-identification. In: NeurIPS (2018)
12. Grafarend, E., Awange, J.: Linear and Nonlinear Models: Fixed Effects, Random Effects, and Total Least Squares. Springer (2012)
13. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. arXiv preprint arXiv:1703.07737 (2017)
14. Hirzer, M., Roth, P.M., Köstinger, M., Bischof, H.: Relaxed pairwise learned metric for person re-identification. In: ECCV. pp. 780–793. Springer (2012)
15. Jia, M., Zhai, Y., Lu, S., Ma, S., Zhang, J.: A similarity inference metric for rgb-infrared cross-modality person re-identification. In: IJCAI (2020)
16. Jin, X., Lan, C., Zeng, W., Chen, Z.: Uncertainty-aware multi-shot knowledge distillation for image-based object re-identification. In: AAAI (2020)
17. Jin, X., Lan, C., Zeng, W., Chen, Z., Zhang, L.: Style normalization and restitution for generalizable person re-identification. In: CVPR (2020)
18. Jin, X., Lan, C., Zeng, W., Wei, G., Chen, Z.: Semantics-aligned representation learning for person re-identification. In: AAAI. pp. 11173–11180 (2020)
19. Kumar, B., Carneiro, G., Reid, I., et al.: Learning local image descriptors with deep siamese and triplet convolutional networks by minimising global loss functions. In: CVPR. pp. 5385–5394 (2016)
20. Li, M., Zhu, X., Gong, S.: Unsupervised person re-identification by deep learning tracklet association. In: ECCV (2018)
21. Li, W., Zhao, R., Tian, L., et al.: Deepreid: Deep filter pairing neural network for person re-identification. In: CVPR (2014)

22. Lin, S., Li, H., Li, C.T., Kot, A.C.: Multi-task mid-level feature alignment network for unsupervised cross-dataset person re-identification. *BMVC* (2018)
23. Liu, J., Zha, Z.J., Chen, D., Hong, R., Wang, M.: Adaptive transfer network for cross-domain person re-identification. In: *CVPR* (2019)
24. Luo, H., Gu, Y., Liao, X., Lai, S., Jiang, W.: Bag of tricks and a strong baseline for deep person re-identification. In: *CVPR workshops* (2019)
25. Qi, L., Wang, L., Huo, J., Zhou, L., Shi, Y., Gao, Y.: A novel unsupervised camera-aware domain adaptation framework for person re-identification. *ICCV* (2019)
26. Qian, X., Fu, Y., Wang, W., et al.: Pose-normalized image generation for person re-identification. In: *ECCV* (2018)
27. Ridgeway, K., Mozer, M.C.: Learning deep disentangled embeddings with the f-statistic loss. In: *NeurIPS*. pp. 185–194 (2018)
28. Ristani, E., Tomasi, C.: Features for multi-target multi-camera tracking and re-identification. In: *CVPR*. pp. 6036–6046 (2018)
29. Song, L., Wang, C., Zhang, L., Du, B., Zhang, Q., Huang, C., Wang, X.: Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognition* p. 107173 (2020)
30. Su, C., Li, J., Zhang, S., et al.: Pose-driven deep convolutional model for person re-identification. In: *ICCV* (2017)
31. Sun, Y., Zheng, L., Yang, Y., Tian, Q., Wang, S.: Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In: *ECCV*. pp. 480–496 (2018)
32. Tang, H., Zhao, Y., Lu, H.: Unsupervised person re-identification with iterative self-supervised domain adaptation. In: *CVPR workshops* (2019)
33. Ustinova, E., Lempitsky, V.: Learning deep embeddings with histogram loss. In: *NeurIPS*. pp. 4170–4178 (2016)
34. Wang, J., Zhu, X., Gong, S., Li, W.: Transferable joint attribute-identity deep learning for unsupervised person re-identification. In: *CVPR* (2018)
35. Wei, L., Zhang, S., Gao, W., Tian, Q.: Person transfer GAN to bridge domain gap for person re-identification. In: *CVPR* (2018)
36. Wojke, N., Bewley, A.: Deep cosine metric learning for person re-identification. In: *WACV*. pp. 748–756 (2018)
37. Yang, F., Li, K., Zhong, Z., Luo, Z., Sun, X., Cheng, H., Guo, X., Huang, F., Ji, R., Li, S.: Asymmetric co-teaching for unsupervised cross domain person re-identification. *AAAI* (2020)
38. Yang, Q., Yu, H.X., Wu, A., Zheng, W.S.: Patch-based discriminative feature learning for unsupervised person re-identification. In: *CVPR* (2019)
39. Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., Hoi, S.C.: Deep learning for person re-identification: A survey and outlook. *arXiv preprint arXiv:2001.04193* (2020)
40. Yu, H.X., Wu, A., Zheng, W.S.: Cross-view asymmetric metric learning for unsupervised person re-identification. In: *ICCV* (2017)
41. Yu, H.X., Wu, A., Zheng, W.S.: Unsupervised person re-identification by deep asymmetric metric embedding. *IEEE TPAMI* (2018)
42. Yu, H.X., Zheng, W.S., Wu, A., Guo, X., Gong, S., Lai, J.H.: Unsupervised person re-identification by soft multilabel learning. In: *CVPR* (2019)
43. Zhai, Y., Lu, S., Ye, Q., Shan, X., Chen, J., Ji, R., Tian, Y.: Ad-cluster: Augmented discriminative clustering for domain adaptive person re-identification. In: *CVPR*. pp. 9021–9030 (2020)
44. Zhai, Y., Ye, Q., Lu, S., Jia, M., Ji, R., Tian, Y.: Multiple expert brainstorming for domain adaptive person re-identification. In: *ECCV* (2020)



45. Zhang, X., Cao, J., Shen, C., You, M.: Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In: ICCV (2019)
46. Zhang, Z., Lan, C., Zeng, W., et al.: Densely semantically aligned person re-identification. In: CVPR (2019)
47. Zhao, H., Tian, M., Sun, S., et al.: Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In: CVPR (2017)
48. Zheng, L., Shen, L., et al.: Scalable person re-identification: A benchmark. In: ICCV (2015)
49. Zheng, L., Yang, Y., Hauptmann, A.G.: Person re-identification: Past, present and future. arXiv preprint arXiv:1610.02984 (2016)
50. Zheng, L., Zhang, H., Sun, S., Chandraker, M., Yang, Y., Tian, Q.: Person re-identification in the wild. In: CVPR. pp. 1367–1376 (2017)
51. Zheng, Z., Zheng, L., Yang, Y.: Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In: ICCV (2017)
52. Zhong, Z., Zheng, L., Cao, D., Li, S.: Re-ranking person re-identification with k-reciprocal encoding. In: CVPR (2017)
53. Zhong, Z., Zheng, L., Li, S., Yang, Y.: Generalizing a person retrieval model hetero- and homogeneously. In: ECCV (2018)
54. Zhong, Z., Zheng, L., Luo, Z., Li, S., Yang, Y.: Invariance matters: Exemplar memory for domain adaptive person re-identification. In: CVPR. pp. 598–607 (2019)
55. Zhou, S., Wang, J., Hou, Q., Gong, Y.: Deep ranking model for person re-identification with pairwise similarity comparison. In: Pacific Rim Conference on Multimedia. pp. 84–94. Springer (2016)
56. Zhu, J.Y., Park, T., Isola, P., et al.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)