

# IDM: An Intermediate Domain Module for Domain Adaptive Person Re-ID

Yongxing Dai<sup>1</sup> Jun Liu<sup>2</sup> Yifan Sun<sup>3</sup> Zekun Tong<sup>4</sup> Chi Zhang<sup>3</sup> Ling-Yu Duan<sup>1,5\*</sup>

<sup>1</sup> Institute of Digital Media (IDM), Peking University, Beijing, China <sup>2</sup> Singapore University of Technology and Design, Singapore

<sup>3</sup> Megvii Technology <sup>4</sup> National University of Singapore, Singapore <sup>5</sup> Peng Cheng Laboratory, Shenzhen, China

{yongxingdai, lingyu}@pku.edu.cn, jun\_liu@sutd.edu.sg, sunyf15@tsinghua.org.cn

## Abstract

Unsupervised domain adaptive person re-identification (UDA re-ID) aims at transferring the labeled source domain’s knowledge to improve the model’s discriminability on the unlabeled target domain. From a novel perspective, we argue that the bridging between the source and target domains can be utilized to tackle the UDA re-ID task, and we focus on explicitly modeling appropriate intermediate domains to characterize this bridging. Specifically, we propose an Intermediate Domain Module (IDM) to generate intermediate domains’ representations on-the-fly by mixing the source and target domains’ hidden representations using two domain factors. Based on the “shortest geodesic path” definition, i.e., the intermediate domains along the shortest geodesic path between the two extreme domains can play a better bridging role, we propose two properties that these intermediate domains should satisfy. To ensure these two properties to better characterize appropriate intermediate domains, we enforce the bridge losses on intermediate domains’ prediction space and feature space, and enforce a diversity loss on the two domain factors. The bridge losses aim at guiding the distribution of appropriate intermediate domains to keep the right distance to the source and target domains. The diversity loss serves as a regularization to prevent the generated intermediate domains from being over-fitting to either of the source and target domains. Our proposed method outperforms the state-of-the-arts by a large margin in all the common UDA re-ID tasks, and the mAP gain is up to 7.7% on the challenging MSMT17 benchmark. Code is available at <https://github.com/SikaStar/IDM>.

## 1. Introduction

Person re-identification (re-ID) [1, 8, 53] aims to identify the person across the non-overlapped cameras. Though fully supervised re-ID methods [41, 45, 3, 60] have achieved




Figure 1. Illustration of our main idea. Assuming that the source and target domains are located in a manifold, there can be some intermediate domains along with the path between the two extreme domains. By utilizing appropriate intermediate domains to bridge the source and target, the source knowledge can be better transferred to the target domain. Different colors and shapes represent different domains and different person identities respectively.

great progress in recent years, it takes much time and effort to label the person images and thus can be difficult to be well applied to some practical scenarios. To overcome this problem, researchers have been focusing on studying unsupervised domain adaptive (UDA) re-ID [48, 13, 38, 2, 9], whose goal is to transfer the knowledge from the labeled source domain to improve the model’s discriminability on the unlabeled target domain. UDA re-ID is a challenging problem because the source and target domains can have two extreme distributions, and there can be no overlap between the two domains’ label space. In this paper, we are devoted to studying the problem of UDA re-ID.

Though many existing methods [48, 13, 16] have made great progress in UDA re-ID, they have not explicitly considered the bridging between the two extreme domains, i.e., what information about the similarity/dissimilarity of the two domains can be utilized to tackle the UDA re-ID task. For example, GAN transferring methods [48, 10] use GANs to translate images’ style across domains and then focus on a single domain’s supervised feature learning. Fine-tuning methods [38, 13, 9] only use the source domain to obtain the pretrained model and then focus on the feature learning in the target domain. Joint training methods [70, 16] just combine the source and target data together for training. From

\*Corresponding Author.

a novel perspective, we explicitly and comprehensively investigate into the bridging between the source and target domains, and argue that using this bridging can help better adapt between two extreme domains in UDA re-ID.

In a spirit of the golden mean, we propose to use the deep model to learn intermediate domains which can bridge the source and target domains and thus ease the UDA re-ID task. The source domain’s knowledge is hard to be directly transferred to the target domain because there can be huge shift between the two extreme domains’ distributions. As shown in Figure 1, assuming the source and target domains are located in a manifold, an appropriate “path” may exist to bridge the two extreme domains along which the source domain’s knowledge can be smoothly transferred to guide the learning of the target domain. Specifically, the appropriate intermediate domains should be located along with this “path” to help the gradual adaptation between the two extreme domains. For example, if an intermediate domain is closer to the source domain, the source reliable labels can be more leveraged. On the contrary, the target domain’s intrinsic distribution can be more exploited.

To model the characteristics of appropriate intermediate domains, we propose a plug-and-play intermediate domain module (IDM). The IDM module can be plugged at any hidden stage of a network, which will generate intermediate domains’ representations on-the-fly to gradually bridge the source and target domains in a joint training framework. Specifically, the IDM module takes some hidden stage’s representations of the source and target domains as the input, and obtains two domain relevance variables that we call domain factors. We use these two domain factors to mix the hidden stage’s representations of both source and target domains, where the mixed representations can represent the characteristics of intermediate domains. All these representations from the source, target and intermediate domains are fed into the network’s next stage on-the-fly.

However, there may be infinite intermediate domains when mixing the source and target domains, and only a part of these domains may be more helpful to smoothly adapt between the two extreme domains (the source and target domains). Inspired by the traditional works [17, 19] that construct intermediate domains using kernel-based methods in a manifold, we first propose the “shortest geodesic path” definition that the appropriate intermediate domains should lie on the shortest geodesic path connecting the source and target domains. From this definition, we further propose two properties where appropriate intermediate domains should satisfy: (1) Keep the right distance to the source and target domains. (2) Be diverse enough to balance the source and target domains’ learning and avoid being over-fitting to either of them. For the first property above, we enforce the bridge losses on both feature and prediction spaces of intermediate domains. For the second property,

we propose a diversity loss by maximizing the standard deviation of the domain factors generated by the IDM module. By enforcing these losses on the learning of the IDM module, we can model the characteristics of appropriate intermediate domains to better transfer the source knowledge to improve the model’s discriminability on the target domain.

Our contributions can be summarized as follows. (1) To the best of our knowledge, we are the first to explicitly consider how to utilize the characteristics of intermediate domains as the bridge to better transfer the source knowledge to the target domain in UDA re-ID. Specifically, we propose a plug-and-play IDM module to generate intermediate domains on-the-fly, which will smoothly bridge the source and target domains to better adapt between the two extremes to ease the UDA re-ID task. (2) To make the IDM module learn more appropriate intermediate domains, we propose two properties of these domains, and design the bridge losses and diversity loss to satisfy these properties. (3) Our method outperforms state-of-the-arts by a large margin on all the common UDA re-ID tasks.

## 2. Related Work

**Unsupervised Domain Adaptation.** A line of works convert the UDA into an adversarial learning task [14, 42, 31, 37]. Another line of works use various metrics to measure and minimize the domain discrepancy, such as MMD [30] or other metrics [39, 72, 32, 24]. Another line of traditional works [17, 7, 19] try to bridge the source and target domains based on intermediate domains. In traditional methods [17, 19], they embed the source and target data into a Grassmann manifold and learn a specific geodesic path between the two domains to bridge the source and target domains, but they are not easily applied to the deep models. In deep methods [18, 6], they either use GANs to generate a domain flow by reconstructing input images on pixel level [18] or learn better domain-invariant features by bridging the learning of the generator and discriminator [6]. However, reconstructing images may not guarantee the high-level domain characteristics in the feature space, and learning domain-invariant features may be not suitable for two extreme domains in UDA re-ID. Different from the above methods, we propose a specific IDM module to model the intermediate domains, which can be easily inserted into the existing deep networks. Instead of hard training for GANs or reconstructing images, our IDM module can be learned in an efficient joint training scheme.

**Unsupervised Domain Adaptive Person Re-ID.** In recent years, many UDA re-ID methods have been proposed and they can be mainly categorized into three types based on their training schemes, *i.e.*, GAN transferring [48, 10, 22, 73], fine-tuning [38, 13, 15, 9, 4, 55, 23, 28, 54], and joint training [69, 70, 44, 16, 11]. GAN transferring methods use GANs to transfer images’ style across domains

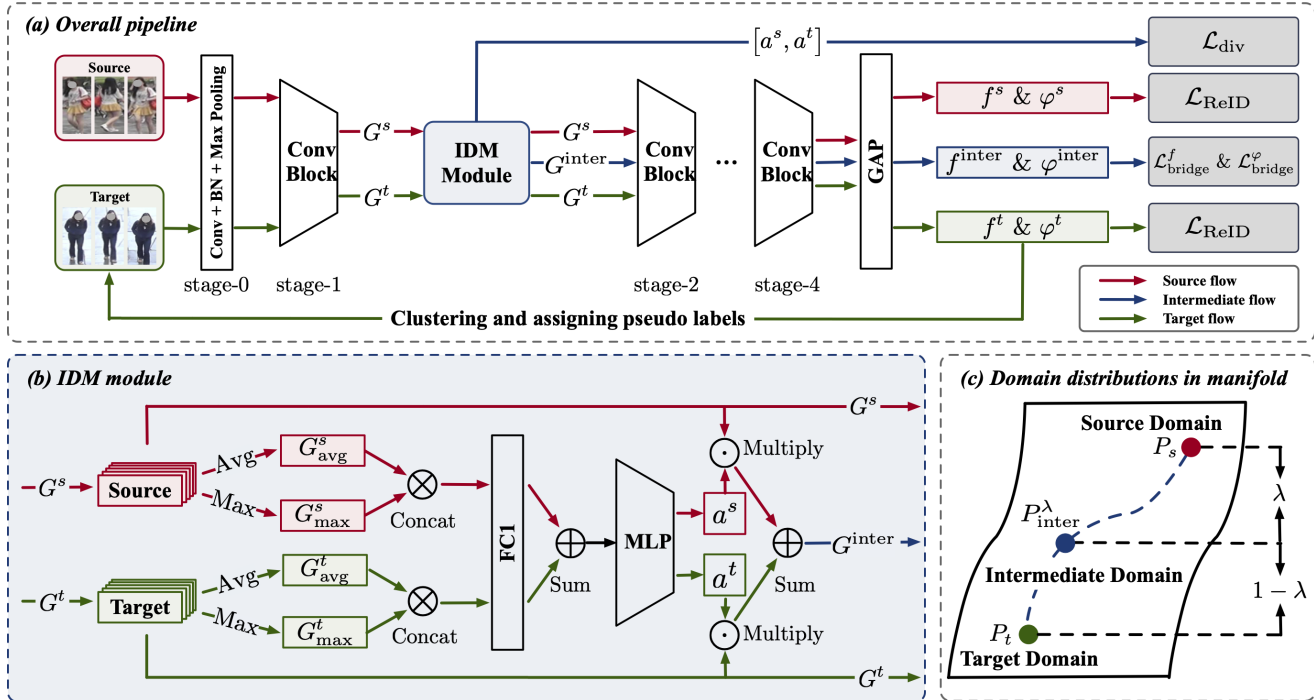


Figure 2. Illustration of our method. (a) In a joint training way, our proposed IDM module can be plugged after any stage (e.g., stage-1) of ResNet-50, where the source and target domains’ hidden representations (i.e.,  $G^s, G^t$ ) are mixed with IDM to generate intermediate domains’ representations  $G^{\text{inter}}$ . The  $G^s, G^t$ , and  $G^{\text{inter}}$  are together fed into the next stage of the same network and the final outputs are their features and predictions:  $(f^s, \varphi^s), (f^t, \varphi^t)$ , and  $(f^{\text{inter}}, \varphi^{\text{inter}})$ . GAP means the global average pooling layer. For simplicity, we omit the classifier. The common ReID loss  $\mathcal{L}_{\text{ReID}}$  (including the classification loss  $\mathcal{L}_{\text{cls}}$  and triplet loss  $\mathcal{L}_{\text{tri}}$ ) are enforced on the source and target domains. To generate more appropriate intermediate domains, we enforce bridge losses ( $\mathcal{L}_{\text{bridge}}^f$  and  $\mathcal{L}_{\text{bridge}}^\varphi$ ) on intermediate domains’  $f^{\text{inter}}$  and  $\varphi^{\text{inter}}$  respectively, and enforce a diversity loss ( $\mathcal{L}_{\text{div}}$ ) to regularize the domain factors  $[a^s, a^t]$  obtained from the IDM module. (b) The detailed structure of the IDM module. “Avg” and “Max” mean average-pooling and max-pooling operations respectively. (c) Assuming different domains as different points in a manifold. An appropriate intermediate domain  $P_{\text{inter}}^\lambda$  should be located along with the shortest geodesic path between the source and target domains (i.e.,  $P_s$  and  $P_t$ ), making  $P_{\text{inter}}^\lambda$  keep the right distance to  $P_s$  and  $P_t$ .

[48, 10] or disentangle features into id-related/unrelated features [73]. For fine-tuning methods, they first train the model with labeled source data and then fine-tune the pre-trained model on target data with pseudo labels. The key component of these methods is how to alleviate the effects of the noisy pseudo labels. However, these methods ignore the labeled source data while fine-tuning on the target data, which will hinder the domain adaptation process because of the catastrophic forgetting in networks. For joint training methods, they combine the source and target data together and train on an ImageNet-pretrained network from scratch. All these joint training methods often utilize the memory bank [51, 50] to improve target domain features’ discriminability. However, these methods just take both the source and target data as the network’s input and train jointly while neglecting the bridging between both domains, i.e., what information of the two domains’ dissimilarities/similarities can be utilized to improve features’ discriminability in UDA re-ID. Different from all the above UDA re-ID methods, we propose to consider the bridging between the source and tar-

get domains by modeling appropriate intermediate domains with a plug-and-paly module, which is helpful for gradually adapting between two extreme domains in UDA re-ID.

**Mixup and Variants.** Mixup [56] is an effective regularization technique to improve the generalization of deep networks by linearly interpolating the image and label pairs, where the interpolating weights are randomly sampled from a Dirichlet distribution. Manifold Mixup [43] extends Mixup to a more general form which can linearly interpolate data at the feature level. Recently, Mixup has been applied to many tasks like point cloud classification [5], object detection [59], and closed-set domain adaptation[52, 49, 34]. Our work differs from these Mixup variants in: (1) All the above methods take Mixup as a data/feature augmentation technique to improve models’ generalization, while we bridge two extreme domains by generating intermediate domains for UDA re-ID. (2) We design an IDM module and enforce specific losses on it to control the bridging process while all the above methods often linearly interpolate data using the random interpolation ratio without constraints.

### 3. Proposed Method

The purpose of our method is to generate appropriate intermediate domains to bridge the source and target domains. By training the source, target, and intermediate domains together at the same time, the network will adaptively characterize the distributions of the source and target domains. This will smoothly adapt between the two extreme domains and better transfer the source knowledge to improve the model’s performance on the target domain.

#### 3.1. Overview

In UDA re-ID, we are often given a labeled source domain dataset  $\{(x_i^s, y_i^s)\}$  and an unlabeled target domain dataset  $\{x_i^t\}$ . The source dataset contains  $N^s$  labeled person images, and the target dataset contains  $N^t$  unlabeled images. Each source image  $x_i^s$  is associated with a person identity  $y_i^s$  and the total number of source domain identities is  $C^s$ . We use ResNet-50 [20] as the backbone network  $f(\cdot)$  and add a hybrid classifier  $\varphi(\cdot)$  after the global average pooling (GAP) layer, where the hybrid classifier is comprised of the batch normalization layer and a  $C^s + C^t$  dimensional fully connected (FC) layer followed by a softmax activation function.  $C^t$  is the cluster number when performing DBSCAN clustering on the target domain features at the beginning of every training epoch like existing clustering-based UDA re-ID methods [13, 38, 16]. After clustering, we assign the pseudo labels  $y_i^t$  for the target data  $x_i^t$ .

Figure 2 (a) shows the pipeline of our method. The overall training scheme is jointly training on the source and target data. A mini-batch including  $n$  source samples and  $n$  target samples is fed to the network. Our proposed Intermediate Domain Module (IDM) can be plugged after any stage of the deep network like ResNet-50. As shown in Figure 2 (a), the IDM module is plugged between the stage-1 and stage-2, where the source and target domains’ hidden representations are mixed with the IDM module to generate new intermediate domains’ representations. Next, all the representations of the source, target, and intermediate domains are fed into the stage-2 until the end of the network. We enforce ReID loss  $\mathcal{L}_{\text{ReID}}$  (including the classification loss [65] and triplet loss [21]) on the source and target domains’ features ( $f^s, f^t$ ) and predictions ( $\varphi^s, \varphi^t$ ). To make the IDM module generate more appropriate intermediate domain representations, we propose two bridge losses and a diversity loss. The bridge losses ( $\mathcal{L}_{\text{bridge}}^f, \mathcal{L}_{\text{bridge}}^\varphi$ ) are respectively enforced on intermediate domains’ features  $f^{\text{inter}}$  and predictions  $\varphi^{\text{inter}}$ . The diversity loss  $\mathcal{L}_{\text{div}}$  is used to enforce on the domain factors obtained from the IDM module, aiming to prevent the module from being overfitting to either of two extreme domains. The IDM module is only used for training and will be discarded in testing.

#### 3.2. Intermediate Domain Module

Our proposed IDM module can be plugged after the hidden stage in the backbone network. The module takes both source and target hidden stage’s representations as the input and generate two domain factors *i.e.*,  $a^s$  and  $a^t$ . With these two domain factors, we can mix the source and target domains’ hidden representations to generate intermediate domains’ representations on-the-fly.

Specifically, we denote backbone network as  $f(x) = f_m(g_m(x))$ , where  $g_m$  represents the part of the network mapping the input data  $x$  to the hidden representation  $g_m(x) \in \mathbb{R}^{h \times w \times c}$  (*i.e.*, feature map) after the  $m$ -th stage and  $f_m$  represents the part of the network mapping the hidden representation  $g_m(x)$  to the 2048-dim feature after the GAP layer. As shown in Figure 2 (b), the IDM module is composed of a fully connected layer FC1, and a MLP (Multi-Layer Perception) followed by a softmax function.

In a mini-batch including  $n$  source and  $n$  target samples, we randomly combine two domains’ samples into  $n$  pairs. For each sample pair  $(x^s, x^t)$ , we can obtain their feature maps at the  $m$ -th hidden stage:  $G^s, G^t \in \mathbb{R}^{h \times w \times c}$ , both of which are performed with the average-pooling and max-pooling operations respectively, generating the  $1 \times 1 \times c$  dimensional features *i.e.*,  $(G_{\text{avg}}^s, G_{\text{max}}^s)$  for the source domain, and  $(G_{\text{avg}}^t, G_{\text{max}}^t)$  for the target domain. We concatenate the *avg* and *max* features for each domain and feed them to FC1. The output feature vectors of FC1 are merged using element-wise summation and fed into MLP to obtain a domain factor vector  $a \in \mathbb{R}^2$ . The procedure to obtain the domain factors in Figure 2 (b) is as follows:

$$a = \delta(\text{MLP}(\text{FC1}([G_{\text{avg}}^s; G_{\text{max}}^s]) + \text{FC1}([G_{\text{avg}}^t; G_{\text{max}}^t])), \quad (1)$$

where  $\delta(\cdot)$  is the softmax function. Next, we define the operation of generating intermediate domain representations after the  $m$ -th hidden stage by mixing the source and target representations with two domain factors as follows:

$$G^{\text{inter}} = a^s \cdot G^s + a^t \cdot G^t, \quad (2)$$

where  $G^s$  and  $G^t$  are the  $m$ -th hidden stage’s representations, and  $a^s$  and  $a^t$  are the two domain factors for the source and target domains respectively ( $[a^s, a^t] = a$ ).

#### 3.3. Modelling Appropriate Intermediate Domains

We can generate infinite intermediate domains by the IDM module in a mixing operation as Eq. (2). However, only a part of these intermediate domains can be appropriate to bridge the source and target domains [19]. Inspired by the traditional works [17, 19] that construct intermediate domains using kernel-based methods, all domains can be assumed to be located in a manifold. Only those intermediate domains located along the shortest geodesic path between the source and target domains can be helpful for the

smooth adaptation procedure [19]. Based on this “shortest geodesic path” definition, we argue that appropriate intermediate domains should satisfy two properties and further propose specifically-designed losses to ensure the two properties in the concept of UDA re-ID. We denote distributions of source, target, and intermediate domains as  $P_s$ ,  $P_t$ , and  $P_{\text{inter}}$ , and use  $d(\cdot)$  to represent geodesic distance in a manifold. As shown in Figure 2 (c), different domains can be seen as different points in a manifold.

**Definition (Shortest geodesic path):** *Appropriate intermediate domains should be located onto the shortest geodesic path between the source and target domains, i.e.,  $d(P_s, P_{\text{inter}}) + d(P_t, P_{\text{inter}}) = d(P_s, P_t)$ .*

Intuitively, for an intermediate domain  $P'_{\text{inter}}$  not on the shortest geodesic path, the triangle inequality ensures that  $d(P_s, P'_{\text{inter}}) + d(P_t, P'_{\text{inter}}) > d(P_s, P_t)$ . As in UDA re-ID, our goal is to reduce the distribution shift  $d(P_s, P_t)$  between the source and target domains and the inappropriate intermediate domain  $P'_{\text{inter}}$  can bring about “extra domain shifts” (i.e., the above inequality) that will degenerate the adaptation procedure.

**Property 1 (Distance should be proportional):** *We use a variable  $\lambda \in [0, 1]$  to control the distance between the source domain to any intermediate domain  $P_{\text{inter}}^{(\lambda)}$  on the shortest geodesic path, i.e.,  $d(P_s, P_{\text{inter}}^{(\lambda)}) = \lambda \cdot d(P_s, P_t)$ . The distances satisfy the proportional relationship:*

$$\frac{d(P_s, P_{\text{inter}}^{(\lambda)})}{d(P_t, P_{\text{inter}}^{(\lambda)})} = \frac{\lambda}{1 - \lambda}. \quad (3)$$

**Sketch of Proof:** From the shortest geodesic path definition, we can derive  $d(P_s, P_{\text{inter}}^{(\lambda)}) + d(P_t, P_{\text{inter}}^{(\lambda)}) = d(P_s, P_t)$ . Then from  $d(P_s, P_{\text{inter}}^{(\lambda)}) = \lambda \cdot d(P_s, P_t)$ , we can get the proportional relationship as Eq. (3).

As shown in Eq. (3), if  $\lambda = 0$  (or  $\lambda=1$ ), the intermediate domain  $P_{\text{inter}}^{(\lambda)}$  is identical to the source (or target) domain  $P_s$  (or  $P_t$ ). In our method, we use the domain factors  $a$  to model the proportional relationship. As shown in Eq. (2), the generated intermediate domains’ representations depend on two domain factors:  $a^s$  and  $a^t$ , where  $a^s, a^t \in [0, 1]$  and  $a^s + a^t = 1$  as we use the softmax function  $\delta$  in Eq. (1). The domain factors  $a$  can be seen as the relevance of the intermediate domain to the other two extreme domains. Thus, the distance relationship (contrary to the relevance relationship) between intermediate domains and other two domains can be formulated as follows:

$$\frac{d(P_s, P_{\text{inter}}^{(a)})}{d(P_t, P_{\text{inter}}^{(a)})} = \frac{a^t}{a^s} = \frac{1 - a^s}{a^s} = \frac{a^t}{1 - a^t}, \quad (4)$$

where  $P_{\text{inter}}^{(a)}$  is the intermediate domain controlled by the domain factors  $a$ . When  $a^s$  gets close to 1, the relevance between  $P_{\text{inter}}^{(a)}$  and  $P_s$  becomes larger while the distance  $d(P_s, P_{\text{inter}}^{(a)})$  becomes smaller.

Based on the above analysis, generating appropriate intermediate domains with domain factors  $a$  becomes finding the points located along the shortest path that should be closest to both  $P_s$  and  $P_t$ , and satisfy Eq. (4). Thus, the problem can be converted into minimizing the loss as:

$$\mathcal{L}_{\text{bridge}} = a^s \cdot d(P_s, P_{\text{inter}}^{(a)}) + a^t \cdot d(P_t, P_{\text{inter}}^{(a)}). \quad (5)$$

If  $a^s > a^t$ ,  $\mathcal{L}_{\text{bridge}}$  will penalize more on  $d(P_s, P_{\text{inter}}^{(a)})$ , otherwise it will penalize more on  $d(P_t, P_{\text{inter}}^{(a)})$ . Intermediate domains can utilize domain factors  $a$  to adaptively balance the minimization of distribution shifts between the source and target domains, resulting in a smooth domain adaptation procedure.

In a deep model for UDA re-ID, we consider to enforce the bridge loss (Eq. (5)) on both the prediction and feature spaces of intermediate domains. For the prediction space, we use the cross-entropy to measure the distribution gap between intermediate domains’ prediction logits and other two extreme domains’ (pseudo) labels (Eq. (6)). For the feature space, we use the L2-norm to measure features’ distance among domains (Eq. (7)). Our proposed two bridge losses are formulated as follows:

$$\mathcal{L}_{\text{bridge}}^{\varphi} = -\frac{1}{n} \sum_{i=1}^n \sum_{k \in \{s,t\}} a_i^k \cdot \left[ y_i^k \log(\varphi(f_m(G_i^{\text{inter}}))) \right], \quad (6)$$

$$\mathcal{L}_{\text{bridge}}^f = \frac{1}{n} \sum_{i=1}^n \sum_{k \in \{s,t\}} a_i^k \cdot \left\| f_m(G_i^k) - f_m(G_i^{\text{inter}}) \right\|_2. \quad (7)$$

In Eq. (6) (7), we use  $k$  to indicate the domain (source or target) and use  $i$  to index the data in a mini-batch.  $G_i^k$  is the  $k$  domain’s representation at the  $m$ -th stage.  $G_i^{\text{inter}}$  is the intermediate domain’s representation at the  $m$ -th hidden stage by mixing  $G_i^s$  and  $G_i^t$  as in Eq. (2). The  $f_m(\cdot)$  is the mapping from the  $m$ -th stage to the features after GAP layer and  $\varphi(\cdot)$  is the classifier.

**Property2 (Diversity):** *Intermediate domains should be as diverse as possible.*

**Sketch of Proof:** Given any intermediate domain  $P_{\text{inter}}^{(\lambda)}$  along the shortest geodesic path from **Property 1**, the intermediate domain will be mainly dominated by the source (or target) domain if  $\lambda \in [0, 0.5)$  (or  $\lambda \in (0.5, 1]$ ). Neither conditions can well bridge the source and target domains, i.e., the shortest geodesic path will be cut off.

To ensure the above property, we propose the diversity loss to enforce on the domain factors  $a$  by maximizing their differences within a mini-batch. This loss is as follows:

$$\mathcal{L}_{\text{div}} = -[\sigma(\{a_i^s\}_{i=1}^n) + \sigma(\{a_i^t\}_{i=1}^n)], \quad (8)$$

where  $\sigma(\cdot)$  means calculating the standard deviation of the values in a mini-batch. By minimizing  $\mathcal{L}_{\text{div}}$ , we can enforce intermediate domains to be as diverse as enough to model the characteristics of the “shortest geodesic path”, which can better bridge the source and target domains.

In conclusion, we use  $\mathcal{L}_{\text{bridge}}^\varphi$  and  $\mathcal{L}_{\text{bridge}}^f$  to enforce generated intermediate domains to be located along the “shortest geodesic path” and use  $\mathcal{L}_{\text{div}}$  to characterize more intermediate domains, which will effectively build a bridge between the two extreme domains. This bridge can adaptively balance the adaptation between the source and target domains when training our network over both the source and target data, *i.e.*, the model can gradually minimize the discrepancy between the two extreme distributions.

### 3.4. Overall Training

The ResNet-50 backbone contains five stages, where stage-0 is comprised of the first Conv, BN and Max Pooling layer and stage-1/2/3/4 correspond to the other four convolutional blocks. We plug our proposed IDM module after the  $m$ -th stage of ResNet-50 (*i.e.*, stage-0/1/2/3/4 as shown in Figure 2). The overall training loss is as follows:

$$\mathcal{L} = \mathcal{L}_{\text{ReID}} + \mu_1 \cdot \mathcal{L}_{\text{bridge}}^\varphi + \mu_2 \cdot \mathcal{L}_{\text{bridge}}^f + \mu_3 \cdot \mathcal{L}_{\text{div}}, \quad (9)$$

where  $\mathcal{L}_{\text{ReID}} = (1 - \mu_1) \cdot \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{tri}}$ , and  $\mu_1, \mu_2, \mu_3$  are the weights to balance losses. The training and testing procedure is shown in Algorithm 1 in the supplementary material.

## 4. Experiments

Five datasets including Market-1501 [64], DukeMTMC-reID [36, 66], MSMT17 [48], PersonX [40], and Unreal [57] are used in our experiments. We use mean average precision (mAP) and Rank-1/5/10 (R1/5/10) of CMC to evaluate performances. In training, we do not use any additional information like temporal consistency in JVTC+ [25]. In testing, there are no post-processing techniques like re-ranking [67] or multi-query fusion [64]. Implementation details can be seen in the supplementary material.

### 4.1. Ablation Study

In this section, we will evaluate the effectiveness of different components of our method. The “Oracle” method in Table 1 means we use the target domain’s ground-truth to train together with the source domain, which can be seen as the upper bound of the performance for UDA re-ID in a joint training pipeline.

**Effectiveness of our proposed IDM module on different baseline methods.** We propose two baseline methods including Naive Baseline (Baseline1) and Strong Baseline (Baseline2). Naive Baseline means only using  $\mathcal{L}_{\text{ReID}}$  to train the source and target data in joint manner as shown in Figure 2. Compared with Naive Baseline, Strong Baseline uses XBM [47] to mine more hard negatives for the triplet loss, which is a variant of the memory bank [51, 50] and easy to implement. Similar to those UDA re-ID methods [70, 16] using the memory bank, we set the memory bank size as the number of all the training data. As shown in Table 1, no matter which baseline we use, our methods can

Table 1. Ablation studies on different components of our method. Baseline1 (Naive Baseline): only using  $\mathcal{L}_{\text{ReID}}$  to train the source and target domains jointly. Baseline2 (Strong Baseline): Baseline1 + XBM [47]. D/M: Duke/Market.

Method	D → M		M → D	
	mAP	R1	mAP	R1
Oracle	83.9	93.2	75.0	86.1
Baseline1	77.0	90.6	63.4	78.4
Baseline1 + Our IDM w/o $\mathcal{L}_{\text{bridge}}^\varphi$	79.4	91.5	66.2	79.8
Baseline1 + Our IDM w/o $\mathcal{L}_{\text{bridge}}^f$	79.2	91.2	65.8	80.4
Baseline1 + Our IDM w/o $\mathcal{L}_{\text{div}}$	78.9	91.1	64.8	79.3
Baseline1 + Our IDM (full)	81.9	92.4	68.6	82.3
Oracle + XBM	86.9	94.8	78.1	88.3
Baseline2	79.1	91.2	65.8	80.1
Baseline2 + Our IDM w/o $\mathcal{L}_{\text{bridge}}^\varphi$	80.7	92.0	67.3	81.8
Baseline2 + Our IDM w/o $\mathcal{L}_{\text{bridge}}^f$	81.6	92.2	69.0	82.0
Baseline2 + Our IDM w/o $\mathcal{L}_{\text{div}}$	79.7	91.6	67.8	81.6
Baseline2 + Our IDM (full)	82.8	93.2	70.5	83.6

obviously outperform the baseline methods by a large margin. Take Market→Duke as an example, mAP/R1 of Baseline2+our IDM (full) is 4.7%/3.5% higher than Baseline2, and mAP/R1 of Baseline1+our IDM (full) is 5.2%/3.9% higher than Baseline1. Because of many state-of-the-arts methods [16, 62, 25, 63] use the memory bank to improve the performance on the target domain, we use Strong Baseline to implement our IDM for fairly comparing with them.

**Effectiveness of the bridge losses for the prediction and feature spaces.** The bridge losses are enforced on both the prediction and feature spaces of intermediate domains. With both losses, we can enforce the generated intermediate domains to be located onto the geodesic path between the source and target domains in a manifold, which are more appropriate to bridge domains to perform more smooth domain adaptation. Take Market→Duke as an example, mAP/R1 of Baseline2+IDM w/o  $\mathcal{L}_{\text{bridge}}^\varphi$  is 3.2%/1.8% lower than Baseline2+IDM, and mAP/R1 of Baseline2+IDM w/o  $\mathcal{L}_{\text{bridge}}^f$  is 1.5%/1.6% lower than Baseline2+IDM. From Table 1, the method w/o  $\mathcal{L}_{\text{bridge}}^f$  is superior to that w/o  $\mathcal{L}_{\text{bridge}}^\varphi$  for Baseline2, while the conclusion is just opposite for Baseline1. The reason may be that Baseline2 uses XBM to mine harder negatives in the feature space, which will affect the effectiveness of  $\mathcal{L}_{\text{bridge}}^f$ .

**Effectiveness of the diversity loss.** We use the diversity loss  $\mathcal{L}_{\text{div}}$  to prevent the IDM module from being over-fitting to either of the source and target domains when generating intermediate domains. Baseline2+IDM w/o  $\mathcal{L}_{\text{div}}$  degenerates much when comparing with the full method. Take Market → Duke as an example (Table 1), mAP of Strong Baseline+IDM w/o  $\mathcal{L}_{\text{div}}$  is 2.7% lower than the full method.

**Effectiveness on which stage to plug the IDM module.** Our IDM module is a plug-and-play module which can be plugged after any stage of the backbone network. In our experiments, we use ResNet-50 as the backbone, which has five stages: stage-0 is comprised of the first Conv, BN

Table 2. Study about which stage of ResNet-50 to plug the IDM module.

Method	Duke → Market		Market → Duke	
	mAP	R1	mAP	R1
Strong Baseline	79.1	91.2	65.8	80.1
IDM after stage-0	<b>82.8</b>	<b>93.2</b>	<b>70.5</b>	<b>83.6</b>
IDM after stage-1	82.3	92.2	70.3	83.2
IDM after stage-2	81.9	92.0	69.2	82.3
IDM after stage-3	81.3	91.8	68.6	82.1
IDM after stage-4	80.2	91.3	68.4	81.5

Table 3. Comparison with different backbones in our method.

Source	Target	w/ ResNet-50		w/ IBN-ResNet-50	
		mAP	R1	mAP	R1
Duke	Market	82.8	93.2	<b>83.9</b>	<b>93.4</b>
Market	Duke	70.5	83.6	<b>71.1</b>	<b>83.9</b>
Duke	MSMT	35.4	63.6	<b>40.8</b>	<b>69.6</b>
Market	MSMT	33.5	61.3	<b>39.3</b>	<b>67.2</b>

and Max Pooling layer and stage-1/2/3/4 correspond to the other four convolutional blocks. We plug our IDM module after different stages to study how different stages will affect the performance of the IDM module. No matter after which stage, our method can all outperform the strong baseline. As shown in Table 2, the performance of the IDM module is gradually declining with the deepening of the network. This phenomenon shows that the domain gap becomes larger and the transferable ability becomes weaker at higher/deeper layers of the network, which satisfies the theory of the domain adaptation [29]. Based on the above analyses, we plug the IDM module after stage-0 of the backbone in all our experiments if not specified.

**Ours vs. traditional mixup methods.** We compare our method with the traditional Mixup [56] and Manifold Mixup (M-Mixup) [43] methods in Figure 3. We use Mixup to randomly mix the source and target domains at the image-level, and use M-Mixup to randomly mix the two domains at the feature-level. The interpolation ratio in Mixup and M-Mixup is randomly sampled from a beta distribution  $\text{Beta}(\alpha, \alpha)$ . Specifically, Mixup and M-Mixup are only the image/feature augmentation technology and we use them to randomly mix the source and target domains. Different from them, our method can adaptively bridge the source and target domains to ensure the two properties that appropriate intermediate domains should satisfy. Thus, our method is superior to these mixup methods in UDA re-ID.

**Scalability for different backbones.** Our proposed IDM module can be easily scalable to other backbones, such as IBN-ResNet-50 [35]. As shown in Table 3, the performance gain is especially obvious on the largest and challenging MSMT dataset. In addition to IBN [35], our method can also be easily implemented with other normalization techniques like AdaBN [26] and CBN [71].

## 4.2. Comparison with the State-of-the-arts

The existing state-of-the-arts (SOTA) UDA re-ID works commonly evaluate the performance on four real → real

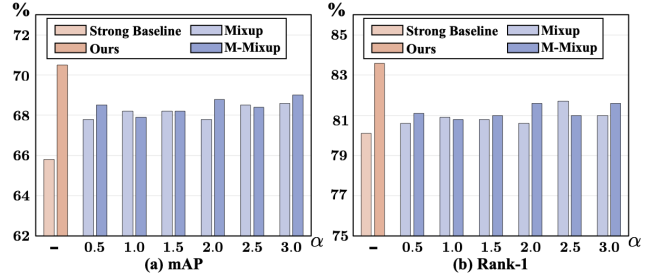


Figure 3. Comparison with traditional mixup methods (Mixup [56], M-Mixup [43]) on Market→Duke. The interpolation ratio in these mixup methods is randomly sampled from a beta distribution  $\text{Beta}(\alpha, \alpha)$ .

tasks [13, 70]. Recently, more challenging synthetic → real tasks [16, 57] are proposed, where they use the synthetic dataset PersonX [40] or Unreal [57] as the source domain and test on other three real re-ID datasets. Using ResNet-50 as the backbone, all the results in Table 4, 5 show that our method can outperform the SOTAs significantly.

**Comparisons on real → real UDA re-ID tasks.** The existing UDA re-ID methods evaluated on real → real UDA tasks can be mainly divided into three categories based on their training schemes. (1) GAN transferring methods include PTGAN [48], SPGAN+LMP [10], [68], and PDA-Net [27]. (2) Fine-tuning methods include PUL [12], PCB-PAST [58], SSG [13], AD-Cluster [54], MMT [15], NRMT [61], MEB-Net [55], Dual-Refinement [9], UNRN [62], and GLT [63]. (3) Joint training methods commonly use the memory bank [51], including ECN [69], MMCL [44], ECN-GPP [70], JVTC+ [25], and SpCL [16]. However, all these methods neglect the significance of intermediate domains, which can smoothly bridge the domain adaptation between the source and target domains to better transfer the source knowledge to the target domain. Instead, we propose an IDM module to generate the appropriate intermediate domains to better improve the performance of UDA re-ID. As shown in Table 4, our method can outperform the second best UDA re-ID methods by a large margin on all these benchmarks.

**Comparisons on synthetic → real UDA re-ID tasks.** Compared with the real → real UDA re-ID tasks, the synthetic → real UDA tasks are more challenging because the domain gap between the synthetic and real images are often larger than that between the real and real images. As shown in Table 5, our method can outperform the SOTA methods by a large margin where mAP of our method is higher than SpCL [16] by 7.5%, 1.3%, and 7.6% when testing on Market-1501, DukeMTMC-reID, and MSMT17 respectively. All these significant results of our method have shown the superiority of our proposed IDM module that can generate the appropriate intermediate domains to better improve the performance on UDA re-ID.

Table 4. Comparison with the state-of-the-art UDA re-ID methods on real  $\rightarrow$  real tasks.

Methods	Reference	DukeMTMC-ReID $\rightarrow$ Market-1501				Market-1501 $\rightarrow$ DukeMTMC-ReID			
		mAP	R1	R5	R10	mAP	R1	R5	R10
PUL [12]	TOMM 2018	20.5	45.5	60.7	66.7	16.4	30.0	43.4	48.5
TJ-AIDL [46]	CVPR 2018	26.5	58.2	74.8	81.1	23.0	44.3	59.6	65.0
SPGAN+LMP [10]	CVPR 2018	26.7	57.7	75.8	82.4	26.2	46.4	62.3	68.0
HHL [68]	ECCV 2018	31.4	62.2	78.8	84.0	27.2	46.9	61.0	66.7
ECN [69]	CVPR 2019	43.0	75.1	87.6	91.6	40.4	63.3	75.8	80.4
PDA-Net [27]	ICCV 2019	47.6	75.2	86.3	90.2	45.1	63.2	77.0	82.5
PCB-PAST [58]	ICCV 2019	54.6	78.4	-	-	54.3	72.4	-	-
SSG [13]	ICCV 2019	58.3	80.0	90.0	92.4	53.4	73.0	80.6	83.2
MMCL [44]	CVPR 2020	60.4	84.4	92.8	95.0	51.4	72.4	82.9	85.0
ECN-GPP [70]	TPAMI 2020	63.8	84.1	92.8	95.4	54.4	74.0	83.7	87.4
JVTC+ [25]	ECCV 2020	67.2	86.8	95.2	97.1	66.5	80.4	89.9	92.2
AD-Cluster [54]	CVPR 2020	68.3	86.7	94.4	96.5	54.1	72.6	82.5	85.5
MMT [15]	ICLR 2020	71.2	87.7	94.9	96.9	65.1	78.0	88.8	92.5
CAIL [33]	ECCV 2020	71.5	88.1	94.4	96.2	65.2	79.5	88.3	91.4
NRMT [61]	ECCV 2020	71.7	87.8	94.6	96.5	62.2	77.8	86.9	89.5
MEB-Net [55]	ECCV 2020	76.0	89.9	96.0	97.5	66.1	79.6	88.3	92.2
SpCL [16]	NeurIPS 2020	76.7	90.3	96.2	97.7	68.8	82.9	90.1	92.5
Dual-Refinement [9]	TIP 2021	78.0	90.9	96.4	97.7	67.7	82.1	90.1	92.5
UNRN [62]	AAAI 2021	78.1	91.9	96.1	97.8	69.1	82.0	90.7	93.5
GLT [63]	CVPR 2021	79.5	92.2	96.5	97.8	69.2	82.0	90.2	92.8
IDM (Ours)	ICCV 2021	<b>82.8</b>	<b>93.2</b>	<b>97.5</b>	<b>98.1</b>	<b>70.5</b>	<b>83.6</b>	<b>91.5</b>	<b>93.7</b>

Methods	Reference	Market-1501 $\rightarrow$ MSMT17				DukeMTMC-reID $\rightarrow$ MSMT17			
		mAP	R1	R5	R10	mAP	R1	R5	R10
ECN [69]	CVPR 2019	8.5	25.3	36.3	42.1	10.2	30.2	41.5	46.8
SSG [13]	ICCV 2019	13.2	31.6	-	49.6	13.3	32.2	-	51.2
ECN-GPP [70]	TPAMI 2020	15.2	40.4	53.1	58.7	16.0	42.5	55.9	61.5
MMCL [44]	CVPR 2020	15.1	40.8	51.8	56.7	16.2	43.6	54.3	58.9
NRMT [61]	ECCV 2020	19.8	43.7	56.5	62.2	20.6	45.2	57.8	63.3
CAIL [33]	ECCV 2020	20.4	43.7	56.1	61.9	24.3	51.7	64.0	68.9
MMT [15]	ICLR 2020	22.9	49.2	63.1	68.8	23.3	50.1	63.9	69.8
JVTC+ [25]	ECCV 2020	25.1	48.6	65.3	68.2	27.5	52.9	70.5	75.9
SpCL [16]	NeurIPS 2020	26.8	53.7	65.0	69.8	26.5	53.1	65.8	70.5
Dual-Refinement [9]	TIP 2021	25.1	53.3	66.1	71.5	26.9	55.0	68.4	73.2
UNRN [62]	AAAI 2021	25.3	52.4	64.7	69.7	26.2	54.9	67.3	70.6
GLT [63]	CVPR 2021	26.5	56.6	67.5	72.0	27.7	59.5	70.1	74.2
IDM (Ours)	ICCV 2021	<b>33.5</b>	<b>61.3</b>	<b>73.9</b>	<b>78.4</b>	<b>35.4</b>	<b>63.6</b>	<b>75.5</b>	<b>80.2</b>

Table 5. Comparison with the state-of-the-art UDA re-ID methods on synthetic  $\rightarrow$  real tasks.

Methods	Reference	PersonX $\rightarrow$ Market-1501				PersonX $\rightarrow$ DukeMTMC-reID				PersonX $\rightarrow$ MSMT17			
		mAP	R1	R5	R10	mAP	R1	R5	R10	mAP	R1	R5	R10
MMT [15]	ICLR 2020	71.0	86.5	94.8	97.0	60.1	74.3	86.5	90.5	17.7	39.1	52.6	58.5
SpCL [16]	NeurIPS 2020	73.8	88.0	95.3	96.9	67.2	81.8	90.2	92.6	22.7	47.7	60.0	65.5
IDM (Ours)	ICCV 2021	<b>81.3</b>	<b>92.0</b>	<b>97.4</b>	<b>98.2</b>	<b>68.5</b>	<b>82.6</b>	<b>91.2</b>	<b>93.4</b>	<b>30.3</b>	<b>58.4</b>	<b>70.7</b>	<b>75.5</b>

Methods	Reference	Unreal $\rightarrow$ Market-1501				Unreal $\rightarrow$ DukeMTMC-reID				Unreal $\rightarrow$ MSMT17			
		mAP	R1	R5	R10	mAP	R1	R5	R10	mAP	R1	R5	R10
JVTC [25]	ECCV 2020	78.3	90.8	-	-	66.1	81.2	-	-	25.0	53.7	-	-
IDM (Ours)	ICCV 2021	<b>83.2</b>	<b>92.8</b>	<b>97.3</b>	<b>98.2</b>	<b>72.4</b>	<b>84.6</b>	<b>92.0</b>	<b>94.0</b>	<b>38.3</b>	<b>67.3</b>	<b>78.4</b>	<b>82.6</b>

## 5. Conclusion

This paper proposes a plug-and-play Intermediate Domain Module (IDM) to tackle the problem of UDA re-ID. From a novel perspective that intermediate domains can bridge the source and target domains, our proposed IDM module can generate appropriate intermediate domain representations to better transfer the source knowledge to improve the model’s discriminability on the target domain. The intermediate domains’ distribution is controlled by the two domain factors generated by the IDM module. Specifically, we propose the bridge losses to enforce the intermedi-

ate domains to be located onto the appropriate path between the source and target domains in a manifold. Besides, we also propose a diversity loss to constrain the domain factors to prevent the intermediate domains from being over-fitting to either of the source and target domains. Extensive experiments have shown the effectiveness of our method.

**Acknowledgements:** This work was supported by the National Natural Science Foundation of China under Grant 62088102, and in part by the PKU-NTU Joint Research Institute (JRI) sponsored by a donation from the Ng Teng Fong Charitable Foundation, and was partially supported by SUTD Project PIE-SGP-AI2020-02.



## References

- [1] Yan Bai, Jile Jiao, Wang Ce, Jun Liu, Yihang Lou, Xuetao Feng, and Ling-Yu Duan. Person30k: A dual-meta generalization network for person re-identification. In *CVPR*, pages 2123–2132, 2021. 1
- [2] Yan Bai, Ce Wang, Yihang Lou, Jun Liu, and Ling-Yu Duan. Hierarchical connectivity-centered clustering for unsupervised domain adaptation on person re-identification. *IEEE TIP*, 2021. 1
- [3] Guangyi Chen, Chunze Lin, Liangliang Ren, Jiwen Lu, and Jie Zhou. Self-critical attention learning for person re-identification. In *ICCV*, pages 9637–9646, 2019. 1
- [4] Guangyi Chen, Yuhao Lu, Jiwen Lu, and Jie Zhou. Deep credible metric learning for unsupervised domain adaptation person re-identification. In *ECCV*, pages 643–659, 2020. 2
- [5] Yunlu Chen, Vincent Tao Hu, Efstratios Gavves, Thomas Mensink, Pascal Mettes, Pengwan Yang, and Cees GM Snoek. Pointmixup: Augmentation for point clouds. In *ECCV*, 2020. 3
- [6] Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Qi Tian. Gradually vanishing bridge for adversarial domain adaptation. In *CVPR*, pages 12455–12464, 2020. 2
- [7] Zhen Cui, Wen Li, Dong Xu, Shiguang Shan, Xilin Chen, and Xuelong Li. Flowing on riemannian manifold: Domain adaptation by shifting covariance. *IEEE transactions on cybernetics*, 44(12):2264–2273, 2014. 2
- [8] Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. Generalizable person re-identification with relevance-aware mixture of experts. In *CVPR*, pages 16145–16154, 2021. 1
- [9] 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. *IEEE TIP*, 2021. 1, 2, 7, 8
- [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*, pages 994–1003, 2018. 1, 2, 3, 7, 8
- [11] Yuhang Ding, Hehe Fan, Mingliang Xu, and Yi Yang. Adaptive exploration for unsupervised person re-identification. *ACM TOMM*, 16(1):1–19, 2020. 2
- [12] Hehe Fan, Liang Zheng, Chenggang Yan, and Yi Yang. Unsupervised person re-identification: Clustering and fine-tuning. *ACM TOMM*, 14(4):1–18, 2018. 7, 8
- [13] Yang Fu, Yunchao Wei, Guanshuo Wang, Yuqian Zhou, Honghui Shi, and Thomas S Huang. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *ICCV*, pages 6112–6121, 2019. 1, 2, 4, 7, 8
- [14] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. 2
- [15] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020. 2, 7, 8
- [16] Yixiao Ge, Dapeng Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *NeurIPS*, 2020. 1, 2, 4, 6, 7, 8
- [17] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, pages 2066–2073, 2012. 2, 4
- [18] Rui Gong, Wen Li, Yuhua Chen, and Luc Van Gool. Dlow: Domain flow for adaptation and generalization. In *CVPR*, pages 2477–2486, 2019. 2
- [19] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Unsupervised adaptation across domain shifts by generating intermediate data representations. *IEEE TPAMI*, 36(11):2288–2302, 2013. 2, 4, 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4
- [21] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 4
- [22] Yan Huang, Qiang Wu, JingSong Xu, and Yi Zhong. Sbsgan: Suppression of inter-domain background shift for person re-identification. In *ICCV*, pages 9527–9536, 2019. 2
- [23] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Global distance-distributions separation for unsupervised person re-identification. In *ECCV*, pages 735–751, 2020. 2
- [24] Guoliang Kang, Liang Zheng, Yan Yan, and Yi Yang. Deep adversarial attention alignment for unsupervised domain adaptation: the benefit of target expectation maximization. In *ECCV*, pages 401–416, 2018. 2
- [25] Jianing Li and Shiliang Zhang. Joint visual and temporal consistency for unsupervised domain adaptive person re-identification. In *ECCV*, pages 483–499, 2020. 6, 7, 8
- [26] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. *arXiv preprint arXiv:1603.04779*, 2016. 7
- [27] 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*, pages 7919–7929, 2019. 7, 8
- [28] Yutian Lin, Lingxi Xie, Yu Wu, Chenggang Yan, and Qi Tian. Unsupervised person re-identification via softened similarity learning. In *CVPR*, pages 3390–3399, 2020. 2
- [29] Mingsheng Long, Yue Cao, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Transferable representation learning with deep adaptation networks. *IEEE TPAMI*, 41(12):3071–3085, 2018. 7
- [30] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, pages 97–105, 2015. 2
- [31] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, 2018. 2
- [32] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217, 2017. 2

- [33] Chuanchen Luo, Chunfeng Song, and Zhaoxiang Zhang. Generalizing person re-identification by camera-aware invariance learning and cross-domain mixup. In *ECCV*, volume 2, page 7, 2020. 8
- [34] Jaemin Na, Heechul Jung, Hyung Jin Chang, and Wonjun Hwang. Fixbi: Bridging domain spaces for unsupervised domain adaptation. In *CVPR*, pages 1094–1103, 2021. 3
- [35] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, pages 464–479, 2018. 7
- [36] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV*, pages 17–35, 2016. 6
- [37] Paolo Russo, Fabio M Carlucci, Tatiana Tommasi, and Barbara Caputo. From source to target and back: symmetric bi-directional adaptive gan. In *CVPR*, pages 8099–8108, 2018. 2
- [38] Liangchen Song, Cheng Wang, Lefei Zhang, Bo Du, Qian Zhang, Chang Huang, and Xinggong Wang. Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognition*, 102:107173, 2020. 1, 2, 4
- [39] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450, 2016. 2
- [40] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *CVPR*, pages 608–617, 2019. 6, 7
- [41] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, pages 480–496, 2018. 1
- [42] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 2
- [43] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, pages 6438–6447, 2019. 3, 7
- [44] Dongkai Wang and Shiliang Zhang. Unsupervised person re-identification via multi-label classification. In *CVPR*, pages 10981–10990, 2020. 2, 7, 8
- [45] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *ACM MM*, pages 274–282, 2018. 1
- [46] 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. 8
- [47] Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R Scott. Cross-batch memory for embedding learning. In *CVPR*, pages 6388–6397, 2020. 6
- [48] 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. 1, 2, 3, 6, 7
- [49] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *ECCV*, pages 540–555, 2020. 3
- [50] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018. 3, 6
- [51] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. Joint detection and identification feature learning for person search. In *CVPR*, pages 3415–3424, 2017. 3, 6, 7
- [52] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *AAAI*, volume 34, pages 6502–6509, 2020. 3
- [53] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE TPAMI*, 2021. 1
- [54] 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*, pages 9021–9030, 2020. 2, 7, 8
- [55] Yunpeng Zhai, Qixiang Ye, Shijian Lu, Mengxi Jia, Rongrong Ji, and Yonghong Tian. Multiple expert brainstorming for domain adaptive person re-identification. *ECCV*, 2020. 2, 7, 8
- [56] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 3, 7
- [57] Tianyu Zhang, Lingxi Xie, Longhui Wei, Zijie Zhuang, Yongfei Zhang, Bo Li, and Qi Tian. Unrealperson: An adaptive pipeline towards costless person re-identification. In *CVPR*, pages 11506–11515, 2021. 6, 7
- [58] Xinyu Zhang, Jiewei Cao, Chunhua Shen, and Mingyu You. Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In *ICCV*, pages 8222–8231, 2019. 7, 8
- [59] Zhi Zhang, Tong He, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of freebies for training object detection neural networks. *arXiv preprint arXiv:1902.04103*, 2019. 3
- [60] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Xin Jin, and Zhibo Chen. Relation-aware global attention for person re-identification. In *CVPR*, pages 3186–3195, 2020. 1
- [61] 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*, pages 526–544, 2020. 7, 8
- [62] Kecheng Zheng, Cuiling Lan, Wenjun Zeng, Zhizheng Zhan, and Zheng-Jun Zha. Exploiting sample uncertainty for domain adaptive person re-identification. In *AAAI*, 2021. 6, 7, 8
- [63] Kecheng Zheng, Wu Liu, Lingxiao He, Tao Mei, Jiebo Luo, and Zheng-Jun Zha. Group-aware label transfer for domain adaptive person re-identification. In *CVPR*, pages 5310–5319, 2021. 6, 7, 8

- [64] 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
- [65] Liang Zheng, Yi Yang, and Alexander G Hauptmann. Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984*, 2016. 4
- [66] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*, pages 3754–3762, 2017. 6
- [67] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. In *CVPR*, pages 1318–1327, 2017. 6
- [68] Zhun Zhong, Liang Zheng, Shaozi Li, and Yi Yang. Generalizing a person retrieval model hetero-and homogeneously. In *ECCV*, pages 172–188, 2018. 7, 8
- [69] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *CVPR*, pages 598–607, 2019. 2, 7, 8
- [70] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Learning to adapt invariance in memory for person re-identification. *IEEE TPAMI*, 2020. 1, 2, 6, 7, 8
- [71] Zijie Zhuang, Longhui Wei, Lingxi Xie, Tianyu Zhang, Hengheng Zhang, Haozhe Wu, Haizhou Ai, and Qi Tian. Rethinking the distribution gap of person re-identification with camera-based batch normalization. In *ECCV*, pages 140–157, 2020. 7
- [72] Junbao Zhuo, Shuhui Wang, Weigang Zhang, and Qingming Huang. Deep unsupervised convolutional domain adaptation. In *ACM MM*, pages 261–269, 2017. 2
- [73] Yang Zou, Xiaodong Yang, Zhiding Yu, BVK Kumar, and Jan Kautz. Joint disentangling and adaptation for cross-domain person re-identification. In *ECCV*, 2020. 2, 3