



# Bridging the Source-to-Target Gap for Cross-Domain Person Re-identification with Intermediate Domains

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

Received: 13 December 2023 / Accepted: 28 June 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

## Abstract

Cross-domain person re-identification (re-ID), such as unsupervised domain adaptive re-ID (UDA re-ID), aims to transfer the identity-discriminative knowledge from the source to the target domain. Existing methods commonly consider the source and target domains are isolated from each other, i.e., no intermediate status is modeled between the source and target domains. Directly transferring the knowledge between two isolated domains can be very difficult, especially when the domain gap is large. This paper, from a novel perspective, assumes these two domains are not completely isolated, but can be connected through a series of intermediate domains. Instead of directly aligning the source and target domains against each other, we propose to align the source and target domains against their intermediate domains so as to facilitate a smooth knowledge transfer. To discover and utilize these intermediate domains, this paper proposes an Intermediate Domain Module (IDM) and a Mirrors Generation Module (MGM). IDM has two functions: (1) it generates multiple intermediate domains by mixing the hidden-layer features from source and target domains and (2) it dynamically reduces the domain gap between the source/target domain features and the intermediate domain features. While IDM achieves good domain alignment effect, it introduces a side effect, i.e., the mix-up operation may mix the identities into a new identity and lose the original identities. Accordingly, MGM is introduced to compensate the loss of the original identity by mapping the features into the IDM-generated intermediate domains without changing their original identity. It allows to focus on minimizing domain variations to further promote the alignment between the source/target domain and intermediate domains, which reinforces IDM into IDM++. We extensively evaluate our method under both the UDA and domain generalization (DG) scenarios and observe that IDM++ yields consistent (and usually significant) performance improvement for cross-domain re-ID, achieving new state of the art. For example, on the challenging MSMT17 benchmark, IDM++ surpasses the prior state of the art by a large margin (e.g., up to 9.9% and 7.8% rank-1 accuracy) for UDA and DG scenarios, respectively. Code is available at <https://github.com/SikaStar/IDM>.

**Keywords** Person re-identification · Unsupervised domain adaptation · Domain generalization · Intermediate domains · Domain bridge

---

Communicated by Zhun Zhong.

---

✉ Ling-Yu Duan  
lingyu@pku.edu.cn

Yongxing Dai  
yongxing.dai@outlook.com

Yifan Sun  
sunyf15@tsinghua.org.cn

Jun Liu  
j.liu81@lancaster.ac.uk

Zekun Tong  
zekuntong@u.nus.edu

## 1 Introduction

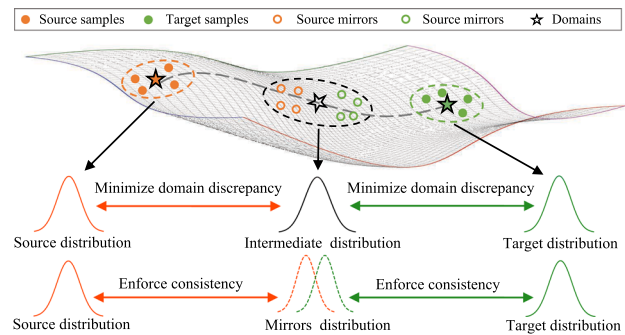
Person re-identification (re-ID) (Zheng et al., 2016; Leng et al., 2019; Ye et al., 2021b) aims to identify the same person across non-overlapped cameras. A critical challenge in

- 1 National Engineering Research Center of Visual Technology, School of Computer Science, Peking University, Beijing, China
- 2 Peng Cheng Laboratory, Shenzhen, China
- 3 Baidu Inc., Beijing, China
- 4 School of Computing and Communications, Lancaster University, Lancaster, UK
- 5 National University of Singapore, Singapore, Singapore

a realistic re-ID system is the cross-domain problem, i.e., the training data and testing data are from different domains. Since annotating the person identities is notoriously expensive, it is of great value to transfer the identity-discriminative knowledge from the source to the target domain without incurring additional annotations. To improve the cross-domain re-ID, there are two popular approaches, i.e., the unsupervised domain adaptation (UDA) and domain generalization (DG). These two approaches are closely related yet have an important difference: during training, UDA has unlabeled target-domain data, while DG is not accessible to the target domain and is thus more challenging. This paper mainly challenges the UDA re-ID and provides compatibility to DG re-ID.

Existing UDA re-ID methods commonly consider the source and target domains are isolated from each other, i.e., there is no intermediate status between the source and target domains. Directly transferring knowledge between two isolated domains can be very difficult, especially when the domain gap is large. Specifically, there are two approaches, i.e., style transfer (Wei et al., 2018; Deng et al., 2018) and pseudo-label-based training (Song et al., 2020; Fu et al., 2019; Dai et al., 2021b; Ge et al., 2020b). The style-transfer methods usually use GANs (Zhu et al., 2017) to transfer the target-domain style onto the labeled source-domain images, so that the deep model can learn from the target-stylized images. Since the source and target domains are isolated and far away from each other, the style transfer procedure can be viewed as jumping from the source to the target domain, which can be difficult. The pseudo-label-based training methods require a clustering procedure as the prerequisite for obtaining the pseudo labels. The domain gap between the isolated source and target domains compromises the clustering accuracy and incurs noisy pseudo labels. Therefore, we conjecture that directly mitigating the domain gap between two isolated (source and target) domains is difficult for both the style-transfer and the pseudo-label-based UDA approaches.

From a novel perspective, this paper considers that the source and target domains are not isolated, but are potentially connected through a series of intermediate domains. In other words, some intermediate domains underpin a path that bridges the gap between the source and target domains. Specifically, we assume that the source and target domains are located in a manifold, as shown in Fig. 1. There is an appropriate “path” connecting these two isolated domains. The source and target domains lie on the two extreme points of this path, while some intermediate domains exist along this path, characterizing the inter-domain connections. This viewpoint leads to an intuition: to align the two extreme points, comparing them against these intermediate points can be more feasible than directly comparing the two extreme points against each other, given that they are far from each other due to the significant domain gap.



**Fig. 1** Illustration of our main idea. Assuming that the source and target domains (in UDA re-ID) are located in a manifold, there can be some intermediate domains along with the path to bridge the two extreme domains. With the generation of intermediate domains, the source and target domains can be smoothly aligned with them. To further preserve the source/target identity information during alignment, we map source/target identities into the intermediate domains to obtain source/target mirrors

Motivated by this, we propose an intermediate domain module (IDM) for UDA re-ID. Instead of directly aligning the source and target domain against each other (i.e., source  $\rightarrow$  target), IDM aligns the source and target domain against a shared set of intermediate domains (i.e., source  $\rightarrow$  intermediate and target  $\rightarrow$  intermediate), as illustrated in Fig. 1. To this end, IDM first synthesizes multiple intermediate domains and then minimizes the distance between each intermediate domain and the source/target domain. Specifically, IDM mixes the source-domain and target-domain features, which are output from a hidden layer of the deep network. The correspondingly-synthesized features exhibit the characteristics of the intermediate domain and are then fed into the sub-sequential hidden layer. The consequential output features of the intermediate domains are compared against the original source and target domain features for minimizing the domain discrepancy. We note that the IDM can be plugged after any hidden layers and is thus a plug-and-play module. There are two critical techniques for constructing the IDM:

- *Modeling the intermediate domain.* We use the Mixup (Zhang et al., 2018) strategy to mix the source and target hidden features (e.g., “Stage-0” in ResNet-50) and forward the mixed hidden feature until the last deep embedding. Importantly, we consider the mix ratio in the hidden layer determines the position of the output feature in the deeply-learned embedding space. More concretely, if the mixup operation in the hidden feature space assigns a larger proportion to the source domain, the output feature should be closer to the source domain than to the target domain (and vice versa). Correspondingly, when we minimize its distance to the source and target domain in parallel, we will emphasize both domains proportionally. In a word, the correspondence between the mix ratio

and the relative distance on the path connecting the source and target domain models the intermediate domain.

- *Enhancing the intermediate domain diversity.* The path between the source and target domains is jointly depicted by all the intermediate domains on this path. Therefore, it is important to generate the intermediate domains densely. Since the mix ratio determines the position of the mixed features on the path, varying the mix ratio results in different intermediate domains. IDM adaptively generates the mix ratios and encourages the mix ratios to be as diverse as possible for different input hidden features.

Based on the IDM in our previous conference version (Dai et al., 2021c), this paper makes an important extension, i.e., reinforcing IDM into IDM++ by further promoting the alignment from source/target domain to intermediate domains. In IDM, when a source and a target-domain feature are mixed up, their identities are accordingly mixed into a new identity (Sect. 3.2.1). This mixup has a side effect, i.e., there are no original identities in the intermediate domains. Therefore, the cross-domain alignment by IDM may focus more on the mixed-identity variation than the diverse cross-domain styles that are irrelevant to the identity. To compensate this, we propose a Mirrors Generation Module (MGM) to map original identities into the already-generated intermediate domains. MGM utilizes a popular statistic-based style transfer approach AdaIN (Huang & Belongie, 2017), which replaces the mean and standard deviation value of the source/target features with those from intermediate features. Consequently, each feature in the source / target domain has a “same-identity different-style” mirror in every intermediate domain, as shown in Fig. 1. It enables IDM++ to focus on the cross-domain variation when aligning the domains (as detailed in Sect. 3.4). Thus, IDM++ brings another round of significant improvement based on IDM.

We note that although the statistic-based style transfer is not new and is widely adopted by many recent out-of-distribution literature (Zhou et al., 2021; Nuriel et al., 2021; Tang et al., 2021), integrating MGM implemented with style transfer into IDM++ brings novelty and significant benefits. First, the style transfer in prior methods is between off-the-shelf domains, while our MGM is the first to conduct style transfer towards on-the-fly synthesized domains (i.e., the IDM-generated intermediate domains). Second, on UDA re-ID, we show that MGM and IDM mutually reinforce each other and jointly bring complementary benefits for IDM++. On the one hand, adding MGM significantly improves IDM (e.g., 8.6% improvement of rank-1 accuracy on Market-1501→MSMT17). On the other hand, IDM is a critical prerequisite for the effectiveness of MGM, because we find that using MGM without IDM only brings slight improvement. Third, MGM endows IDM++ with the extra capacity of improving domain generalization, allowing IDM++ to be

the first framework to integrate both UDA and DG capacity for cross-domain re-ID, achieving state of the art results on both scenarios.

In addition to the reinforcement from IDM to IDM++, this paper makes two more extensions *w.r.t.* the applied scenarios and the experimental evaluation. Overall, compared with our previous conference version (Dai et al., 2021c), the extensions are from three aspects: (1) *Method.* We reinforce IDM into IDM++ by adding a novel MGM (Sect. 3.4). MGM maps the training identities from the source/target domain as their mirrors in the IDM-generated intermediate domains. It allows IDM++ to directly minimize the variation of the same identity across various domains. This new advantage improves IDM by a large margin. (2) *Applied scenarios.* We investigate IDM++ under two popular cross-domain re-ID scenarios, i.e., unsupervised domain adaptation (UDA) and the more challenging domain generalization (DG). We show that IDM++ provides a unified framework for these two scenarios and brings general improvement. (3) *Experimental evaluation.* We provide more comprehensive experimental evaluation, including experiments on more detailed ablation studies (Sect. 4.3), analysis on parameter sensitivity (Sect. 4.4), two challenging protocols of DG re-ID (Sect. 4.5), investigations on key designs of IDM++ (Sect. 4.6), and discussions on the distribution along the domain bridge and alignment of domains (Sect. 4.7).

Our contributions can be summarized as follows.

- We innovatively propose to explore intermediate domains for the cross-domain re-ID task. We consider that some intermediate domains can be densely distributed between the source/target domain and jointly depict a smooth bridge across the domain gap. Minimizing the domain discrepancy along this bridge improves domain alignment.
- We propose an IDM module to discover the desired intermediate domains by on-the-fly mixing the source and target domain features with varying mix-ratios. Afterwards, IDM mitigates the gap from source/target to the intermediate domains so as to align the source and target domains.
- We propose an MGM module to further promote the source/target to intermediate domain alignment. In compensation for a loss of IDM (i.e., the feature mix-up loses the original training identities), MGM maps the original training identities into the intermediate domains. Incorporating MGM and IDM, IDM++ brings another round of substantial improvement.
- We show that our IDM++ provides a unified framework for two popular cross-domain re-ID scenarios, i.e., UDA and DG. Extensive experiments under twelve UDA re-ID benchmarks and two DG re-ID protocols validate that

IDM++ brings general improvement and sets new state of the art for both scenarios.

## 2 Related Work

**Unsupervised Domain Adaptation** A line of works converts the UDA into an adversarial learning task (Ganin & Lempit-sky, 2015; Tzeng et al., 2017; Long et al., 2018b; Russo et al., 2018). Another line of works uses various metrics to measure and minimize the domain discrepancy, such as MMD (Long et al., 2015) or other metrics (Sun & Saenko, 2016; Zhuo et al., 2017; Long et al., 2017; Kang et al., 2018). Another line of traditional works (Gong et al., 2012; Cui et al., 2014; Gopalan et al., 2013) tries to bridge the source and target domains based on intermediate domains. In traditional methods (Gong et al., 2012; Gopalan et al., 2013), 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. Still, they are not easily applied to the deep models. In deep methods (Gong et al., 2019; Cui et al., 2020), they either use GANs to generate a domain flow by reconstructing input images on pixel level (Gong et al., 2019) or learn better domain-invariant features by bridging the learning of the generator and discriminator (Cui et al., 2020). However, reconstructing images may not guarantee the high-level domain characteristics in the feature space or introduce unnecessary noise in the pixel space. It will become harder to adapt between the two isolated domains especially when the large domain gap. Recently, the concept of intermediate modalities has been explored in the context of cross-modal re-ID (Li et al., 2020; Wei et al., 2021; Zhang et al., 2021b; Ye et al., 2021a; Huang et al., 2022; Yu et al., 2023), which also motivates us to utilize intermediate domains to tackle the cross-domain re-ID.

Unlike the above methods, we propose a lightweight 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 can be learned in an efficient joint training scheme. Besides, this paper also proposes a new module (MGM) to encourage the model to focus on minimizing the cross-domain variation when aligning the source and target domains.

**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 (Wei et al., 2018; Deng et al., 2018; Huang et al., 2019; Zou et al., 2020; Huang et al., 2021), fine-tuning (Song et al., 2020; Fu et al., 2019; Ge et al., 2020a; Dai et al., 2021b; Chen et al., 2020b; Zhai et al., 2020b; Jin et al., 2020a; Lin et al., 2020; Wang et al., 2022; Zhai et al., 2020a, 2023), and joint train-

ing (Zhong et al., 2019, 2020b; Wang & Zhang, 2020; Ge et al., 2020b; Ding et al., 2020; Zheng et al., 2021c; Isobe et al., 2021). GAN transferring methods use GANs to transfer images' style across domains (Wei et al., 2018; Deng et al., 2018) or disentangle features into id-related/unrelated features (Zou et al., 2020). 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 (Xiao et al., 2017; Wu et al., 2018) 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 bridge 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 bridge between the source and target 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. Besides, we also propose a novel statistic-based feature augmentation module to exploit the intermediate domains' styles, which can encourage the model to focus on learning identity-discriminative features when aligning the domains. The above exploitation of intermediate domains has not been investigated in this field as far as we know.

**Mixup and Variants** Mixup (Zhang et al., 2018) 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 (Verma et al., 2019) 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 (Chen et al., 2020a), object detection (Zhang et al., 2019b), and closed-set domain adaptation (Xu et al., 2020; Wu et al., 2020; Na et al., 2021). 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 cross-domain 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.

**Domain Generalization** Domain generalization (DG) aims to improve the model's generalization on one or more target domains by only using one or more source domains for training. Different from UDA, target data is not accessible for training DG models. For more comprehensive surveys of DG, please refer to Zhou et al. (2021), Wang et al. (2021), and Shen et al. (2021). Following the taxonomy in a DG survey (Zhou et al., 2021), we briefly review the two categories that are most related to our work: (1) domain alignment and (2) data augmentation. The key idea of the existing domain alignment methods is to align different source domains by only minimizing their discrepancy. Different from these works, we align the source and target domains by aligning them with the synthetic intermediate domain respectively. The motivation of our work is to ease the procedure of domain alignment using intermediate domains, which fundamentally differs from the existing works. Data augmentation works in DG aim to improve models' generalization by designing image-based (Volpi et al. 2018; Qiao et al. 2020; Zhou et al. 2020b, a) or feature-based (Mancini et al. 2020; Xu et al. 2021; Zhou et al. 2021; Nuriel et al. 2021; Tang et al. 2021; Du et al. 2022; Zhou et al. 2023; Li et al. 2024; Zhao et al. 2023) augmentation methods, which can avoid over-fitting to source domains. For example, several feature-based augmentation researches (Zhou et al. 2021; Nuriel et al. 2021; Tang et al. 2021) focus on manipulating feature statistics based on the observation that CNN feature statistics (i.e., mean and standard deviation) can represent an image's style (Huang & Belongie, 2017). Motivated by them, we propose a new augmentation module (i.e., MGM) implemented by the arbitrary-style-transfer technique AdaIN (Huang & Belongie, 2017) to augment source/target features with statistics of the already-generated intermediate domains. Different from the existing augmentation works, in this paper, we exploit the diverse inter-domain variation to augment source/target identities with different domain styles, and we utilize this augmentation mechanism to complement with the procedure of generating intermediate domains by our IDM.

**Domain Generalizable Person Re-ID** Recently, a newly proposed cross-domain re-ID task called domain generalizable re-ID (DG re-ID) (Song et al., 2019; Jia et al., 2019; Jin et al., 2020b; Zhao et al., 2021b; Choi et al., 2021; Dai et al., 2021a; Pu et al., 2023; Zhang et al., 2022; Tan et al., 2023; Xiang et al., 2023; Ni et al., 2022) has gained much interest of researchers. DG re-ID aims to improve the generalization on the target domains that are not accessible during training, where only one or more source domains can be used for training. Different from the classification task in DG, the target domain does not share the label space with the

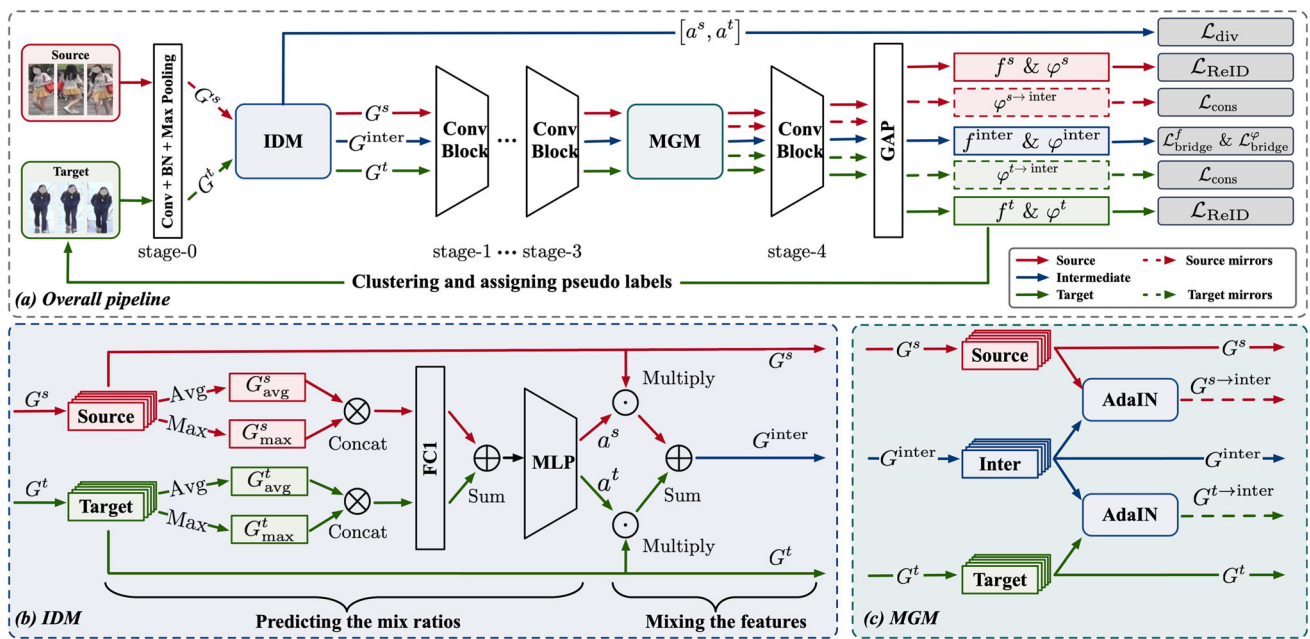
source domains, which poses a new challenge over conventional DG settings. The existing DG re-ID works mainly fall into two categories: (1) designing specific modules that can learn better domain-invariant representations (Song et al., 2019; Jia et al., 2019; Jin et al., 2020b), and (2) adopting the meta-learning paradigm to simulate the domain bias by splitting training domains into pseudo-seen and pseudo-unseen domains (Zhao et al., 2021b; Choi et al., 2021; Dai et al., 2021a). For example, SNR (Jin et al., 2020b) designs a Style Normalization and Restitution module to disentangle object representations into identity-irrelevant and identity-relevant ones. RaMoE (Dai et al., 2021a) introduces meta-learning into a novel mixture-of-experts paradigm via an effective voting-based mixture mechanism, which can learn to dynamically aggregate multi-source domains information. MDA (Ni et al., 2022) is an innovative approach that seeks to align the distributions of source and target domains with a known prior distribution. This is achieved through a meta-learning strategy that facilitates generalization and supports fast adaptation to unseen domains. Different from the above works, we aim to exploit the bridge among different source domains to tackle the problem of DG re-ID, which may provide a novel perspective in solving this problem.

### 3 Methodology

This paper explores intermediate domains for cross-domain re-ID. We first focus on the unsupervised domain adaptation (UDA) scenario (from Sects. 3.1 to 3.4) and then extend our method to the domain generalization (DG) scenario (Sect. 3.5). Specifically, for UDA, we first introduce the overall pipeline of IDM++ in Sect. 3.1, and then elaborate the key component IDM for discovering the intermediate domains in Sect. 3.2. Given the learned intermediate domains, we disentangle the source-to-target domain alignment into a more feasible one, i.e., aligning source/target to intermediate domains in Sect. 3.3. Afterwards, we further improve the source / target to intermediate domain alignment with MGM and summarize the overall training for IDM++ in Sect. 3.4. Section 3.5 illustrates how to apply IDM++ for DG re-ID.

#### 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 adopt the pseudo-label-based pipeline to perform clustering to assign the pseudo labels for the target samples. Similar to the pseudo-label-based UDA re-ID methods (Fu et al., 2019;



**Fig. 2** The overall of IDM++ (a), which is comprised of IDM (b) and MGM (c). In (a), IDM++ combines the source and target domain for joint training. It first plugs IDM after a bottom layer (e.g., stage-0 in ResNet-50 (He et al., 2016)) and then plugs MGM after a top layer (e.g., stage-3). IDM mixes the hidden-layer features from the source and target domain (i.e.,  $G^s$ ,  $G^t$ ) to generate intermediate-domain features  $G^{inter}$ .  $G^s$ ,  $G^t$ , and  $G^{inter}$  continue to proceed along the network until the final outputs, i.e., the deep embedding ( $f^s$ ,  $f^t$  and  $f^{inter}$ ) and the softmax prediction ( $\varphi^s$ ,  $\varphi^t$  and  $\varphi^{inter}$ ). Since the intermediate-domain features are initially generated after stage-0, we consider the following stages already provide intermediate domains. Given the intermediate-domains features at the end of stage-3, MGM maps the features in the source and

target domain into the intermediate domains, i.e.,  $G^{s \rightarrow inter}$  and  $G^{t \rightarrow inter}$ , which are finally output as  $\varphi^{s \rightarrow inter}$  and  $\varphi^{t \rightarrow inter}$  in the softmax prediction layer. Overall, there are four losses for learning IDM++: 1) Given the output of the source and target domain, we enforce the popular ReID loss  $\mathcal{L}_{ReID}$  (including the classification loss  $\mathcal{L}_{cls}$  and triplet loss  $\mathcal{L}_{tri}$ ). 2) Given the output of IDM, i.e.,  $f^{inter}$  and  $\varphi^{inter}$  in the intermediate domain, we enforce two bridge losses i.e.,  $\mathcal{L}_{bridge}^f$  and  $\mathcal{L}_{bridge}^\varphi$ . (3) Given the output of MGM, i.e.,  $\varphi^{s \rightarrow inter}$  and  $\varphi^{t \rightarrow inter}$  in the intermediate domain, we enforce the consistency loss  $\mathcal{L}_{cons}$ . (4) To generate more intermediate domains, we use  $\mathcal{L}_{div}$  to enlarge the variance of the mix ratio

Song et al., 2020; Ge et al., 2020b), we perform DBSCAN clustering on the target domain features at the beginning of every training epoch and assign the cluster id (a total of  $C^t$  clusters) as the pseudo identity  $y_i^t$  of the target sample  $x_i^t$ . We use ResNet-50 (He et al., 2016) 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.

Figure 2a shows the overall framework of our method. Our IDM++ adopts the source-target joint training pipeline, i.e., each mini-batch consists of  $n$  source samples and  $n$  target samples. IDM++ has two key components, i.e., IDM and MGM. Both IDM and MGM can be seamlessly plugged into the backbone network (e.g., Resnet-50), while IDM should be placed on earlier layers because it is a prerequisite for MGM. An optimized configuration is to plug IDM between the stage-0 and stage-1 and plug MGM between the stage-3 and stage-4. The IDM generates intermediate domains by mixing the source-domain and target-domain feature with

two corresponding mix ratios, i.e.,  $a^s$  and  $a^t$ . The mixed features, along with the original source-domain and target-domain features in the hidden layer continue to proceed until the deep embedding, i.e., the features after the global average pooling (GAP), and correspondingly become  $f^{inter}$ ,  $f^s$  and  $f^t$ . With the classifier, these features are mapped to  $\varphi^{inter}$ ,  $\varphi^s$  and  $\varphi^t$  respectively. The MGM maps the original source/target identities into the IDM-generated intermediate domains to obtain the mirrors. After the classifier, we obtain the source and target mirrors' predictions:  $\varphi^{s \rightarrow inter}$  and  $\varphi^{t \rightarrow inter}$ .

### 3.2 Intermediate Domain Module

In this section, we first illustrate how to utilize IDM to generate intermediate domains' features in the hidden stage of the backbone (Sect. 3.2.1). Next, we provide the manifold assumption to exploit two properties (proportional distance relationship in Sect. 3.2.2 and diversity in Sect. 3.2.3) that the intermediate domains should satisfy in the deeply-learned output space.

### 3.2.1 Modelling the Intermediate Domains

We denote backbone network as  $f(x) = f_m(G)$ , where  $G$  is the hidden feature map  $G \in \mathbb{R}^{h \times w \times c}$  after the  $m$ -th stage and  $f_m$  represents the part of the network mapping the hidden representation  $G$  after the  $m$ -th hidden stage to the 2048-dim feature after the GAP layer. As shown in Fig. 2b, IDM contains two steps: (1) **predicting the mix ratios** and (2) **mixing the features**. Specifically, IDM is plugged after a bottom (earlier) stage in the backbone network and takes the hidden-layer features from both the source and target domain ( $G^s$  and  $G^t$ ) as its input to predict the two mix ratios (i.e.,  $a^s$  and  $a^t$ ). Next, IDM mixes  $G^s$  and  $G^t$  with two corresponding mix ratios to generate intermediate domains hidden-layer features  $G^{\text{inter}}$ . The  $G^s$ ,  $G^t$  and  $G^{\text{inter}}$  forward into the next hidden stage until the final output of the network.

**Predicting the Mix Ratios** As shown in Fig. 2b, in a mini-batch comprised of  $n$  source and  $n$  target samples, we randomly assign all the samples into  $n$  pairs so that each pair contains a source-domain and a target-domain sample. For each sample pair  $(x^s, x^t)$ , the network obtains two feature maps at the  $m$ -th hidden stage:  $G^s, G^t \in \mathbb{R}^{h \times w \times c}$ . IDM uses an average pooling and a max pooling operation to transform each single feature into two  $1 \times 1 \times c$  dimensional features. Therefore, for each source-target pair, we have  $(G_{avg}^s, G_{max}^s)$  for the source domain, and  $(G_{avg}^t, G_{max}^t)$  for the target domain. IDM concatenates the *avg*-pooled and *max*-pooled features for each domain and feed them into a fully-connected layer, i.e., FC1 in Fig. 2b. The output feature vectors of FC1 are merged using element-wise summation and fed into a multi-layer perception (MLP) to obtain a mix ratio vector  $a = [a^s, a^t] \in \mathbb{R}^2$ . The above step for predicting the mix ratios is formulated as:


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

where  $\delta(\cdot)$  is the softmax function to ensure  $a^s + a^t = 1$ .

**Mixing the Features** Given the predicted mix ratios  $a^s$  and  $a^t$ , IDM mixes the source-domain and target-domain features ( $G^s$  and  $G^t$ ) using weighted mean operation, as illustrated in Fig. 2 (b). The mixing step is formulated as:

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

These hidden features (i.e.,  $G^s$ ,  $G^t$ , and  $G^{\text{inter}}$ ) proceed until the outputs (e.g., embeddings after the GAP layer or logits after the classifier in re-ID) of the network. We note a side-effect accompany the above mixing step, i.e., the mixing of identities. Concretely, we consider the mixed feature  $G^{\text{inter}}$  has a new identity, i.e.,  $y^{\text{inter}} = a^s \cdot y^s + a^t \cdot y^t$ .



**Fig. 3** Assuming domains as points in a manifold. The “shortest geodesic path” connecting the source and target domains (i.e.,  $P_s$  and  $P_t$ ) is bridged by dense intermediate domains (i.e.,  $P_{\text{inter}}$ ). We use the mix ratios (i.e.,  $a^s$  and  $a^t$ ) to approximate the distance relationship between domains in the manifold

### 3.2.2 Proportional Distance Relationship

A series of smooth transformations from the source to the target domain, can be conceptualized as a form of a ‘manifold’ in the high-dimensional feature space. Though it is not very mathematically rigorous, the concept of ‘manifold’ can still serve as an effective tool to analyze the geometric relationships in the deep embedding space. We assume that the source and target domains are located on a manifold, i.e., a topological space that locally resembles Euclidean space. On a manifold, the distance between two isolated and far-away points cannot be measured through Euclidean distance. Instead, the distance should be measure through the “shortest geodesic distance” (Gong et al., 2012; Gopalan et al., 2013), i.e., the curve representing the shortest path bridging two isolated points in a manifold surface (as shown in Fig. 3)

**Definition (Shortest Geodesic Distance)** To estimate such “shortest geodesic distance”, we assume there exist many intermediate domains that are close to both the source and target domains. In other words, we assume that the distance between the intermediate domain and the source/target domain is small enough that they may be viewed as lying in a local Euclidean space. Therefore, we sum up the source-to-intermediate Euclidean distance and the target-to-intermediate Euclidean distance as the shortest geodesic distance between the source and target domain samples, which is formulated as:

$$d(P_s, P_t) = d(P_s, P_{\text{inter}}) + d(P_t, P_{\text{inter}}). \quad (3)$$

where the  $P_s$ ,  $P_t$  and  $P_{\text{inter}}$  are the source domain, the target domain and the intermediate domain, respectively. In IDM, when the mixed feature  $G^{\text{inter}}$  (Eq. (2)) proceeds along the sub-sequential layers, we view all the corresponding features in the sub-sequential layers are within the intermediate domain. Based on the above definition, the objective of aligning the source against target domain can be transformed into simultaneously aligning the source and target domain against the intermediate domains.

In Eq. (2), the mix ratios  $a^s$  and  $a^t$  jointly control the relevance of the intermediate domains to the source and target. If the source ratio  $a^s$  is larger than the target ratio  $a^t$ , the synthetic intermediate domain should be closer to the source domain than to the target domain. Therefore, we utilize the mix ratio to approximately estimate the distance relationship between the intermediate domains' outputs and other two isolated domains' outputs. Specifically, we formulate such distance relationship of these domains' outputs as the following **Property-1**.

**Property-1 (Distance Should be Proportional):** We use the mix ratios (i.e.,  $a^s$  and  $a^t$ ) to approximate the distance relationship between the intermediate domains' output  $P_{\text{inter}}$  and the other two isolated domains' output (i.e.,  $P_s$  and  $P_t$ ). The distance relationship is formulated as follows:

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

Using the proportional distance relationship in Eq. (4), we can approximately locate the position of the intermediate domains in the deeply-learned embedding space or prediction space.

### 3.2.3 Promoting the Diversity of Intermediate Domains

As shown in the “shortest geodesic distance” definition, the two isolated domains (i.e., source and target) cannot be directly measured in the Euclidean space if they are far away from each other. The distance of two distant points must be measured by the points existing between them. To estimate the “shortest geodesic distance” more precisely, we need as many as possible intermediate points. Compared with sparse intermediate points, more dense points can characterize the “shortest geodesic path” bridging the source and target domains more comprehensively. Therefore, we propose the **Property-2** as follows.

**Property-2 (Diversity)** Intermediate domains should be as diverse as possible.

Based on the above property, the distance of the source and target domains (i.e.,  $d(P_s, P_t)$ ) can be correctly measured in the manifold by approximating the distance between the dense intermediate domains and the other two isolated domains (i.e.,  $d(P_s, P_{\text{inter}})$  and  $d(P_t, P_{\text{inter}})$ ). To satisfy such property when generating intermediate domains in training, we propose a diversity loss by maximizing the differences of the mix ratios ( $a^s$  and  $a^t$ ) within a mini-batch. This loss is formulated as follows:

$$\mathcal{L}_{\text{div}} = -[\sigma(\{a_i^s\}_{i=1}^n) + \sigma(\{a_i^t\}_{i=1}^n)], \quad (5)$$

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.

### 3.3 Aligning by Intermediate Domains

Instead of directly aligning the source and target domains (i.e., source  $\rightarrow$  target), we utilize the IDM to align the source and target domains against the synthetic intermediate domains (i.e., source  $\rightarrow$  intermediate and target  $\rightarrow$  intermediate), as illustrated in Fig. 1. Given the “shortest geodesic distance” definition and **Property-1** in Sect. 3.2.1, we propose the bridge loss to adaptively minimize the discrepancy between the IDM-generated intermediate domains' output (i.e.,  $P_{\text{inter}}$ ) and the source and target domains' output (i.e.,  $P_s$  and  $P_t$ ). The general form of the bridge loss is formulated as follows:

$$\mathcal{L}_{\text{bridge}} = a^s \cdot d(P_s, P_{\text{inter}}) + a^t \cdot d(P_t, P_{\text{inter}}). \quad (6)$$

The intuition of the  $\mathcal{L}_{\text{bridge}}$  is based on **Property-1**. As shown in **Property-1**, we utilize the proportional relationship (i.e., Eq. (4)) to estimate the location of the IDM-generated intermediate domains along the domain bridge. If the intermediate domain  $P_{\text{inter}}$  is closer to the source  $P_s$  than the target  $P_t$  (i.e.,  $a^s$  is larger than  $a^t$ ), the objective of minimizing the discrepancy of “source  $\rightarrow$  target” (i.e.,  $d(P_s, P_{\text{inter}})$ ) is easier than minimizing the discrepancy of “target  $\rightarrow$  intermediate” (i.e.,  $d(P_t, P_{\text{inter}})$ ). Thus, we force the minimization objective to penalize more on  $d(P_s, P_{\text{inter}})$  than  $d(P_t, P_{\text{inter}})$  by multiplying  $d(P_s, P_{\text{inter}})$  by  $a^s$  and multiplying  $d(P_t, P_{\text{inter}})$  by  $a^t$ . By utilizing the mix ratios as the weighting coefficients in Eq. (6), the alignment between the intermediate domains and the source and target domains can be adaptively conducted.

In a deep model for UDA re-ID, we consider to enforce the bridge loss (Eq. (6)) on the feature and prediction space of the network to align the source and target by the intermediate domains. For the prediction space, we use the cross-entropy to measure the discrepancy between intermediate domains' prediction logits and other two extreme domains' (pseudo) labels (Eq. (7)). For the feature space, we use the L2-norm to measure features' distance among domains (Eq. (8)). Our proposed two bridge losses are formulated as follows:

$$\mathcal{L}_{\text{bridge}}^{\varphi} = -\frac{1}{n} \sum_{i=1}^n \sum_{k \in \{s,t\}} a_i^k \cdot \left[ y_i^k \log(\varphi(f_m(G_i^{\text{inter}}))) \right], \quad (7)$$

$$\mathcal{L}_{\text{bridge}}^f = \frac{1}{n} \sum_{i=1}^n \sum_{k \in \{s,t\}} a_i^k \cdot \left\| f_m(G_i^k) - f_m(G_i^{\text{inter}}) \right\|_2. \quad (8)$$



In Eqs. (7) (8), 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 hidden stage to the features after GAP layer and  $\varphi(\cdot)$  is the classifier.

In conclusion, we first utilize IDM to generate the intermediate domains by mixing up the representations of the source and target (with two corresponding mix ratios ( $a^s$  and  $a^t$ )) at the  $m$ -th hidden stage. Second, we propose the diversity loss  $\mathcal{L}_{\text{div}}$  by maximizing the diversity of the mix ratios, in order to generate intermediate domains bridging the source and target as diverse as possible. Third, we propose the bridge losses (i.e.,  $\mathcal{L}_{\text{bridge}}^\varphi$  and  $\mathcal{L}_{\text{bridge}}^f$ ) to adaptively minimize the discrepancy of “source  $\rightarrow$  intermediate” and the discrepancy of “target  $\rightarrow$  intermediate”.

### 3.4 Mirrors Generation Module

We recall that when IDM mixes a source-domain and a target-domain feature to obtain an intermediate-domain feature, it has an side-effect of mixing their identities (Sect. 3.2.1). Consequentially, the intermediate domains do not contain the original training identities. We argue that these original identities are also beneficial for domain alignment, because anchoring the identity allows focusing on the cross-domain variance and thus benefits the domain alignment. To this end, Mirrors Generation Module (MGM) (1) maps the original source and target-domain identities into the intermediate domains and obtain their mirrors and (2) uses a cross-domain consistency loss between the mirrors and their original features for domain alignment.

#### Mapping the Features into the Intermediate Domains

Given a sample’s feature map  $G \in \mathbb{R}^{H \times W \times C}$  at the  $l$ -th hidden stage of the network, its feature statistics are denoted as the mean  $\mu(G) \in \mathbb{R}^C$  and the standard deviation  $\sigma(G) \in \mathbb{R}^C$ , where each channel is calculated across spatial dimensions independently as follows:

$$\mu_c(G) = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W G_{hwc}, \tag{9}$$

$$\sigma_c(G) = \sqrt{\frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W (G_{hwc} - \mu_c(G))^2}. \tag{10}$$

Motivated by recent studies on out-of-distribution (Zhou et al., 2021; Nuriel et al., 2021; Tang et al., 2021), feature statistics (i.e.,  $\mu$ ,  $\sigma$ ) are utilized to represent a sample’s domain characteristics and the normalized content (i.e.,  $(G - \mu)/\sigma$ ) is utilized to represent a sample’s identity-

specific information. For a sample’s feature map  $G^a$  from the domain “ $a$ ” and a sample’s feature map  $G^b$  from the domain “ $b$ ”, we utilize AdaIN (Huang & Belongie, 2017) to map  $G^a$  into the domain “ $b$ ” as follows:

$$G^{a \rightarrow b} = \text{AdaIN}(G^a, G^b) = \sigma(G^b) \left( \frac{G^a - \mu(G^a)}{\sigma(G^a)} \right) + \mu(G^b). \tag{11}$$

By utilizing  $\text{AdaIN}(G^a, G^b)$ , we can easily characterize the sample’s feature  $G^a$  with the style of the domain “ $b$ ” to get the mirror feature  $G^{a \rightarrow b}$ . Compared with  $G^a$ ,  $G^{a \rightarrow b}$  can represent the same person identity but different domain style.

Similar to the formulation in Sect. 3.2, we can obtain source, target, and intermediate domains’ feature maps  $\{(G_i^s, G_i^t, G_i^{\text{inter}})\}_{i=1}^n$  at the  $l$ -th hidden stage of the network. Then we utilize Eq. (11) to characterize the source and target feature maps with intermediate domains’ statistics as follows:

$$\begin{cases} G_i^{s \rightarrow \text{inter}} = \text{AdaIN}(G_i^s, G_i^{\text{inter}}) \\ G_i^{t \rightarrow \text{inter}} = \text{AdaIN}(G_i^t, G_i^{\text{inter}}) \end{cases}. \tag{12}$$

With Eq. (12), we can obtain the source mirrors  $\{G_i^{s \rightarrow \text{inter}}\}_{i=1}^n$  and target mirrors  $\{G_i^{t \rightarrow \text{inter}}\}_{i=1}^n$  characterized by intermediate domain styles. We denote the operation of Eq. (12) as generating mirrors in MGM in Fig. 2c.

**Cross-Domain Consistency** MGM uses cross-domain consistency to make the original identity-discriminative information consistent across different domains. Specifically, MGM enforces a consistency loss between the prediction space  $\varphi$  of samples and the  $\varphi$  of their mirrors. For simplicity, we denote predictions of  $n$  samples in a domain as  $\{\varphi_i\}_{i=1}^n$ , representing an identity’s distribution characteristic.

In the prediction space, we conduct the softmax function on the logits after the classifier to model the prediction distribution. Given the feature map  $G$  at the  $l$ -th hidden stage of the network, and the (pseudo) label  $y$ , the prediction distribution is formulated as follows:

$$P(y | G) = \frac{\exp(\varphi_y(f_l(G))/\tau)}{\sum_{i=1}^{C^s+C^t} \exp(\varphi_i(f_l(G))/\tau)}, \tag{13}$$

where  $f_l(\cdot)$  denotes the part of the network mapping  $G$  after the  $l$ -th hidden stage to the 2048-dim feature after the GAP layer,  $\varphi_i$  denotes the logit of the classifier for class  $i$ , and  $\tau > 0$  is the temperature scaling parameter. Thus, we propose a class-wise consistency loss that enforces consistency on predictions of samples and their mirrors in intermediate domains. Given a sample’s feature map  $G^a$  from domain “ $a$ ” and its mirror (sharing the same identity of  $G^a$ )  $G^{a \rightarrow b}$  stylized by domain “ $b$ ”, their discrepancy in the prediction space is calculated as follows:

**Algorithm 1** The overall training procedure

---

**Input:** Source dataset  $\{(x_i^s, y_i^s)\}$  and target dataset  $\{x_i^t\}$ ;  
**Output:** The trained backbone network  $f(\cdot)$  and classifier  $\varphi(\cdot)$ ;

- 1 Initialize the ImageNet-pretrained backbone network  $f(\cdot)$ ;
- 2 Initialize the XBM memory as an empty queue  $M$ ;
- 3 Plug our IDM and MGM at the  $m$ -th and  $l$ -th stage respectively.
- 4 **for**  $epoch = 1$  to  $MaxEpochs$  **do**
- 5     Use  $f(\cdot)$  to extract features  $\{f_i^t\}$  for the target dataset  $\{x_i^t\}$ ;
- 6     Assign pseudo labels  $\{y_i^t\}$  for target samples  $\{x_i^t\}$  by performing DBSCAN clustering on their features  $\{f_i^t\}$ ;
- 7     **for**  $iter = 1$  to  $MaxIters$  **do**
- 8         Sample a mini-batch including  $n$  source samples  $\{(x_i^s, y_i^s)\}_{i=1}^n$  and  $n$  target samples  $\{(x_i^t, y_i^t)\}_{i=1}^n$ ;
- 9         Using IDM to obtain  $f^{inter}$  and  $\varphi^{inter}$ ;
- 10         Using MGM to obtain  $\varphi^{s \rightarrow inter}$  and  $\varphi^{t \rightarrow inter}$ ;
- 11         **if using XBM then**
- 12             Enqueue( $M, \{(f^t, y^t), \{(f^s, y^s)\}$ );
- 13             **if**  $M$  is full **then**
- 14                 Dequeue( $M$ );
- 15             **end**
- 16             Use all entries in  $M$  for the hard negatives mining in the triplet loss in  $\mathcal{L}_{ReID}$  from Eq. (16);
- 17         **end**
- 18         Optimizing the network by the gradients of Eq. (16);
- 19     **end**
- 20 **end**

---

$$D_\varphi(G^a, G^{a \rightarrow b}) = \text{KL} \left( P(y | G^a) \parallel P(y | G^{a \rightarrow b}) \right) + \text{KL} \left( P(y | G^{a \rightarrow b}) \parallel P(y | G^a) \right), \quad (14)$$

where KL is the Kullback–Leibler (KL) divergence. With Eq. (14), the cross-domain consistency loss is formulated as follows:

$$\mathcal{L}_{\text{cons}} = \tau^2 \cdot \left( D_\varphi(G^s, G^{s \rightarrow inter}) + D_\varphi(G^t, G^{t \rightarrow inter}) \right). \quad (15)$$

Following the knowledge distillation method (Hinton et al., 2015), we multiple the square of the temperature  $\tau^2$ . By minimizing Eq. (15), we can encourage the model to focus on minimizing the diverse domain variation during the domain alignment.

**Overall Training for IDM++** Commonly, 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 IDM at the  $m$ -th stage of ResNet-50 (e.g., stage-0 as shown in Fig. 2) to generate intermediate domains, and plug the MGM at the  $l$ -th stage (e.g., stage-3 in Fig. 2). The overall training loss is as follows:

$$\mathcal{L} = \mathcal{L}_{\text{ReID}} + \mu_1 \cdot \mathcal{L}_{\text{bridge}}^\varphi + \mu_2 \cdot \mathcal{L}_{\text{bridge}}^f + \mu_3 \cdot \mathcal{L}_{\text{div}} + \mu_4 \cdot \mathcal{L}_{\text{cons}}, \quad (16)$$

where  $\mathcal{L}_{\text{ReID}} = (1 - \mu_1) \cdot \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{tri}}$ , and  $\mu_1, \mu_2, \mu_3, \mu_4$  are the weights to balance losses. The training procedure is shown in Algorithm 1.

### 3.5 Extension to Domain Generalizable Person Re-ID

Compared with UDA re-ID, domain generalizable re-ID (DG re-ID) is more challenging because target domain data can not be accessible during training. Different from UDA re-ID, the keynote of DG re-ID is to make the deep model robust to the variation of multiple source domains, so that the learned features can be well generalized to any unseen target domain. For example, the existing DG re-ID works (Song et al., 2019; Jin et al., 2020b; Zhao et al., 2021a) also call these features as domain-invariant features, or identity-relevant features.

While our IDM++ is initially designed for UDA, we may apply it to DG scenario through slight modifications. Specifically, we merge multiple domains (without using domain labels) into a hybrid training domain. During training, each mini-batch contains features from multiple domains. For a mini-batch of  $n$  samples, we first obtain their feature maps  $\{G_i^s\}_{i=1}^n$  at the  $m$ -th hidden stage and randomly shuffle them along the batch dimension to obtain out-of-order feature maps  $\{\widetilde{G}_i^s\}_{i=1}^n$ . Next, we consider  $\{\widetilde{G}_i^s\}_{i=1}^n$  as pseudo target domains' feature maps and use IDM to generate intermediate domains' representations with Eq. (2). When mapping source identities into intermediate domains with the MGM, we just use AdaIN( $\cdot, \cdot$ ) to change statistics values of the source-domain features with statistics values of the intermediate-domain features and obtain source mirrors  $\{G_i^{s \rightarrow inter}\}_{i=1}^n$ . Next, we only enforce the consistency loss  $\mathcal{L}_{\text{cons}}$  between  $\{G_i^{s \rightarrow inter}\}_{i=1}^n$  and  $\{G_i^s\}_{i=1}^n$ . The learning objective of DG re-ID is the same as the objective of UDA re-ID (i.e., Eq. (16)).

Through the above training procedure, IDM++ improves the robustness against domain variation. There are two main reasons for the improvement. First, IDM adds more source domains for training, because the synthetic intermediate domains may be viewed as extra source domains for training. Second, MGM further enforces consistency among multiple domains and promotes learning domain-invariant features. Therefore, IDM and MGM jointly helps IDM++ to enhance the feature robustness against domain variation, which benefits generalization towards novel unseen target domains.

## 4 Experiments

We introduce the datasets and evaluation protocols in Sect. 4.1. Section 4.2 describes the implementation details of IDM++ for both UDA and DG re-ID, respectively. Section 4.3 conducts ablation study to analyze the improvement of IDM++, as well as its two key components (i.e., IDM and

MGM). Section 4.4 analyzes the impact of some important hyper-parameters. Section 4.5 compares our IDM++ with the state-of-the-art methods on twelve UDA re-ID benchmarks and two protocols of DG re-ID. Section 4.6 validates some key designs within IDM and MGM, while Sect. 4.7 investigates the mechanism of IDM++ and reveals how the intermediate domains improve the domain alignment.

#### 4.1 Datasets and Evaluation Protocols

**Datasets** A total of six person re-ID datasets are used in our experiments, including Market-1501 (Zheng et al., 2015), DukeMTMC-reID (Ristani et al., 2016; Zheng et al., 2017), CUHK03 (Li et al., 2014), MSMT17 (Wei et al., 2018), PersonX (Sun & Zheng, 2019), and Unreal (Zhang et al., 2021a). Among these datasets, PersonX and Unreal are synthetic datasets, and all the remaining datasets are from the real-world. We adopt CUHK03-NP (Zhong et al., 2017) when testing on CUHK03.

**Evaluation Protocols** 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+ (Li & Zhang, 2020). In testing, there are no post-processing techniques like re-ranking (Zhong et al., 2017) or multi-query fusion (Zheng et al., 2015). We follow the mainstream evaluation protocols to evaluate our method in both UDA re-ID and DG re-ID tasks. (1) UDA re-ID protocols: During training, we train on a labeled source dataset and an unlabeled target dataset. We use mAP and CMC to test the model on the target dataset. Specifically, we evaluate on two kinds of UDA re-ID tasks, i.e., real  $\rightarrow$  real and synthetic  $\rightarrow$  real, where the source datasets are real and synthetic respectively. (2) DG re-ID protocols: Following recent DG re-ID works (Zhao et al., 2021b; Dai et al., 2021a), we adopt the leave-one-out setting. Given a pool of datasets (Market-1501, DukeMTMC-reID, CUHK03, and MSMT17), we use a single dataset for testing and use all the other datasets for training. Specifically, we use two protocols, the “DG-partial” (Zhao et al., 2021b) and the “DG-full” (Dai et al., 2021a). DG-partial uses only the training set while DG-full combines the training and testing set for training.

#### 4.2 Implementation Details

ResNet-50 (He et al., 2016) pretrained on ImageNet is adopted as the backbone network. Following (Ge et al., 2020b), domain-specific BNs (Chang et al., 2019) are also used in the backbone network to narrow domain gaps. Following (Luo et al., 2019), we resize the image size to  $256 \times 128$  and apply some common image augmentation techniques, including random flipping, random cropping, and random erasing (Zhong et al., 2020a).

For UDA re-ID, we perform DBSCAN (Ester et al., 1996) clustering on the unlabeled target data to assign pseudo labels before each training epoch, which is consistent with (Fu et al., 2019; Song et al., 2020; Ge et al., 2020b). The mini-batch size is 128, including 64 source images of 16 identities and 64 target images of 16 pseudo identities. We totally train 50 epochs and each epoch contains 400 iterations. The initial learning rate is set as  $3.5 \times 10^{-4}$  which will be divided by 10 at the 20th and 40th epoch respectively. The Adam optimizer with weight decay  $5 \times 10^{-4}$  and momentum 0.9 is adopted in our training. We implement Strong Baseline of UDA re-ID with XBM (Wang et al., 2020) to mine more hard negatives for the triplet loss.

For DG re-ID, the training setting follows BoT (Luo et al., 2019). However, directly enforcing the cross-domain consistency loss (Eq. (15)) on the domain samples and their mirrors may cause the feature collapse problem (Grill et al., 2020). Feature collapse means the learned features become uninformative and indistinguishable from one another. The collapse problem can occur when enforcing consistency loss in domain generalization, as it might lead the model to learn features that are too smooth and lack the necessary discriminability to distinguish between different identities. As a result, to prevent learning identity-discriminative features from getting into collapse solutions when enforcing the consistency loss in DG re-ID, we use the twin-branch network structure (duplicating the branch from the  $l$ -th stage to the classifier), and feed the original samples and mirrors into different branches. Similar to BYOL (Grill et al., 2020), we enforce the consistency loss on the outputs of the twin-branches.

If not specified, the loss weights  $\mu_1, \mu_2, \mu_3, \mu_4$  are set as 0.7, 0.1, 1.0, 1.0 respectively, and the temperature  $\tau$  is set as 0.5 in all our experiments. In the IDM module, the FC1 layer is parameterized by  $W_1 \in \mathbb{R}^{c \times 2c}$  and MLP is composed of two fully connected layers (followed by batch normalization) which are parameterized by  $W_2 \in \mathbb{R}^{(c/r) \times c}$  and  $W_3 \in \mathbb{R}^{2 \times (c/r)}$  respectively, where  $c$  is the representations' channel number after the  $m$ -th stage and  $r$  is the reduction ratio. If not specified, we plug the IDM module at the stage-0 of ResNet-50 and set  $r$  as 2, and plug the MGM at the stage-3. The IDM and MGM modules are only used in training and will be discarded in testing. Our method is implemented with Pytorch, and four Nvidia Tesla V100 GPUs are used for training and only one GPU is used for testing.


#### 4.3 Ablation Study

This section evaluates the effectiveness of each component in IDM++ and the results are provided in Table 1. The “Oracle” method in Table 1 uses the ground-truth label for the target (and source) domain, marking the upper bound of UDA re-ID accuracy.

**Table 1** Ablation studies on different components of our method

Method	IDM			MGM	DukeMTMC→Market-1501		Market-1501→DukeMTMC	
	$\mathcal{L}_{\text{bridge}}^{\phi}$	$\mathcal{L}_{\text{bridge}}^f$	$\mathcal{L}_{\text{div}}$	$\mathcal{L}_{\text{cons}}$	mAP	R1	mAP	R1
Oracle					83.9	93.2	75.0	86.1
Baseline1					77.0	90.6	63.4	78.4
IDM		✓	✓		79.4	91.5	66.2	79.8
	✓		✓		79.2	91.2	65.8	80.4
	✓	✓			78.9	91.1	64.8	79.3
	✓	✓	✓		81.9	92.4	68.6	82.3
IDM++		✓	✓	✓	81.2	92.3	68.5	81.8
	✓		✓	✓	81.4	91.8	69.5	82.1
	✓	✓		✓	80.9	92.0	69.9	82.4
	✓	✓	✓	✓	82.4	92.7	71.6	83.4
Oracle					86.9	94.8	78.1	88.3
Baseline2					79.1	91.2	65.8	80.1
IDM		✓	✓		80.7	92.0	67.3	81.8
	✓		✓		81.6	92.2	69.0	82.0
	✓	✓			79.7	91.6	67.8	81.6
	✓	✓	✓		82.8	93.2	70.5	83.6
IDM++		✓	✓	✓	83.4	93.1	70.0	83.0
	✓		✓	✓	84.3	93.4	72.2	83.9
	✓	✓		✓	83.8	93.3	72.5	84.4
	✓	✓	✓	✓	85.3	94.2	73.2	85.5

Baseline1 (Naive Baseline): only using  $\mathcal{L}_{\text{ReID}}$  to train the source and target domains jointly. Baseline2 (Strong Baseline): Baseline1 + XBM (Wang et al., 2020)



**Fig. 4** Analysis of different values of hyper-parameters:  $\mu_1$ ,  $\mu_2$ ,  $\mu_3$ ,  $\mu_4$ , and  $\tau$  on the target domain when transferring from Market-1501 to DukeMTMC-reID

**IDM Improves the Baseline** IDM is inserted in the hidden stage of the network to generate intermediate domains' features that can be utilized to better align the source and target domains. In IDM, there are two functions: (1) using the diversity loss (i.e.,  $\mathcal{L}_{\text{div}}$ ) to generate dense intermediate domains as many as possible and (2) using the bridge losses (i.e.,  $\mathcal{L}_{\text{bridge}}^{\phi}$  and  $\mathcal{L}_{\text{bridge}}^f$ ) to dynamically minimize the discrepancy between the intermediate domains and the source and target domains. We evaluate the effectiveness of both functions in Table 1. (1) The effectiveness of  $\mathcal{L}_{\text{div}}$ : “Baseline2 + IDM w/o  $\mathcal{L}_{\text{div}}$ ” degenerates much when comparing with the full method. The Rank-1 of “Baseline2 + IDM++ w/o  $\mathcal{L}_{\text{div}}$ ” is 1.1% lower than the full method on Market → Duke. (2) The effectiveness of  $\mathcal{L}_{\text{bridge}}^{\phi}$  and

$\mathcal{L}_{\text{bridge}}^f$ : Taking Market→Duke as an example, mAP/R1 of “Baseline2 + IDM++ w/o  $\mathcal{L}_{\text{bridge}}^{\phi}$ ” is 3.2%/2.5% lower than “Baseline2 + IDM++(full)”. The mAP/R1 of “Baseline2 + IDM++ w/o  $\mathcal{L}_{\text{bridge}}^f$ ” is 1.0%/1.6% lower than “Baseline2 + IDM++(full)”. Compared with  $\mathcal{L}_{\text{bridge}}^f$ ,  $\mathcal{L}_{\text{bridge}}^{\phi}$  is more important based on Baseline2. The reason may be that Baseline2 uses XBM (Wang et al., 2020) to mine harder negatives in the feature space, which will affect the effectiveness of  $\mathcal{L}_{\text{bridge}}^f$ . The above analyses show that all the functions in the IDM are important for learning appropriate intermediate domains that can help to bridge the source and target domains.

**MGM Achieves Further Improvement Based on IDM** Compared with our previous work (IDM (Dai et al., 2021c)),

this paper proposes a novel MGM to map source/target features into intermediate domains, which can help learn better identity-discriminative representations across domains. We first use the MGM to map the source and target domains' feature maps to intermediate domains at the  $l$ -th (if not mentioned, we set  $l$  as 3) stage and then we use the class-wise domain consistency loss  $\mathcal{L}_{\text{cons}}$  to enforce the consistency on predictions of source and target samples together with their mirrors. In Table 1, we evaluate the effectiveness of the newly proposed MGM compared with our previous IDM. Taking Market→Duke as an example, mAP/R1 of “Baseline2 + IDM++(full)” is 2.7%/1.9% higher than “Baseline2 + IDM(full)”.

**IDM++ integrating IDM and MGM Brings General Improvement on Multiple Baselines** 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 Fig. 2. Compared with Naive Baseline, Strong Baseline uses XBM (Wang et al., 2020) to mine more hard negatives for the triplet loss, which is a variant of the memory bank (Xiao et al., 2017; Wu et al., 2018) and easy to implement. Similar to those UDA re-ID methods (Zhong et al., 2020b; Ge et al., 2020b) 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 obviously outperform the baseline methods by a large margin. On Market → Duke, mAP/R1 of “Baseline2 + our IDM++ (full)” is 7.4%/5.4% higher than Baseline2, and mAP/R1 of “Baseline1 + our IDM++ (full)” is 8.2% and 5.0% higher than Baseline1. Because of many state-of-the-arts methods (Ge et al. 2020b; Li and Zhang 2020; Zheng et al. 2021a,b) 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.

#### 4.4 Parameter Analysis

In Fig. 4, we provide the visualization on the sensitivity of hyper-parameters. Specifically, we evaluate the loss weights:  $\mu_1, \mu_2, \mu_3$ , and  $\mu_4$  in Eq. (16), and the temperature value  $\tau$  in Eq. (13). We conduct the experiments on Market-1501 → DukeMTMC-reID and evaluate the mAP and Rank-1 on the target domain: DukeMTMC-reID. As shown in Fig. 4, these hyper-parameters are not very sensitive. The performance achieves the best when we set  $\mu_1, \mu_2, \mu_3, \mu_4$ , and  $\tau$  as 0.7, 0.1, 1.0, 1.0, and 0.5 respectively. If not specified, we utilize the above setting of hyper-parameters in all other experiments in this paper.

#### 4.5 Comparison with the State-of-the-arts

The existing state-of-the-art (SOTA) UDA re-ID works commonly evaluate the performance on six real → real tasks (Fu et al., 2019; Wu et al., 2019; Zhong et al., 2020b). Recently, more challenging synthetic → real tasks (Ge et al., 2020b; Zhang et al., 2021a) are proposed, where they use the synthetic dataset PersonX (Sun & Zheng, 2019) or Unreal (Zhang et al., 2021a) as the source domain and test on other three real re-ID datasets. Besides, the existing SOTA DG re-ID works usually evaluate the performance under two challenging “leave-one-out” protocols (Zhao et al., 2021b; Dai et al., 2021a), i.e., “DG-partia” and “DG-full”. In the “DG-partial” protocol (Zhao et al., 2021b), the train-sets of three datasets among the four datasets are used for training and the test-set of the remaining dataset is used for testing. In the “DG-full” protocol (Dai et al., 2021a), all data (including train-sets and test-sets) of three datasets among the four datasets are used for training and the test-set of the remaining dataset is used for testing. To compare with the SOTAs fairly, we also take ResNet-50 as the backbone. All the results in Tables 2, and 3 show that our method can outperform the UDA re-ID SOTAs by a large margin, and the results in Table 4 show that our method is superior to the DG re-ID 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 (Wei et al., 2018), SPGAN+LMP (Deng et al., 2018), (Zhong et al., 2018), and PDA-Net (Li et al., 2019). (2) Fine-tuning methods include PUL (Fan et al., 2018), PCB-PAST (Zhang et al., 2019a), SSG (Fu et al., 2019), AD-Cluster (Zhai et al., 2020a), MMT (Ge et al., 2020a), NRMT (Zhao et al., 2020), MEB-Net (Zhai et al., 2020b), Dual-Refinement (Dai et al., 2021b), UNRN (Zheng et al., 2021a), and GLT (Zheng et al., 2021b). (3) Joint training methods commonly use the memory bank (Xiao et al., 2017), including ECN (Zhong et al., 2019), MMCL (Wang & Zhang, 2020), ECN-GPP (Zhong et al., 2020b), JVTC+ (Li & Zhang, 2020), SpCL (Ge et al., 2020b), HCD (Zheng et al., 2021c), CCL (Isobe et al., 2021). 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. In our previous work IDM (Dai et al., 2021c), we propose an IDM module to generate appropriate intermediate domains to better improve the performance of UDA re-ID. As an extension of IDM (Dai et al., 2021c), we further propose IDM++ by integrating MGM and IDM to focus on the cross-domain variation when aligning the source and target domains. As shown in Table 2, our method can outperform the second best UDA re-ID methods by a large

**Table 2** Comparison with the state-of-the-art UDA re-ID methods on real → real tasks

Methods	Reference	DukeMTMC-ReID → Market-1501				Market-1501 → DukeMTMC-ReID			
		mAP	R1	R5	R10	mAP	R1	R5	R10
PUL (Fan et al., 2018)	TOMM 2018	20.5	45.5	60.7	66.7	16.4	30.0	43.4	48.5
TJ-AIDL (Wang et al., 2018)	CVPR 2018	26.5	58.2	74.8	81.1	23.0	44.3	59.6	65.0
SPGAN+LMP (Deng et al., 2018)	CVPR 2018	26.7	57.7	75.8	82.4	26.2	46.4	62.3	68.0
HHL (Zhong et al., 2018)	ECCV 2018	31.4	62.2	78.8	84.0	27.2	46.9	61.0	66.7
ECN (Zhong et al., 2019)	CVPR 2019	43.0	75.1	87.6	91.6	40.4	63.3	75.8	80.4
PDA-Net (Li et al., 2019)	ICCV 2019	47.6	75.2	86.3	90.2	45.1	63.2	77.0	82.5
PCB-PAST (Zhang et al., 2019a)	ICCV 2019	54.6	78.4	—	—	54.3	72.4	—	—
SSG (Fu et al., 2019)	ICCV 2019	58.3	80.0	90.0	92.4	53.4	73.0	80.6	83.2
MMCL (Wang & Zhang, 2020)	CVPR 2020	60.4	84.4	92.8	95.0	51.4	72.4	82.9	85.0
ECN-GPP (Zhong et al., 2020b)	TPAMI 2020	63.8	84.1	92.8	95.4	54.4	74.0	83.7	87.4
JVTC+ (Li & Zhang, 2020)	ECCV 2020	67.2	86.8	95.2	97.1	66.5	80.4	89.9	92.2
AD-Cluster (Zhai et al., 2020a)	CVPR 2020	68.3	86.7	94.4	96.5	54.1	72.6	82.5	85.5
MMT (Ge et al., 2020a)	ICLR 2020	71.2	87.7	94.9	96.9	65.1	78.0	88.8	92.5
CAIL (Luo et al., 2020)	ECCV 2020	71.5	88.1	94.4	96.2	65.2	79.5	88.3	91.4
NRMT (Zhao et al., 2020)	ECCV 2020	71.7	87.8	94.6	96.5	62.2	77.8	86.9	89.5
MEB-Net (Zhai et al., 2020b)	ECCV 2020	76.0	89.9	96.0	97.5	66.1	79.6	88.3	92.2
SpCL (Ge et al., 2020b)	NeurIPS 2020	76.7	90.3	96.2	97.7	68.8	82.9	90.1	92.5
Dual-Refinement (Dai et al., 2021b)	TIP 2021	78.0	90.9	96.4	97.7	67.7	82.1	90.1	92.5
UNRN (Zheng et al., 2021a)	AAAI 2021	78.1	91.9	96.1	97.8	69.1	82.0	90.7	93.5
GLT (Zheng et al., 2021b)	CVPR 2021	79.5	92.2	96.5	97.8	69.2	82.0	90.2	92.8
HCD (Zheng et al., 2021c)	ICCV 2021	80.0	91.5	—	—	70.1	82.2	—	—

Table 2 continued

Methods	Reference	DukeMTMC-ReID → Market-1501				Market-1501 → DukeMTMC-ReID			
		mAP	R1	R5	R10	mAP	R1	R5	R10
CCL (Isobe et al., 2021)	ICCV 2021	<u>83.4</u>	<b>94.2</b>	–	–	<u>70.8</u>	83.5	–	–
Our IDM (Dai et al., 2021c)	ICCV 2021	82.8	<u>93.2</u>	<u>97.5</u>	<u>98.1</u>	70.5	<u>83.6</u>	<u>91.5</u>	<u>93.7</u>
Our IDM++	This paper	<b>85.3</b>	<b>94.2</b>	<b>97.7</b>	<b>98.5</b>	<b>73.2</b>	<b>85.5</b>	<b>92.1</b>	<b>94.3</b>
Methods	Reference	MSMT17 → Market-1501				MSMT17 → DukeMTMC-reID			
		mAP	R1	R5	R10	mAP	R1	R5	R10
CASCL (Wu et al., 2019)	ICCV 2019	35.5	65.4	80.6	86.2	37.8	59.3	73.2	77.8
MAR (Yu et al., 2019)	CVPR 2019	40.0	67.7	81.9	87.3	48.0	67.1	79.8	84.2
PAUL (Yang et al., 2019)	CVPR 2019	40.1	68.5	82.4	87.4	53.2	72.0	82.7	86.0
DG-Net++ (Zou et al., 2020)	ECCV 2020	64.6	83.1	91.5	94.3	58.2	75.2	73.6	86.9
D-MMD (Mekhazni et al., 2020)	ECCV 2020	50.8	72.8	88.1	92.3	51.6	68.8	82.6	87.1
MMT-dbscan (Ge et al., 2020a)	ICLR 2020	75.6	89.3	95.8	97.5	63.3	77.4	88.4	91.7
SpCL (Ge et al., 2020b)	NeurIPS 2020	77.5	89.7	96.1	97.6	69.3	82.9	91.0	93.0
HCD (Zheng et al., 2021c)	ICCV 2021	80.2	91.4	–	–	71.2	83.1	–	–
Our IDM (Dai et al., 2021c)	ICCV 2021	<u>82.1</u>	<u>92.4</u>	<u>97.5</u>	<u>98.4</u>	<u>71.9</u>	<u>83.6</u>	<u>91.5</u>	<u>93.4</u>
Our IDM++	This paper	<b>85.2</b>	<b>94.1</b>	<b>97.7</b>	<b>98.5</b>	<b>73.6</b>	<b>84.6</b>	<b>92.2</b>	<b>94.1</b>

Table 2 continued

Methods	Reference	MSMT17 → Market-1501					MSMT17 → DukeMTMC-reID				
		mAP	R1	R5	R10	R10	mAP	R1	R5	R10	R10
		Market-1501 → MSMT17					DukeMTMC-reID → MSMT17				
	Reference	mAP	R1	R5	R10	R10	mAP	R1	R5	R10	R10
ECN (Zhong et al., 2019)	CVPR 2019	8.5	25.3	36.3	42.1	46.8	10.2	30.2	41.5	46.8	
SSG (Fu et al., 2019)	ICCV 2019	13.2	31.6	–	49.6	51.2	13.3	32.2	–	51.2	
ECN-GPP (Zhong et al., 2020b)	TPAMI 2020	15.2	40.4	53.1	58.7	61.5	16.0	42.5	55.9	61.5	
MMCL (Wang & Zhang, 2020)	CVPR 2020	15.1	40.8	51.8	56.7	58.9	16.2	43.6	54.3	58.9	
NRMT (Zhao et al., 2020)	ECCV 2020	19.8	43.7	56.5	62.2	63.3	20.6	45.2	57.8	63.3	
CAIL (Luo et al., 2020)	ECCV 2020	20.4	43.7	56.1	61.9	68.9	24.3	51.7	64.0	68.9	
MMT (Ge et al., 2020a)	ICLR 2020	22.9	49.2	63.1	68.8	69.8	23.3	50.1	63.9	69.8	
JVTC+ (Li & Zhang, 2020)	ECCV 2020	25.1	48.6	65.3	68.2	75.9	27.5	52.9	70.5	75.9	
SpCL (Ge et al., 2020b)	NeurIPS 2020	26.8	53.7	65.0	69.8	70.5	26.5	53.1	65.8	70.5	
Dual-Refinement (Dai et al., 2021b)	TIP 2021	25.1	53.3	66.1	71.5	73.2	26.9	55.0	68.4	73.2	
UNRN (Zheng et al., 2021a)	AAAI 2021	25.3	52.4	64.7	69.7	70.6	26.2	54.9	67.3	70.6	
GLT (Zheng et al., 2021b)	CVPR 2021	26.5	56.6	67.5	72.0	74.2	27.7	59.5	70.1	74.2	
HCD (Zheng et al., 2021c)	ICCV 2021	28.4	54.9	–	–	–	29.3	56.1	–	–	
CCL (Isobe et al., 2021)	ICCV 2021	<u>35.8</u>	<u>65.8</u>	–	–	–	<u>36.3</u>	<u>66.6</u>	–	–	
Our IDM (Dai et al., 2021c)	ICCV 2021	33.5	61.3	<u>73.9</u>	78.4	<u>80.2</u>	35.4	63.6	<u>75.5</u>	<u>80.2</u>	
Our IDM++	This paper	<b>40.2</b>	<b>69.9</b>	<b>80.2</b>	<b>83.7</b>	<b>84.0</b>	<b>40.5</b>	<b>69.5</b>	<b>80.3</b>	<b>84.0</b>	

The best results are highlighted with bold and the second best results are marked with underline



**Table 3** Comparison with the state-of-the-art UDA person re-ID methods on synthetic → real tasks

Methods	Reference	PersonX → Market-1501			PersonX → DukeMTMC-reID			PersonX → MSMT17			
		mAP	R1	R5	mAP	R1	R5	mAP	R1	R5	
MMT (Ge et al., 2020a)	ICLR 2020	71.0	86.5	94.8	60.1	74.3	86.5	90.5	39.1	52.6	58.5
SpCL (Ge et al., 2020b)	NeurIPS 2020	73.8	88.0	95.3	67.2	81.8	90.2	92.6	47.7	60.0	65.5
CCL (Isobe et al., 2021)	ICCV 2021	79.6	<u>92.5</u>	-	-	-	-	-	53.2	-	-
Our IDM (Dai et al., 2021c)	ICCV 2021	81.3	92.0	<u>97.4</u>	<u>68.5</u>	82.6	<u>91.2</u>	<u>93.4</u>	<u>58.4</u>	<u>70.7</u>	<u>75.5</u>
Our IDM++	This paper	<b>84.3</b>	<b>93.7</b>	<b>97.5</b>	<b>71.5</b>	<b>84.1</b>	<b>91.9</b>	<b>94.0</b>	<b>68.3</b>	<b>79.1</b>	<b>82.8</b>
Methods	Reference	Unreal → Market-1501			Unreal → DukeMTMC-reID			Unreal → MSMT17			
JVTC (Li & Zhang, 2020)	ECCV 2020	mAP	R1	R5	mAP	R1	R5	R10	mAP	R1	R5
Our IDM (Dai et al., 2021c)	ICCV 2021	78.3	90.8	-	66.1	81.2	-	-	25.0	53.7	-
Our IDM++	This paper	<b>85.2</b>	<b>93.6</b>	<b>97.5</b>	<b>73.2</b>	<b>85.0</b>	<b>92.4</b>	<b>94.6</b>	<b>41.5</b>	<b>71.0</b>	<b>81.2</b>
		83.2	<u>92.8</u>	<u>97.3</u>	<u>72.4</u>	<u>84.6</u>	<u>92.0</u>	<u>94.0</u>	<u>38.3</u>	<u>67.3</u>	<u>78.4</u>
		<b>98.2</b>	<b>98.2</b>	<b>98.3</b>	<b>94.6</b>	<b>94.6</b>	<b>94.6</b>	<b>94.6</b>	<b>81.2</b>	<b>81.2</b>	<b>84.7</b>

The best results are highlighted with bold and the second best results are marked with underline

**Table 4** Comparison with the state-of-the-art DG person re-ID methods under the protocols of DG-partial and DG-full

DG-partial Method	Source	IDs	Images	Target: Market-1501 mAP	R1	Source	IDs	Images	Target: DukeMTMC mAP	R1
QACov <sub>50</sub> (Liao & Shao, 2020)	D+C3+MT	2510	56,508	39.5	68.6	M+C3+MT	2559	52,922	43.4	64.9
M <sup>3</sup> L (Zhao et al., 2021b)				<u>51.1</u>	<u>76.5</u>				<u>48.2</u>	<u>67.1</u>
BoT (Luo et al., 2019)				46.3	75.1				46.8	65.3
Our IDM (Dai et al., 2021c)				47.8	75.8				48.1	66.3
Our IDM++				<b>52.6</b>	<b>79.2</b>				<b>52.7</b>	<b>70.9</b>
Method	Source	IDs	Images	Target: CUHK03 mAP	R1	Source	IDs	Images	Target: MSMT17 mAP	R1
QACov <sub>50</sub> (Liao & Shao, 2020)	M+D+MT	2494	62,079	19.2	22.9	M+D+C3	2220	36,823	10.0	29.9
M <sup>3</sup> L (Zhao et al., 2021b)				<u>30.9</u>	<u>31.9</u>				13.1	32.0
BoT (Luo et al., 2019)				26.0	24.9				12.6	31.7
Our IDM (Dai et al., 2021c)				27.8	27.7				<u>14.8</u>	<u>36.5</u>
Our IDM++				<b>32.7</b>	<b>33.6</b>				<b>19.5</b>	<b>44.3</b>
DG-full										
Method	Source	IDs	Images	Target: Market-1501 mAP	R1	Source	IDs	Images	Target: DukeMTMC mAP	R1
RaMoE (Dai et al., 2021a)	D+C3+MT	7380	176,949	<u>56.5</u>	<u>82.0</u>	M+C3+MT	7069	169,957	<u>56.9</u>	<u>73.6</u>
BoT (Luo et al., 2019)				50.2	77.0				52.6	68.4
Our IDM (Dai et al., 2021c)				52.2	77.3				53.1	70.0
Our IDM++				<b>58.8</b>	<b>82.2</b>				<b>58.3</b>	<b>74.5</b>
Method	Source	IDs	Images	Target: CUHK03 mAP	R1	Source	IDs	Images	Target: MSMT17 mAP	R1
RaMoE (Dai et al., 2021a)	M+D+MT	7414	192,271	35.5	<u>36.6</u>	M+D+C3	4780	79,927	13.5	34.1
BoT (Luo et al., 2019)				30.5	29.9				12.0	29.9
Our IDM (Dai et al., 2021c)				<u>35.7</u>	36.0				<u>15.3</u>	<u>36.7</u>
Our IDM++				<b>39.9</b>	<b>40.9</b>				<b>20.1</b>	<b>44.3</b>

The best results are highlighted with bold and the second best results are marked with underline

margin on all these benchmarks. Especially when taking the most challenging MSMT17 as the target domain, our method outperforms the SOTA method CCL (Isobe et al., 2021) by 4.4% mAP on Market-1501  $\rightarrow$  MSMT17 and 4.2% mAP on DukeMTMC-reID  $\rightarrow$  MSMT17. Adding MGM can significantly improve IDM, e.g., the performance gain is up to 6.7% mAP and 8.6% Rank-1 on Market-1501  $\rightarrow$  MSMT17.

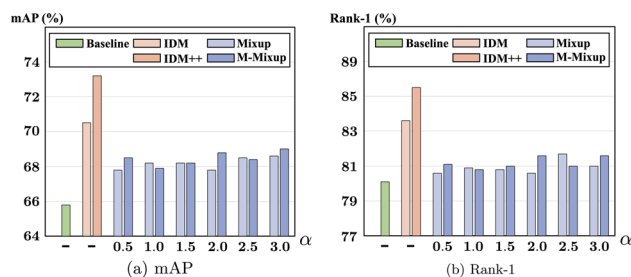
#### Comparisons on synthetic $\rightarrow$ real UDA re-ID tasks.

Compared with the real  $\rightarrow$  real UDA re-ID tasks, the synthetic  $\rightarrow$  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 3, our method can outperform the SOTA methods by a large margin where mAP of our method is higher than CCL (Isobe et al., 2021) by 4.7% and 8.8% when transferring from PersonX to Market-1501 and MSMT17 respectively. All these significant performance gains in Table 3 have shown the superiority of our method that can better improve the performance on UDA re-ID by extending IDM to IDM++.

**Comparisons on DG Re-ID Tasks** In DG re-ID tasks, we utilize BoT (Luo et al., 2019) as the baseline method and it is competitive to the existing methods (Liao & Shao, 2020; Zhao et al., 2021a; Dai et al., 2021a). Based on this baseline method, we implement our method IDM++ and our previous work IDM. As shown in Table 4, we compare our method with the SOTAs under two common protocols: i.e., DG-partial (Zhao et al., 2021a) and DG-full (Dai et al., 2021a). Under the “DG-partial” protocol, mAP of our IDM++ outperforms the SOTA method M<sup>3</sup>L (Zhao et al., 2021a) by 1.5%, 4.5%, 1.8%, and 6.4% when testing on Market-1501, DukeMTMC, CUHK03, and MSMT17 respectively. Under the “DG-full” protocol, our method outperforms the SOTA method RaMoE (Dai et al., 2021a) significantly as well. Compared with our previous work (IDM (Dai et al., 2021c)), the performance gain of our IDM++ is significant. Taking the protocol of “DG-ful” as an example, mAP of IDM++ outperforms IDM by 6.6%, 5.2%, 4.2%, and 4.8% when testing on Market-1501, DukeMTMC, CUHK03, and MSMT17 respectively. In DG re-ID, the performance gain mainly comes from our MGM. The IDM method is superior to the baseline BoT because mixing source and target domains to generate intermediate domains can be seen as only considering inter-domain knowledge. By reinforcing IDM into IDM++ by adding MGM, features are learned to be more identity-discriminative and more robust to cross-domain variations.

#### 4.6 Investigation on Key Designs of IDM++

**Predicting Mixing Ratios by IDM is More Essential than Random** We compare our method with the traditional Mixup (Zhang et al., 2018) and Manifold Mixup (M-Mixup) (Verma



**Fig. 5** Comparison with traditional mixup methods (Mixup (Zhang et al., 2018), M-Mixup (Verma et al., 2019)) on Market $\rightarrow$ Duke. The interpolation ratio in these mixup methods is randomly sampled from a beta distribution Beta( $\alpha$ ,  $\alpha$ )

et al., 2019) methods in Fig. 5. 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 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. Unlike them, our method can adaptively generate the mix ratios to mix up the source and target domains representations to generate intermediate domain representations.

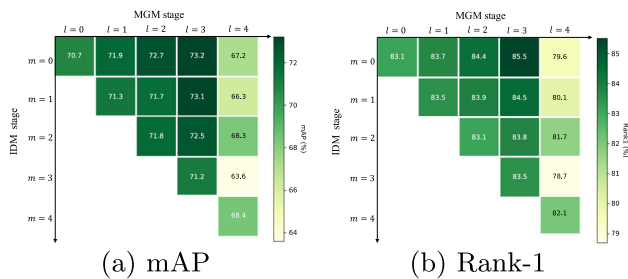
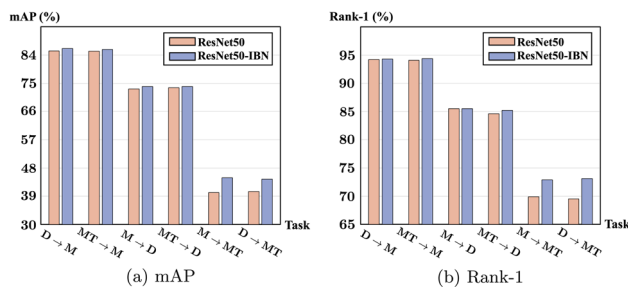
We attribute the reasons why the predictive mixing ratio is better than random as follows:

- *Adaptivity* This adaptivity of predicting mixing ratio is crucial because the optimal mix ratio that minimizes the domain gap may not be the same for all samples or during different training stages. In contrast, random sampling from a beta distribution does not take into account the data distribution or the training dynamics, potentially leading to a less effective domain adaptation process.
- *Optimization* By making the mix ratio learnable, we enable the model to optimize it as part of the training objective. This optimization process is guided by the overall goal of domain alignment and identity preservation, which is directly relevant to the task at hand. Randomly sampling mix ratios, on the other hand, treats the ratio as a hyperparameter that is not influenced by the learning process, potentially missing out on the benefits of end-to-end learning.
- *Fine-grained control* The learnable mix ratio provides fine-grained control over the generation of intermediate domains. This control is essential for capturing the nuances of domain-specific information and aligning the intermediate domains more effectively with both the source and target domains. Random sampling lacks this level of control and may not explore the full potential of intermediate representations.

**Table 5** Comparisons on different mirrors generation mechanisms

Mirrors generation mechanism	Market → Duke	
	mAP	R1
IDM+MGM w/ mirrors-I	73.2	85.5
IDM+MGM w/ mirrors-ST	72.5	83.8
MGM w/ mirrors-ST	68.0	81.7

The mirrors-I means using MGM to map source and target samples into intermediate domains. The mirrors-ST means mapping source and target samples into each other's domain

**Fig. 6** Evaluation on which stage (stage- $m$ ) to plug our IDM and which stage (stage- $l$ ) to plug our MGM. We provide the mAP and Rank-1 on Market→Duke**Fig. 7** Comparisons between different backbones: ResNet50 and ResNet50-IBN. We denote Market-1501, DukeMTMC-reID, and MSMT17 as M, D, and MT respectively

**The Mirrors Generation Mechanism in MGM** We compare with different mirrors generation mechanisms in Table 5. In our IDM++, we use the mechanism “IDM+MGM w/ mirrors-I” as shown in Sect. 3.4. If we use the MGM to map source and target features into each other's domain (“mirrors-ST”), the inputs of the MGM (as shown in Fig. 2) are only the source and target hidden representations (i.e.,  $G^S$  and  $G^T$ ). Thus, we can obtain the source mirrors  $G^{S \rightarrow T}$  and the target mirrors  $G^{T \rightarrow S}$  that are stylized by the target and source domains respectively. In Table 5, the mechanism “MGM w/ mirrors-ST” means directly aligning between the source and target domains, which is inferior to aligning by intermediate domains. The mechanism “IDM+MGM w/ mirrors-ST” means utilizing the IDM to generate intermediate domains but mapping source/target identities into each other's domain to obtain mirrors. From Table 5, we can see the significance of our proposed mirrors generation mechanism.

**Plugging IDM and MGM at Which Stage** Our IDM and MGM are two plug-and-play modules which can be seamlessly plugged into 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 and Max Pooling layer and stage-1/2/3/4 correspond to the other four convolutional blocks. We plug our IDM at the  $m$ -th stage and plug our MGM at the  $l$ -th stage to study how different combinations of both kinds of stages will affect the performance of our method. The stage- $l$  of MGM should not be smaller than the stage- $m$  of IDM because IDM is a prerequisite for MGM (i.e., MGM needs to map source/target identities into intermediate domains generated by IDM). As shown in Fig. 6, we provide the visualization of the target domain's mAP/Rank-1 on different combinations of the stage- $m$  and the stage- $l$  when transferring from Market-1501 to DukeMTMC. The Fig. 6 shows that the performance is gradually declining with the deepening of the network (i.e.,  $m$  ranges from 0 to 4). 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 (Long et al., 2018a). When the stage- $l$  of MGM ranges from 0 to 3, the performance becomes better, which shows that our proposed MGM can complement with the IDM in deeper layers that have weak transferable ability. We observe that when  $l$  is 4, the model will be degenerated because features after stage-4 will be directly conducted with global average pooling to learn with the triplet loss. Compared with shallow layers' features, the little manipulation on stage-4 features will make a significant effect on the deep embedding learning. Based on the above analysis, we plug the IDM at the stage-0 and plug the MGM at the stage-3 in all our experiments if not specified.

**Scalability for Different Backbones** Our proposed IDM++ can be easily extended to other backbones which are equipped with different normalization layers (Pan et al., 2018; Li et al., 2016; Zhuang et al., 2020). We take the commonly used IBN (Pan et al., 2018) as an example and use the ResNet50-IBN (Pan et al., 2018) as a stronger backbone. As shown in Fig. 7, our method can achieve the consistent performance gains with the stronger backbone and the performance gain is especially obvious on the largest and most challenging MSMT dataset.

## 4.7 Analysis and Discussions

**Analysis on Pseudo Labels Quality** Our method can also be included into the pseudo-label-based UDA re-ID methods (Dai et al., 2021b; Ge et al., 2020b) where the quality of pseudo labels is important during training, because target training data is supervised by pseudo labels. Following the existing methods (Dai et al., 2021b; Yang et al., 2020;




Fig. 8 Analysis on the F-score, NMI score, and Rank-1 score when training on Market-1501 and testing on MSMT17

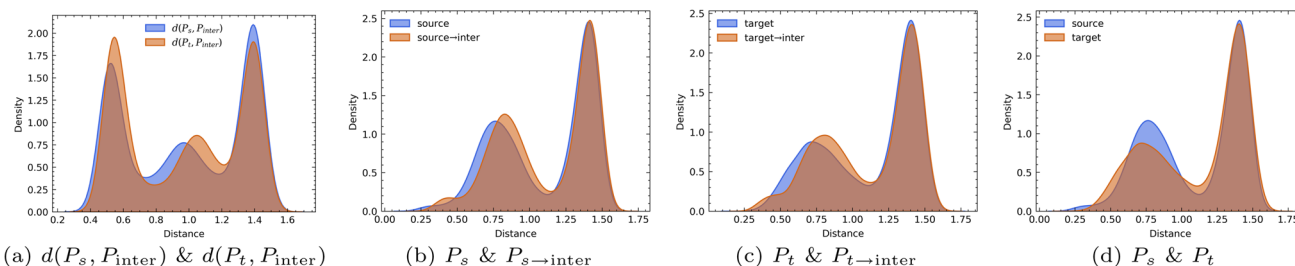


Fig. 9 Distribution of samples' pair-wise distance (Euclidean distance of the deep embedding) on the target dataset DukeMTMC-reID when taking Market-1501 as the source dataset. **a** Distance between source and intermediate domains  $d(P_s, P_{inter})$ , and distance between target and intermediate domains  $d(P_t, P_{inter})$ . **b** Distribution of source samples  $P_s$  and their mirrors  $P_{s \rightarrow inter}$  in intermediate domains. **c** Distribution of target samples  $P_t$  and their mirrors  $P_{t \rightarrow inter}$  in intermediate domains. **d** Distribution of source and target samples

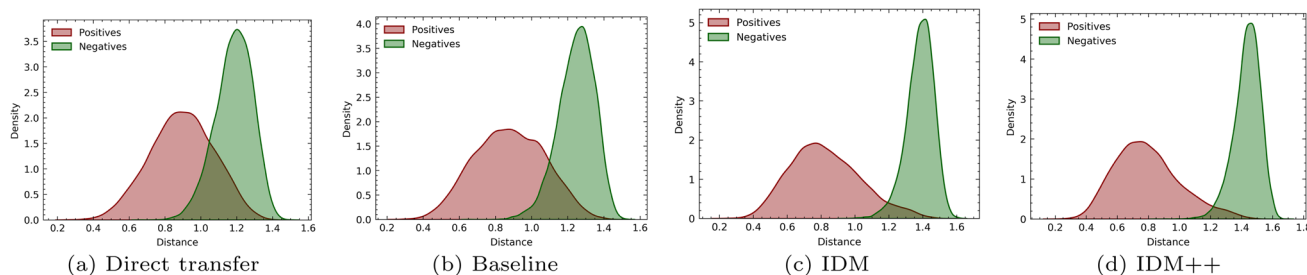


Fig. 10 Target domain's distribution of different methods when transferring from Market-1501 to DukeMTMC-reID

Ge et al., 2020b), we use both the BCubed F-score (Amigó et al., 2009) and the Normalized Mutual Information (NMI) score to evaluate the quality of pseudo labels. For F-score and NMI score, the score towards 1.0 implies better clustering quality and less label noise. As shown in Fig. 8, we visualize the F-score, NMI score, and Rank-1 score when training on Market-1501 and testing on MSMT17. It shows that our IDM++ can outperform our IDM and Baseline method by a large margin no matter the quality of pseudo labels or on the Rank-1 performance.

**Analysis on the bridge and alignment of domains' distribution.** In Fig. 9, we provide the visualization on the bridge between intermediate domains and the source/target domain, and the distribution alignment between domains

when transferring from the source domain Market-1501 to the target domain DukeMTMC-reID. In Fig. 9 (a), the distance distribution  $d(P_s, P_{inter})$  between the source and intermediate domains is almost aligned with  $d(P_t, P_{inter})$  between the target and intermediate domains. The little discrepancy between  $d(P_s, P_{inter})$  and  $d(P_t, P_{inter})$  shows the distribution of learned intermediate domains is diverse enough to bridge source and target domains. If the learned intermediate domains are not diverse or are dominated by either domain, the discrepancy of the distribution between  $d(P_s, P_{inter})$  and  $d(P_t, P_{inter})$  will be large. To evaluate the effectiveness of our proposed cross-domain consistency loss that enforces the consistency between source/target samples and their mirrors generated by the MGM, we randomly

sample 20,000 sample pairs in each domain and calculate the Euclidean distances of their L2-normalized features to visualize the distributions in Fig. 9b, c. As shown in Fig. 9b, the distribution of source samples and the distribution of their mirrors in intermediate domains can be well aligned. For target samples and their mirrors, their distributions are also well aligned as shown in Fig. 9c. It shows that our proposed MGM can well preserve the original source/target identity-discriminative distribution in intermediate domains. By aligning source and target domains to intermediate domains respectively, distributions of source and target domains can be well aligned as shown in Fig. 9d.

**Analysis on the discriminability of the target domain's features.** In Fig. 10, we visualize on the distribution of target domain's positives and negatives and compare among four methods including Direct transfer, Baseline, our IDM, and our IDM++. Specifically, we randomly sample 10,000 positive pairs and 10,000 negative pairs in the test-set of the target domain when transferring from Market-1501 to DukeMTMC-reID. When the overlap between positives' and negatives' distributions is smaller, the model will be more discriminative. As shown in Fig. 10, when training with our IDM or IDM++, the discriminability gain is significant than the Direct transfer or the Baseline method.

## 5 Conclusion

This paper provides a new perspective for cross-domain re-ID: the source and target domain are not isolated, but are connected through a series of intermediate domains. This perspective motivates us to align the source/target domain against these shared intermediate domains, instead of the popular source-to-target alignment. To this end, we propose an IDM and an MGM module and integrate them into IDM++. IDM first discovers the intermediate domains by mixing the source-domain and target-domain features within the deep network and then minimizes the source-to-intermediate and target-to-intermediate discrepancy. MGM further reinforces the domain alignment by solving a side effect of IDM. Specifically, the intermediate domains generated by IDM lack the original identities in the source/target domain, while MGM compensates for this lack by mapping the original identities into the IDM-generated intermediate domains. Both IDM and MGM can be seamlessly plugged into the backbone network and facilitate mutual benefits. Experimental results show that IDM significantly improves UDA baseline and MGM brings another round of substantial improvement. Integrating IDM and MGM, IDM++ achieves the new state of the art on both UDA and DG scenarios for cross-domain re-ID.

**Acknowledgements** This work was supported by the National Natural Science Foundation of China under Grant 62088102, in part by the PKU-NTU Joint Research Institute (JRI) sponsored by a donation from the Ng Teng Fong Charitable Foundation and in part by AI Joint Lab of Future Urban Infrastructure sponsored by Fuzhou Chengtuo New Infrastructure Group and Boyun Vision Co. Ltd.

**Data availability statement** The datasets generated and analysed during the current study are available from the corresponding author on reasonable request.

## References

- Amigó, E., Gonzalo, J., Artiles, J., & Verdejo, F. (2009). A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information Retrieval*, 12(4), 461–486.
- Chang, W.-G., You, T., Seo, S., Kwak, S., & Han, B. (2019). Domain-specific batch normalization for unsupervised domain adaptation. In *Proceedings of the CVPR* (pp. 7354–7362).
- Chen, Y., Hu, V. T., Gavves, E., Mensink, T., Mettes, P., Yang, P., & Snoek, C. G. (2020a). Pointmixup: Augmentation for point clouds. In *Proceedings of the ECCV*.
- Chen, G., Lu, Y., Lu, J., & Zhou, J. (2020b). Deep credible metric learning for unsupervised domain adaptation person re-identification. In *Proceedings of the ECCV* (pp. 643–659).
- Choi, S., Kim, T., Jeong, M., Park, H., & Kim, C. (2021). Meta batch-instance normalization for generalizable person re-identification. In *Proceedings of the CVPR* (pp. 3425–3435).
- Cui, S., Wang, S., Zhuo, J., Su, C., Huang, Q., & Tian, Q. (2020). Gradually vanishing bridge for adversarial domain adaptation. In *Proceedings of the CVPR* (pp. 12455–12464).
- Cui, Z., Li, W., Xu, D., Shan, S., Chen, X., & Li, X. (2014). Flowing on Riemannian manifold: Domain adaptation by shifting covariance. *IEEE Transactions on Cybernetics*, 44(12), 2264–2273.
- Dai, Y., Li, X., Liu, J., Tong, Z., & Duan, L.-Y. (2021a). Generalizable person re-identification with relevance-aware mixture of experts. In *Proceedings of the CVPR* (pp. 16145–16154).
- Dai, Y., Liu, J., Bai, Y., Tong, Z., & Duan, L.-Y. (2021b). Dual-refinement: Joint label and feature refinement for unsupervised domain adaptive person re-identification. In *Proceedings of the IEEE TIP*.
- Dai, Y., Liu, J., Sun, Y., Tong, Z., Zhang, C., & Duan, L.-Y. (2021c). IDM: An intermediate domain module for domain adaptive person re-ID. In *Proceedings of the ICCV* (pp. 11864–11874).
- Deng, W., Zheng, L., Ye, Q., Kang, G., Yang, Y., & Jiao, J. (2018). Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *Proceedings of the CVPR* (pp. 994–1003).
- Ding, Y., Fan, H., Xu, M., & Yang, Y. (2020). Adaptive exploration for unsupervised person re-identification. *ACM TOMM*, 16(1), 1–19.
- Du, D., Chen, J., Li, Y., Ma, K., Wu, G., Zheng, Y., & Wang, L. (2022). Cross-domain gated learning for domain generalization. *International Journal of Computer Vision*, 130(11), 2842–2857.
- Ester, M., Krieger, H.-P., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the KDD* (Vol. 96, pp. 226–231).
- Fan, H., Zheng, L., Yan, C., & Yang, Y. (2018). Unsupervised person re-identification: Clustering and fine-tuning. *ACM TOMM*, 14(4), 1–18.
- Fu, Y., Wei, Y., Wang, G., Zhou, Y., Shi, H., & Huang, T. S. (2019). Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *Proceedings of the ICCV* (pp. 6112–6121).

- Ganin, Y., & Lempitsky, V. (2015). Unsupervised domain adaptation by backpropagation. In *Proceedings of the ICML* (pp. 1180–1189).
- Ge, Y., Chen, D., & Li, H. (2020a). Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *Proceedings of the ICLR*.
- Ge, Y., Chen, D., Zhu, F., Zhao, R., & Li, H. (2020b). Self-paced contrastive learning with hybrid memory for domain adaptive object re-ID. In *Proceedings of the NeurIPS*.
- Gong, B., Shi, Y., Sha, F., & Grauman, K. (2012). Geodesic flow kernel for unsupervised domain adaptation. In *Proceedings of the CVPR* (pp. 2066–2073).
- Gong, R., Li, W., Chen, Y., & Gool, L. V. (2019). DLOW: Domain flow for adaptation and generalization. In *Proceedings of the CVPR* (pp. 2477–2486).
- Gopalan, R., Li, R., & Chellappa, R. (2013). Unsupervised adaptation across domain shifts by generating intermediate data representations. *IEEE TPAMI*, 36(11), 2288–2302.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., kavukcuoglu, K., Munos, R., & Valko, M. (2020). Bootstrap your own latent—A new approach to self-supervised learning. In *Proceedings of the NeurIPS* (Vol. 33, pp. 21271–21284).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the CVPR* (pp. 770–778).
- Hinton, G., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network. arXiv preprint [arXiv:1503.02531](https://arxiv.org/abs/1503.02531)
- Huang, X., & Belongie, S. (2017). Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the ICCV* (pp. 1501–1510).
- Huang, Y., Wu, Q., Xu, J., & Zhong, Y. (2019). SBSGAN: Suppression of inter-domain background shift for person re-identification. In *Proceedings of the ICCV* (pp. 9527–9536).
- Huang, Y., Wu, Q., Xu, J., Zhong, Y., & Zhang, Z. (2021). Unsupervised domain adaptation with background shift mitigating for person re-identification. *International Journal of Computer Vision*, 129(7), 2244–2263.
- Huang, Z., Liu, J., Li, L., Zheng, K., & Zha, Z.-J. (2022). Modality-adaptive mixup and invariant decomposition for RGB-infrared person re-identification. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, pp. 1034–1042).
- Isobe, T., Li, D., Tian, L., Chen, W., Shan, Y., & Wang, S. (2021). Towards discriminative representation learning for unsupervised person re-identification. In *Proceedings of the ICCV* (pp. 8526–8536).
- Jia, J., Ruan, Q., & Hospedales, T. M. (2019). Frustratingly easy person re-identification: Generalizing person re-ID in practice. In *Proceedings of the BMVC*.
- Jin, X., Lan, C., Zeng, W., & Chen, Z. (2020a). Global distance-distributions separation for unsupervised person re-identification. In *Proceedings of the ECCV* (pp. 735–751).
- Jin, X., Lan, C., Zeng, W., Chen, Z., & Zhang, L. (2020b). Style normalization and restitution for generalizable person re-identification. In *Proceedings of the CVPR* (pp. 3143–3152).
- Kang, G., Zheng, L., Yan, Y., & Yang, Y. (2018). Deep adversarial attention alignment for unsupervised domain adaptation: The benefit of target expectation maximization. In *Proceedings of the ECCV* (pp. 401–416).
- Leng, Q., Ye, M., & Tian, Q. (2019). A survey of open-world person re-identification. *IEEE TCSVT*, 30(4), 1092–1108.
- Li, J., & Zhang, S. (2020). Joint visual and temporal consistency for unsupervised domain adaptive person re-identification. In *Proceedings of the ECCV* (pp. 483–499).
- Li, D., Wei, X., Hong, X., & Gong, Y. (2020). Infrared-visible cross-modal person re-identification with an x modality. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 34, pp. 4610–4617).
- Li, W., Zhao, R., Xiao, T., & Wang, X. (2014). DeepReID: Deep filter pairing neural network for person re-identification. In *Proceedings of the CVPR* (pp. 152–159).
- Li, Y., Wang, N., Shi, J., Liu, J., & Hou, X. (2016). Revisiting batch normalization for practical domain adaptation. arXiv preprint [arXiv:1603.04779](https://arxiv.org/abs/1603.04779)
- Li, Y., Zhang, D., Keuper, M., & Khoreva, A. (2024). Intra- & extra-source exemplar-based style synthesis for improved domain generalization. *International Journal of Computer Vision*, 132(2), 446–465.
- Li, Y.-J., Lin, C.-S., Lin, Y.-B., & Wang, Y.-C. F. (2019). Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation. In *Proceedings of the ICCV* (pp. 7919–7929).
- Liao, S., & Shao, L. (2020). Interpretable and generalizable person re-identification with query-adaptive convolution and temporal lifting. In *Proceedings of the ECCV* (pp. 456–474).
- Lin, Y., Xie, L., Wu, Y., Yan, C., & Tian, Q. (2020). Unsupervised person re-identification via softened similarity learning. In *Proceedings of the CVPR* (pp. 3390–3399).
- Long, M., Cao, Y., Wang, J., & Jordan, M. (2015). Learning transferable features with deep adaptation networks. In *Proceedings of the ICML* (pp. 97–105).
- Long, M., Cao, Y., Cao, Z., Wang, J., & Jordan, M. I. (2018a). Transferable representation learning with deep adaptation networks. *IEEE TPAMI*, 41(12), 3071–3085.
- Long, M., Cao, Z., Wang, J., & Jordan, M. I. (2018b). Conditional adversarial domain adaptation. In *Proceedings of the NeurIPS*.
- Long, M., Zhu, H., Wang, J., & Jordan, M. I. (2017). Deep transfer learning with joint adaptation networks. In *Proceedings of the ICML* (pp. 2208–2217).
- Luo, C., Song, C., & Zhang, Z. (2020). Generalizing person re-identification by camera-aware invariance learning and cross-domain mixup. In *Proceedings of the ECCV* (pp. 2, 7).
- Luo, H., Jiang, W., Gu, Y., Liu, F., Liao, X., Lai, S., & Gu, J. (2019). A strong baseline and batch normalization neck for deep person re-identification. *IEEE TMM*, 22(10), 2597–2609.
- Mancini, M., Akata, Z., Ricci, E., & Caputo, B. (2020). Towards recognizing unseen categories in unseen domains. In *Proceedings of the ECCV* (pp. 466–483). Springer: Berlin.
- Mekhazni, D., Bhuiyan, A., Ekladios, G., & Granger, E. (2020). Unsupervised domain adaptation in the dissimilarity space for person re-identification. In *Proceedings of the ECCV* (pp. 159–174).
- Na, J., Jung, H., Chang, H. J., & Hwang, W. (2021). FixBi: Bridging domain spaces for unsupervised domain adaptation. In *Proceedings of the CVPR* (pp. 1094–1103).
- Ni, H., Song, J., Luo, X., Zheng, F., Li, W., & Shen, H. T. (2022). Meta distribution alignment for generalizable person re-identification. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2487–2496).
- Nuriel, O., Benaim, S., & Wolf, L. (2021). Permuted ADAIN: Reducing the bias towards global statistics in image classification. In *Proceedings of the CVPR* (pp. 9482–9491).
- Pan, X., Luo, P., Shi, J., & Tang, X. (2018). Two at once: Enhancing learning and generalization capacities via IBN-net. In *Proceedings of the ECCV* (pp. 464–479).
- Pu, N., Zhong, Z., Sebe, N., & Lew, M. S. (2023). A memorizing and generalizing framework for lifelong person re-identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11), 13567–13585.
- Qiao, F., Zhao, L., & Peng, X. (2020). Learning to learn single domain generalization. In *Proceedings of the CVPR* (pp. 12556–12565).
- Ristani, E., Solera, F., Zou, R., Cucchiara, R., & Tomasi, C. (2016). Performance measures and a data set for multi-target, multi-camera tracking. In *Proceedings of the ECCV* (pp. 17–35).

- Russo, P., Carlucci, F. M., Tommasi, T., & Caputo, B. (2018). From source to target and back: Symmetric bi-directional adaptive GAN. In *Proceedings of the CVPR* (pp. 8099–8108).
- Shen, Z., Liu, J., He, Y., Zhang, X., Xu, R., Yu, H., & Cui, P. (2021). Towards out-of-distribution generalization: A survey. arXiv preprint [arXiv:2108.13624](https://arxiv.org/abs/2108.13624)
- Song, J., Yang, Y., Song, Y.-Z., Xiang, T., & Hospedales, T. M. (2019). Generalizable person re-identification by domain-invariant mapping network. In *Proceedings of the CVPR* (pp. 719–728).
- Song, L., Wang, C., Zhang, L., Du, B., Zhang, Q., Huang, C., & Wang, X. (2020). Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognition*, 102, 107173.
- Sun, B., & Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. In *Proceedings of the ECCV* (pp. 443–450).
- Sun, X., & Zheng, L. (2019). Dissecting person re-identification from the viewpoint of viewpoint. In *Proceedings of the CVPR* (pp. 608–617).
- Tan, W., Ding, C., Wang, P., Gong, M., & Jia, K. (2023). Style interleaved learning for generalizable person re-identification. *IEEE Transactions on Multimedia*.
- Tang, Z., Gao, Y., Zhu, Y., Zhang, Z., Li, M., & Metaxas, D. N. (2021). Crossnorm and selfnorm for generalization under distribution shifts. In *Proceedings of the ICCV* (pp. 52–61).
- Tzeng, E., Hoffman, J., Saenko, K., & Darrell, T. (2017). Adversarial discriminative domain adaptation. In *Proceedings of the CVPR* (pp. 7167–7176).
- Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Lopez-Paz, D., & Bengio, Y. (2019). Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the ICML* (pp. 6438–6447).
- Volpi, R., Namkoong, H., Sener, O., Duchi, J., Murino, V., & Savarese, S. (2018). Generalizing to unseen domains via adversarial data augmentation. In *Proceedings of the NeurIPS*.
- Wang, D., & Zhang, S. (2020). Unsupervised person re-identification via multi-label classification. In *Proceedings of the CVPR* (pp. 10981–10990).
- Wang, J., Lan, C., Liu, C., Ouyang, Y., Zeng, W., & Qin, T. (2021). Generalizing to unseen domains: A survey on domain generalization. arXiv preprint [arXiv:2103.03097](https://arxiv.org/abs/2103.03097)
- Wang, J., Zhu, X., Gong, S., & Li, W. (2018). Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *Proceedings of the CVPR* (pp. 2275–2284).
- Wang, W., Zhao, F., Liao, S., & Shao, L. (2022). Attentive waveblock: Complementarity-enhanced mutual networks for unsupervised domain adaptation in person re-identification and beyond. In *Proceedings of the IEEE TIP*.
- Wang, X., Zhang, H., Huang, W., & Scott, M. R. (2020). Cross-batch memory for embedding learning. In *Proceedings of the CVPR* (pp. 6388–6397).
- Wei, L., Zhang, S., Gao, W., & Tian, Q. (2018). Person transfer GAN to bridge domain gap for person re-identification. In *Proceedings of the CVPR* (pp. 79–88).
- Wei, Z., Yang, X., Wang, N., & Gao, X. (2021). Syncretic modality collaborative learning for visible infrared person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 225–234).
- Wu, A., Zheng, W.-S., & Lai, J.-H. (2019). Unsupervised person re-identification by camera-aware similarity consistency learning. In *Proceedings of the ICCV* (Vol. 6922–6931).
- Wu, Y., Inkpen, D., & El-Roby, A. (2020). Dual mixup regularized learning for adversarial domain adaptation. In *Proceedings of the ECCV* (pp. 540–555).
- Wu, Z., Xiong, Y., Yu, S. X., & Lin, D. (2018). Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the CVPR* (pp. 3733–3742).
- Xiang, S., Gao, J., Guan, M., Ruan, J., Zhou, C., Liu, T., Qian, D., & Fu, Y. (2023). Learning robust visual-semantic embedding for generalizable person re-identification. arXiv preprint [arXiv:2304.09498](https://arxiv.org/abs/2304.09498)
- Xiao, T., Li, S., Wang, B., Lin, L., & Wang, X. (2017). Joint detection and identification feature learning for person search. In *Proceedings of the CVPR* (pp. 3415–3424).
- Xu, M., Zhang, J., Ni, B., Li, T., Wang, C., Tian, Q., & Zhang, W. (2020). Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI* (Vol. 34, pp. 6502–6509).
- Xu, Z., Liu, D., Yang, J., Raffel, C., & Niethammer, M. (2021). Robust and generalizable visual representation learning via random convolutions. In *Proceedings of the ICLR*.
- Yang, F., Li, K., Zhong, Z., Luo, Z., Sun, X., Cheng, H., Guo, X., Huang, F., Ji, R., & Li, S. (2020). Asymmetric co-teaching for unsupervised cross-domain person re-identification. In *Proceedings of the AAAI* (Vol. 34, pp. 12597–12604).
- Yang, Q., Yu, H.-X., Wu, A., & Zheng, W.-S. (2019). Patch-based discriminative feature learning for unsupervised person re-identification. In *Proceedings of the CVPR* (pp. 3633–3642).
- Ye, M., Ruan, W., Du, B., & Shou, M. Z. (2021a). Channel augmented joint learning for visible-infrared recognition. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 13567–13576).
- Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., & Hoi, S. C. (2021b). Deep learning for person re-identification: A survey and outlook. *IEEE TPAMI*, 44, 2872–2893.
- Yu, H., Cheng, X., Peng, W., Liu, W., & Zhao, G. (2023). Modality unifying network for visible-infrared person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 11185–11195).
- Yu, H.-X., Zheng, W.-S., Wu, A., Guo, X., Gong, S., & Lai, J.-H. (2019). Unsupervised person re-identification by soft multilabel learning. In *Proceedings of the CVPR* (pp. 2148–2157).
- Zhai, Y., Peng, P., Jia, M., Li, S., Chen, W., Gao, X., & Tian, Y. (2023). Population-based evolutionary gaming for unsupervised person re-identification. *International Journal of Computer Vision*, 131(1), 1–25.
- Zhai, Y., Lu, S., Ye, Q., Shan, X., Chen, J., Ji, R., & Tian, Y. (2020a). Ad-cluster: Augmented discriminative clustering for domain adaptive person re-identification. In *Proceedings of the CVPR* (pp. 9021–9030).
- Zhai, Y., Ye, Q., Lu, S., Jia, M., Ji, R., & Tian, Y. (2020b). *Multiple expert brainstorming for domain adaptive person re-identification*. In *Proceedings of the ECCV*.
- Zhang, H., Cisse, M., Dauphin, Y. N., & Lopez-Paz, D. (2018). Mixup: Beyond empirical risk minimization. In *Proceedings of the ICLR*.
- Zhang, P., Dou, H., Yu, Y., & Li, X. (2022). Adaptive cross-domain learning for generalizable person re-identification. In: *European conference on computer vision* (pp. 215–232). Berlin: Springer.
- Zhang, T., Xie, L., Wei, L., Zhuang, Z., Zhang, Y., Li, B., & Tian, Q. (2021a). Unrealperson: An adaptive pipeline towards costless person re-identification. In *Proceedings of the CVPR* (pp. 11506–11515).
- Zhang, X., Cao, J., Shen, C., & You, M. (2019a). Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In *Proceedings of the ICCV* (pp. 8222–8231).
- Zhang, Z., He, T., Zhang, H., Zhang, Z., Xie, J., & Li, M. (2019b). Bag of freebies for training object detection neural networks. arXiv preprint [arXiv:1902.04103](https://arxiv.org/abs/1902.04103)
- Zhang, Y., Yan, Y., Lu, Y., & Wang, H. (2021b). Towards a unified middle modality learning for visible-infrared person re-identification. In *Proceedings of the 29th ACM international conference on multimedia* (pp. 788–796).
- Zhao, F., Liao, S., Xie, G.-S., Zhao, J., Zhang, K., & Shao, L. (2020). Unsupervised domain adaptation with noise resistible mutual



- training for person re-identification. In *Proceedings of the ECCV* (pp. 526–544).
- Zhao, J., Zhao, Y., Li, J., Yan, K., & Tian, Y. (2021a). Heterogeneous relational complement for vehicle re-identification. In *Proceedings of the ICCV* (pp. 205–214).
- Zhao, Y., Zhong, Z., Yang, F., Luo, Z., Lin, Y., Li, S., & Sebe, N. (2021b). Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *Proceedings of the CVPR* (pp. 6277–6286).
- Zhao, Y., Zhong, Z., Zhao, N., Sebe, N., & Lee, G. H. (2023). Style-hallucinated dual consistency learning: A unified framework for visual domain generalization. *International Journal of Computer Vision*, *132*, 837–853.
- Zheng, K., Lan, C., Zeng, W., Zhan, Z., & Zha, Z.-J. (2021a). Exploiting sample uncertainty for domain adaptive person re-identification. In *Proceedings of the AAAI*.
- Zheng, K., Liu, W., He, L., Mei, T., Luo, J., & Zha, Z.-J. (2021b). Group-aware label transfer for domain adaptive person re-identification. In *Proceedings of the CVPR* (pp. 5310–5319).
- Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., & Tian, Q. (2015). Scalable person re-identification: A benchmark. In *Proceedings of the ICCV* (pp. 1116–1124).
- Zheng, Y., Tang, S., Teng, G., Ge, Y., Liu, K., Qin, J., Qi, D., & Chen, D. (2021c). Online pseudo label generation by hierarchical cluster dynamics for adaptive person re-identification. In *Proceedings of the ICCV* (pp. 8371–8381).
- Zheng, L., Yang, Y., & Hauptmann, A. G. (2016). Person re-identification: Past, present and future. arXiv preprint [arXiv:1610.02984](https://arxiv.org/abs/1610.02984)
- Zheng, Z., Zheng, L., & Yang, Y. (2017). Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the ICCV* (pp. 3754–3762).
- Zhong, Z., Zheng, L., Cao, D., & Li, S. (2017). Re-ranking person re-identification with k-reciprocal encoding. In *Proceedings of the CVPR* (pp. 1318–1327).
- Zhong, Z., Zheng, L., Kang, G., Li, S., & Yang, Y. (2020a). Random erasing data augmentation. In *Proceedings of the AAAI* (Vol. 34, pp. 13001–13008).
- Zhong, Z., Zheng, L., Li, S., & Yang, Y. (2018). Generalizing a person retrieval model hetero-and homogeneously. In *Proceedings of the ECCV* (pp. 172–188).
- Zhong, Z., Zheng, L., Luo, Z., Li, S., & Yang, Y. (2019). Invariance matters: Exemplar memory for domain adaptive person re-identification. In *Proceedings of the CVPR* (pp. 598–607).
- Zhong, Z., Zheng, L., Luo, Z., Li, S., & Yang, Y. (2020b). Learning to adapt invariance in memory for person re-identification. In *Proceedings of the IEEE TPAMI*.
- Zhou, K., Liu, Z., Qiao, Y., Xiang, T., & Loy, C. C. (2021). Domain generalization in vision: A survey. arXiv preprint [arXiv:2103.02503](https://arxiv.org/abs/2103.02503)
- Zhou, K., Yang, Y., Hospedales, T., & Xiang, T. (2020a). Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI* (pp. 13025–13032).
- Zhou, K., Yang, Y., Hospedales, T., & Xiang, T. (2020b). Learning to generate novel domains for domain generalization. In *Proceedings of the ECCV* (pp. 561–578). Berlin: Springer.
- Zhou, K., Yang, Y., Qiao, Y., & Xiang, T. (2021). Domain generalization with mixstyle. In *Proceedings of the ICLR*.
- Zhou, K., Yang, Y., Qiao, Y., & Xiang, T. (2023). Mixstyle neural networks for domain generalization and adaptation. *International Journal of Computer Vision*, *132*, 822–836.
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the ICCV* (pp. 2223–2232).
- Zhuang, Z., Wei, L., Xie, L., Zhang, T., Zhang, H., Wu, H., Ai, H., & Tian, Q. (2020). Rethinking the distribution gap of person re-identification with camera-based batch normalization. In *Proceedings of the ECCV* (pp. 140–157).
- Zhuo, J., Wang, S., Zhang, W., & Huang, Q. (2017). Deep unsupervised convolutional domain adaptation. *Proceedings of the ACM MM* (pp. 261–269).
- Zou, Y., Yang, X., Yu, Z., Kumar, B., & Kautz, J. (2020). Joint disentangling and adaptation for cross-domain person re-identification. In *Proceedings of the ECCV*.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.