

This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Towards Discriminative Representation Learning for Unsupervised Person Re-identification

Takashi Isobe^{1,2}, Dong Li¹, Lu Tian¹, Weihua Chen³, Yi Shan¹, Shengjin Wang^{2*} ¹Xilinx Inc., Beijing, China. ²Department of Electronic Engineering, Tsinghua University ³Machine Intelligence Technology Lab, Alibaba Group

{dongl, lutian, yishan}@xilinx.com jbj18@mails.tsinghua.edu.cn
wgsg@tsinghua.edu.cn kugang.cwh@alibaba-inc.com

Abstract

In this work, we address the problem of unsupervised domain adaptation for person re-ID where annotations are available for the source domain but not for target. Previous methods typically follow a two-stage optimization pipeline, where the network is first pre-trained on source and then fine-tuned on target with pseudo labels created by feature clustering. Such methods sustain two main limitations. (1) The label noise may hinder the learning of discriminative features for recognizing target classes. (2) The domain gap may hinder knowledge transferring from source to target. We propose three types of technical schemes to alleviate these issues. First, we propose a cluster-wise contrastive learning algorithm (CCL) by iterative optimization of feature learning and cluster refinery to learn noise-tolerant representations in the unsupervised manner. Second, we adopt a progressive domain adaptation (PDA) strategy to gradually mitigate the domain gap between source and target data. Third, we propose Fourier augmentation (FA) for further maximizing the class separability of re-ID models by imposing extra constraints in the Fourier space. We observe that these proposed schemes are capable of facilitating the learning of discriminative feature representations. Experiments demonstrate that our method consistently achieves notable improvements over the state-of-the-art unsupervised re-ID methods on multiple benchmarks, e.g., surpassing MMT largely by 8.1%, 9.9%, 11.4% and 11.1% mAP on the Market-to-Duke, Duke-to-Market, Market-to-MSMT and Duke-to-MSMT tasks, respectively.



Figure 1. Given labeled source data and unlabeled target data, our goal is to learn feature representations for recognizing target classes. For this unsupervised domain adaptation task, we propose three technical schemes to learn discriminative target features: (a) clusterwise contrastive learning, (b) progressive domain adaptation (c) Fourier augmentation.

1. Introduction

Person re-identification (re-ID) is an important task in intelligent surveillance, which aims at identifying the person across different camera views. Recent person re-ID methods have achieved impressive performance owing to the advancement of deep convolutional neural networks (CNNs) [52, 67, 51, 31, 6, 88, 27]. However, the success is mainly attributed to supervised learning over massive human-labeled data. The need of time-consuming manual annotations substantially limits the scalability of re-ID models. Besides, directly applying a pre-trained re-ID model to other new domains may cause significant performance drop due to the inherent data distribution shift across different surveillance cameras. Recently, unsupervised domain adaptation (UDA)

^{*}Corresponding author

has thus attracted much attention to adapt the model learned on a labeled source domain to an unlabeled target domain. Prior unsupervised re-ID methods typically rely on iterative training with pseudo labels generated by clustering algorithms on the target domain [79, 70, 17, 35]. These existing methods have shown promising results but still sustain two main limitations. (1) The label noise may mislead an unexpected optimization direction for network training with the unlabeled target data. (2) The knowledge (i.e. the model ability of distinguishing person identities) learned on source can not be sufficiently transferred to target by simply fine-tuning the source model.

To alleviate these problems, we investigate three aspects to facilitating the learning of discriminative feature representations for better recognizing target classes by (1) reducing the label noise on the unlabeled target data, (2) better transferring knowledge learned from source to target, and (3) adding extra training constraints. Accordingly, we propose a unified framework to achieve these goals. First, inspired by the recent contrastive learning algorithm for unsupervised visual representation learning [21, 8, 66], we propose a cluster-wise contrastive learning algorithm to learn noisetolerant representations on the unlabeled target data (Figure 1 (a)). Specifically, we construct a momentum-based movingaverage (MMA) feature encoder and build a dynamic queue to provide sufficient negative samples across multiple minibatches for training. Unlike the instance-wise supervision [21, 8, 66], we incorporate cluster-wise supervision generated by clustering, which is amenable for the high-level re-ID task. Our contrastive learning and feature clustering is performed in an alternating way so that the noise of pseudo labels can be largely reduced. Second, most of existing methods apply a two-stage training process in which the network is first pre-trained on source and then fine-tuned on target. Instead of directly fine-tuning, we adopt a collaborative learning mechanism on both domains with a shared feature encoder (Figure 1 (b)). By gradually decreasing the training weights on source and increasing weights on target, we can better transfer the model ability of distinguishing person identities from source to target. Third, we propose to impose extra constraints in the Fourier space for maximizing the class separability of re-ID models (Figure 1 (c)). We view the amplitude spectrum feature as a kind of nonlinear transformations and compute additional loss functions (e.g. cross-entropy loss) to augment training. In summary, with the proposed method, we can learn better discriminative feature representations and further improve the state-of-the-art performance of unsupervised re-ID.

The main contributions of this paper are summarized as follows. (1) We propose a cluster-wise contrastive learning algorithm to learn noise-tolerant feature representations on the unlabeled target data. The label noise can be largely reduced in the iterative optimization procedure of feature clustering and learning. (2) We propose a progressive domain adaptation strategy to gradually transfer the knowledge learned by the labeled source domain into the unlabeled target domain for unsupervised re-ID. (3) We propose to impose Fourier constraints to further maximize the class separability of the model. We observe the frequency spectrum features can be complementary to the spatial features and beneficial for improving the re-ID performance. (4) Empirical evaluations demonstrate that our method consistently outperforms prior state-of-the-art methods on multiple benchmarks by a large margin. Particularly, using the same ResNet-50 backbone, we surpass MMT [17] by 8.1%, 9.9%, 11.4% and 11.1% mAP on the Market-to-Duke, Duke-to-Market, Market-to-MSMT and Duke-to-MSMT tasks, respectively.

2. Related Work

Unsupervised Visual Representation Learning. Unsupervised visual representation learning aims to learn rich feature representations from large-scale unlabeled images, which is also related to our work. The key idea to perform unsupervised learning is constructing pretext tasks with free supervision. Typical methods include recovering the input image by auto-encoders [57, 43, 77], predicting spatial context [13, 40], clustering features [4, 3], tracking [61] or segmenting objects [42] in videos and discriminating the instance-wise samples [66, 21, 8]. Similar to [21], we also build a contrastive self-supervised learning framework to learn representations on the unlabeled target domain. However, in contrast to relying on the instance-wise supervision by maximizing agreement between differently augmented views of the same instance, we incorporate cluster-wise supervision generated by iterative clustering into contrastive learning. We observe that the class information is more suitable to learn discriminative representations for the re-ID task. Moreover, rather than collecting all the samples from the queue [21], we filter out those having the same pseudo-class with the anchor to ensure the quality of negative samples. Contrastive learning is also widely used in many supervised learning methods where training samples are off-the-shelf with labels. In this work, we focus on how to collect meaningful pairs and reduce the label noise in the unsupervised case.

Unsupervised Domain Adaptation. Generic unsupervised domain adaptation (UDA) methods address the closedset problem where the target domain shares the same semantic classes with the source domain. Typical UDA methods focus on reducing the domain discrepancy by aligning data distribution between source and target domains [69, 47, 74, 46, 48], training adversarial domain-classifiers to encourage features of source and target domains to be indistinguishable [78, 55, 28], or learning domain-specific properties [2, 19, 38, 78, 5]. In this work, we tackle the more challenging open-set problem of UDA for re-ID, where the classes between the source and target domains are not shared. In fact, our method does not rely on any assumption on the classes. The classes between source and target can be exactly the same, totally different or partially shared.

Unsupervised Cross-Domain Person Re-ID. Although supervised person re-ID methods have achieved great performance on the trained data domain [23, 52, 67, 41, 51, 31, 6, 88, 24, 25, 80], the accuracy often drops significantly when directly testing on a different domain. Recently, unsupervised cross-domain person re-ID methods [7, 12, 49, 34, 60, 83] have attracted much attention to address the problem. Typical approaches [79, 70, 44, 17, 71] take a pre-trained model on the labeled source domain as the initialized feature encoder, and further optimize it on the unlabeled target domain by metric learning or unsupervised clustering. Instead of directly fine-tuning the source model, we progressively transfer the knowledge from source to target. Some approaches [17, 86, 73] apply soft labels for training on target, which can reduce the effect of noise to some extent in the optimization process. NRMT [81] introduces a collaborative clustering to fit to noisy instances. Another line of recent work [75] attempts to learn domain-invariant features from style-transferred images. DG-Net++ [89] disentangles feature space from each domain into id-related and id-unrelated components However, the model performance heavily counts on the image generation quality and how to optimize the class separability of learned representations is often neglected. Recent work of [18, 82] jointly optimizes both source and target domains to produce reliable pseudo labels. Our work is related to [18] in the aspect of contrastive learning. The main differences are three-fold. (1) For the contrastive loss, [18] integrates instance-level, cluster-level and class-level supervision on both domains while we only employ cluster-level supervision on the target data. Besides, we do not rely on additional tricks to select clusters (e.g., independence or compactness used in [18]). (2) [18] also performs joint learning of source and target domains. Differently, we focus on progressive training by gradually decreasing the source weights and increasing the target weights. (3) [18] only relies on conventional spatial features while we propose to add Fourier constraints for improving the class separability of re-ID models.

Learning in Fourier Space. The discrete Fourier transform (DFT) converts a finite sequence of values into components of different frequencies, which is a classical mathematical transform method and has many practical applications such as digital signal processing and image processing. Its fast computation algorithm of fast Fourier transform (FFT) and the variant of discrete cosine transform (DCT) have been widely used in data compression. Recently, Fourier transform has also been explored for compressing CNNs by grouping the frequency coefficients of kernel weights



Figure 2. Illustration of the proposed unified framework for unsupervised cross-domain person re-ID.

into hash buckets [9], discarding the low-energy frequency coefficients [63] or training with band-limited frequency spectra [14]. Another line of recent work attempts to learn straight from the compressed representations in the Fourier space [54, 20, 15, 68] for efficient training and inference. To reduce the discrepancy between the source and target distributions, FDA [72] adopts style transfer in Fourier space by swapping the low-frequency spectrum of source with the target, which shows promising results on semantic segmentation. However, image-level perceptual changes can cause significant deterioration of the performance of person re-ID since it relies heavily on the appearance characteristics of persons. Different from prior methods, we apply 1D FFT for converting the output of network from spatial space to Fourier space, and then combine the losses that computed on the FFT features for network optimization.

3. Methodology

We denote the source domain data as $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$, where x_i^s and y_i^s indicate the *i*-th training samples and its corresponding class label. The target domain is denoted as $\mathcal{D}_t = \{x_i^t\}_{i=1}^{N_t}$, where class labels are not available. The goal of unsupervised cross-domain person re-ID is to learn a mapping function $f_{\theta}(\cdot)$ to identify the class label (i.e., person identity) for each target image, where θ is the parameters to be learned. The general optimization target can be formulated as:

$$\mathcal{L}(\theta) = \lambda_{s}(e) \cdot \mathcal{L}^{s}(\theta) + \lambda_{t}(e) \cdot \mathcal{L}^{t}(\theta), \qquad (1)$$

where \mathcal{L}^s and \mathcal{L}^t indicate the optimization targets on the source and target domains, respectively. $\lambda_s(e)$ and $\lambda_t(e)$ are variables that change over time (*e* means epoch) to control training on the source and target domains, respectively.



Figure 3. For a fair comparison, we first pre-train the networks on the labeled source data, and then fine-tune these networks with different contrastive learning methods on target. "B", "M" and "C" represents "Baseline", "MoCo" and "CCL", respectively.

Most of existing methods adopt a two-stage optimization pipeline to address this task. That is, the model is first pretrained on source using the ground-truth labels and then finetuned on target using pseudo labels generated by clustering. We formulate the two-stage baseline method as:

$$\mathcal{L}(\theta) = \lambda_{\rm s}(e) \cdot (\mathcal{L}_{\rm ce}(\theta; y^s) + \mathcal{L}_{\rm tri}(\theta; y^s)) + \lambda_{\rm t}(e) \cdot (\mathcal{L}_{\rm ce}(\theta; \hat{y}^t) + \mathcal{L}_{\rm tri}(\theta; \hat{y}^t)),$$
(2)

$$\lambda_{s}(e) = \begin{cases} 1 & e \in (0, e_{1}] \\ 0 & e \in (e_{1}, e_{2}] \end{cases}, \ \lambda_{t}(e) = \begin{cases} 0 & e \in (0, e_{1}] \\ 1 & e \in (e_{1}, e_{2}] \end{cases}$$
(3)

where \mathcal{L}_{ce} and \mathcal{L}_{tri} represents the cross-entropy classification and triplet loss [26], respectively. y^s and \hat{y}^t means the ground-truth class labels on source and pseudo labels on target, respectively.

However, the pseudo labels generated by clustering inevitably contain noise (i.e., wrong labels), which may cause wrong optimization directions during network training.

3.1. Cluster-wise Contrastive Learning

In order to reduce label noise on the unlabeled target data, we propose a cluster-wise contrastive learning algorithm, which is inspired by the recent success of unsupervised feature learning [21, 62]. In detail, we perform feature clustering and contrastive learning in an alternating manner, towards refining the noisy pseudo labels and updating network weights iteratively. In each round of alternating training, we first employ unsupervised feature clustering (e.g., DBSCAN) to generate pseudo labels and design a cluster-wise contrastive loss to train the network:

$$\mathcal{L}_{\rm ccl}^t(\theta) = -\log \frac{\exp\left(f_{\theta}(x^t) f_{\theta}(x_p^t)/\tau\right)}{\sum_{x_n \in \mathbb{N}_{\rm past}} \exp\left(f_{\theta}(x^t) f_{\hat{\theta}}(x_n^t)/\tau\right)}, \quad (4)$$

where x_p^t and x_n^t indicate the positive (i.e., same pseudo class) and negative (i.e., different pseudo classes) samples to x^t , respectively. $\hat{\theta}$ represents a momentum-based movingaveraging (MMA) feature encoder to maintain consistency



Figure 4. Illustration of the proposed progressive domain adaptation (PDA) in term of the re-ID and clustering performance. (a) the re-ID performance (mAP and rank-1 accuracy) on the test set. (b) the clustering performance (NMI [64] and F-measure [1]) on the training set.

between past and current features, which is updated as $\hat{\theta} \leftarrow$ $m\hat{\theta} + (1-m)\theta$ during training. Here, m is a momentum coefficient which is set to 0.99 in our method. We only update the parameters of the regular encoder θ by backpropagation and constrain no gradient back-propagating for the MMA encoder $\hat{\theta}$. For the anchor x^t , we select its positive sample within the current batch and collect past negative samples from previous batches. To implement cross-batch sampling, we build a dynamic queue to memorize the past features \mathbb{N}_{past} . However, it is unreasonable to simply take all the past features as negative samples because previous batches may contain positive samples of the anchor. Hence, based on pseudo labels generated in the current round, we drop the instances which have the same class with the anchor from the queue to ensure the quality of negative samples. By collecting sufficient negative samples across multiple batches for training, the network can be better optimized compared to a naive contrastive loss with very limited training samples within a single batch.

We summarize the main differences between CCL with MoCo [21] as follows. First, MoCo relies on the instancewise supervision by maximizing agreement between differently augmented views of the same instance, while we extend it to a cluster-wise version by exploiting the pseudo labels generated by clustering to construct pairs for learning. We observe that such class information is more suitable for the re-ID task. Second, in contrast to keep updating the queue throughout the training process, we refresh the queue when pseudo labels are updated by a new round of feature clustering, owing to the fact of the class label of a specific training sample is not consistent across different clustering procedures.

Figure 3 reinforces our intuition that cluster-wise supervision is crucial for the re-ID performance. MoCo almost fails on both benchmarks. By combining with the baseline (Eq. 2), MoCo still yields inferior results. One can reason that instance-level supervision and its optimization target of MoCo is different from the re-ID task. Training with instance-wise pairs may hinder the learning of feature representations to distinguish different classes. We also find that CCL can boost the baseline by a large margin on both benchmarks, validating the non-trivial design and effectiveness of our method.

3.2. Progressive Domain Adaptation

Individually training each domain (Eq. 3) is not optimal for knowledge transfer, especially when there is a large discrepancy between source and target domains. Moreover, when the number of labeled source images is limited, it can easily induce an overfitting trap and hamper the knowledge transfer from source to target.

To alleviate this problem, we propose a progressive domain adaptation strategy to gradually optimize $\mathcal{L}(\theta)$ from source to target. Specifically, we decrease the source training weights and increase the target training weights over time. Unlike the two-stage training baseline (Eq. 3), we can formulate λ_s and λ_t as:

$$\lambda_{s}(e) = \begin{cases} 1, & e \in (0, e_{1}] \\ w(e), e \in (e_{1}, e_{2}] , \lambda_{t}(e) = \begin{cases} 0, & e \in (0, e_{1}] \\ 1 - w(e), e \in (e_{1}, e_{2}] \\ 0, & e \in (e_{2}, e_{3}] \end{cases}$$
(5)

where w(e) defines a decay policy. For example, a multistep policy can be illustrated in Figure 2. The training process is divided into three phases according to Eq. 5. First, we follow common practice in supervised re-ID to pre-train the model on source $(\mathcal{L}(\theta) = \mathcal{L}^{s}(\theta))$ as an initialization for the subsequent optimization. Second, we jointly train the network on both source and target domains. For the labeled source data, the optimization objective remains the same as the pre-training phase. For the unlabeled target data, the optimization objective is combination of CCL, cross-entropy and triplet losses based on pseudo labels. Third, since our goal is to accurately predict the target classes as possible, we continue to train the network with the target data only in the final phase $(\mathcal{L}(\theta) = \mathcal{L}^t(\theta))$. Figure 4 shows the clustering and re-ID performance throughout training. According to the clustering performance, the results imply that our method can gradually reduce label noise and yield cleaner clusters compared to the two-stage baseline. According to the re-ID performance, the results imply that our method can learn better features gradually and achieve higher recognition performance than the baseline.

3.3. Fourier Augmentation

Inspired by [53], we consider to impose extra optimization constraints in Fourier space. Specifically, we first apply Fast Fourier Transform (FFT) to compute the real and imaginary components of the 1D CNN output features. We then exploit the amplitude spectrum $\mathcal{M}(x) = ||\mathcal{F}(f_{\theta}(x))||$ to



Figure 5. Analysis of the proposed Fourier augmentation (FA) in feature space using t-SNE [56] visualization. (Zoom-in for best view)

compute the cross-entropy loss for the target data. In order to better understand the proposed Fourier augmentation scheme, we provide analysis in the following aspects. (1) $\mathcal{M}(\cdot)$ can be viewed as a kind of nonlinear mappings. Thus, joint training on the spatial and Fourier features implies that loss functions are computed for different nonlinear features of a training image. We empirically find that such amplitude spectrum features performs better than an extra single MLP layer. (2) Figure 5 visualizes the CNN feature distribution with and without our Fourier augmentation. he qualitative results show that by imposing these extra constrains for training, different classes can be better distinguished. The qualitative results show that by imposing these extra constrains for training, different classes can be better distinguished. (3) In mathematics, Parseval's Theorem (a special case of Plancherel Theorem) states the relation between a signal and its Fourier transform. In our case, the relation becomes:

$$||f_{\theta}(x)||^{2} = \frac{1}{D} ||\mathcal{M}(f_{\theta}(x))||^{2}$$
(6)

where D means the feature length. According to this nature, the triplet losses based on Euclidean distance will be equivalent for the spatial and Fourier features. Hence, we only add the cross entropy loss in the Fourier space.

3.4. Overall Optimization

The overall optimization objective of our method can be defined as:

$$\mathcal{L}(\theta) = \lambda_{\rm s}(e) \cdot \mathcal{L}^s + \lambda_{\rm t}(e) \cdot (\delta \mathcal{L}_{\rm ccl}^t + \gamma \mathcal{L}_{\rm spa}^t + (1 - \gamma) \mathcal{L}_{\rm fre}^t),$$
(7)

where γ is a loss weight that balances the spatial and Fourier losses. δ controls the effect of cluster-wise contrastive learning. We compute both of the cross-entropy and triplet losses for \mathcal{L}^s and \mathcal{L}^t_{spa} and only compute the cross-entropy loss for \mathcal{L}^t_{fre} .

Table 1. Performance comparisons on multiple benchmarks for unsupervised cross-domain person re-ID. The supervised baseline is obtained by training with cross-entropy and triplet losses using the ground-truth labels of target data. More stronger supervised baseline is obtained by combining the loss in Fourier space. [†] means our re-implementation of [17] with the DBSCAN clustering algorithm for fair comparisons. The results of "Ours*" are obtained by combining the proposed method with soft cross-entropy loss, soft triplet loss and mutual learning strategy introduced by [17].

Method		Market	-to-Duke	Duke-to	o-Market	Market-t	o-MSMT	Duke-to-MSMT	
		mAP	rank-1	mAP	rank-1	mAP	rank-1	mAP	rank-1
PUL [16]	TOMM'18	16.4	30.0	20.5	45.5	-	-	-	-
HHL [85]	ECCV'18	27.2	46.9	31.4	62.2	-	-	-	-
PTGAN [65]	CVPR'18	-	27.4	-	38.6	2.9	10.2	3.3	11.8
TJ-AIDL [59]	CVPR'18	23.0	44.3	26.5	58.2	-	-	-	-
ARN [33]	CVPRW'18	33.4	60.2	39.4	70.3	-	-	-	-
MMFA [34]	BMVC'19	24.7	45.3	27.4	56.7				
PDA-Net [32]	ICCV'19	45.1	63.2	47.6	75.2	-	-	-	-
PCB-PAST [79]	ICCV'19	54.3	72.4	54.6	78.4	-	-	-	-
SSG [70]	ICCV'19	53.4	73.0	58.3	80.0	13.2	31.6	13.3	32.2
CR-GAN [10]	ICCV'19	48.6	84.7	54.0	77.7	-	-	-	-
ECN++ [87]	TPAMI'20	54.4	74.0	63.8	84.1	15.2	40.4	16.0	42.5
MMCL [58]	CVPR'20	51.4	72.4	60.4	84.4	15.1	40.8	16.2	43.6
SNR [29]	CVPR'20	58.1	76.3	61.7	82.8	-	-	-	-
DG-Net++ [89]	ECCV'20	63.8	78.9	61.7	82.1	-	-	-	-
NRMT [81]	ECCV'20	62.2	77.8	71.7	87.8	19.8	43.7	20.6	45.2
MEB-Net [76]	ECCV'20	66.1	79.6	76.0	89.9	-	-	-	-
MMT (k-means) [17]	ICLR'20	65.1	78.0	71.2	87.7	22.9	49.2	23.5	50.1
MMT(DBSCAN) [†] [17]	ICLR'20	62.7	76.8	73.5	89.7	24.4	50.7	25.2	53.2
SpCL [18]	NeurIPS'20	-	-	-	-	26.8	53.7	-	-
Baseline	Ours	53.7	69.9	63.6	82.5	14.5	33.3	17.1	38.4
Baseline + CCL	Ours	59.6	75.0	71.1	87.8	20.1	42.7	22.9	48.4
Baseline + CCL + PDA	Ours	67.3	80.9	80.3	92.5	30.7	59.0	30.1	59.5
Baseline + CCL + PDA + FA	Ours	69.4	82.7	82.2	93.6	32.9	61.8	32.7	62.7
Baseline + CCL + PDA + FA	Ours*	70.8	83.5	83.4	94.2	35.8	65.8	36.3	66.6
Supervised baselin	e	72.3	84.4	82.8	93.6	44.7	72.4	44.7	72.4
Supervised baseline +	FA	74.4	86.0	84.5	94.8	47.1	75.2	47.1	75.2

Table 2. Comparison with the state-of-the-art unsupervised re-ID methods under the synthetic-to-real setting.

Method	PersonX-	to-Market	PersonX-	to-MSMT
	mAP	rank-1	mAP	rank-1
MMT [17]	70.7	86.2	18.2	39.5
SpCL [18]	73.8	88.0	22.7	47.7
Ours	78.4	91.3	26.2	50.1
Ours*	79.6	92.5	28.9	53.2

Table 3. Comparison with other unsupervised domain adaptation methods for Vehicle re-ID tasks. The results of MMT are taken from [18].

Method	VehicleII	D-to-VeRi	Vehicle	VehicleX-to-VeRi		
	mAP	rank-1	mAP	rank-1		
MMT [17]	35.3	74.6	35.6	76.0		
SpCL [18]	38.4	79.9	38.3	82.1		
Ours	41.2	83.6	41.4	85.3		
Ours*	42.7	84.7	42.5	86.5		

4. Experiments

4.1. Datasets

We evaluate the proposed PDA algorithm on three realworld person re-ID datasets: DukeMTMC-reID [45], Market-1501 [84] and MSMT17 [65]. The DukeMTMC-reID dataset

Table 4. Performance comparisons with other contrastive learning methods and our method. "†" means our implementation based on the official code. The cross-entropy and triplet losses are not used for all the experiments here.

Method	Market	-to-Duke	Duke-to	o-Market
	mAP	rank-1	mAP	rank-1
SupCon [†] [30]	66.0	79 .4	75.4	88.1
InstDisc [†] [66] MoCo [†] [21]	1.9 10.3	4.1 17.7	2.4 11.7	5.9 26.2
CCL (Ours)	56.8	71.9	67.5	84.2

contains 1,812 identities with 36,411 images captured by eight cameras, which splits 702 identities with 16,522 images for training and the remaining images for testing. The Market-1501 dataset consists of 32,688 images of 1,501 identities captured by six cameras, where the training set contains 12,936 images of 751 identities, and the test set contains 19,732 images of 750 identities. The MSMT17 dataset is a large-scale person re-ID dataset, which consists of 126,441 bounding boxes of 4,101 identities captured by fifteen cameras, for which 32,621 images of 1,041 identities are used for training. We report performance on four real-world unsupervised domain adaptation tasks: Duke-to-Market, Marketto-Duke, Duke-to-MSMT and Market-to-MSMT, where the ground-truth labels are provided on source only. We also conduct domain adaptation experiments under the syntheticto-real setting, where PersonX [50] is used as the synthetic source domain. We use the standard mean average precision (mAP) and cumulative matching characteristics (CMC) at rank-1 accuracy as evaluation metrics.

4.2. Implementation Details

We use the network (e.g., ResNet-50 [22]) pre-trained on ImageNet [11] as the initial feature encoder. We first train the network for $e_1 = 20$ epochs on source. For progressive training on both source and target domains, we adopt a kstep policy where the loss weights of source and target are decreased and increased by k times, respectively. We train 30 epochs (i.e., $e_2 = 50$) in this phase. To learn discriminative features as possible to distinguish target classes, we continue to optimize the model with the target data only in another 30 epochs (i.e., $e_3 = 80$). The output of MMA encoder is used for inference. For hyper-parameters δ and γ , we conduct parameter analysis to obtain the best choices ($\delta =$ 0.1, $\gamma = 0.7$) on the Market-to-Duke benchmark and fix them on the other benchmarks. For cluster-wise contrastive learning, we set the temperature parameter τ as 0.07 [21] and set the maximum length of the queue as 1024. The training data are cropped to 256×128 and augmented by flipping and rotating with a probability of 0.5. The network is optimized by Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$ and weight decay of 5×10^{-4} . We set a constant learning rate of 3.5×10^{-3} for the entire training process. All of our experiments are conducted on PyTorch 1.1 with 4 TITAN XP GPUs. It costs 8, 10, 15, 15 hours for training our models on Duke-to-Market, Market-to-Duke, Market-to-MSMT, Duketo-MSMT, respectively.

4.3. Comparisons to the State-of-the-Arts

We compare the proposed algorithm with the state-of-theart methods on multiple real-world benchmarks for unsupervised cross-domain person re-ID in Table 1. Our method consistently outperforms the existing methods by a large margin and achieves the best mAP performance on all the four benchmarks. Specifically, with the same backbone (ResNet-50) and clustering algorithm (DBSCAN), our method surpasses MMT [17] by 8.1%, 9.9%, 11.4% and 11.1% mAP on the Market-to-Duke, Duke-to-Market, Market-to-MSMT and Duke-to-MSMT benchmarks, respectively. Compared to the other recent unsupervised re-ID methods (e.g., by alternatively training and clustering [70, 75, 81] or generating synthetic training data by GAN [10]), our PDA method can also obtain superior performance in terms of both mAP and rank-1 accuracy. Compared to the [18] which also uses contrastive learning, our method achieves notable gains on the Market-to-MSMT benchmark, e.g., +9.0% of mAP and +12.1% of rank-1 accuracy. We implement two supervised

Table 5. Ablation studies of the proposed cluster-wise contrastive learning (CCL) algorithm on Market-to-Duke and Duke-to-Market benchmarks. The cross-entropy and triplet losses are used for all the experiments here.

Method		Market-to-Duke			Duke-to-Market		
		mAP	rank-1		mAP	rank-1	
(i). w/o cluster-wise pairs (ii). w/o past negatives (iii). w/o instance drop (iv). $ \mathbb{N}_{past} = 512$ (v). $ \mathbb{N}_{past} = 1024$ (vi) $ \mathbb{N}_{past} = 2048$		37.3 66.8 68.1 68.5 69.4 67.6	50.7 80.0 81.4 82.0 82.7 80.8		46.1 78.6 81.0 81.6 82.2 80.7	66.5 92.7 93.0 93.4 93.6 92.9	

baselines with the standard cross-entropy and triplet losses using the ground-truth labels of target data. Our method can achieve similar or comparable results in the challenging unsupervised learning scenario. We also compare with existing unsupervised re-ID methods under the synthetic-to-real setting in Table 2. Our method consistently surpasses prior state-of-the-art methods on both benchmarks with a large margin, e.g., outperforming [18] by 5.8% of mAP and 6.2% rank-1 accuracy on PersonX-to-Market. We also evaluate our method on other object re-ID benchmarks [37, 36, 39] in Table 3. Our method achieves superior performance than MMT and SpCL by a large margin, e.g., +4.3% and +4.2% mAP over SpCL on VehicleID-to-VeRi and VehicleX-to-VeRi, repectively.

4.4. Ablation Study

Contributions from Algorithmic Components. Table 1 also shows the relative contributions from each algorithmic component. Our cluster-wise contrastive learning algorithm brings significant improvement over the baseline, e.g., 53.7% vs. 59.6% mAP on Market-to-Duke. With the proposed progressive domain adaptation strategy, we obtain another remarkable performance gains, e.g., 59.6% vs. 67.3%. By adding the extra loss for training in Fourier space, we can obtain consistent improvements by around 2% mAP for all the four benchmarks. By combining our method with other training strategies proposed by MMT (e.g., soft losses and mutual learning), we can further obtain improved performance.

Ablation studies on contrastive learning. First, we compare the proposed CCL algorithm with three contrastive learning methods in Table 4. For a fair comparison, we first pre-train the networks on the labeled source data, and then fine-tune these networks with different contrastive losses on target. For the supervised baseline [30], we use the GT labels of target data to generate the training pairs. For other unsupervised contrastive learning methods [66, 21], we find they almost fail on these benchmarks. This is because instance-level supervision is often used to learn general feature representations and its optimization target is different from the re-ID task. Directly applying such instance-wise pairs for

Table 6. Ablation studies of the proposed progressive domain adaptation (PDA) strategy on Market-to-Duke and Duke-to-Market benchmarks.

Method	Market	-to-Duke	Duke-to-Market		
	mAP	rank-1	mAP	rank-1	
(i). source only	31.1	48.8	33.7	62.3	
(ii). two-stage training	62.1	77.7	74.0	89.2	
(iii). static weights (0.8;0.2)	51.9	67.2	62.7	79.1	
(iv). static weights (0.5;0.5)	56.2	73.8	66.8	84.2	
(v). static weights (0.2;0.8)	58.8	74.9	68.8	85.5	
(vi). 2-step policy	67.8	81.2	80.8	92.4	
(vii). 3-step policy	69.4	82.7	82.2	93.6	
(viii). 4-step policy	68.2	81.6	81.1	92.9	
(ix). linear policy	67.6	81.0	80.7	92.5	

training may hinder the discriminability of model to distinguish different high-level classes (i.e., person identities). Second, we conduct detailed ablation experiments on our cluster-wise contrastive learning algorithm in Table 5. Without cluster-wise pairs (i.e., using instance-wise pairs as in [21]), the performance drops significantly on both benchmarks (e.g, only 37.3% mAP on Market-to-Duke and 46.1% mAP on Duke-to-Market). Without collecting negative samples in the past iterations (i.e., collecting them from the current batch only) or without dropping positive instances from past features, both experiments obtain degraded performance. We also test different sizes of the queue to store the past features for training and find $|(N)_{past}| = 1024$ performs best in our setting.

Effect of progressive weights. We conduct ablation experiments to show the effectiveness of our progressive domain adaptation strategy in Table 6. The results of (i) show that the performance is poor by directly testing the pretrained source model on target without training. This is not surprising since no knowledge is transferred to the unlabeled target domain. Two-stage training (i.e., first pre-training on source and then fine-tuning on target) obtain inferior results compared to our progressive training strategy. We test different combinations of static loss weights for joint training source and target. The best choice ($\lambda_s = 0.2, \lambda_t = 0.8$) is still worse than our progressive weights. We also investigate different multi-step policies as well as the linear policy¹ and find that 3-step policy performs best in our experiments.

Fourier space vs. Spatial space. The motivation of our Fourier augmentation to exploit extra feature space to facilitate the network training. We compare the proposed FA with other alternative nonlinear mappings in Table 7. The proposed FA outperforms a single MLP layer (FC + ReLU) on both Market-to-Duke and Duke-to-Market benchmarks, which validates the superiority of the proposed method. With the spatial features or the Fourier features only for training, we achieve similar results on both benchmarks (e.g., 67.3% vs. 67.6% mAP on Market-to-Duke). By joint training in

Table 7. Comparisons with different nonlinear mappings.

Method	Market-to-Duke			Duke-to-Market			
	mAP	rank-1		mAP	rank-1		
MLP	68.2	81.5		81.3	93.2		
Spatial Fourier Spatial + Fourier	67.3 67.6 69.4	80.9 81.3 82.7		80.3 80.8 82.2	92.5 92.6 93.6		

Table 8. Ablation studies on loss weights δ and γ .

Loss v	veights	Market	Market-to-Duke		Duke-to-Market		
δ	γ	mAP	rank-1	mAP	rank-1		
	0	66.7	80.9	79.4	91.2		
0.01	0.5	67.7	81.7	80.1	92.1		
0.01	0.7	68.2	82.4	80.9	92.8		
	1	66.1	80.2	78.5	90.6		
	0	67.6	81.3	80.8	92.6		
0.1	0.5	69.2	82.4	81.8	93.6		
0.1	0.7	69.4	82.7	82.2	93.6		
	1	67.3	80.9	80.3	92.5		
	0	66.0	79.8	77.3	91.3		
	0.5	66.8	80.5	78.1	91.9		
1	0.7	67.2	80.9	78.7	92.5		
	1	65.4	79.1	76.5	90.8		

the spatial and Fourier space, we can further improve the performance.

Hyper-parameter analysis. To investigate the importance of loss weights δ and γ , we conduct experiments by changing γ from 0 to 1 under a fixed δ . Table 8 shows that $\delta = 0.1$ and $\gamma = 0.7$ perform best on both Market-to-Duke and Duke-to-Market benchmarks.

5. Conclusion

In this work, we propose a unified framework by incorporating three technical schemes to address the challening unsupervised cross-domain re-ID problem. To learn noisetolerant feature representations, we propose a cluster-wise contrastive learning algorithm by iterative optimization of feature learning and clustering. Instead of simply fine-tuning the pre-trained source model, we adopt a progressive training mechanism to gradually transfer the knowledge from source to target. Furthermore, We impose extra training constraints on the Fourier space for further maximizing the class separability of re-ID models. Our method consistently outperforms prior unsupervised re-ID methods on multiple benchmarks by a large margin. We believe that an extension of this work is to address large variations (e.g., large poses, partial occlusion) in the unsupervised setting.

Acknowledgement We truly thank Yuwing Tai and Xin Tao for helpful discussion. This work was supported by the National Natural Science Foundation of China under Grant No. 61771288, Cross-Media Intelligent Technology Project of Beijing National Research Center for Information Science and Technology (BNRist) under Grant No. BNR2019TD01022 and the research fund under Grant No. 2019GQG0001 from the Institute for Guo Qiang, Tsinghua University.

¹linear policy: $w(e) = \frac{1}{e_1 - e_2} \cdot e + \frac{e_2}{e_2 - e_1}$

References

- Enrique Amigó, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information retrieval*, 12(4):461–486, 2009.
- [2] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *NeurIPS*, 2016.
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In ECCV, 2018.
- [4] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curated data. In *ICCV*, 2019.
- [5] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *CVPR*, 2019.
- [6] Xiaobin Chang, Timothy M Hospedales, and Tao Xiang. Multi-level factorisation net for person re-identification. In *CVPR*, 2018.
- [7] Xiaobin Chang, Yongxin Yang, Tao Xiang, and Timothy M Hospedales. Disjoint label space transfer learning with common factorised space. In AAAI, 2019.
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [9] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Q Weinberger, and Yixin Chen. Compressing convolutional neural networks in the frequency domain. In *KDD*, 2016.
- [10] Yanbei Chen, Xiatian Zhu, and Shaogang Gong. Instanceguided context rendering for cross-domain person reidentification. In *ICCV*, 2019.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [12] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *CVPR*, 2018.
- [13] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, 2015.
- [14] Adam Dziedzic, John Paparrizos, Sanjay Krishnan, Aaron Elmore, and Michael Franklin. Band-limited training and inference for convolutional neural networks. In *ICML*, 2019.
- [15] Max Ehrlich and Larry S Davis. Deep residual learning in the jpeg transform domain. In *ICCV*, 2019.
- [16] Hehe Fan, Liang Zheng, Chenggang Yan, and Yi Yang. Unsupervised person re-identification: Clustering and fine-tuning. *TOMM*, 14(4):83:1–83:18, 2018.
- [17] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual meanteaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020.
- [18] Yixiao Ge, Dapeng Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *NeurIPS*, 2020.

- [19] Boqing Gong, Kristen Grauman, and Fei Sha. Connecting the dots with landmarks: Discriminatively learning domaininvariant features for unsupervised domain adaptation. In *ICML*, 2013.
- [20] Lionel Gueguen, Alex Sergeev, Ben Kadlec, Rosanne Liu, and Jason Yosinski. Faster neural networks straight from jpeg. In *NeurIPS*, 2018.
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [23] Lingxiao He, Jian Liang, Haiqing Li, and Zhenan Sun. Deep spatial feature reconstruction for partial person reidentification: Alignment-free approach. In CVPR, 2018.
- [24] Lingxiao He, Yinggang Wang, Wu Liu, He Zhao, Zhenan Sun, and Jiashi Feng. Foreground-aware pyramid reconstruction for alignment-free occluded person re-identification. In *CVPR*, 2019.
- [25] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object reidentification. arXiv preprint arXiv:2102.04378, 2021.
- [26] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *CoRR*, abs/1703.07737, 2017.
- [27] Takashi Isobe, Jian Han, Fang Zhuz, Yali Liy, and Shengjin Wang. Intra-clip aggregation for video person reidentification. In *ICIP*, 2020.
- [28] Takashi Isobe, Xu Jia, Shuaijun Chen, Jianzhong He, Yongjie Shi, Jianzhuang Liu, Huchuan Lu, and Shengjin Wang. Multitarget domain adaptation with collaborative consistency learning. In CVPR, 2021.
- [29] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *CVPR*, 2020.
- [30] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020.
- [31] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In CVPR, 2018.
- [32] Yu-Jhe Li, Ci-Siang Lin, Yan-Bo Lin, and Yu-Chiang Frank Wang. Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation. In *ICCV*, 2019.
- [33] Yu-Jhe Li, Fu-En Yang, Yen-Cheng Liu, Yu-Ying Yeh, Xiaofei Du, and Yu-Chiang Frank Wang. Adaptation and reidentification network: An unsupervised deep transfer learning approach to person re-identification. In CVPRW, 2018.
- [34] Shan Lin, Haoliang Li, Chang-Tsun Li, and Alex Chichung Kot. Multi-task mid-level feature alignment network for unsupervised cross-dataset person re-identification. In *BMVC*, 2018.
- [35] Yutian Lin, Xuanyi Dong, Liang Zheng, Yan Yan, and Yi Yang. A bottom-up clustering approach to unsupervised person re-identification. In AAAI, 2019.
- [36] Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In CVPR, 2016.

- [37] Xinchen Liu, Wu Liu, Tao Mei, and Huadong Ma. A deep learning-based approach to progressive vehicle reidentification for urban surveillance. In *ECCV*, 2016.
- [38] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- [39] Milind Naphade, Shuo Wang, David C Anastasiu, Zheng Tang, Ming-Ching Chang, Xiaodong Yang, Liang Zheng, Anuj Sharma, Rama Chellappa, and Pranamesh Chakraborty. The 4th ai city challenge. In CVPRW, 2020.
- [40] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In ECCV, 2016.
- [41] Sakrapee Paisitkriangkrai, Chunhua Shen, and Anton Van Den Hengel. Learning to rank in person re-identification with metric ensembles. In CVPR, 2015.
- [42] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017.
- [43] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [44] Lei Qi, Lei Wang, Jing Huo, Luping Zhou, Yinghuan Shi, and Yang Gao. A novel unsupervised camera-aware domain adaptation framework for person re-identification. In *ICCV*, 2019.
- [45] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In ECCV, 2016.
- [46] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In CVPR, 2018.
- [47] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *AAAI*, 2018.
- [48] Kihyuk Sohn, Wenling Shang, Xiang Yu, and Manmohan Chandraker. Unsupervised domain adaptation for distance metric learning. In *ICLR*, 2019.
- [49] Liangchen Song, Cheng Wang, Lefei Zhang, Bo Du, Qian Zhang, Chang Huang, and Xinggang Wang. Unsupervised domain adaptive re-identification: Theory and practice. *CoRR*, abs/1807.11334, 2018.
- [50] Xiaoxiao Sun and Liang Zheng. Dissecting person reidentification from the viewpoint of viewpoint. In *CVPR*, 2019.
- [51] Yifan Sun, Liang Zheng, Weijian Deng, and Shengjin Wang. Svdnet for pedestrian retrieval. In *ICCV*, 2017.
- [52] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 2018.
- [53] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020.

- [54] Robert Torfason, Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Towards image understanding from deep compression without decoding. In *ICLR*, 2018.
- [55] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017.
- [56] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [57] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, 2008.
- [58] Dongkai Wang and Shiliang Zhang. Unsupervised person re-identification via multi-label classification. In CVPR, 2020.
- [59] Jingya Wang, Xiatian Zhu, Shaogang Gong, and Wei Li. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In CVPR, 2018.
- [60] Menglin Wang, Baisheng Lai, Jianqiang Huang, Xiaojin Gong, and Xian-Sheng Hua. Camera-aware proxies for unsupervised person re-identification. In AAAI, 2020.
- [61] Xiaolong Wang and Abhinav Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015.
- [62] Xun Wang, Haozhi Zhang, Weilin Huang, and Matthew R Scott. Cross-batch memory for embedding learning. In CVPR, 2020.
- [63] Yunhe Wang, Chang Xu, Chao Xu, and Dacheng Tao. Packing convolutional neural networks in the frequency domain. *TPAMI*, 41(10):2495–2510, 2018.
- [64] Zhongdao Wang, Liang Zheng, Yali Li, and Shengjin Wang. Linkage based face clustering via graph convolution network. In CVPR, 2019.
- [65] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, 2018.
- [66] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In CVPR, 2018.
- [67] Jing Xu, Rui Zhao, Feng Zhu, Huaming Wang, and Wanli Ouyang. Attention-aware compositional network for person re-identification. In *CVPR*, 2018.
- [68] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In CVPR, 2020.
- [69] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, 2017.
- [70] Fu Yang, Wei Yunchao, Wang Guanshuo, Zhou Yuqian, Shi Honghui, and Huang Thomas. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *ICCV*, 2019.
- [71] Fengxiang Yang, Zhun Zhong, Zhiming Luo, Yuanzheng Cai, Yaojin Lin, Shaozi Li, and Nicu Sebe. Joint noise-tolerant learning and meta camera shift adaptation for unsupervised person re-identification. In AAAI, 2021.

- [72] Yanchao Yang and Stefano Soatto. Fda: Fourier domain adaptation for semantic segmentation. In CVPR, 2020.
- [73] Hong-Xing Yu, Wei-Shi Zheng, Ancong Wu, Xiaowei Guo, Shaogang Gong, and Jian-Huang Lai. Unsupervised person re-identification by soft multilabel learning. In CVPR, 2019.
- [74] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Central moment discrepancy (cmd) for domain-invariant representation learning. In *ICLR*, 2017.
- [75] Yunpeng Zhai, Shijian Lu, Qixiang Ye, Xuebo Shan, Jie Chen, Rongrong Ji, and Yonghong Tian. Ad-cluster: Augmented discriminative clustering for domain adaptive person re-identification. In *CVPR*, 2020.
- [76] Yunpeng Zhai, Qixiang Ye, Shijian Lu, Mengxi Jia, Rongrong Ji, and Yonghong Tian. Multiple expert brainstorming for domain adaptive person re-identification. In *ECCV*, 2020.
- [77] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In ECCV, 2016.
- [78] Weichen Zhang, Wanli Ouyang, Wen Li, and Dong Xu. Collaborative and adversarial network for unsupervised domain adaptation. In *CVPR*, 2018.
- [79] Xinyu Zhang, Jiewei Cao, Chunhua Shen, and Mingyu You. Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In *ICCV*, 2019.
- [80] Yuqi Zhang, Qian Qi, Chong Liu, Weihua Chen, Fan Wang, Hao Li, and Rong Jin. Graph convolution for re-ranking in person re-identification. arXiv preprint arXiv:2107.02220, 2021.
- [81] Fang Zhao, Shengcai Liao, Guo-Sen Xie, Jian Zhao, Kaihao Zhang, and Ling Shao. Unsupervised domain adaptation with noise resistible mutual-training for person re-identification. In ECCV, 2020.
- [82] Kecheng Zheng, Cuiling Lan, Wenjun Zeng, Zhizheng Zhan, and Zheng-Jun Zha. Exploiting sample uncertainty for domain adaptive person re-identification. In AAAI, 2021.
- [83] Kecheng Zheng, Wu Liu, Lingxiao He, Tao Mei, Jiebo Luo, and Zheng-Jun Zha. Group-aware label transfer for domain adaptive person re-identification. In *CVPR*, 2021.
- [84] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015.
- [85] Zhun Zhong, Liang Zheng, Shaozi Li, and Yi Yang. Generalizing a person retrieval model hetero-and homogeneously. In *ECCV*, 2018.
- [86] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *CVPR*, 2019.
- [87] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Learning to adapt invariance in memory for person re-identification. *TPAMI*, 2020.
- [88] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. Camera style adaptation for person re-identification. In *CVPR*, 2018.
- [89] Yang Zou, Xiaodong Yang, Zhiding Yu, BVK Kumar, and Jan Kautz. Joint disentangling and adaptation for cross-domain person re-identification. In *ECCV*, 2020.