

DASK: Distribution Rehearsing via Adaptive Style Kernel Learning for Exemplar-Free Lifelong Person Re-Identification

Kunlun Xu¹, Chenghao Jiang¹, Peixi Xiong², Yuxin Peng¹, Jiahuan Zhou^{1*}

¹Wangxuan Institute of Computer Technology, Peking University, Beijing, China

²Intel Labs

xkl@stu.pku.edu.cn, cj533@scarletmail.rutgers.edu, peixi.xiong@intel.com, {pengyuxin, jiahuanzhou}@pku.edu.cn

Abstract

Lifelong person re-identification (LReID) is an important but challenging task that suffers from catastrophic forgetting due to significant domain gaps between training steps. Existing LReID approaches typically rely on data replay and knowledge distillation to mitigate this issue. However, data replay methods compromise data privacy by storing historical exemplars, while knowledge distillation methods suffer from limited performance due to the cumulative forgetting of undistilled knowledge. To overcome these challenges, we propose a novel paradigm that models and rehearses the distribution of the old domains to enhance knowledge consolidation during the new data learning, possessing a strong anti-forgetting capacity without storing any exemplars. Specifically, we introduce an exemplar-free LReID method called Distribution Rehearsing via Adaptive Style Kernel Learning (DASK). DASK includes a Distribution Rehearser Learning mechanism that learns to transform arbitrary distribution data into the current data style at each learning step. To enhance the style transfer capacity, an Adaptive Kernel Prediction network is explored to achieve an instance-specific distribution adjustment. Additionally, we design a Distribution Rehearsing-driven LReID Training module, which rehearses old distribution based on the new data via the old AKPNet model, achieving effective knowledge accumulation. Experimental results show our DASK outperforms the existing methods by 3.6%-6.8% and 4.5%-6.5% on seen and unseen domains, respectively.

Code — <https://github.com/zhouliahuan1991/AAAI2025-LReID-DASK>

Introduction

Person Re-identification (ReID) aims to identify a person of interest across different camera viewpoints (Shi et al. 2023). Extensive research has shown that ReID models trained on a specific dataset often perform poorly when applied to new datasets due to significant domain gaps (Zhao et al. 2021b; Zhou, Su, and Wu 2022). This limitation has spurred a growing interest in the Lifelong Person Re-identification (LReID) task, which focuses on training models with non-stationary datasets step by step, to enhance their adaptability across various domains (Pu et al. 2021). A primary challenge in

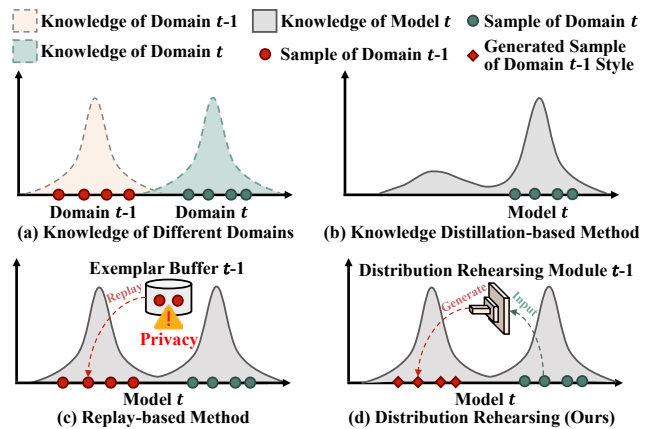


Figure 1: (a) Cross-domain distribution gap entails knowledge difference and catastrophic forgetting. (b) Knowledge distillation suffers inadequate anti-forgetting capacity. (c) Replay-based methods require an exemplar buffer, compromising data privacy. (d) Our method rehearses old distribution via old-style data generation, possessing a strong **anti-forgetting capacity** without storing **historical exemplars**.

LReID is catastrophic forgetting, where the model’s performance on previously learned datasets deteriorates drastically after being trained on new data (Li et al. 2024).

Recently, various LReID methods have emerged, which can be broadly categorized into two main branches: data replay-based and knowledge distillation-based. Data replay-based methods (Yu et al. 2023) rely on storing historical exemplars to rehearse old knowledge while training new models (Fig. 1 (c)). Despite their promising anti-forgetting capabilities, these methods pose significant challenges regarding data privacy and impose a substantial storage burden, limiting their applicability in real-world scenarios (Pu et al. 2022). On the other hand, knowledge distillation-based methods (Pu et al. 2021) focus on maintaining feature consistency between old and new models to preserve historical knowledge. However, due to the pronounced domain gap between old and new data (Fig. 1 (a)), the features learned from historical data may not be adequately represented in the new data (Xu, Zou, and Zhou 2024). Consequently, the

*Corresponding author

corresponding discriminative knowledge is often overwritten by the new model, resulting in insufficient anti-forgetting capacity (Fig. 1 (b)). Additionally, as undistilled historical knowledge becomes inaccessible in subsequent training steps, the forgetting effect accumulates over time, leading to progressively degraded performance on older domains.

To harness the benefits of strong anti-forgetting capabilities and privacy-friendliness from the above LReID solutions, we propose a novel paradigm that regenerates historical data distributions for old knowledge rehearsal without storing any exemplars. Our core motivation, illustrated in Fig. 1 (d), is to develop a distribution rehearsing model capable of transforming data from the new domain to the old domain. Consequently, when new domain data is acquired, old-style data can be synthesized, enabling the model to achieve performance comparable to data replay-based methods.

To achieve this, we propose a novel LReID method named Distribution Rehearse via Adaptive Style Kernel Learning (DASK), which decomposes the LReID learning process into two key components: Distribution Rehearser Learning (DRL) and Distribution Rehearsing-driven LReID Training (DRRT). DRL aims to learn to transform data from arbitrary distributions to the current data style in a self-supervised manner. Considering that different instances possess unique offsets from the target distribution, we propose predicting a specific distribution transfer kernel for each instance utilizing an adaptive kernel prediction network (AKPNet). Then this kernel is exploited to transform the input image into a target domain style. DRRT aims to jointly utilize the new data and the AKPNet model obtained in the previous learning step to consolidate the new and old knowledge into a unified model. Specifically, an old AKPNet model is utilized to transform real new data into old-style data. Then, a joint knowledge consolidation module leverages the new and generated old-style data simultaneously, where a joint learning loss is employed to accumulate knowledge from both kinds of data. Extensive experimental results validate the superiority of our DASK in knowledge consolidation performance, compared to state-of-the-art exemplar-free LReID methods.

To summarize, the contributions of this paper are as follows: (1) We introduce a new LReID paradigm that rehearses the distribution of historical domains to mitigate catastrophic forgetting. (2) A novel non-exemplar LReID method named DASK is proposed, where an instance-adaptive distribution adjustment mechanism is developed to achieve high-quality old-style data generation, and a joint knowledge consolidation module is designed to achieve effective knowledge accumulation. (3) Extensive experimental results demonstrate that DASK achieves state-of-the-art performance in both anti-forgetting and generalization capacity.

Related Work

Person Re-Identification

Previous person re-identification (ReID) works focused on a closed setting where the test domain is identical to the training domain (Shi et al. 2025; Yin et al. 2024). However, such a setting is often impractical since the domains are variable due to the dynamic environment and camera view-

points (Zhou et al. 2017; Zou et al. 2020). To settle this issue, some works investigated the Domain Adaptation (DA) and Domain Generalization (DG) tasks (Zheng et al. 2021; Shi et al. 2024). However, both DA and DG only consider utilizing static training data, neglecting the training data occurring gradually as time changes. Such a practical scenario is denoted Lifelong Person Re-Identification (LReID).

Lifelong Person Re-Identification

The catastrophic forgetting problem (Liu, Peng, and Zhou 2024) is the main challenge of LReID. To settle this, existing LReID works (Ge et al. 2022; Pu et al. 2023) focused on data replay and knowledge distillation strategies.

The data replay strategy focuses on storing and replaying exemplars from historical domains (Wu and Gong 2021; Ge et al. 2022; Yu et al. 2023; Chen, Lagadec, and Bremond 2022; Huang et al. 2022). Although promising anti-forgetting capacity has been shown, such a strategy is impractical due to data privacy (Xu et al. 2024a; Li et al. 2024). Therefore, in this paper, we investigate the scenario where no historical exemplar is accessible.

Knowledge distillation (KD) (Pu et al. 2021; Sun and Mu 2022; Xu et al. 2024b; Cui et al. 2024) aims to mitigate the semantic drift by constraining the output consistency between the new and old model (Yang et al. 2023; Xu, Zou, and Zhou 2024). However, since the model output can only reflect the shared features between the new and old domains, and the unique features of the historical domains are inevitably forgotten as the model updates (Xu, Zou, and Zhou 2024), the anti-forgetting capacity of KD is usually limited. In this paper, to address the drawbacks of KD, we propose rehearsing the historical features from the input level.

Distribution Rehearsing

Existing distribution rehearsing works in ReID (Deng et al. 2018; Zhou et al. 2024; Nguyen et al. 2024) focused on domain adaptation and multi-domain generalization where the data of multiple training domains are available at once (Li et al. 2023; Tan et al. 2023; Wei et al. 2018). However, in LReID, the datasets of different domains are provided step by step. Therefore, the existing methods that rely on directly utilizing multi-domain inputs are infeasible in LReID. In this paper, we instead investigate transferring the data of arbitrary distributions to the known style, thereby the old distribution can be rehearsed without storing historical exemplars.

The Proposed Method

Problem Definition

In the LReID task, a stream of T training datasets $\mathcal{D} = \{D_t\}_{t=1}^T$ is given step by step, where $D_t = \{(x_i, y_i)\}_{i=1}^{n_t}$ contains n_t pairs of image x_i and corresponding identity label y_i . When D_t is given, the previous $t - 1$ datasets are inaccessible (Xu et al. 2024b). After the t -th training step, the trained model is denoted as \mathbf{M}_t . The final model \mathbf{M}_T is evaluated on T testing datasets $\mathcal{D}^{te} = \{D_t^{te}\}_{t=1}^T$ which are collected from T training domains. Besides, additional U datasets $\mathcal{D}^{un} = \{D_t^{un}\}_{t=1}^U$ collected from novel domains are used to evaluate the generalization capacity \mathbf{M}_T .

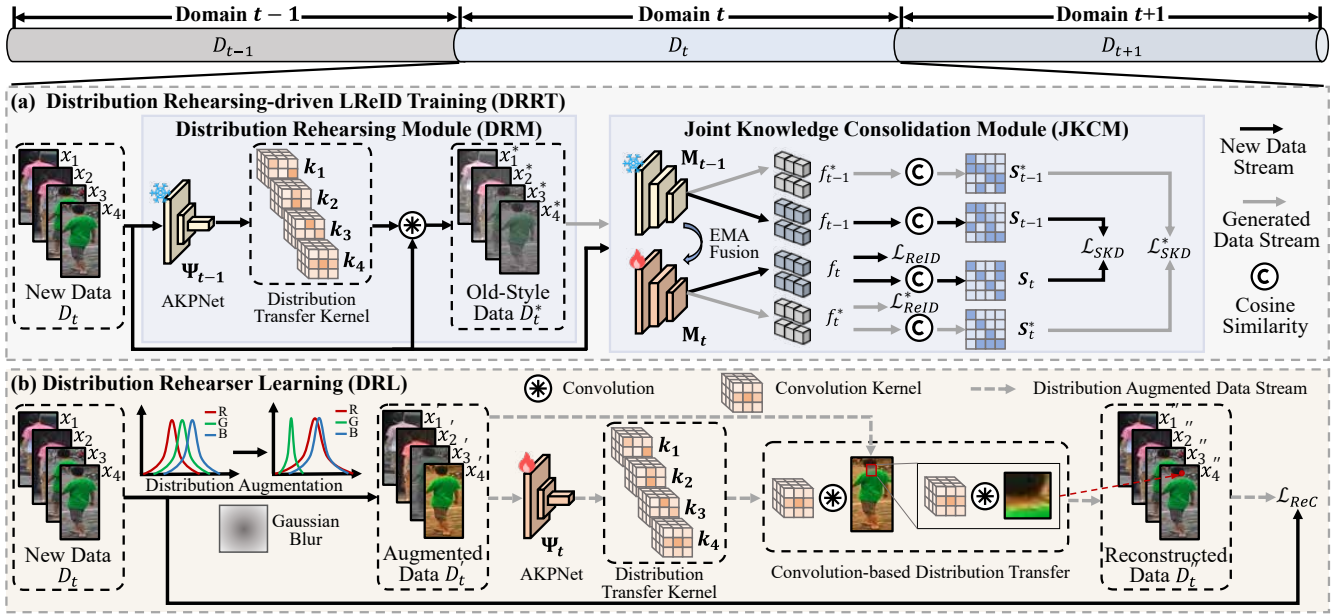


Figure 2: The overall components of our DASK method. At each training step, a new domain dataset D_t is given. (a) DRRT scheme generates the old-style data to enhance knowledge consolidation as the LReID model M_t learns from D_t . (b) DRL mechanism trains an Adaptive Kernel Prediction Network (AKPNet) to achieve instance-specific distribution adjustment, aiming to transform the data of arbitrary domains to domain t , preparing for D_t -style data generation in subsequent training steps.

Overview

As shown in Fig. 2, the proposed method contains a Distribution Rehearsing-driven LReID Training (DRRT) scheme and an additional Distribution Rehearsal Learning (DRL) mechanism. DRRT aims to learn a ReID model M_t based on the old model M_{t-1} , an new AKPNet model Ψ_{t-1} and the new data D_t . DRL aims to obtain an AKPNet model Ψ_t of the current training step based on the new data D_t .

Distribution Rehearsing-driven LReID Training

The main function of DRRT is to generate old-style data and jointly utilize the real new data and generated data to achieve new knowledge learning and mitigate catastrophic forgetting. As shown in Fig. 2 (a), DRRT primarily contains a Distribution Rehearsing Module (DRM) and a Joint Knowledge Consolidation Module (JKCM).

Distribution Rehearsing Module: Given the new data D_t , DRM utilizes an AKPNet model Ψ_{t-1} to generate a distribution transfer kernel $k_i \in \mathbb{R}^{C \times C \times k \times k}$ for each image $x_i \in \mathbb{R}^{W \times H \times C}$, i.e. $k_i = \Psi_{t-1}(x_i)$, where H , W , and C are the image height, width, and channel, respectively. k_i is actually a convolution kernel and k is the kernel size. Then, k_i is utilized to process x_i by:

$$x_i^* = k_i \otimes x_i, \quad (1)$$

where \otimes denotes the convolution process and $x_i^* \in \mathbb{R}^{W \times H \times C}$ is a generated image that contains the distribution information of the historical domain. Then, x_i^* is assigned the same identity label y_i with x_i .

Joint Knowledge Consolidation Module: When the real new data D_t and the generated data D_t^* are obtained, a

frozen old ReID model M_{t-1} and a learnable new ReID model M_t are utilized to process D_t and D_t^* , where the parameters of M_t are initialized with M_{t-1} .

For the new data D_t , given a batch of B input images $x = \{x_i\}_{i=1}^B$, the ReID models transform each image into a d -dimensional feature. Specifically, the features extracted by M_{t-1} and M_t are denoted as $f_{t-1} = \{f_{t-1}^{(i)} \in \mathbb{R}^d\}_{i=1}^B$ and $f_t = \{f_t^{(i)} \in \mathbb{R}^d\}_{i=1}^B$, separately. Then, the cross-instance similarity matrices $S_{t-1} \in \mathbb{R}^{B \times B}$ and $S_t \in \mathbb{R}^{B \times B}$ are calculated by :

$$\begin{cases} S_{t-1} = \frac{F_{t-1}}{\|F_{t-1}\|} \cdot \frac{F_{t-1}^\top}{\|F_{t-1}^\top\|} \\ S_t = \frac{F_t}{\|F_t\|} \cdot \frac{F_t^\top}{\|F_t^\top\|} \end{cases}, \quad (2)$$

where $F_{t-1} \in \mathbb{R}^{B \times d}$ and $F_t \in \mathbb{R}^{B \times d}$ are the matrix forms of f_{t-1} and f_t , respectively. $\frac{F_{t-1}}{\|F_{t-1}\|}$ is a row-wise L2 normalization process to facilitate model convergence (Xu et al. 2024b). Following (Xu, Zou, and Zhou 2024), the cross-instance similarity knowledge distillation loss is adopted:

$$\mathcal{L}_{SKD} = \frac{1}{B} \sum_{i=1}^B KL((S_{t-1})_i || (S_t)_i), \quad (3)$$

where $(S_t)_i$ denotes the i -th row of matrix S_t . \mathcal{L}_{SKD} is an anti-forgetting loss that aims to maintain the predicted cross-instance similarity between the old and new models.

To ensure the new knowledge learning, the classical ReID loss \mathcal{L}_{ReID} is adopted which consists of a Triplet loss

\mathcal{L}_{Tri} (He et al. 2021) and a Cross-Entropy loss \mathcal{L}_{CE} . Following the previous LReID works (Xu et al. 2024b), \mathcal{L}_{ReID} is calculated by:

$$\mathcal{L}_{ReID} = \mathcal{L}_{Tri} + \mathcal{L}_{CE}. \quad (4)$$

Besides, for the generated data D_t^* , M_{t-1} and M_t are adopted to extract the features $f_{t-1}^* = \{(f_{t-1}^{(i)})^* \in \mathbb{R}^d\}_{i=1}^B$ and $f_t^* = \{(f_t^{(i)})^* \in \mathbb{R}^d\}_{i=1}^B$ from the input batch $x^* = \{x_i^*\}_{i=1}^B$. Then, Eq. 2 and Eq. 3 are adopted to obtain rehearsed SKD loss \mathcal{L}_{SKD}^* based on f_t^* and f_{t-1}^* . Then, the rehearsed ReID loss \mathcal{L}_{ReID}^* is calculated following Eq. 4.

Note that \mathcal{L}_{ReID}^* ensures the new model M_t learns with old-style data under accurate new identity labels, which not only enhance historical discriminative knowledge consolidation but also guide the model to mine the domain-irrelevant knowledge, thus improving the model’s generalization capability. Besides, \mathcal{L}_{SKD}^* aims to maintain the cross-instance similarity structure under the historical data style, improving the structural knowledge transfer from M_{t-1} to M_t , which is complementary to the function of \mathcal{L}_{ReID}^* .

ReID Model Training and Inference: When jointly utilizing D_t^* and D_t to training the new model M_t , the overall loss is calculated by:

$$\mathcal{L}_{Total} = \mathcal{L}_{ReID} + \alpha \mathcal{L}_{SKD} + \beta (\mathcal{L}_{ReID}^* + \alpha \mathcal{L}_{SKD}^*), \quad (5)$$

where α and β are hyperparameters to balance the new knowledge learning and historical knowledge forgetting. Besides, following previous works (Li et al. 2024; Xu et al. 2024b), the Exponential Moving Average (EMA) strategy is adopted to fuse the new and old model:

$$M_t \leftarrow \lambda M_{t-1} + (1 - \lambda) M_t, \quad (6)$$

where λ is a weight to balance the new and old parameters.

During inference, only the final model M_T is utilized to extract image features for person matching. Thus our method does not introduce additional computational overhead beyond a feature extractor backbone at the test stage.

Distribution Rehearser Learning

Due to data privacy, the data of different domains can not be obtained simultaneously, thus the existing style transfer methods that rely on the data of multiple accessible domains are infeasible (Deng et al. 2018; Zhou et al. 2024; Nguyen et al. 2024). To settle this, in our DRL mechanism, we first introduce a **Distribution Augmentation** strategy to generate random domain data for cross-domain distribution transfer learning. Then, a **Distribution Reconstruction** design is adopted to guide the model to learn to transfer the distribution augmented data into the current domain.

Distribution Augmentation: Firstly, given the new domain data D_t , we obtain the mean $\mu_i = \{\mu_i^r, \mu_i^g, \mu_i^b\}$ and standard deviation $\sigma_i = \{\sigma_i^r, \sigma_i^g, \sigma_i^b\}$ of the R, G, B channels in each image x_i . Then, the standard deviations of μ_i and σ_i across the images in D_t are also obtained, which are denoted as $\bar{\sigma}_\mu = \{\bar{\sigma}_\mu^r, \bar{\sigma}_\mu^g, \bar{\sigma}_\mu^b\}$ and $\bar{\sigma}_\sigma = \{\bar{\sigma}_\sigma^r, \bar{\sigma}_\sigma^g, \bar{\sigma}_\sigma^b\}$, respectively. Note that μ_i and σ_i reflect the instance-specific color distribution, while $\bar{\sigma}_\mu$ and $\bar{\sigma}_\sigma$ reflect the overall color distribution within domain t .

Then, for an input image x_i , its distribution is augmented by Gaussian Sampling. Specifically, the augmented mean $\mu_i^{r'}$ and standard deviation $\sigma_i^{r'}$ of the R channel are sampled according to $\mu_i^{r'} \sim \mathcal{N}(\mu_i^r, \bar{\sigma}_\mu^{r2})$ and $\sigma_i^{r'} \sim \mathcal{N}(\sigma_i^r, \bar{\sigma}_\sigma^{r2})$, respectively. Furthermore, the augmented R channel data $x_i^{r'} \in \mathbb{R}^{W \times H}$ of the image x_i is obtained by:

$$(x_i^{r'})_{m,n} = \frac{(x_i^r)_{m,n} - \mu_i^r + \mu_i^{r'}}{\sigma_i^r} \cdot \sigma_i^{r'}, \quad (7)$$

where x_i^r is the original R channel data of x_i , (m, n) is the coordinate of a pixel. Similarly, the augmented G channel $x_i^{g'}$ and B channel $x_i^{b'}$ are generated. Once the color distribution augmentation is applied, the random Gaussian blur is also employed to further improve the texture diversity of the augmented image x_i' .

Distribution Reconstruction: To address the factor that different instances possess unique distribution drifts, we introduce the adaptive kernel prediction network (AKPNet) which consists of a lightweight CNN backbone and a linear layer. Specifically, AKPNet generates a Distribution Transfer Kernel k_i for the input x_i' , as shown in Fig. 2 (b). Then, k_i serves as a convolution kernel to process x_i' and the obtained reconstructed image is denoted as x_i'' . A self-supervised reconstruction loss \mathcal{L}_{ReC} is applied as follows:

$$\mathcal{L}_{ReC} = \|x_i - x_i''\|, \quad (8)$$

where $x_i'' = \Psi_t(x_i') \otimes x_i'$. Therefore, Ψ_t is forced to recover the original distribution of x_i from the augmented data x_i' . Thus when the new data D_{t+1} from an arbitrary domain is given, the distribution of D_t can be rehearsed.

Discussion and Analysis

In this section, we discuss the effectiveness of DASK in learning to transform x_i' into the distribution of x_i , in comparison with existing distribution rehearsing solutions.

Compared to Statistics Predicting-based Method: The existing method, CoP (Gu et al. 2023) considers the reverse process of Eq. 7, where the intensity of the R channel of a pixel $(x_i^r)_{m,n}$ can be represented as:

$$(x_i^r)_{m,n} = \frac{\sigma_i^r}{\sigma_i^{r'}} (x_i^{r'})_{m,n} + \mu_i^r - \mu_i^{r'}, \quad (9)$$

which is a hypothetical model that only involves the linear transformation of pixels. However, directly predicting $\frac{\sigma_i^r}{\sigma_i^{r'}}$ and $\mu_i^r - \mu_i^{r'}$ lacks the content-level constraints, leading to unrealistic reconstruction as shown in Fig. 5 (b). Instead, in our DASK, the reconstruction for each pixel $(x_i^r)_{m,n}$ can be deduced as a neighborhood-weighting process:

$$(x_i^r)_{m,n} = \sum_{j \in \{r,g,b\}} \sum_{p=-1}^1 \sum_{q=-1}^1 k_{i,p,q}^j (x_i^{r'})_{m+p,n+q} + b_i^r, \quad (10)$$

where $k_{i,p,q}^j$ is a weight element of k_i , and b_i^r is the bias of channel R . Note that Eq. 9 is a special case of Eq. 10 when $k_{i,p,q}^j = 0, w.r.t. p \neq 0, q \neq 0$. Thus, the proposed instance-adaptive transfer kernel prediction can handle the instance-specific linear hypothesis and model the adjacent relations simultaneously, obtaining better reconstructions.

	Method	Publication	Market-1501		CUHK-SYSU		DukeMTMC		MSMT17		CUHK03		Seen-Avg		UnSeen-Avg	
			mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1
	JointTrain	-	78.9	90.9	86.7	88.2	71.2	82.9	36.2	61.2	61.2	63.4	66.8	77.3	59.4	52.6
CIL	LwF	T-PAMI 2017	56.3	77.1	72.9	75.1	29.6	46.5	6.0	16.6	36.1	37.5	40.2	50.6	47.2	42.6
	SPD	ICCV 2019	35.6	61.2	61.7	64.0	27.5	47.1	5.2	15.5	42.2	44.3	34.4	46.4	40.4	36.6
	PRD	ICML 2023	7.3	18.0	33.5	35.6	3.7	7.6	0.8	2.4	33.8	33.8	15.8	19.5	23.0	17.7
	PRAKA	ICCV 2023	37.4	61.3	69.3	71.8	35.4	55.0	10.7	27.2	54.0	55.6	41.3	54.2	47.7	41.6
	FCS	CVPR 2024	58.3	78.5	75.1	76.2	42.6	59.8	10.2	24.3	35.3	34.9	44.3	54.7	52.1	44.2
LReID	CRL	WACV 2021	58.0	78.2	72.5	75.1	28.3	45.2	6.0	15.8	37.4	39.8	40.5	50.8	47.8	43.5
	AKA	CVPR 2021	51.2	72.0	47.5	45.1	18.7	33.1	16.4	37.6	27.7	27.6	32.3	43.1	44.3	40.4
	PatchKD	ACM MM 2022	68.5	85.7	75.6	78.6	33.8	50.4	6.5	17.0	34.1	36.8	43.7	53.7	49.1	45.4
	MEGE	T-PAMI 2023	39.0	61.6	73.3	76.6	16.9	30.3	4.6	13.4	36.4	37.1	34.0	43.8	47.7	44.0
	ConRFL	PR 2023	59.2	78.3	82.1	84.3	45.6	61.8	12.6	30.4	51.7	53.8	50.2	61.7	57.4	52.3
	CKP	ACM MM 2024	53.8	76.0	81.2	83.0	49.7	67.0	18.4	40.8	44.1	45.8	49.4	62.5	58.0	51.0
	LSTKC	AAAI 2024	54.7	76.0	81.1	83.4	49.4	66.2	20.0	43.2	44.7	46.5	50.0	63.1	57.0	49.9
	DKP	CVPR 2024	60.3	80.6	83.6	85.4	51.6	68.4	19.7	41.8	43.6	44.2	51.8	64.1	59.2	51.6
DASK	This Paper	61.2	82.3	81.9	83.7	58.5	74.6	29.1	57.6	46.2	48.1	55.4	69.3	65.3	58.4	

Table 1: Comparison results of Seen-domain anti-forgetting and UnSeen-domain generalization on Training Order-1.

Compared to Shared Generation Network-based Methods: When training a shared generation network for the images, if a shallow network is adopted, the parameter $k_{i,p,q}^j$ in Eq. 10 is fixed across all inputs. Thus the reconstruction loss forces the model to learn an average style of the target domain (as the illustration of Shared Conv in Fig. 5 (b)). While deep networks can model more complex distributions, they demand large amounts of training data, which is impractical for lifelong learning (Wang et al. 2020). Besides, deep generation networks introduce higher computational overhead (Wang et al. 2023) compared to our lightweight AKPNet. Therefore, our approach offers an effective and efficient means to achieve high-quality reconstruction.

Experiment

Datasets and Evaluation Metrics

Datasets: All the experiments are conducted on the LReID benchmark (Pu et al. 2021), which contains five training subsets (Market-1501 (Zheng et al. 2015), DukeMTMC-reID (Ristani et al. 2016), CUHK-SYSU (Xiao et al. 2016), MSMT17-V2 (Wei et al. 2018), and CUHK03 (Li et al. 2014)) and seven novel subsets (CUHK01 (Li, Zhao, and Wang 2012), CUHK02 (Li and Wang 2013), VIPeR (Gray and Tao 2008), PRID (Hirzer et al. 2011), i-LIDS (Branch 2006), GRID (Loy, Xiang, and Gong 2010), and SenseReID (Zhao et al. 2017)). The training and novel subsets are used for anti-forgetting and generalization evaluation, respectively. Following the previous works (Sun and Mu 2022), two training orders are adopted to imitate varying domain gaps^{1 2}.

Evaluation Metrics: Following previous LReID works (Pu et al. 2021; Sun and Mu 2022), the mean Average Precision (mAP) and Rank@1 (R@1) accuracy are calculated on

¹(Order-1) Market-1501→CUHK-SYSU→DukeMTMC-reID→MSMT17→CUHK03

²(Order-2) DukeMTMC-reID→MSMT17→Market-1501→CUHK-SYSU→CUHK03

the subsets to show the model’s adaptation to each domain. Besides, the average mAP and R@1 are also calculated to evaluate the overall anti-forgetting and generalization capacity on the training (Seen) and novel (Unseen) domains.

Implementation Details

To ensure a fair comparison with previous works (Xu et al. 2024b), we adopt ResNet-50 as the backbone of our ReID model. The backbone of the AKPNet is the Mobilenet-v3 (Howard et al. 2019). When training the ReID model, the first dataset D_1 is trained for 80 epochs and the subsequent $t - 1$ datasets are trained for 60 epochs. As for the distribution rehearsing learning, each dataset is trained for 50 epochs. The input images are resized to 256×128 with random cropping, erasing, and horizontal flipping augmentation. The hyperparameters α , β , and λ are set to 1.0, 4.5, and 0.5 respectively. The training pipeline shown in Fig. 2 (a) DRRT and (b) DRL are conducted independently. To prevent the continuous growth of storage overhead, **only one old AKPNet model**, Ψ_{t-1} , is retained for DRRT at the t -th training step. All experiments are conducted on a single NVIDIA 3090 GPU.

Compared Methods

We compare the proposed method with two streams of methods: *Class Incremental Learning (CIL)* methods (LwF (Li and Hoiem 2017), SPD (Tung and Mori 2019), PRAKA (Shi and Ye 2023), and FCS (Li, Peng, and Zhou 2024)) and *Non-exemplar LReID* methods (CRL (Zhao et al. 2021a), AKA (Pu et al. 2021), PatchKD (Sun and Mu 2022), ConRFL (Huang et al. 2023), MEGE (Pu et al. 2023), LSTKC (Xu, Zou, and Zhou 2024), CKP (Xu et al. 2024a) and DKP (Xu et al. 2024b)). CIL is the most popular lifelong learning task and the CIL methods can be directly applied to the LReID task. Besides, the *JointTrain* that indicates collecting the training data of all seen domains to train the model is also reported, which is commonly regarded as the upper bound of LReID methods (Pu et al. 2023). All

Method	Publication	DukeMTMC		MSMT17		Market-1501		CUHK-SYSU		CUHK03		Seen-Avg		UnSeen-Avg		
		mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	mAP	R@1	
JointTrain	-	71.2	82.9	36.2	61.2	78.9	90.9	86.7	88.2	61.2	63.4	66.8	77.3	59.4	52.6	
CIL	LwF	T-PAMI 2017	42.7	61.7	5.1	14.3	34.4	58.6	69.9	73.0	34.1	34.1	37.2	48.4	44.0	40.1
	SPD	ICCV 2019	28.5	48.5	3.7	11.5	32.3	57.4	62.1	65.0	43.0	45.2	33.9	45.5	39.8	36.3
	PRD	ICML 2023	3.6	8.2	0.6	1.8	8.9	22.3	34.6	36.1	35.4	35.3	16.6	20.7	20.7	15.0
	PRAKA	ICCV 2023	31.2	48.7	6.6	19.1	47.8	69.8	70.4	73.0	54.9	56.6	42.2	53.4	48.4	41.1
	FCS	CVPR 2024	53.6	70.0	9.5	23.5	48.7	69.9	76.2	78.2	37.1	38.4	45.0	56.0	52.7	45.1
	LReID	CRL	WACV 2021	43.5	63.1	4.8	13.7	35.0	59.8	70.0	72.8	34.5	36.8	37.6	49.2	45.3
AKA		CVPR 2021	32.5	49.7	-	-	-	-	-	-	-	-	-	-	40.8	37.2
PatchKD		ACM MM 2022	58.3	74.1	6.4	17.4	43.2	67.4	74.5	76.9	33.7	34.8	43.2	54.1	48.6	44.1
MEGE		T-PAMI 2023	21.6	35.5	3.0	9.3	25.0	49.8	69.9	73.1	34.7	35.1	30.8	40.6	44.3	41.1
ConRFL		PR 2023	34.4	51.3	7.6	20.1	61.6	80.4	82.8	85.1	49.0	50.1	47.1	57.4	57.9	53.4
CKP		ACM MM 2024	49.4	67.0	14.5	33.8	56.0	77.6	83.2	84.9	45.3	47.1	49.7	62.1	57.2	50.0
LSTKC		AAAI 2024	49.9	67.6	14.6	34.0	55.1	76.7	82.3	83.8	46.3	48.1	49.6	62.1	57.6	49.6
DKP		CVPR 2024	53.4	70.5	14.5	33.3	60.6	81.0	83.0	84.9	45.0	46.1	51.3	63.2	59.0	51.6
DASK	This Paper	55.7	74.4	25.2	51.9	71.6	87.7	84.8	86.2	48.4	49.8	57.1	70.0	65.5	57.9	

Table 2: Comparison results of Seen-domain anti-forgetting and UnSeen-domain generalization on Training Order-2.

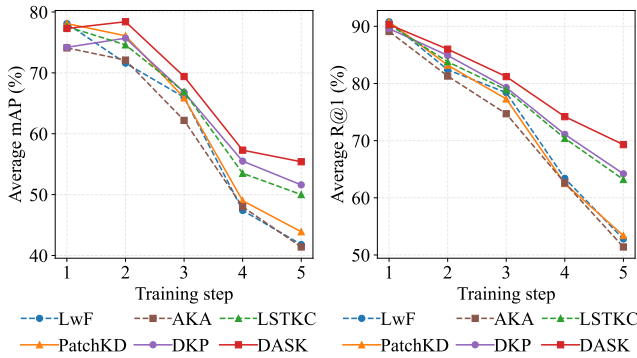


Figure 3: Anti-forgetting tendency on seen domains.

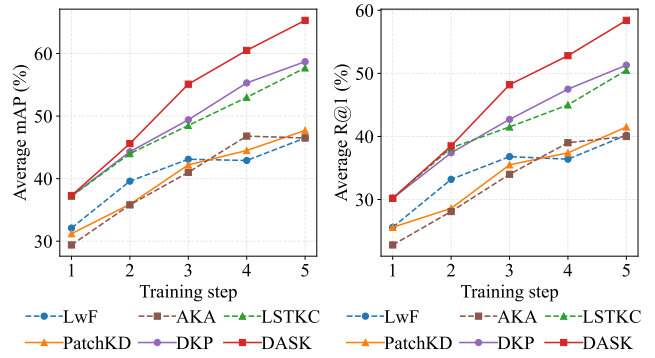


Figure 4: Generalization tendency on unseen domains.

experimental results are reported following the official publications or implemented with the official codes.

We present the performance of different methods on each seen domain, and the average performance across all seen domains (Seen-Avg) and unseen domains (UnSeen-Avg) in Tab. 1 and Tab. 2, corresponding to Training Order-1 and Training Order-2 respectively. The best and second best results are highlighted in **Red** and **Blue**, separately.

Seen-Domain Performance Evaluation

Compared to CIL Methods: In Tab. 1 and Tab. 2, our DASK outperforms all CIL methods on the first four subsets, obtaining **11.1%/14.6%** and **12.1%/14.0%** Seen-Avg mAP/R@1 improvements on two training orders. This is because CIL methods focus on improving the recognition capacity of known classes, whereas the identity classes of the test data are usually unseen during training. DASK obtains inferior results compared to PRAKA on the last subset CUHK03 due to the anti-forgetting designs of DASK partly limiting the learning of new data. The significantly superior average performance of DASK underscores that the proposed distribution rehearsing design can boost the model

to achieve a better balance between learning and forgetting.

Compared to LReID Methods: In Tab. 1 and Tab. 2, our DASK achieves leading performance on most subsets, with **3.6%/5.2%** and **5.8%/6.8%** Seen-Avg mAP/R@1 improvements compared to the state-of-the-art DKP. DASK achieves competitive results compared to PatchKD and ConRFL on the initial and last subsets, respectively. This is because PatchKD adopts strong anti-forgetting constraints and ConRFL adopts looser ones, resulting in over-emphasized performance on the first and last domains. In contrast, our method achieves a better balance between acquisition and forgetting under the proposed distribution rehearsing mechanism. As a result, significantly superior results on the middle subsets and Seen-Avg metrics are obtained.

Seen Domain Performance Tendency: Fig. 3 shows the seen domain evaluation results across the training steps. Our DASK achieves similar performance with existing methods initially and consistently outperforms them from the second training step. These results arise because DASK adopts the same baseline as existing methods but the distribution rehearsing design improves the knowledge consolidation capacity of our model during subsequent learning steps.

Method	Seen-Avg		UnSeen-Avg	
	mAP	R@1	mAP	R@1
Baseline	48.7	60.7	55.5	48.2
Style Augmentation	48.6	63.1	59.9	52.3
Shared Convolution	48.6	61.2	55.7	48.8
Statistical Prediction (CoP)	51.2	65.1	61.3	53.9
DASK	55.4	69.3	65.3	58.4

Table 3: Ablation on the input training data of ReID model.

Baseline	\mathcal{L}_{ReID}^*	\mathcal{L}_{SKD}^*	Seen-Avg		UnSeen-Avg	
			mAP	R@1	mAP	R@1
✓	✗	✗	48.7	60.7	55.5	48.2
✓	✓	✗	55.0	68.8	63.4	56.0
✓	✗	✓	49.9	61.5	56.3	49.2
✓	✓	✓	55.4	69.3	65.3	58.4

Table 4: Ablation on the rehearsing-guided loss components.

Unseen-Domain Generalization Evaluation

Compared to CIL and LReID Methods: As shown in the UnSeen-Avg terms in Tab. 1 and Tab. 2, our DASK outperforms all CIL and LReID methods by at least **6.1%/6.1%** and **6.5%/4.5%** mAP/R@1 on two training orders, which is a more significant improvement compared to the seen domains. This can be attributed to the distribution rehearsing design that enriches the data domains at each training step, guiding the model to learn both intra-domain and inter-domain discriminative knowledge, thus improving the model’s generalization capacity to unknown conditions.

UnSeen Domain Performance Tendency: In Fig. 4, our DASK gradually demonstrates significant superiority in unseen domain performance compared to existing methods as the training steps increase. This is because our distribution rehearsing design enables the model to jointly learn from multiple seen domains to mine generalizable knowledge.

Ablation Studies

We conduct ablation studies on the proposed components under Training Order-1 of the LReID benchmark as follows.

Ablation on Distribution Rehearsing Strategy: In Tab. 3, additional distribution rehearsing choices are compared with our DASK: (1) Baseline denotes using only the new data for training. (2) Style Augmentation denotes adopting the distribution-augmented data in Fig. 2 (b). (3) Shared Convolution denotes using a shared convolution kernel for all images. (4) Statistical Prediction indicates predicting only the mean and standard deviation of the transferred image, as conducted by CoP (Gu et al. 2023). The results show our DASK significantly outperforms the other solutions. Note that Style Augmentation can generate unreasonably hard samples that disturb model learning, Shared Convolution applies the same transformation to all images, neglecting instance-level distribution alignment, and Statistical Prediction ignores informative contextual clues. As a result, the distribution rehearsing strategy of DASK generates higher-quality old-style data and improves lifelong learning due to the adaptive distribution transfer kernel design.

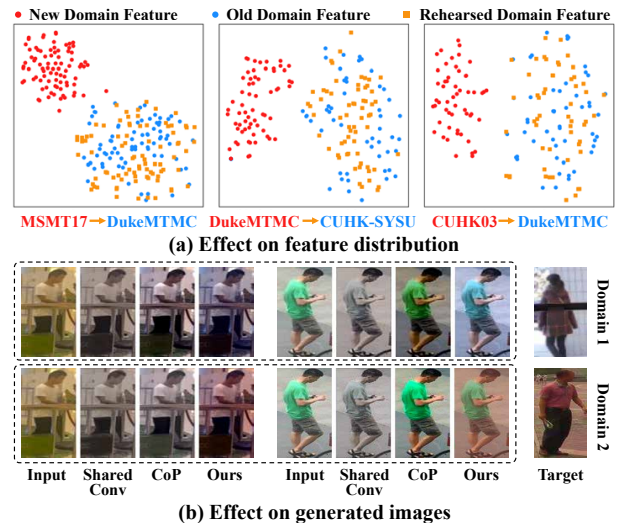


Figure 5: Visualization of distribution rehearsing effects.

Effectiveness of Rehearsing-guided Loss Components:

In Tab. 4, Baseline denotes training with \mathcal{L}_{ReID} and \mathcal{L}_{SKD} . By progressively incorporating the \mathcal{L}_{ReID}^* and \mathcal{L}_{SKD}^* , the results demonstrate marked improvements in LReID performance. Specifically, \mathcal{L}_{ReID}^* facilitates the network in reinforcing discriminative features from old domains. Concurrently, \mathcal{L}_{SKD}^* enhances structural consistency, which mitigates the forgetting caused by parameter overwriting.

Visualization results: Fig. 5 (a) is the t-SNE visualization of the features of the new, old, and rehearsed (transferred from new) domains. There are obvious gaps between new and old feature distributions, while the distributions of rehearsed features and old features overlap significantly. These results demonstrate that our domain rehearsing strategy can effectively generate old-style data. Besides, Fig. 5 (b) visualizes generated images of Shared Convolution, CoP and ours. It is evident that our generations more closely resemble the styles of various target domains, supporting the analysis in Sec. *Ablation on Distribution Rehearsing Strategy*.

Conclusion

In this paper, we propose DASK, a novel exemplar-free LReID method. The key idea is a paradigm that models and rehearses the distribution of the old domains to enhance lifelong knowledge consolidation. To achieve this, a Distribution Rehearser Learning mechanism is introduced, where an adaptive kernel prediction network (AKPNet) is designed to generate an instance-specific domain transfer kernel, achieving a fine-grained style adjustment. Besides, a Distribution Rehearsing-driven LReID Training scheme is introduced to accomplish new knowledge learning and old knowledge accumulation. Our results underscore that the input distribution is closely relevant to the learned knowledge in LReID, and the image level distribution rehearsing is effective for addressing catastrophic forgetting. In future work, we will investigate extending our distribution rehearsing design to more lifelong learning tasks.

Acknowledgments

This work was supported by grants from the National Natural Science Foundation of China (62376011, 61925201, 62132001, 62432001) and Beijing Natural Science Foundation (L247006).

References

- Branch, H. O. S. D. 2006. Imagery Library for Intelligent Detection Systems (i-lids). In *2006 IET Conf. on Crime and Security*, 445–448. IET.
- Chen, H.; Lagadec, B.; and Bremond, F. 2022. Unsupervised Lifelong Person Re-identification via Contrastive Rehearsal. *arXiv:2203.06468*.
- Cui, Z.; Zhou, J.; Wang, X.; Zhu, M.; and Peng, Y. 2024. Learning Continual Compatible Representation for Re-indexing Free Lifelong Person Re-identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 16614–16623. IEEE.
- Deng, W.; Zheng, L.; Ye, Q.; Kang, G.; Yang, Y.; and Jiao, J. 2018. Image-Image Domain Adaptation with Preserved Self-Similarity and Domain-Dissimilarity for Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 994–1003. IEEE.
- Ge, W.; Du, J.; Wu, A.; Xian, Y.; Yan, K.; Huang, F.; and Zheng, W.-S. 2022. Lifelong Person Re-identification by Pseudo Task Knowledge Preservation. In *Proc. AAAI Conf. Artif. Intell.*, 688–696.
- Gray, D.; and Tao, H. 2008. Viewpoint Invariant Pedestrian Recognition with an Ensemble of Localized Features. In *Proc. Eur. Conf. Comput. Vis.*, 262–275. Springer.
- Gu, J.; Luo, H.; Wang, K.; Jiang, W.; You, Y.; and Zhao, J. 2023. Color Prompting for Data-free Continual Unsupervised Domain Adaptive Person Re-Identification. *arXiv preprint arXiv:2308.10716*.
- He, S.; Luo, H.; Wang, P.; Wang, F.; Li, H.; and Jiang, W. 2021. TransReID: Transformer-based Object Re-Identification. In *Proc. IEEE Int. Conf. Comput. Vis.*, 14993–15002. IEEE.
- Hirzer, M.; Belezni, C.; Roth, P. M.; and Bischof, H. 2011. Person Re-Identification by Descriptive and Discriminative Classification. In *Image Analysis*, 91–102. Springer.
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for Mobilenetv3. In *Proc. IEEE Int. Conf. Comput. Vis.*, 1314–1324. IEEE.
- Huang, J.; Yu, X.; An, D.; Wei, Y.; Bai, X.; Zheng, J.; Wang, C.; and Zhou, J. 2023. Learning Consistent Region Features for Lifelong Person Re-Identification. *Pattern Recognition*, 144: 109837.
- Huang, Z.; Zhang, Z.; Lan, C.; Zeng, W.; Chu, P.; You, Q.; Wang, J.; Liu, Z.; and Zha, Z.-j. 2022. Lifelong Unsupervised Domain Adaptive Person Re-identification with Coordinated Anti-forgetting and Adaptation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 14288–14297. IEEE.
- Li, Q.; Peng, Y.; and Zhou, J. 2024. FCS: Feature Calibration and Separation for Non-Exemplar Class Incremental Learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 28495–28504. IEEE.
- Li, Q.; Xu, K.; Peng, Y.; and Zhou, J. 2024. Exemplar-Free Lifelong Person Re-identification via Prompt-Guided Adaptive Knowledge Consolidation. *Int. J. Comput. Vis.*, 1–16.
- Li, W.; and Wang, X. 2013. Locally Aligned Feature Transforms across Views. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 3594–3601. IEEE.
- Li, W.; Zhao, R.; and Wang, X. 2012. Human Reidentification with Transferred Metric Learning. In *ACCV*, 31–44. Springer.
- Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deep-ReID: Deep Filter Pairing Neural Network for Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 152–159. IEEE.
- Li, Y.; Song, J.; Ni, H.; and Shen, H. T. 2023. Style-Controllable Generalized Person Re-Identification. In *Proc. ACM Int. Conf. Multimedia*, 7912–7921.
- Li, Z.; and Hoiem, D. 2017. Learning without Forgetting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(12): 2935–2947.
- Liu, Z.; Peng, Y.; and Zhou, J. 2024. Compositional Prompting for Anti-Forgetting in Domain Incremental Learning. *Int. J. Comput. Vis.*, 1–18.
- Loy, C. C.; Xiang, T.; and Gong, S. 2010. Time-Delayed Correlation Analysis for Multi-Camera Activity Understanding. *Int. J. Comput. Vis.*, 90(1): 106–129.
- Nguyen, V. D.; Mirza, S.; Zakeri, A.; Gupta, A.; Khaldi, K.; Aloui, R.; Mantini, P.; Shah, S. K.; and Merchant, F. 2024. Tackling Domain Shifts in Person Re-Identification: A Survey and Analysis. In *CVPR*, 4149–4159. IEEE.
- Pu, N.; Chen, W.; Liu, Y.; Bakker, E. M.; and Lew, M. S. 2021. Lifelong Person Re-Identification via Adaptive Knowledge Accumulation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 7897–7906. IEEE.
- Pu, N.; Liu, Y.; Chen, W.; Bakker, E. M.; and Lew, M. S. 2022. Meta Reconciliation Normalization for Lifelong Person Re-Identification. In *Proc. ACM Int. Conf. Multimedia*, 541–549.
- Pu, N.; Zhong, Z.; Sebe, N.; and Lew, M. S. 2023. A Memorizing and Generalizing Framework for Lifelong Person Re-Identification. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Ristani, E.; Solera, F.; Zou, R.; Cucchiara, R.; and Tomasi, C. 2016. Performance Measures and A Data Set for Multi-Target, Multi-Camera tracking. In *Proc. Eur. Conf. Comput. Vis.*, 17–35. Springer.
- Shi, J.; Yin, X.; Chen, Y.; Zhang, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2025. Multi-memory matching for unsupervised visible-infrared person re-identification. In *Proc. Eur. Conf. Comput. Vis.*, 456–474. Springer.
- Shi, J.; Yin, X.; Zhang, Y.; Xie, Y.; Qu, Y.; et al. 2024. Learning commonality, divergence and variety for unsupervised visible-infrared person re-identification. In *Proc. Int. Conf. Neural Inf. Process. Syst.*

- Shi, J.; Zhang, Y.; Yin, X.; Xie, Y.; Zhang, Z.; Fan, J.; Shi, Z.; and Qu, Y. 2023. Dual Pseudo-Labels Interactive Self-Training for Semi-Supervised Visible-Infrared Person Re-Identification. In *Proc. IEEE Int. Conf. Comput. Vis.*, 11218–11228.
- Shi, W.; and Ye, M. 2023. Prototype Reminiscence and Augmented Asymmetric Knowledge Aggregation for Non-Exemplar Class-Incremental Learning. In *Proc. IEEE Int. Conf. Comput. Vis.*, 1772–1781. IEEE.
- Sun, Z.; and Mu, Y. 2022. Patch-Based Knowledge Distillation for Lifelong Person Re-Identification. In *Proc. ACM Int. Conf. Multimedia*, 696–707.
- Tan, W.; Ding, C.; Wang, P.; Gong, M.; and Jia, K. 2023. Style Interleaved Learning for Generalizable Person Re-Identification. *IEEE Trans. Multimedia*, 26: 1600–1612.
- Tung, F.; and Mori, G. 2019. Similarity-Preserving Knowledge Distillation. In *Proc. IEEE Int. Conf. Comput. Vis.*, 1365–1374. IEEE.
- Wang, H.; Li, Y.; Wang, Y.; Hu, H.; and Yang, M.-H. 2020. Collaborative Distillation for Ultra-Resolution Universal Style Transfer. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 1860–1869. IEEE.
- Wang, Z.; Zhao, L.; Zuo, Z.; Li, A.; Chen, H.; Xing, W.; and Lu, D. 2023. MicroAST: towards super-fast ultra-resolution arbitrary style transfer. In *Proc. AAAI Conf. Artif. Intell.*, volume 37, 2742–2750.
- Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person Transfer GAN to Bridge Domain Gap for Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 79–88. IEEE.
- Wu, G.; and Gong, S. 2021. Generalising without Forgetting for Lifelong Person Re-identification. In *Proc. AAAI Conf. Artif. Intell.*, 2889–2897.
- Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2016. End-to-End Deep Learning for Person Search. arXiv:1604.01850.
- Xu, K.; Zhang, H.; Li, Y.; Peng, Y.; and Zhou, J. 2024a. Mitigate Catastrophic Remembering via Continual Knowledge Purification for Noisy Lifelong Person Re-Identification. In *ACM Multimedia 2024*.
- Xu, K.; Zou, X.; Peng, Y.; and Zhou, J. 2024b. Distribution-Aware Knowledge Prototyping for Non-Exemplar Lifelong Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 16604–16613. IEEE.
- Xu, K.; Zou, X.; and Zhou, J. 2024. LSTKC: Long Short-Term Knowledge Consolidation for Lifelong Person Re-identification. In *Proc. AAAI Conf. Artif. Intell.*, volume 38, 16202–16210.
- Yang, Z.; Wu, D.; Zhang, W.; Li, B.; and Wang, W. 2023. Handling Label Uncertainty for Camera Incremental Person Re-Identification. In *Proc. ACM Int. Conf. Multimedia*, 6253–6263. ISBN 9798400701085.
- Yin, X.; Shi, J.; Zhang, Y.; Lu, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024. Robust pseudo-label learning with neighbor relation for unsupervised visible-infrared person re-identification. In *Proc. ACM Int. Conf. Multimedia*, 2242–2251.
- Yu, C.; Shi, Y.; Liu, Z.; Gao, S.; and Wang, J. 2023. Lifelong Person Re-Identification via Knowledge Refreshing and Consolidation. In *Proc. AAAI Conf. Artif. Intell.*, 3295–3303.
- Zhao, B.; Tang, S.; Chen, D.; Bilen, H.; and Zhao, R. 2021a. Continual Representation Learning for Biometric Identification. In *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 1197–1207. IEEE.
- Zhao, H.; Tian, M.; Sun, S.; Shao, J.; Yan, J.; Yi, S.; Wang, X.; and Tang, X. 2017. Spindle Net: Person Re-Identification with Human Body Region Guided Feature Decomposition and Fusion. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 907–915. IEEE.
- Zhao, Y.; Zhong, Z.; Yang, F.; Luo, Z.; Lin, Y.; Li, S.; and Sebe, N. 2021b. Learning to Generalize Unseen Domains via Memory-Based Multi-Source Meta-Learning for Person Re-Identification. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 6277–6286. IEEE.
- Zheng, K.; Lan, C.; Zeng, W.; Zhang, Z.; and Zha, Z.-J. 2021. Exploiting Sample Uncertainty for Domain Adaptive Person Re-identification. In *Proc. AAAI Conf. Artif. Intell.*, volume 35, 3538–3546.
- Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable Person Re-Identification: A Benchmark. In *Proc. IEEE Int. Conf. Comput. Vis.*, 1116–1124. IEEE.
- Zhou, J.; Su, B.; and Wu, Y. 2022. Discriminative self-paced group-metric adaptation for online visual identification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(4): 4368–4383.
- Zhou, J.; Yu, P.; Tang, W.; and Wu, Y. 2017. Efficient online local metric adaptation via negative samples for person re-identification. In *Proc. IEEE Int. Conf. Comput. Vis.*, 2420–2428.
- Zhou, K.; Yang, Y.; Qiao, Y.; and Xiang, T. 2024. Mixstyle Neural Networks for Domain Generalization and Adaptation. *Int. J. Comput. Vis.*, 132(3): 822–836.
- Zou, Y.; Yang, X.; Yu, Z.; Kumar, B. V.; and Kautz, J. 2020. Joint Disentangling and Adaptation for Cross-Domain Person Re-Identification. In *Proc. Eur. Conf. Comput. Vis.*, 87–104. Springer.