



SCORE: SubDistribution-Aware Collaborative Knowledge Reinforcing for Cloth-Hybrid Lifelong Person Re-identification

Kunlun Xu¹, Liangyu Ma¹, Jiangmeng Li², Xin Tong³, Xiaode Liu³,
Yufei Guo³✉, and Jiahuan Zhou¹✉

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

{xkl,2300012949}@stu.pku.edu.cn, jiahuanzhou@pku.edu.cn

² University of Chinese Academy of Sciences, Beijing, China

jiangmeng2019@iscas.ac.cn

³ Intelligent Science and Technology Academy, CASIC, Beijing 100041, China

{xin.tong,lxde,yfguo}@pku.edu.cn

Abstract. Lifelong Person Re-Identification (LReID) aims to train a unified person retrieval model from a non-stationary data stream. Existing LReID methods mainly focus on scenarios where the clothing of each person is consistent. Recently, the Cloth-Hybrid LReID (CH-LReID) where cloth-consistent and cloth-changing data alternately occur, has emerged as a more practical and challenging scenario. Due to the conflict between clothing-relevant and clothing-irrelevant knowledge, the well-known catastrophic forgetting problem is significantly exacerbated in this task. To address this issue, we propose a **SubDistribution-aware Collaborative Knowledge REinforcing** (SCORE) framework, where our key idea is explicitly modeling the intra-identity diversity to continually consolidate distinct cloth-consistent and cloth-changing knowledge. Specifically, an Adaptive SubDistribution Modeling mechanism is developed, where a set of distributional subprototypes is assigned to each identity to capture the intra-identity diversity, improving the compatibility between cloth-consistent and cloth-changing knowledge. Then, a Distributional Knowledge Reinforcement scheme is introduced, where the knowledge of old distributional subprototypes is retained in the new ones by a collaborative aligning mechanism. Extensive experiments show that our SCORE achieves the state-of-the-art performance. Our code is available at <https://github.com/zhoujiahuan1991/ECCV2026-SCORE>.

Keywords: Cloth-Hybrid Lifelong Person Re-Identification · SubDistribution Modeling · Knowledge Conflict

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-032-37362-5_6.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
P. Favaro et al. (Eds.): ECCV 2026, LNCS 17049, pp. 93–110, 2026.
https://doi.org/10.1007/978-3-032-37362-5_6

1 Introduction

Person Re-Identification (ReID) [5, 28] aims to match the same person across cameras. Traditional ReID methods focus on static scenarios with stationary cameras and environments [4, 6, 27, 42]. Recently, Lifelong Person Re-Identification (LReID) [23, 31, 35, 47] has attracted increasing research interest by considering that non-stationary data arrives continually, where the catastrophic forgetting of past knowledge is the key challenge. However, existing LReID methods primarily consider that the clothing of each person is consistent [1, 12, 33, 37, 43], as shown in Fig. 1 (a). Since lifelong learning usually involves long-term data collection, changes in people’s clothing inevitably occur from time to time, leading to a practical Cloth-Hybrid LReID (CH-LReID) [2], as shown in Fig. 1 (b).




Fig. 1. (a) Previous LReID methods primarily focus on the cloth-consistent lifelong person re-identification scenario. (b) This paper addresses a more practical and challenging cloth-hybrid lifelong person re-identification problem where the cloth-consistent and cloth-changing training data occur alternatively.

In CH-LReID, the catastrophic forgetting problem is further exacerbated due to knowledge conflicts between training stages. As shown in Fig. 2 (a), at stage $t - 1$, the model learns from a cloth-consistent dataset, in which clothing cues help to identify people. However, when cloth-changing data is introduced at stage t , the inter-cloth matching constraints cause the model to forget the previously learned cloth-consistent knowledge. Later, at stage $t + 1$, the model is trained again on cloth-consistent data, causing the overwriting of cloth-changing knowledge. In addition, the model’s performance on earlier domains remains limited because some old knowledge cannot be reproduced due to domain gaps across stages. This ongoing vanishment of cloth-consistent and cloth-changing knowledge results in severe performance degradation across all past domains.

Although Cloth-Changing ReID methods [7, 9, 11] have been proposed in recent years, the issues in Fig. 2 (a) persist when they are combined with




Fig. 2. (a) Due to the optimization conflict between cloth-consistent and cloth-changing scenarios, the catastrophic forgetting problem emerges as a critical challenge in CH-LReID. (b)(c) We propose to explicitly model the intra-identity distribution diversity, improving the compatibility between the cloth-consistent and cloth-changing knowledge. Thereby, it significantly addresses the catastrophic forgetting problem.

LReID approaches. This is because these methods focus only on improving cloth-changing knowledge acquisition, without addressing the inherent knowledge conflict problem. Moreover, most of them rely on auxiliary information such as clothing labels, gait features, or body shape contours during training. However, obtaining these auxiliary data requires substantial manual effort, which limits their practicality in the lifelong learning scenario, which typically requires efficiency.

To address these limitations, as shown in Fig. 2 (c), we propose to explicitly model the intra-identity distribution diversity, thereby enabling the model to mine and store the cloth-consistent and cloth-changing knowledge simultaneously. As a consequence, inter-stage knowledge conflict can be effectively alleviated, making the catastrophic forgetting problem largely mitigated (as shown in Fig. 2 (b)).

To achieve this, we introduce a SubDistribution-aware **C**ollaborative Knowledge **R**einforcing (SCORE) framework. Firstly, an Adaptive SubDistribution

Modeling (ASD) mechanism is developed, where the learnable distributional subprototypes are introduced to model the intra-identity subdistributions. Then, three optimization paradigms, Prototype Separating, Instance Assembling, and Distribution Purifying, are introduced to guide the prototypes to capture inter-identity subdistributions. To mitigate the forgetting of abundant knowledge embedded in subdistributions from historical stages, a Distributional Knowledge Reinforcement (DKR) scheme is introduced to guide the new data distribution modeling by exploiting historical subdistributions, where a novel subprototype-level knowledge transfer loss is proposed to achieve interactive subdistribution modeling. To better retain the knowledge within the cross-instance representational relations, an Instance-oriented Structural Knowledge Retention (ISK) module is developed to further improve the overall anti-forgetting capacity. Extensive experimental results on the CH-LReID benchmark demonstrate that our approach outperforms existing approaches by large margins. In summary, our contributions are threefold:

(1) To address the cyclic knowledge conflict and amplified catastrophic forgetting in CH-LReID, we propose a novel approach SCORE that explicitly models the intra-identity variations to enhance the compatibility of the cloth-consistent and cloth-changing knowledge.

(2) An Adaptive SubDistribution Modeling mechanism is designed to automatically capture intra-identity subdistributions without relying on auxiliary data. Besides, the Distributional Knowledge Reinforcement scheme is developed to utilize historical subdistributional knowledge to guide new subdistribution learning, effectively achieving knowledge integration.

(3) Extensive experiments demonstrate the superiority of our proposed method in jointly consolidating the cloth-consistent and cloth-changing knowledge compared to the state-of-the-art approaches.

2 Related Work

2.1 Lifelong Person Re-identification

Lifelong Person Re-Identification (LReID) aims to learn from non-stationary data [16, 23, 26], where catastrophic forgetting [15, 29, 32] is the primary challenge. Existing LReID approaches can be divided into two streams. Data replay-based methods [31, 43] mitigate the forgetting of old knowledge by revisiting stored historical data during the learning of new domains. Although this method has achieved notable progress, growing concerns over data privacy [24, 40] have drawn increasing attention to non-exemplar-based approaches. Knowledge distillation-based methods [24, 29, 34] reduce the discrepancy between old and new domains by enforcing output consistency. Some methods, *e.g.*, DKP [36], introduce identity prototypes to reserve historical information. The above LReID methods typically focus on cloth-consistent scenarios where the cross-domain knowledge shares similar semantic clues [23, 37, 40]. However, when applied to cloth-hybrid settings where the cloth-consistent and cloth-changing data show up alternately, the knowledge conflict caused by semantic clue change across domains would further intensify the catastrophic forgetting issue [37].

2.2 Cloth-Changing Person Re-identification

Cloth-changing ReID is a challenging task where the diversity in clothing styles hinders the learning of robust knowledge. To address this issue, some methods focused on disentangling clothing features during inference. Specifically, [3, 41] leverage clothing labels to force the model to rely less on clothing information. Besides, [8, 19] employ a semantic segmentation model to mask clothing information and introduce a consistency constraint between the original image and the clothing-removed image during training. Alternatively, other methods utilize information from additional modalities to learn clothes-irrelevant features, *e.g.*, gait information [13] and silhouette information [39]. However, their reliance on auxiliary information limits their application in practice. Recently, some methods have made notable advancements without introducing auxiliary information. For example, methods [9, 20] employ data augmentation to diversify the colors and textures of clothing in the dataset, thereby reducing the association between identity and specific clothing. However, when faced with datasets that continuously come in a non-stationary manner, these methods suffer from catastrophic forgetting, where previously learned knowledge is overwritten, leading to degraded overall performance on seen domains.

2.3 Distribution Learning

Existing distribution learning methods primarily focus on modeling data uncertainty to address out-of-distribution (OOD) data [18, 21, 44, 46]. In the LReID task, DKP [36] performs instance-level modeling for each sample and aggregates them into prototype distributions, improving the model’s ability to acquire new knowledge. In this paper, targeting the cloth-hybrid scenario, we model multiple subdistributions within each identity to capture diverse discriminative knowledge adaptively, where learnable subprototypes are introduced to model the subdistributions. By collaboratively utilizing the different subdistributions learned from both cloth-consistent and cloth-changing data across stages, we aim to reduce knowledge conflicts and alleviate catastrophic forgetting.

3 The Proposed Method

3.1 Preliminary

In CH-LReID, a stream of T cloth-hybrid training datasets $\mathcal{D} = \{D_t\}_{t=1}^T$ is provided sequentially. Specifically, given a dataset D_t containing N_t identities, the included data may be either cloth-consistent or cloth-changing, but the specific type is unknown in advance. Besides, for each identity, it is also not known beforehand whether they appear in multiple outfits. Furthermore, previously seen $t - 1$ datasets are not accessible at the training stage t .




Fig. 3. The overview of our SCORE approach. (a) Adaptive SubDistribution Modeling (ASD) captures intra-identity diversity and encodes it into distributional subprototypes. (b) Distributional Knowledge Reinforcement (DKR) performs distributional alignment to guide new distribution learning with historical distribution. (c) Instance-oriented Structural Knowledge Retention (ISK) integrates cross-instance relational knowledge across domains to improve the old knowledge consolidation.

3.2 Overview

As shown in Fig. 3, at the t -th training stage, given the new dataset D_t , we introduce k learnable subprototypes in the feature space for each identity, denoted as $\mathcal{P}_t = \{(p_i^t, y_i^t)\}_{i=1}^{k \times N_t}$. Specifically, p_i^t is a multi-dimensional Gaussian distribution, and y_i^t is its corresponding identity label. The center of p_i^t is denoted as $c_i^t \in \mathbb{R}^d$, and $v_i^t \in \mathbb{R}^d$ indicates the standard deviation vector derived from a diagonal covariance matrix $\Sigma_i^t \in \mathbb{R}^{d \times d}$ of p_i^t . Besides, we define $\mathcal{C}^t = \{c_i^t\}_{i=1}^{k \times N_t}$ and $\mathcal{V}^t = \{v_i^t\}_{i=1}^{k \times N_t}$ as the overall center set and standard deviation set at stage t , respectively.

The Adaptive SubDistribution Modeling (ASD) module aims to capture the intra-identity subdistributions via an adaptive distributional subprototype optimization mechanism. Besides, the Distributional Knowledge Reinforcement (DKR) module utilizes the preserved historical distributional information to guide the new subdistribution learning. Furthermore, the Instance-oriented Structural Knowledge Retention (ISK) module integrates fine-grained cross-instance relational knowledge from the previous stage into the new model to enhance the historical knowledge consolidation capacity.

3.3 Adaptive SubDistribution Modeling

Given a batch of input data $\{(x_i, y_i)\}_{i=1}^B \in D_t$, where B denotes the batch size, suppose there are n distinct identities in this batch, and let the corresponding identity label set be $\mathcal{Y} = \{y_i\}_{i=1}^n$. The model backbone $\mathcal{M}_t$ extracts visual features from the inputs, producing a set of feature representations $\mathcal{F} = \{f_i\}_{i=1}^B$.

To enhance the discriminative capability of the model, we employ the cross-entropy loss L_{ce} [37] and the triplet loss L_{tri} [29, 34] as the base supervision:

$$\begin{cases} \mathcal{L}_{ce-s} = -\sum_{i=1}^B y_i \log \rho(\mathbf{W}_t f_i) \\ \mathcal{L}_{tri-s} = \log \left(1 + \exp \left(\|\tilde{f}_a - \tilde{f}_p\|_2^2 - \|\tilde{f}_a - \tilde{f}_n\|_2^2 \right) \right) \end{cases}, \quad (1)$$

where $\mathbf{W}_t$ denotes the linear projection parameter to map the features to logits, ρ represents the softmax function, $\tilde{f}$ is the L_2 -normalized version of the feature, and $< a, p, n >$ is a triplet set. Therefore, the basic loss function $\mathcal{L}_{base}$ in LReID is formulated as:

$$\mathcal{L}_{base} = \mathcal{L}_{ce-s} + \mathcal{L}_{tri-s}. \quad (2)$$

In the following, we introduce the mechanism of subdistribution Optimization:

Prototype Separating: Prototype Separating aims to ensure the clustering of subprototypes of the same identity and separate the subprototypes of different identities. Specifically, for each input batch, we retrieve the corresponding prototypes based on $\mathcal{Y}$, denoted as $\mathcal{P}^* = \{(p_i^t, y_i^t)\}_{i=1}^{k \times n}$, where $y_i^t \in \mathcal{Y}$. To ensure the coherence and the distinction among prototypes of the same and different identities respectively, we apply the cross-entropy loss and the triplet loss to $\mathcal{P}^*$, obtaining a prototype separating loss:

$$\mathcal{L}_p = \mathcal{L}_{ce-p} + \mathcal{L}_{tri-p}, \quad (3)$$

Instance Assembling: Instance Assembling aims to assign each instance to a subprototype and encourage the instance feature to approach the corresponding prototype. Since different outfits of the same identity may vary significantly in clothing style, color, and appearance cues, the knowledge encoded by each outfit can differ substantially. To capture this intra-identity diversity, each subprototype is modeled as a distribution, enabling different subprototypes to represent varying degrees of intra-identity variability. To measure how well an instance feature matches each subprototype, we compute the maximum likelihood of the instance feature under every subprototype. The resulting likelihood matrix $\mathcal{M} \in \mathbb{R}^{B \times kN_t}$ is formulated as:

$$\mathcal{M}_{i,j} = \frac{1}{(2\pi)^{d/2} |\Sigma_j^t|^{1/2}} \exp \left(-\frac{1}{2} (f_i - c_j^t)^T (\Sigma_j^t)^{-1} (f_i - c_j^t) \right). \quad (4)$$

where $i \in [1, B]$ and $j \in [1, kN_t]$.

Then, we pull each instance feature toward the most similar subprototype of the same identity. Accordingly, the instance assembling loss $\mathcal{L}_a$ is formulated as:

$$\mathcal{L}_a = \frac{1}{B} \sum_{i=1}^B (f_i - c_i^*)^2, \quad (5)$$

where c_i^* denotes the center of the most similar subprototype that shares the same identity label with f_i .

Distribution Purifying: Distribution Purifying aims to ensure that each subdistribution exclusively models instances belonging to its corresponding identity. To achieve this, we first obtain the likelihood of the most compatible subprototypes from the likelihood matrix $\mathcal{M}$ for each instance, denoted as $\mathbf{m} \in \mathbb{R}^B$. A corresponding correctness indicator $\mathbf{z} \in \{0, 1\}^B$ is also obtained which reflects whether the matched subprototype shares the same identity label:

$$\begin{cases} \mathbf{m}_i = \max_j \mathcal{M}_{i,j} \\ \mathbf{z}_i = \mathbb{I}(y_i = y_{\arg \max_j \mathcal{M}_{i,j}}^t) \end{cases}, \quad (6)$$

where $\mathbb{I}$ is an indicator function. Then, we design a distribution purifying loss $\mathcal{L}_d$ which is computed as follows:

$$\mathcal{L}_d = -\frac{1}{B} \sum_{i=1}^B [\mathbf{z}_i \log(\rho(\mathbf{m}_i)) + (1 - \mathbf{z}_i) \log(1 - \rho(\mathbf{m}_i))]. \quad (7)$$

Overall SubDistribution Optimization Loss: Prototype Separating, Instance Assembling, and Distribution Purifying mechanisms are complementary to each other since they perform inter-prototype, instance-centric, and prototype-centric SubDistribution evolution, respectively. Therefore, we compute the overall subdistribution optimization Loss as follows:

$$\mathcal{L}_{sub} = \mathcal{L}_p + \mathcal{L}_a + \mathcal{L}_d. \quad (8)$$

By incorporating the above constraints, our method encodes intra-identity differentiated knowledge in diverse subprototypes. Note that even for the cloth-consistent dataset, the subdistributions are applicable since all the distributions could be approximated as Gaussian mixtures [22].

3.4 Distributional Knowledge Reinforcement

In this section, we propose a collaborative aligning mechanism, which leverages the accumulated discriminative knowledge from the previous stage to guide the new subdistribution learning, thereby improving the coexistence of cloth-consistent and cloth-changing knowledge.

Specifically, as shown in Fig. 3, when training in the new t -th stage, we sample s subprototype features $\mathcal{G}_{t-1} = \{g_i^{t-1}\}_{i=1}^{N_s^{t-1}}$ from all subprototypes in the $\mathcal{P}_{t-1}$ based on the distribution $\mathcal{N}(c_i^{t-1}, (v_i^{t-1})^2)$, where $N_s^{t-1} = s \times k \times N_{t-1}$. It is worth noting that, compared to directly using the subprototype centers, sampling from subdistributions allows for better utilization of the abundant distributional knowledge learned by each subprototype.

For the $\mathcal{P}_t$ that are currently adaptively learnable, we only select the subprototypes that can be correctly matched according to Eq. 6 and also sample s

subprototype features $\mathcal{G}_t = \{g_i^t\}_{i=1}^{N_s^t}$, where $N_s^t = s \times n_c$ and n_c denotes the number of correctly matched subprototypes. This is because the correctly matched subprototypes reflect the information that can currently accurately describe the identity, whereas the incorrectly matched subprototypes indicate confusion and bias in the knowledge they represent.

We convert $\mathcal{G}_{t-1}$ and $\mathcal{G}_t$ into matrix forms $\mathbf{G}_{t-1} \in \mathbb{R}^{N_s^{t-1} \times d}$ and $\mathbf{G}_t \in \mathbb{R}^{N_s^t \times d}$, respectively. Then, a knowledge shift matrix $\mathbf{K}_s \in \mathbb{R}^{N_s^t \times N_s^{t-1}}$ for the subprototypes is obtained by:

$$\mathbf{K}_s = \boldsymbol{\rho}(\mathbf{G}_t \mathbf{G}_{t-1}^\top / \lambda_1), \quad (9)$$

where the softmax function $\boldsymbol{\rho}$ is applied across each row, and λ_1 serves as a temperature parameter to adjust the sharpness of the matrix values. To reduce the knowledge conflict caused by cloth-consistent and cloth-changing data across different stages, we perform distillation on the knowledge shift matrix $\mathbf{K}_s$ to collaboratively leverage both new and old knowledge. Accordingly, the subprototype-level knowledge transfer loss is defined as follows:

$$\mathcal{L}_t = \mathcal{KL}(\boldsymbol{\rho}(\mathbf{K}_s \mathbf{K}_s^\top / \lambda_2) \parallel \boldsymbol{\rho}(\mathbf{G}_t \mathbf{G}_t^\top / \lambda_2)), \quad (10)$$

where $\mathcal{KL}$ denotes the Kullback-Leibler (KL) divergence and λ_2 is another temperature parameter.

3.5 Instance-Oriented Structural Knowledge Retention

The DKR module facilitates the collaborative use of overall knowledge from both the new and old stages from the perspective of distributions. However, it lacks the integration of cross-instance fine-grained knowledge. To address this, inspired by the existing works [37], we propose an Instance-oriented Structural Knowledge Retention module. First, we use the old backbone $\phi_{t-1}(\cdot)$, which was already trained in the $t - 1$ stage, to extract features from the current input batch, obtaining the representations in the old feature space denoted as $\mathbf{F}_o \in \mathbb{R}^{B \times d}$. Then, to preserve fine-grained old knowledge, we maximize the structural consistency of the input between the new and old feature spaces as follows:

$$\mathcal{L}_s = \mathcal{KL}(\boldsymbol{\rho}(\mathbf{F}_o \mathbf{F}_o^\top / \lambda_2) \parallel \boldsymbol{\rho}(\mathbf{F} \mathbf{F}^\top / \lambda_2)), \quad (11)$$

where $\mathbf{F} \in \mathbb{R}^{B \times d}$ is the matrix form of $\mathcal{F}$. Note that $\boldsymbol{\rho}(\mathbf{F}_o \mathbf{F}_o^\top / \lambda_2)$ and $\boldsymbol{\rho}(\mathbf{F} \mathbf{F}^\top / \lambda_2)$ both reflect the cross-instance representational relations, preserving the structured knowledge in the feature space. By guiding the model to focus on this structural consistency, the consolidation of fine-grained knowledge is significantly improved.

3.6 Training and Inference

During training, the overall loss $\mathcal{L}$ for SCORE can be computed as follows:

$$\mathcal{L} = \mathcal{L}_{base} + \alpha \mathcal{L}_{sub} + \beta (\mathcal{L}_t + \mathcal{L}_s), \quad (12)$$

where α and β are hyperparameters. Since both $\mathcal{L}_t$ and $\mathcal{L}_s$ involves old knowledge, they share the same weight for simplicity.

After training at the t -th stage, we fuse the newly learned model $\mathbf{M}_t$ with the old model $\mathbf{M}_{t-1}$ as a post-processing step following [36]:

$$\mathbf{M}_t \leftarrow \frac{\mathbf{M}_t + \mathbf{M}_{t-1}}{2}. \quad (13)$$

During the inference phase, only the backbone is exploited to extract image features for person retrieval.

4 Experiments

4.1 Datasets and Evaluation Protocol

Following previous works [2], we evaluate the effectiveness of our method on the CH-LReID benchmark [2], which consists of three cloth-consistent datasets (Market-1501 [45], MSMT17-V2 [30] and CUHK03 [17]), as well as two cloth-changing datasets (LTCC [25] and PRCC [39]). Following the previous works [2], we conduct training with two distinct dataset orders: (Order-1) Market $\rightarrow$ LTCC $\rightarrow$ PRCC $\rightarrow$ MSMT17 $\rightarrow$ CUHK03 and (Order-2) MSMT17 $\rightarrow$ PRCC $\rightarrow$ Market $\rightarrow$ CUHK03 $\rightarrow$ LTCC.

Following previous works [2, 23, 36], we adopt Rank-1 accuracy (R@1) and mean Average Precision (mAP) as evaluation metrics. To comprehensively verify the model’s performance, both cloth-consistent and cloth-changing setting are evaluated.

4.2 Implementation Details

To ensure a fair comparison with existing methods [2, 36, 37], we adopt ResNet50 [10] as the backbone network. For both training orders, the first dataset is trained for 80 epochs, while each of the subsequent datasets is trained for 60 epochs. We use a mini-batch size of 128, constructed by sampling 32 identities with 4 images per identity. Each image is resized to 256×128 during training. The model is optimized using SGD with a learning rate of 0.008 and a weight decay of 0.0001. Additionally, the temperature parameters λ_1 and λ_2 are set to 0.8. All experiments are conducted on a single NVIDIA 4090 GPU.

4.3 Comparison with State-of-the-Art Methods

We compare our proposed SCORE method against the current state-of-the-art CH-LReID method DKC [2]. Besides, we also compare with LReID methods, including LwF [14], AKA [23], PatchKD [29], LSTKC [37], USP [38], DKP [36] and DASK [33]. In addition, we incorporate the state-of-the-art cloth-changing ReID method DSIFLF [19], with the anti-forgetting strategy of LReID. The

Table 1. Performance comparison on Training Order-1: Market $\rightarrow$ LTCC $\rightarrow$ PRCC $\rightarrow$ MSMT17 $\rightarrow$ CUHK03.

Method	Cloth-Consistent											Cloth-Changing				Overall				
	Market	LTCC	PRCC	MSMT17	CUHK03	Average	LTCC	PRCC	Average	Average	LTCC	PRCC	Average	Average						
	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	R@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1						
JointTrain	64.1	82.5	42.6	62.1	94.6	98.7	18.4	40.8	44.4	46.4	52.8	66.1	10.1	23.0	32.7	33.8	21.4	28.4	43.8	55.3
SFT	28.5	52.0	28.5	49.3	92.5	97.3	7.0	19.6	44.0	45.6	40.1	52.8	6.9	15.3	21.8	21.2	14.4	18.3	32.7	42.9
LwF[14]	44.5	65.8	21.6	40.0	87.4	91.3	4.0	11.6	25.5	25.0	36.6	46.7	5.9	12.5	25.9	26.7	15.9	19.6	30.7	39.0
AKA[23]	48.0	69.5	25.4	45.1	88.1	93.3	4.2	12.0	31.2	31.2	39.4	50.2	6.5	12.8	26.5	26.7	16.5	19.8	32.8	41.5
PatchKD[29]	68.0	85.5	30.8	54.7	93.5	96.5	5.7	15.6	33.2	32.9	46.2	57.0	7.2	17.9	26.1	26.0	16.7	22.0	37.8	47.0
LSTKC[37]	39.9	63.4	39.6	65.4	95.9	98.9	11.5	29.2	48.1	50.1	47.0	61.4	8.3	19.4	24.0	22.9	16.2	21.2	38.2	49.9
USP[38]	<u>66.3</u>	74.9	37.2	49.6	88.8	92.6	6.6	14.6	38.7	33.9	47.5	53.1	7.8	17.3	25.7	23.6	16.8	20.5	38.7	43.8
DKP[36]	51.3	73.0	44.2	68.6	98.4	<u>99.4</u>	13.5	31.6	39.5	40.6	49.4	62.6	9.3	21.4	35.5	34.8	22.4	28.1	41.7	52.8
DASK[33]	56.8	78.6	32.7	<u>72.4</u>	<u>98.5</u>	<u>99.4</u>	25.5	52.8	<u>42.4</u>	44.4	51.2	<u>69.5</u>	11.3	26.5	33.6	31.6	22.5	29.0	43.0	<u>57.9</u>
DSIFLF [†] [19]	47.5	69.2	42.1	<u>65.5</u>	97.0	97.8	11.0	26.4	39.7	41.6	47.5	60.1	8.9	17.9	34.0	31.7	21.5	24.8	40.0	50.0
DKC[2]	57.9	78.5	<u>49.2</u>	<u>72.4</u>	98.6	99.5	16.2	36.7	40.7	42.4	<u>52.5</u>	65.9	10.1	25.5	36.9	<u>35.5</u>	<u>23.5</u>	<u>30.5</u>	<u>44.2</u>	<u>55.8</u>
SCORE	63.1	<u>82.9</u>	58.2	78.4	98.7	99.5	<u>23.5</u>	<u>50.2</u>	44.0	<u>45.2</u>	57.5	71.2	<u>11.0</u>	<u>25.8</u>	<u>36.2</u>	36.4	23.6	31.1	47.8	59.7

Table 2. Performance comparison on Training Order-2: MSMT17 $\rightarrow$ PRCC $\rightarrow$ Market-1501 $\rightarrow$ CUHK03 $\rightarrow$ LTCC.

Method	Cloth-Consistent											Cloth-Changing					Overall			
	MSMT17	PRCC	Market	CUHK03	LTCC	Average	PRCC	LTCC	Average	Average										
	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1	mAPR@1										
JointTrain	18.4	40.8	94.6	98.7	64.1	82.5	44.4	46.4	42.6	62.1	52.8	66.1	32.7	33.8	10.1	23.0	21.4	28.4	43.8	55.3
SFT	2.4	9.1	84.5	95.9	16.9	37.2	8.2	8.8	37.9	57.6	30.0	41.7	21.9	24.3	9.1	20.9	15.5	22.6	25.8	36.3
LwF[14]	12.7	28.8	91.3	95.5	25.2	49.0	8.6	8.0	41.9	63.3	35.9	48.9	27.9	28.3	8.7	18.4	18.3	23.4	30.9	41.6
AKA[23]	15.0	32.4	93.0	97.0	26.6	51.0	13.4	11.6	43.4	64.0	38.3	51.2	28.2	27.7	8.4	15.8	18.3	21.8	32.6	42.8
PatchKD[29]	<u>22.6</u>	47.9	95.0	98.5	38.2	64.2	21.8	23.1	45.8	70.0	44.7	60.7	28.1	27.6	9.3	21.7	18.7	24.7	37.3	50.4
LSTKC[37]	6.9	20.4	92.8	97.0	39.8	64.3	25.2	26.0	51.1	65.5	43.2	54.6	27.3	27.0	10.9	22.4	19.1	24.7	36.3	46.1
USP[38]	22.7	37.2	92.5	95.9	37.8	53.5	17.3	14.6	<u>63.5</u>	72.4	46.8	54.7	28.7	29.4	11.9	26.0	20.3	27.7	39.2	47.0
DKP[36]	11.8	29.0	96.7	98.8	46.0	69.8	29.9	30.9	55.8	75.8	48.0	60.9	34.1	34.6	11.6	24.2	22.9	29.4	40.8	51.9
DASK[33]	13.3	35.0	<u>98.2</u>	<u>99.4</u>	51.3	<u>74.3</u>	26.1	25.5	35.3	72.4	44.8	61.3	34.8	31.9	<u>13.5</u>	<u>27.8</u>	24.1	29.8	38.9	52.3
DSIFLF [†] [19]	8.6	22.4	96.7	97.2	50.0	71.4	25.8	25.7	57.4	74.8	47.7	58.3	32.8	30.3	12.8	23.2	22.1	26.8	40.6	49.3
DKC[2]	15.4	35.3	98.1	99.0	<u>51.8</u>	72.7	<u>31.9</u>	<u>32.1</u>	57.2	77.9	<u>50.9</u>	<u>63.4</u>	38.2	37.9	13.0	27.0	25.6	<u>32.5</u>	<u>43.7</u>	<u>54.6</u>
SCORE	<u>19.9</u>	<u>44.5</u>	98.9	99.5	55.3	77.3	34.5	35.1	65.0	81.1	54.7	67.5	<u>36.5</u>	<u>35.4</u>	<u>14.4</u>	32.7	<u>25.4</u>	34.0	46.4	57.9

resulting model is named DSIFLF[†]. Moreover, we include both sequential fine-tuning and joint training settings. The best and second-best results are highlighted in **bold** and underline, respectively.

Compared on the Cloth-Consistent Data: As shown in Table 1 and Table 2, our SCORE outperforms all existing methods under the Cloth-Consistent. Specifically, on Training Order-1, SCORE surpasses the state-of-the-art DASK and DKC with **5.0%** and **1.7%** Average mAP and R@1 improvements, respectively. Furthermore, on Training Order-2, SCORE outperforms the state-of-the-art DKC with **3.8%/4.1%** Average mAP/R@1 improvements. These results demonstrate the effectiveness of our method in the CH-LReID scenario, where it can distinguish individuals by effectively integrating both clothing-related and clothing-irrelevant knowledge learned across different training stages, rather than suppressing the use of clothing information.

These advantages arise from our SubDistribution-aware Collaborative Knowledge Reinforcing mechanism, which effectively alleviates the conflict between cloth-changing and cloth-consistent knowledge, thereby improving the cross-domain knowledge accumulation capacity. In contrast, existing LReID works, *e.g.*, DASK, suffer from performance limitations as they discard clothing information in the feature space for each identity, which inevitably hinders the effective use of valuable clothing-related cues. The existing CH-LReID method, DKC, though considered estimating the subdistributions, it only exploited the statistical distribution and neglected the dynamic modeling subdistributions. Thus, it suffered from limited representation and utilization of inter-identity diversity.

Compared on the Cloth-Changing Data: As shown in Table 1, SCORE outperforms state-of-the-art DKC by **0.1%/0.6%** under Average mAP/R1 on the Cloth-Changing data under Training-Order-1. Besides, under Training-Order-2, our SCORE surpasses DKC by **1.5%** Average R@1 and obtains comparable performance on Average mAP. This indicates that, for cloth-changing data, our approach effectively captures and preserves the identity-specific differentiated knowledge introduced by variations in clothing. By collaboratively leveraging diverse discriminative knowledge, it mitigates the conflict induced by data type switching.

Compared on the Overall Performance: As shown in Table 1 and Table 2, our SCORE achieves state-of-the-art overall performance compared to the existing approaches. Specifically, SCORE outperforms the existing methods by a minimal improvement of **3.6%/1.8%** and **2.7%/3.3%** under Training Order-1 and Training-Order-2, respectively. These results arise from the proposed SubDistribution-aware Collaborative Knowledge Reinforcing paradigm that effectively improves the compatibility between the Cloth-Changing and Cloth-Consistent knowledge.




Fig. 4. Performance tendency at different training stages.

Performance Tendency in Cloth-Hybrid Scenario: Figure 4 illustrates the overall performance tendency of different methods on seen domains. The results show that our SCORE, although exhibiting comparable performance with existing methods, shows a consistent improvement over existing approaches. This highlights that our method is capable of consistent learning in cloth-hybrid scenarios, and achieves superior anti-forgetting performance by integrating diverse discriminative knowledge to alleviate knowledge conflicts.

Table 3. Ablation study of different components.

.9Base	.9ASD	.9ISK	.9DKR	Cloth-Consist.		Cloth-Chang.	
				mAP	R@1	mAP	R@1
✓				40.1	52.8	14.4	18.3
✓	✓			50.1	63.9	22.3	27.6
✓	✓	✓		52.8	66.3	23.0	29.6
✓	✓	✓	✓	57.5	71.2	23.6	31.1

4.4 Ablation Study

Ablations on Model Components. As shown in Table 3, we adopt Sequential Finetuning (SFT) as the baseline and incrementally incorporate the ASD, ISK, and DKR modules. When ASD is adopted, the model achieves **10.0%/11.1%** and **7.9%/9.3%** improvements in mAP/R@1 on Cloth-Consistent and Cloth-Changing data, respectively. This improvement is attributed to the SubDistribution modeling mechanism, which enhances the compatibility between Cloth-Consistent and Cloth-Changing knowledge, thereby mitigating the catastrophic forgetting issue. Furthermore, as ISK is introduced, additional **2.7%/2.4%** and **0.7%/2.0%** improvements under mAP/R@1 metrics on Cloth-Consistent and Cloth-Changing data are achieved. This is because ISK transfers the inter-instance structure knowledge, which is complementary to the clothing knowledge. Finally, incorporating DKR yields additional gains of **4.7%/4.9%** and **0.6%/1.5%**, resulting in total improvements of **17.4%/18.4%** and **9.2%/12.8%**. These gains stem from explicitly leveraging historical distributions to guide new model learning, further alleviating conflicts between Cloth-Consistent and Cloth-Changing knowledge.

Note that the performance improvement of the Cloth-Changing data is less significant on the Cloth-Consistent data. This is because the Cloth-Changing scenario is inherently more challenging, making performance gains more difficult to achieve.

Ablations on Hyperparameters. We investigate the influence of the hyperparameters α, β, s, k , and present the results in Fig. 6 and Fig. 5. α and β are the loss weights to emphasize the subdistribution modeling and historical knowledge utilization, respectively. The proper α and β can facilitate model learning, while too large values can lead to overfitting and limited robustness. Similarly, a proper predefined subdistribution number per class k and instance sample number s can promote model learning, while a too large k and s can cause optimization instability and overfitting. In this paper, we empirically set α, β, k, s to 0.75, 0.50, 5 and 1 by default.




Fig. 5. Ablation studies on feature sample number s .




Fig. 6. Ablation studies on hyperparameters under training Order-1. The default values are highlighted by the dashed lines.

4.5 Visualization Results

We visualize the features of Cloth-Consistent data (MSMT17) and Cloth-Changing data (LTCC) in comparison with the state-of-the-art DKC method, as shown in Fig. 7. The results indicate that our approach achieves better feature separation across identities on both types of data. The results arise because DKC adopts unlearnable distributions, suffering from limited knowledge representation capacity. In contrast, our SCORE model introduces an adaptive subdistribution learning and reinforcement paradigm, effectively enhancing the consolidation of Cloth-Consistent and Cloth-Changing knowledge.




Fig. 7. The t-SNE visualization of cloth-consistent and cloth-changing training data.

5 Conclusion

In this paper, we focus on the challenging Cloth-Hybrid Lifelong Person Re-Identification (CH-LReID) task and introduce a novel SubDistribution-aware Collaborative Knowledge REinforcing (SCORE) approach. To address the deteriorated forgetting caused by the cloth-consistent and cloth-changing knowledge conflict, we propose to explicitly model the intra-identity diversity to improve the compatibility of both knowledge. To achieve this, an Adaptive SubDistribution Modeling mechanism is developed, where a set of distributional subprototypes is assigned to each identity to capture the intra-identity diversity. Besides, a Distributional Knowledge Reinforcement scheme is introduced, where the knowledge of old distributional subprototypes is retained in the new ones by a collaborative aligning mechanism. Extensive experiments demonstrate that our SCORE achieves the state-of-the-art performance. Our approach underscores the potential of automatic subdistribution modeling in improving conflict knowledge integration.

Acknowledgements. This work was a research achievement of National Engineering Research Center of New Electronic Publishing Technologies, and was supported by the National Natural Science Foundation of China (62376011, 62406313), the National Key R&D Program of China (2024YFA1410000), and the Postdoctoral Fellowship Program of China Postdoctoral Science Foundation, Grant No. YJB20250283.

References

1. Cui, Z., Peng, Y., Wang, X., Zhu, M., Zhou, J.: Continual vision-language retrieval via dynamic knowledge rectification. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, pp. 11704–11712 (2024)
2. Cui, Z., Zhou, J., Peng, Y.: DKC: differentiated knowledge consolidation for cloth-hybrid lifelong person re-identification. In: Proceedings of the Computer Vision and Pattern Recognition Conference, pp. 3573–3582 (2025)
3. Cui, Z., Zhou, J., Peng, Y., Zhang, S., Wang, Y.: DCR-reid: deep component reconstruction for cloth-changing person re-identification. IEEE Trans. Circuits Syst. Video Technol. **33**(8), 4415–4428 (2023)

4. Gong, Y., Huang, L., Chen, L.: Person re-identification method based on color attack and joint defence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4313–4322 (2022)
5. Gong, Y., Zhong, Z., Qu, Y., Luo, Z., Ji, R., Jiang, M.: Cross-modality perturbation synergy attack for person re-identification. In: The Thirty-Eighth Annual Conference on Neural Information Processing Systems
6. Gong, Y., Zhong, Z., Qu, Y., Luo, Z., Ji, R., Jiang, M.: Cross-modality perturbation synergy attack for person re-identification. In: NeurIPS, vol. 37, pp. 23352–23377 (2024)
7. Gu, X., Chang, H., Ma, B., Bai, S., Shan, S., Chen, X.: Clothes-changing person re-identification with RGB modality only. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1060–1069 (2022)
8. Guo, P., Liu, H., Wu, J., Wang, G., Wang, T.: Semantic-aware consistency network for cloth-changing person re-identification. In: Proceedings of the 31st ACM International Conference on Multimedia, pp. 8730–8739 (2023)
9. Han, K., Gong, S., Huang, Y., Wang, L., Tan, T.: Clothing-change feature augmentation for person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 22066–22075 (2023)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
11. Huang, Y., Wu, Q., Xu, J., Zhong, Y., Zhang, Z.: Clothing status awareness for long-term person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11895–11904 (2021)
12. Huang, Z., et al.: Lifelong unsupervised domain adaptive person re-identification with coordinated anti-forgetting and adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14288–14297 (2022)
13. Jin, X., et al.: Cloth-changing person re-identification from a single image with gait prediction and regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14278–14287 (2022)
14. Li, D., Chen, X., Zhang, Z., Huang, K.: Learning deep context-aware features over body and latent parts for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 384–393 (2017)
15. Li, Q., Peng, Y., Zhou, J.: Progressive prototype evolving for dual-forgetting mitigation in non-exemplar online continual learning. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp. 2477–2486 (2024)
16. Li, Q., Xu, K., Peng, Y., Zhou, J.: Exemplar-free lifelong person re-identification via prompt-guided adaptive knowledge consolidation. *Int. J. Comput. Vision* **132**(11), 4850–4865 (2024)
17. Li, W., Zhao, R., Xiao, T., Wang, X.: Deepreid: deep filter pairing neural network for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 152–159 (2014)
18. Li, Y., Duan, Y., Kuang, Z., Chen, Y., Zhang, W., Li, X.: Uncertainty estimation via response scaling for pseudo-mask noise mitigation in weakly-supervised semantic segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, pp. 1447–1455 (2022)
19. Li, Y., Cheng, D., Fang, C., Jiao, C., Wang, N., Gao, X.: Disentangling identity features from interference factors for cloth-changing person re-identification. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp. 2252–2261 (2024)

20. Liu, F., Ye, M., Du, B.: Cloth-aware augmentation for cloth-generalized person re-identification. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp. 4053–4062 (2024)
21. Lu, F., Zhu, K., Zhai, W., Zheng, K., Cao, Y.: Uncertainty-aware optimal transport for semantically coherent out-of-distribution detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3282–3291 (2023)
22. McLachlan, G.J., Rathnayake, S.: On the number of components in a gaussian mixture model. *Wiley Interdisc. Rev. Data Min. Knowl. Discov.* **4**(5), 341–355 (2014)
23. Pu, N., Chen, W., Liu, Y., Bakker, E.M., Lew, M.S.: Lifelong person re-identification via adaptive knowledge accumulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7901–7910 (2021)
24. Pu, N., Liu, Y., Chen, W., Bakker, E.M., Lew, M.S.: Meta reconciliation normalization for lifelong person re-identification. In: Proceedings of the 30th ACM International Conference on Multimedia, pp. 541–549 (2022)
25. Qian, X., et al.: Long-term cloth-changing person re-identification. In: Proceedings of the Asian Conference on Computer Vision (2020)
26. Shi, J., et al.: Multi-memory matching for unsupervised visible-infrared person re-identification. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) ECCV 2024. LNCS, vol. 15076, pp. 456–474. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-72649-1_26
27. Shi, J., Yin, X., Zhang, Y., Xie, Y., Qu, Y., et al.: Learning commonality, divergence and variety for unsupervised visible-infrared person re-identification. *Adv. Neural. Inf. Process. Syst.* **37**, 99715–99734 (2024)
28. Shi, J., et al.: Dual pseudo-labels interactive self-training for semi-supervised visible-infrared person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11218–11228 (2023)
29. Sun, Z., Mu, Y.: Patch-based knowledge distillation for lifelong person re-identification. In: Proceedings of the 30th ACM International Conference on Multimedia, pp. 696–707 (2022)
30. Wei, L., Zhang, S., Gao, W., Tian, Q.: Person transfer GAN to bridge domain gap for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 79–88 (2018)
31. Wu, G., Gong, S.: Generalising without forgetting for lifelong person re-identification. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, pp. 2889–2897 (2021)
32. Xu, K., Cheng, H., Li, J., Zou, X., Zhou, J.: Vision-language attribute disentanglement and reinforcement for lifelong person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 40397–40406 (2026)
33. Xu, K., Jiang, C., Xiong, P., Peng, Y., Zhou, J.: Dask: distribution rehearsing via adaptive style kernel learning for exemplar-free lifelong person re-identification. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, pp. 8915–8923 (2025)
34. Xu, K., Zhang, H., Li, Y., Peng, Y., Zhou, J.: Mitigate catastrophic remembering via continual knowledge purification for noisy lifelong person re-identification. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp. 5790–5799 (2024)

35. Xu, K., Zhuo, F., Li, J., Zou, X., Zhou, J.: Self-reinforcing prototype evolution with dual-knowledge cooperation for semi-supervised lifelong person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3564–3574 (2025)
36. Xu, K., Zou, X., Peng, Y., Zhou, J.: Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16604–16613 (2024)
37. Xu, K., Zou, X., Zhou, J.: LSTKC: long short-term knowledge consolidation for lifelong person re-identification. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, pp. 16202–16210 (2024)
38. Yan, Y., Yu, H., Wang, Y., Song, S., Huang, W., Jin, J.: Unified stability and plasticity for lifelong person re-identification in cloth-changing and cloth-consistent scenarios. *IEEE Trans. Circuits Syst. Video Technol.* (2024)
39. Yang, Q., Wu, A., Zheng, W.S.: Person re-identification by contour sketch under moderate clothing change. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**(6), 2029–2046 (2019)
40. Yang, Z., Wu, D., Zhang, W., Li, B., Wang, W.: Handling label uncertainty for camera incremental person re-identification. In: Proceedings of the 31st ACM International Conference on Multimedia, pp. 6253–6263 (2023)
41. Yang, Z., Lin, M., Zhong, X., Wu, Y., Wang, Z.: Good is bad: causality inspired cloth-debiasing for cloth-changing person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1472–1481 (2023)
42. Yin, X., et al.: Robust pseudo-label learning with neighbor relation for unsupervised visible-infrared person re-identification. In: Proceedings of the 32nd ACM International Conference on Multimedia, pp. 2242–2251 (2024)
43. Yu, C., Shi, Y., Liu, Z., Gao, S., Wang, J.: Lifelong person re-identification via knowledge refreshing and consolidation. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, pp. 3295–3303 (2023)
44. Yu, T., Li, D., Yang, Y., Hospedales, T.M., Xiang, T.: Robust person re-identification by modelling feature uncertainty. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 552–561 (2019)
45. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: a benchmark. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1116–1124 (2015)
46. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *Int. J. Comput. Vision* **129**(4), 1106–1120 (2021)
47. Zhou, J., Xu, K., Zhuo, F., Zou, X., Peng, Y.: Distribution-aware knowledge aligning and prototyping for non-exemplar lifelong person re-identification. *IEEE Trans. Pattern Anal. Mach. Intell.* (2025)