

Self-Reinforcing Prototype Evolution with Dual-Knowledge Cooperation for Semi-Supervised Lifelong Person Re-Identification

Kunlun Xu¹ Fan Zhuo¹ Jiangmeng Li² Xu Zou³ Jiahuan Zhou^{1*}

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

² University of Chinese Academy of Sciences, Beijing, China

³ School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan, China

xkl@stu.pku.edu.cn fanzhuo@stu.xidian.edu.cn jiangmeng2019@iscas.ac.cn zoux@hust.edu.cn jiahuanzhou@pku.edu.cn

Abstract

Current lifelong person re-identification (LReID) methods predominantly rely on fully labeled data streams. However, in real-world scenarios where annotation resources are limited, a vast amount of unlabeled data coexists with scarce labeled samples, leading to the Semi-Supervised LReID (Semi-LReID) problem where LReID methods suffer severe performance degradation. Existing LReID methods, even when combined with semi-supervised strategies, suffer from limited long-term adaptation performance due to struggling with the noisy knowledge occurring during unlabeled data utilization. In this paper, we pioneer the investigation of Semi-LReID, introducing a novel Self-Reinforcing **PR**ototype **E**volution with **D**ual-Knowledge **C**ooperation framework (SPRED). Our key innovation lies in establishing a self-reinforcing cycle between dynamic prototype-guided pseudo-label generation and new-old knowledge collaborative purification to enhance the utilization of unlabeled data. Specifically, learnable identity prototypes are introduced to dynamically capture the identity distributions and generate high-quality pseudo-labels. Then, the dual-knowledge cooperation scheme integrates current model specialization and historical model generalization, refining noisy pseudo-labels. Through this cyclic design, reliable pseudo-labels are progressively mined to improve current-stage learning and ensure positive knowledge propagation over long-term learning. Experiments on the established Semi-LReID benchmarks show that our SPRED achieves state-of-the-art performance. Our source code is available at <https://github.com/zhoujiahuan1991/ICCV2025-SPRED>

1. Introduction

Person re-identification (ReID) [7, 15, 16, 69] has been extensively studied in stationary scenarios [12, 19, 25, 60].

*Corresponding author




Figure 1. (a) In Semi-LReID, the training data is continuously obtained with partial labeling. (b) As the labeling rate decreases, the performance of existing LReID methods suffers severe degradation. (c) Our SPRED improves the knowledge acquisition of unlabeled data significantly.

Since real-world applications often involve dynamic environments where new data continuously arrives [13, 50, 56, 57, 64], lifelong person re-identification (LReID) [8, 52, 53] has emerged as a research frontier, where a significant challenge is the issue of catastrophic forgetting [9, 54]. Traditional LReID methods typically assume that all training samples are labeled [33, 52, 59], an assumption that is often impractical due to the high labor costs of ReID annotation [12, 38, 61]. To address this limitation, we explore a more realistic setting, *semi-supervised lifelong person re-identification* (Semi-LReID), where only a subset of the data stream is labeled, as depicted in Fig. 1 (a). In such a scenario, the scarcity of labeled samples leads to significant performance degradation in existing LReID methods, as shown in Fig. 1 (b). Consequently, the core challenge in




Figure 2. (a) Given the labeled and unlabeled data of the new domain, we aim to learn actual ID distributions. (b) SSL tends to approximate ID distributions by emphasizing regions near the labeled data. (c) Old knowledge recognizes samples of the same ID from historical experience (old model). (d) The proposed dual-knowledge cooperation mechanism jointly utilizes the newly learned and complementary historical knowledge to improve ID assignment for unlabeled samples.

Semi-LReID is to effectively leverage the unlabeled data while mitigating catastrophic forgetting during long-term learning.

A typical approach to tackle Semi-LReID is to adapt existing semi-supervised lifelong learning (SSLL) techniques. However, current SSLL methods [10, 22, 46], usually built on an existing semi-supervised learning baseline [41], predominantly develop replay-based anti-forgetting mechanisms. Given the privacy-sensitive nature of human images, storing historical samples for replay is often impractical [33, 53]. Recently proposed replay-free SSLL methods [14] integrate self-supervised pre-training constraints to improve unlabeled data utilization. However, since ReID is an identity-level discrimination task where identity-irrelevant visual cues dominate [29, 37], self-supervised pre-training constraints often exhibit limited robustness [39].

An alternative approach is to integrate semi-supervised learning (SSL) techniques [37, 49, 67] into LReID frameworks. While it can enhance unlabeled data exploitation and alleviate catastrophic forgetting, the overall performance remains limited. As shown in Fig. 2 (b), SSL methods tend to approximate identity distributions by overemphasizing regions near labeled samples, resulting in erroneous pseudo-labels [1, 41]. These noisy labels, in turn, propagate errors during training and hinder the adapta-

tion of the model to new data [52], ultimately diminishing knowledge acquisition capacity in long-term learning.

To address these limitations, we propose a novel Semi-LReID framework, Self-Reinforcing Prototype Evolution with Dual-Knowledge Cooperation (SPRED). As shown in Fig. 2 (b), our core idea is to jointly leverage newly learned knowledge (Fig. 2 (b)) and complementary historical knowledge (Fig. 2 (c)) to form a Dual-Knowledge Cooperation mechanism (Fig. 2 (d)) that enhances reliable pseudo-label generation. Specifically, SPRED comprises two key components: Dual-Knowledge Cooperation-driven Pseudo-label Purification (DKCP) and Self-Reinforcing Prototype Evolution (SPE). SPE employs learnable identity prototypes to dynamically model identity distributions from pseudo-labels and limited ground-truth annotations. It incorporates a neighbor prototype labeling strategy where ambiguous predictions are discarded and high-quality pseudo-labels are generated. DKCP clusters new data using both the current and historical models and selects reliable pseudo-labels based on clustering consistency among unlabeled samples. These selected pseudo-labels, in turn, motivate the identity prototypes learning in SPE, forming a self-reinforcing cycle that iteratively improves pseudo-label quality and prototype learning. Besides, a prototype structure-based knowledge distillation loss is developed to mitigate catastrophic forgetting, further improving the long-term knowledge consolidation capacity. Extensive experiments on established Semi-LReID benchmarks demonstrate that SPRED significantly enhances unlabeled data utilization and long-term knowledge accumulation compared to existing solutions.

In summary, the contributions of this work are three-fold: (1) We present a pioneering study on the Semi-LReID task, identifying the limitations of current methods in this setting and introducing a dedicated Semi-LReID benchmark for future research. (2) A novel Semi-LReID method, SPRED, is proposed, forming a cyclically evolved unlabeled data utilization mechanism. The learnable identity prototypes are designed to dynamically enhance pseudo-label generation, while the complementary knowledge from the old model is integrated with new knowledge to purify pseudo-labels. (3) Extensive experiments demonstrate the effectiveness of SPRED in both unlabeled data utilization and long-term knowledge accumulation within the challenging Semi-LReID scenarios.

2. Related Work

2.1. Lifelong Person Re-Identification

Lifelong person re-identification [33, 50, 55] aims to develop a person retrieval model capable of adapting to multiple domains by learning from non-stationary data. The primary challenge in LReID is the catastrophic forgetting




Figure 3. The overview of our SPRED. Given a dataset D_t , two modules operate cyclically throughout each training epoch e . (a) The DKCP mechanism refines the pseudo-labels assigned to unlabeled samples, which are initially generated by (b) the SPE scheme. This scheme introduces learnable prototypes to facilitate LReID training using both artificial labels and pseudo-labels while leveraging the updated prototypes to assign pseudo-labels in the subsequent epoch.

problem [34], where previously acquired discriminative knowledge is overwritten as new data is introduced [35, 53, 54, 62]. Existing LReID methods typically assume access to abundant labeled data during training. However, when labeled data is limited, their performance deteriorates significantly, highlighting the need for effective strategies to leverage unlabeled data during lifelong learning.

2.2. Semi-Supervised Learning

Semi-supervised learning (SSL) is widely applied in classification and person re-identification tasks where vast images are available, but only a small fraction of the samples are labeled [37]. Pseudo-label methods aim to predict labels for unannotated data, thereby expanding the labeled training set [6, 37, 43, 48, 49]. However, these methods are prone to overfitting artificial labels, leading to noisy pseudo-labels. Such noisy data can introduce erroneous knowledge, which not only degrades the performance of new data but also exacerbates catastrophic forgetting [52]. Self-supervised-based methods employ both labeled and unlabeled data, using supervised and self-supervised supervision, respectively [11, 63]. These approaches typically require large-scale datasets, making them less suitable for lifelong learning scenarios where training data is often limited [33]. Consistency regularization methods enforce consistency in the model’s predictions across different perturbations of the same input [1, 24, 30, 36, 41]. These approaches focus on intra-instance consistency but neglect inter-instance relation learning. However, in Semi-LReID, the inter-instance relations are crucial during testing. As a result, these methods tend to yield sub-optimal results in Semi-LReID scenarios.

2.3. Semi-Supervised Lifelong Learning

Semi-supervised Lifelong learning (SSLL) considers a more realistic scenario where only a limited number of samples are annotated in the continual data stream [10]. Thus, the new knowledge learning and catastrophic forgetting problems are both challenging. Existing SSLL works primarily focus on the classification [2, 32, 40, 45] task. They usually exploit the SSL technique to learn from new data and adopt a replay buffer to store historical exemplars, which are used to mitigate forgetting [10, 22]. However, since human images are privacy-sensitive, these technologies are impractical in real scenarios [59]. Recently, a replay-free SSLL method, HDC [48], was proposed, where the self-supervised pre-training constraints are adopted to utilize unlabeled data. Since ReID is an identity-specific discrimination task, and identity-irrelevant visual cues usually dominate image information. Self-supervised pre-training constraints often exhibit limited robustness [29, 37]. Therefore, in this paper, we provide a pioneering investigation on the semi-supervised lifelong person re-identification task and propose an exemplar-free solution for realistic applications.

3. The Proposed Method

3.1. Problem Definition

In Semi-LReID, a set of T datasets $\mathcal{D} = \{D_t\}_{t=1}^T$ are given in order. Each dataset D_t contains an unlabeled subset X_t^u and a labeled subset (X_t^l, Y_t^l) . The identity labels between different training datasets are disjoint, i.e., $Y_t^l \cap Y_s^l = \emptyset$.

when $t \neq s$. During the t -th training step, only D_t is available, the previous and later datasets are inaccessible. Besides, a set of T test datasets $\mathcal{D}^{te} = \{D_t^{te}\}_{t=1}^T$ is collected to evaluate the long-term knowledge accumulation capacity of the model. In addition, U datasets $\mathcal{D}^{un} = \{D_i^{un}\}_{i=1}^U$ from novel domains are exploited to evaluate the generalizability of the model.

3.2. Overview

As shown in Fig. 3, the proposed SPRED method comprises two main components: a Dual-Knowledge Cooperation-driven Pseudo-label Purification (DKCP) mechanism and a Self-Reinforcing Prototype Evolution (SPE) scheme. SPE contains a Dynamic Prototype-guided LReID Learning module (DPL) and a Neighbor Prototype Labeling module (NPL), which aim to conduct dynamic prototype modeling and prototype-based pseudo-label generation, respectively. DKCP consists of an Old Knowledge Aligning-based Clustering module (OKAC) and a New Knowledge Clustering module (NKC), which adopt the new and old models to generate new sample clusters, respectively. The clustering results are exploited to filter unreliable pseudo-labels.

3.3. Self-Reinforcing Prototype Evolution

In this section, we first present the learnable prototype modeling process in the Semi-LReID scenario, followed by a detailed explanation of the prototype-based pseudo-label generation strategy.

Dynamic Prototype-guided LReID Learning: At training step t , given L_t identities, a set of learnable prototypes $\mathcal{P}_t = \{p_i^t\}_{i=1}^{L_t}$ is adopted, where each $p_i^t \in \mathbb{R}^d$ is a d -dimensional vector corresponding to a specific identity. Additionally, we exploit a set of valid pseudo labels, denoted Y_t^{pse} , with an associated image set X_t^{pse} . At the e -th epoch, a new LReID model $\mathbf{M}_t^e$ is trained on these data. The final LReID model learned from D_t is denoted as $\mathbf{M}_t$.

New Knowledge Learning: Given an input image x which is sampled from $X_t^l \cup X_t^{pse}$, $\mathbf{M}_t^e$ is adopted to extract the feature f_t^e . Then a prototype-orient identity learning loss is introduced, which is calculated as follows:

$$\mathcal{L}_p = -\log \frac{e^{\langle f_t^e, p_i^x \rangle}}{\sum_{i=1}^{L_t} e^{\langle f_t^e, p_i^x \rangle}}, \quad (1)$$

where p_i^x denotes the prototype sharing the same label as x . As illustrated in Fig. 3 (b), Eq. (1) encourages the instance feature to align more closely with its corresponding identity prototype while distancing itself from other prototypes. Since both $\mathcal{P}_t$ and $\mathbf{M}_t^e$ are learnable, the prototypes and model parameters are updated dynamically during training.

To further enhance feature learning in LReID, we also integrate the widely-used Triplet loss $\mathcal{L}_{Tri}$ [54]. The total loss for new knowledge learning is thus formulated as:

$$\mathcal{L}_{base} = \mathcal{L}_p + \mathcal{L}_{Tri}. \quad (2)$$

Old Knowledge Non-Forgetting: Since $\mathcal{L}_p$ guides the model to learn one-instance vs. all-prototype similarity structures, we aim to preserve the structural knowledge learned by the old model to mitigate catastrophic forgetting. Specifically, let $f_{t-1} \in \mathbb{R}^d$ denote the feature extracted by $\mathbf{M}_{t-1}$. The prototype-based old structure is represented as:

$$S(f_{t-1}, \mathcal{P}_{t-1}) = [g(f_{t-1}, \mathcal{P}_{t-1}, 1), \dots, g(f_{t-1}, \mathcal{P}_{t-1}, N_{t-1})] \quad (3)$$

where $g(f_{t-1}, \mathcal{P}_{t-1}, j) = \frac{e^{\langle f_{t-1}, p_{t-1}^j \rangle}}{\sum_{i=1}^{N_{t-1}} e^{\langle f_{t-1}, p_{t-1}^i \rangle}}$ represents the affinity between input image x and each prototype p_{t-1}^j according to the old knowledge.

Similarly, the prototype-based new structure $S(f_t^e, \mathcal{P}_{t-1})$ could be obtained. Then, to alleviate structural drift as the new model evolves, we introduce a prototype-guided feature structure maintaining loss:

$$\mathcal{L}_s = \mathcal{L}_{KL}(S(f_{t-1}, \mathcal{P}_{t-1}) || S(f_t^e, \mathcal{P}_{t-1})), \quad (4)$$

where $\mathcal{L}_{KL}$ is the Kullback Leibler divergence [20].

Neighbor Prototype Labeling: Existing pseudo-label generation methods generally rely on classification scores and a threshold to create pseudo-labels [5, 11], where the correct prediction score is susceptible to multiple neighbor classes, causing many informative samples to be discarded.

To address this issue, we propose to focus on the top-2 nearest classes to obtain the pseudo-labels and ignore the less-related classes, as shown in Fig. 3 (b). Specifically, given the image feature $f_t^{e-1} \in \mathbb{R}^d$ extracted by $\mathbf{M}_t^{e-1}$, we obtain the top-2 nearest prototypes, where the nearest and the second nearest prototypes are represented as p_A and p_B , respectively. The classification score of the p_A class is obtained by:

$$s_A = \frac{e^{\langle f_t^{e-1}, p_A \rangle}}{e^{\langle f_t^{e-1}, p_A \rangle} + e^{\langle f_t^{e-1}, p_B \rangle}}. \quad (5)$$

Then, we select the unlabeled samples with a threshold T_p . Therefore, the predicted label l_i of a sample x_i can be obtained by:

$$l_i = \begin{cases} y_i & x_i \in X_t^l \\ l_A & x_i \notin X_t^l, s_A > T_p \\ N/A & x_i \notin X_t^l, s_A \leq T_p \end{cases} \quad (6)$$

where y_i denotes the artificial label of x_i , l_A represents the corresponding label of p_A , and N/A means the pseudo-label is unavailable according to the prototypes.

3.4. Dual-Knowledge Cooperation-driven Pseudo-label Purification

The pseudo-labels generated in Sec. 3.3 are inevitably noisy due to the prototypes overfitting the labeled data, which can

misguide model training and exacerbate catastrophic forgetting [52]. To settle this, we propose to leverage both new and old knowledge from the current and previous models to filter noisy pseudo-labels, thereby enhancing the effective utilization of unlabeled samples.

Old Knowledge Aligning-based Clustering: Due to the domain shift across training steps, directly applying $\mathbf{M}_{t-1}$ to process D_t leads to suboptimal results. To settle this, we employ a distribution alignment network (DANet) Θ_{t-1} to map the new data D_t to the distribution of D_{t-1} . The transformed data are denoted as $D_t^* = \{x_i^*\}_{i=1}^{n_t}$, where n_t is the number of images in D_t . The architecture and training protocol for DANet is based on [18], with additional details available in our Supplementary Material.

Then, the old model $\mathbf{M}_{t-1}$ is exploited to extract image features $F_{t-1}^* = \{f_{t-1}^{(i)}\}_{i=1}^{|D_t^*|}$ from the transformed data D_t^* . Subsequently, the DBSCAN clustering algorithm [18] is applied to F_{t-1}^* , resulting in N_C^{t-1} clusters, denoted as $C_{t-1} = \{C_1^{t-1}, C_2^{t-1}, \dots, C_{N_C^{t-1}}^{t-1}\}$. Each cluster C_i^{t-1} represents a group of instances exhibiting similar patterns based on the knowledge encoded in $\mathbf{M}_{t-1}$.

New Knowledge Clustering: Given the new data D_t , the new model $\mathbf{M}_t^{e-1}$ is utilized to extract features $F_t^{e-1} = \{f_t^{(i),e-1}\}_{i=1}^{N_t}$. Next, the clustering algorithm is also applied to process F_t^{e-1} and a set of $N_C^{t,e-1}$ clusters $C_t^{e-1} = \{C_1^{e-1}, C_2^{e-1}, \dots, C_{N_C^{t,e-1}}^{e-1}\}$ is obtained. Here, each cluster C_i^{e-1} is composed of instances whose features extracted by the new model are highly similar to each other.

Noisy Pseudo-Label Filtering: As shown in Fig. 3 (a), we aim to utilize the clustering results generated by new and old models to filter the noisy pseudo-labels. Firstly, to bridge the clustering results and pseudo-labels, we split the new samples into sets according to the pseudo-labels generated in Sec. 3.3, resulting in a collection $\mathcal{S}^{e-1} = \{S_1^{e-1}, S_2^{e-1}, \dots, S_{L_t}^{e-1}\}$, where each subset $S_i^{e-1} = \{x_j : x_j \in X_t^l \cup X_t^u, l_j = i\}$ is set of samples with the same pseudo-label.

For an unlabeled instance x , pseudo-label set $\mathcal{S}$ and clustering result $\mathcal{C}$, we define the set-based label confidence score $LC_s(x, \mathcal{S}, \mathcal{C})$ which is calculated by

$$LC_s(x, \mathcal{S}, \mathcal{C}) = \frac{S_x \cap C_x}{S_x \cup C_x}, \quad (7)$$

where S_x and C_x are the elements of $\mathcal{S}$ and $\mathcal{C}$ respectively, which both contain x .

Given the new knowledge clustering result $\mathcal{C}_t^{e-1}$ and an unlabelled sample x , a threshold T_c is adopted to obtain the purified pseudo-label set $Y_{e-1}^{pse} = \{y_{i,e-1}^{pse} : x_i \in X_t^u\}$ where $y_{i,e-1}^{pse}$ is obtained by

$$y_{i,e-1}^{pse} = \begin{cases} l_i & LC_s(x_i, \mathcal{S}^{e-1}, \mathcal{C}_t^{e-1}) > T_c \\ N/A & LC_s(x_i, \mathcal{S}^{e-1}, \mathcal{C}_t^{e-1}) \leq T_c \end{cases} \quad (8)$$

Similarly, given the old knowledge clustering result $\mathcal{C}_{t-1}$, we can obtain another purified pseudo-label set $Y_{t-1}^{pse} = \{y_{i,t-1}^{pse} : x_i \in X_t^u\}$ by introducing a threshold T_o , where $y_{i,t-1}^{pse}$ is obtained following Eq. (8).

Compared to using a purified pseudo-label set alone or using their intersection, we experimentally found that using Y_{e-1}^{pse} and Y_{t-1}^{pse} results in the highest performance, *i.e.*, the pseudo label for LReID model training is obtained by

$$Y_t^{pse} = Y_{t-1}^{pse} \cup Y_{e-1}^{pse}. \quad (9)$$

This underscores that the new and old model clustering results are complementary to each other.

3.5. Training and Inference

When training on the first dataset, only $\mathcal{L}_{base}$ is exploited as the supervision. When training on the subsequent datasets, the overall loss is calculated by:

$$\mathcal{L} = \mathcal{L}_{base} + \alpha \mathcal{L}_s, \quad (10)$$

where α is a hyperparameter to balance the new knowledge learning and old knowledge forgetting.

During inference, following the previous works [44, 53], the image features extracted by the final model $\mathbf{M}_T$ are adopted for person retrieval.

4. Experiments

4.1. Datasets and Evaluation Metrics

In this section, we present a new and large-scale Semi-LReID benchmark for comprehensive model evaluation:

Datasets: 12 ReID datasets are adopted to compose our Semi-LReID benchmark, 5 of which are set as training domains (Market1501 [66], LPW [42], CUHK-SYSU [51], MSMT17-V2 [47], and CUHK03 [28]), and the other 7 smaller datasets are set as test-only domains (CUHK01 [27], CUHK02 [26], VIPeR [17], PRID [21], iLIDS [3], GRID [31], and SenseReID [65]). More details of the datasets are provided in our Supplementary Materials.

Label Rate: Following the previous semi-supervised learning settings [4, 37, 67], diverse labeling rates (10%, 20%, 50%) are adopted to evaluate the effectiveness of the methods under different scenarios, where the labeled data are sampled randomly.

Evaluation Metrics: Following the standard LReID protocols [33, 54], mean Average Precision (mAP) and Rank-1 (R@1) accuracy are adopted to evaluate model performance on each ReID dataset. Moreover, the average mAP and R@1 across datasets are calculated to assess the overall performance on diverse domains.

4.2. Implementation Details

Following the previous works [44, 54], the ResNet-50 is set as the LReID backbone. The model is initially trained

r	Type	Method	Market-1501		CUHK-SYSU		LPW		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
50%	LReID	PatchKD [44]	65.8	84.5	73.0	75.9	25.7	34.4	5.3	14.7	19.3	19.7	37.8	45.8	40.3	34.1
		DKP [53]	49.5	73.9	80.7	83.2	35.4	47.6	14.4	33.9	26.6	26.3	41.3	53.0	52.8	46.6
		LSTKC [54]	48.4	71.4	80.8	83.4	39.3	52.0	16.7	38.2	37.0	38.1	44.4	56.6	53.3	45.7
	SSL+LReID	CDMAD [†] [23]	42.0	70.5	77.6	80.7	21.9	35.4	7.4	23.4	14.5	15.1	32.7	45.0	47.4	40.3
		ShrinkM [†] [58]	51.9	75.9	77.5	80.6	42.6	57.5	21.7	49.3	37.7	39.4	46.3	60.5	57.9	50.6
		SimMV2 [†] [68]	34.6	59.3	63.2	65.9	28.0	41.1	11.5	30.3	29.6	31.4	33.4	45.6	46.9	40.3
		DPIS [†] [37]	49.9	72.1	82.0	84.4	40.8	53.7	17.2	38.8	35.5	35.6	45.1	56.9	55.3	48.5
		HDC [†] [48]	51.8	75.2	79.7	82.2	43.5	57.0	22.2	49.1	40.1	42.3	47.5	61.2	58.2	50.6
	Ours	SPRED	59.8	79.7	82.8	84.6	49.7	60.2	17.4	38.2	40.6	41.3	50.1	60.8	57.8	50.3
		SPRED [‡]	59.4	80.7	79.6	81.5	47.3	60.2	21.6	47.3	48.4	49.4	51.3	63.8	60.7	53.9
20%	LReID	PatchKD [44]	43.6	65.7	67.1	70.5	10.0	15.8	3.0	8.8	7.0	6.2	26.1	33.4	30.4	24.8
		DKP [53]	30.4	55.8	73.5	76.1	21.5	31.0	8.8	24.1	8.3	7.0	28.5	38.8	40.8	34.8
		LSTKC [54]	34.8	58.6	77.1	80.0	28.2	40.4	12.2	29.7	19.4	18.2	34.3	45.4	47.0	40.1
	SSL+LReID	CDMAD [†] [23]	29.8	56.9	74.0	77.5	19.1	31.1	6.0	19.9	9.5	8.6	27.7	38.8	43.6	37.0
		ShrinkM [†] [58]	41.7	68.0	73.3	76.4	34.8	49.2	17.5	42.9	22.8	22.4	38.0	51.8	53.0	45.8
		SimMV2 [†] [68]	34.5	59.1	65.1	68.4	26.1	38.1	10.5	28.6	19.0	19.2	31.0	42.7	45.7	38.3
		DPIS [†] [37]	34.8	59.3	77.9	81.1	29.3	41.1	12.5	30.8	18.5	16.8	34.6	45.8	47.7	40.7
		HDC [†] [48]	39.9	65.5	75.3	78.5	31.8	44.8	16.4	40.3	25.3	24.1	37.7	50.6	52.5	45.7
	Ours	SPRED	54.7	75.7	81.1	83.0	45.0	56.9	16.1	36.3	35.8	35.4	46.5	57.5	55.7	48.4
		SPRED [‡]	54.6	77.1	78.5	80.5	44.9	58.9	21.3	47.5	38.6	38.2	47.6	60.4	58.9	51.5
10%	LReID	PatchKD [44]	20.6	39.0	58.5	62.4	3.5	6.1	1.1	3.7	3.0	2.3	17.3	22.7	20.8	16.5
		DKP [53]	17.9	39.5	67.5	70.3	12.5	19.0	5.7	17.4	4.9	3.2	21.7	29.9	32.9	27.1
		LSTKC [54]	24.0	46.7	72.8	76.8	20.0	29.6	7.6	20.4	14.2	12.1	27.7	37.1	39.5	32.7
	SSL+LReID	CDMAD [†] [23]	23.0	49.8	72.2	75.9	16.6	26.4	5.8	19.3	7.8	6.9	25.1	35.7	42.7	37.1
		ShrinkM [†] [58]	32.2	58.1	70.2	74.0	27.7	42.1	13.0	35.1	18.5	17.9	32.3	45.4	48.0	41.2
		SimMV2 [†] [68]	29.2	54.1	64.3	68.0	20.5	32.0	8.5	24.4	14.9	14.1	27.5	38.5	42.3	35.2
		DPIS [†] [37]	25.4	48.1	74.0	77.4	20.9	30.6	8.6	22.3	14.0	12.2	28.6	38.1	41.4	34.2
		HDC [†] [48]	26.0	48.8	70.2	73.9	20.8	29.7	8.1	23.1	12.4	10.9	27.5	37.3	45.6	39.2
	Ours	SPRED	45.9	69.4	79.3	81.6	40.0	51.1	14.2	33.0	36.7	37.1	43.2	54.4	51.7	44.6
		SPRED [‡]	49.5	74.0	76.5	79.1	40.9	55.4	19.4	45.3	37.5	38.4	44.8	58.4	55.9	48.7

Table 1. Results comparison under different label rate r . [†] indicates the state-of-the-art SSL method is integrated with the anti-forgetting designs of LSTKC. [‡] represents adopting the data augmentation strategy of ShrinkM.

on the first dataset for 80 epochs, with subsequent datasets trained for 60 epochs each. For optimization, an SGD optimizer with an initial learning rate of 8×10^{-3} is used. To ensure fair comparisons with existing LReID methods, standard LReID data augmentations [53], including random cropping, erasing, and horizontal flipping, are applied to obtain the SPRED model. Additionally, we observed that the strong-weak augmentation strategy commonly used in SSL [58], can boost generalization in Semi-LReID. Therefore, we incorporate this SSL augmentation strategy, with further details provided in the supplementary material, to obtain SPRED[‡]. The experiments are conducted on the label rate $r=10\%$ by default.

4.3. Compared Methods

We compare the proposed method with two categories of approaches: A. *LReID* methods, including PatchKD [44],

LSTKC [54], and DKP [53]. B. *SSL* methods, including SSL-Classification methods, CDMAD [23], ShrinkM [58], SimMV2 [68] and SSL-ReID methods, DPIS [37], HDC [48]. Since SSL methods do not incorporate an anti-forgetting mechanism, the anti-forgetting strategy of the state-of-the-art LReID method (LSTKC [54]) is integrated into these SSL approaches. The resulting methods are referred to as *SSL+LReID* approaches. All experimental results are obtained using the official implementation or following the official paper, with a parameter grid search conducted to ensure optimal performance.

The performance of different methods on each seen domain is shown in Tab. 1. Besides, we also report the average performance across all seen domains (Seen-Avg) and unseen domains (UnSeen-Avg). The best and second best results are highlighted in **Red** and **Blue**, separately.




Figure 4. Seen domain performance curves under 10% label rate.

4.4. Seen-Domain Performance Evaluation

Compared to LReID Methods: In Tab. 1, our SPRED consistently outperforms existing LReID methods in terms of Seen-Avg mAP/R@1 across different labeling rates. Specifically, compared to the state-of-the-art LSTKC, SPRED achieves notable improvements of **5.7%/4.2%**, **12.2%/12.1%** and **15.5%/17.3%** in Seen-Avg mAP/R@1 as the labeling rate decreases from 50% to 10%. These results verify the effectiveness of SPRED in label-limited settings, primarily due to its dynamic prototype learning and dual-knowledge cooperation mechanisms, which enhance the utilization of unlabeled data.

Compared to SSL+LReID Methods: As shown in Tab. 1, our SPRED surpasses existing SSL+LReID approaches, achieving at least **8.5%/5.7%** and **10.9%/9.0%** improvements in Seen-Avg mAP/R@1 under 20% and 10% labeling rates, respectively. Under a 50% labeling rate, SPRED demonstrates a 2.6% increase in Seen-Avg mAP but a 0.4% decrease in Seen-Avg R@1 compared to HDC. This higher R@1 performance by HDC primarily stems from its data augmentation strategy, which enhances input data diversity to facilitate LReID training.

By incorporating the widely used data augmentation strategy of SSL [58], our enhanced SPRED[‡] model surpasses all SSL+LReID methods in Seen-Avg mAP/R@1 metrics across different labeling rates. Specifically, we observe improvements of **3.8%/2.6%**, **9.6%/8.6%** and **12.5%/13.0%** in Seen-Avg mAP/R@1 at 50%, 20%, and 10% label rates, respectively. This superior performance is attributed to our self-reinforcing prototype evolution framework, which enhances pseudo-label quality and boosts the capacity of the model for new knowledge acquisition.

Seen Domain Performance Tendency: In Fig. 4, we evaluate model performance on already encountered domains after each training step. The results indicate that both SPRED and SPRED[‡] consistently outperform all existing methods. This advantage is attributed to our prototype-based anti-forgetting loss and effective pseudo-label purification mechanisms, which enhance knowledge consolida-




Figure 5. Generalization curves under 10% label rate.

Base	DPL	NPL	DKCP	Seen-Avg mAP	Seen-Avg R@1	UnSeen-Avg mAP	UnSeen-Avg R@1
✓				27.7	37.1	39.5	32.7
✓	✓			28.8	38.2	41.7	34.2
✓	✓	✓		38.0	49.8	47.4	40.5
✓	✓	✓	✓	43.2	54.4	51.7	44.6

Table 2. Ablation study of different components.

tion and mitigate the catastrophic forgetting of historical correct knowledge caused by noisy pseudo-labels.

4.5. Unseen-Domain Generalization Evaluation

Compared to LReID and SSL+LReID Methods: As shown in Tab. 1, our SPRED outperforms all LReID methods on the UnSeen-Avg mAP/R@1 performance. Besides, our SPRED[‡] surpasses SSL+LReID approaches with **2.5%/3.3%**, **5.9%/5.7%**, and **7.9%/7.5%** UnSeen-Avg mAP/R@1 across label rate from 50% to 10%. These results underscore our method’s effectiveness in accumulating generalizable knowledge, which is attributed to an efficient pseudo-label generation mechanism that reduces overfitting to the limited labeled data.

UnSeen Domain Performance Tendency: In Fig. 5, we assess the model’s performance on unseen domains after each training step. The results show that our SPRED and SPRED[‡] achieve leading performance consistently, primarily due to enhanced utilization of unlabeled data, which integrates more generalizable knowledge into the model.

4.6. Ablation Studies

Ablations on model components. In Tab. 2, we adopt LSTKC as the Base model and progressively integrate the proposed DPL, NPL, and DKCP modules. The results show that each component substantially enhances model performance. When all components are incorporated, our method surpasses the Base model by **15.5%/17.3%** Seen-Avg mAP/R@1 and **12.2%/11.9%** UnSeen-Avg mAP/R@1.

Ablations on hyperparameters. We investigate the effects of the hyperparameters α , T_c , T_o , and T_p , on the model




Figure 6. Ablation studies on hyperparameters under 10% label rate. The default values are highlighted by the dashed lines.




Figure 7. T-SNE visualization of labeled and unlabeled features in the training dataset MSMT17.

in Fig. 6. The results indicate that model performance is relatively robust to changes in these hyperparameters. In practice, we set α , T_c , T_o , and T_p to 4.0, 0.1, 0.6, and 0.7 by default, respectively.

4.7. Visualization Results

Feature Distribution Visualization: In Fig. 7, we visualize the training data features of different identities (colors) in comparison with existing approaches. The results demonstrate that our method achieves a more cohesive grouping of labeled and unlabeled features while existing methods exhibit substantial labeled-unlabeled feature disparity and limited discriminative ability between identities. These results highlight our superiority in mining knowledge from unlabeled data and boosting the LReID performance.

Identity Prediction Visualization: To demonstrate the




Figure 8. Label prediction accuracy of different methods.

knowledge acquisition capacity of different methods, we visualize the identity prediction accuracy on the challenging Market-1501 and MSMT17 datasets in Fig. 8. Unlike existing methods, whose accuracy plateaus after 25 epochs, our SPRED method continues to improve, ultimately achieving significantly higher final accuracy. These results confirm the effectiveness of our cyclically evolved unlabeled data utilization mechanism, in which pseudo-label prediction and purification iteratively enhance each other.

5. Conclusion

In this paper, a practical Semi-LReID task is investigated, aiming to learn from a data stream with a vast amount of unlabeled data alongside scarce labeled samples. To support further research, a new large-scale Semi-LReID benchmark is introduced. Additionally, a novel Semi-LReID approach, SPRED, is proposed. Specifically, a self-reinforcing cycle is established, wherein the learnable prototype is introduced to dynamically model identity distribution and generate high-quality pseudo-labels. Besides, a dual-knowledge collaborative purification mechanism is designed to filter the underlying unreliable pseudo-labels by exploiting the new knowledge and the complementary old knowledge. Extensive experiments demonstrate that the proposed method significantly outperforms existing approaches, particularly under conditions with limited labeled data.

Acknowledgements: This work was supported by the National Key R&D Program of China (2024YFA1410000), the National Natural Science Foundation of China (62376011), and the China Postdoctoral Science Foundation (2025T180417).

References

- [1] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *NeurIPS*, 32, 2019. 2, 3
- [2] Dhanajit Brahma, Vinay Kumar Verma, and Piyush Rai. Hypernetworks for continual semi-supervised learning. *arXiv preprint arXiv:2110.01856*, 2021. 3
- [3] Home Office Scientific Development Branch. Imagery library for intelligent detection systems (i-lids). In *2006 IET conference on crime and security*, pages 445–448. IET, 2006. 5
- [4] Zhaowei Cai, Avinash Ravichandran, Subhransu Maji, Charles Fowlkes, Zhuowen Tu, and Stefano Soatto. Exponential moving average normalization for self-supervised and semi-supervised learning. In *CVPR*, pages 194–203. IEEE, 2021. 5
- [5] Yongrui Chen, Xinnan Guo, Tongtong Wu, Guilin Qi, Yang Li, and Yang Dong. Learn from yesterday: A semi-supervised continual learning method for supervision-limited text-to-sql task streams. In *AAAI*, pages 12682–12690, 2023. 4
- [6] Yuhao Chen, Xin Tan, Borui Zhao, Zhaowei Chen, Renjie Song, Jiajun Liang, and Xuequan Lu. Boosting semi-supervised learning by exploiting all unlabeled data. In *CVPR*, pages 7548–7557, 2023. 3
- [7] Zhenyu Cui, Jiahuan Zhou, and Yuxin Peng. Dma: Dual modality-aware alignment for visible-infrared person re-identification. *T-IFS*, 19:2696–2708, 2024. 1
- [8] Zhenyu Cui, Jiahuan Zhou, Xun Wang, Manyu Zhu, and Yuxin Peng. Learning continual compatible representation for re-indexing free lifelong person re-identification. In *CVPR*, pages 16614–16623, 2024. 1
- [9] Zhenyu Cui, Jiahuan Zhou, and Yuxin Peng. Dkc: Differentiated knowledge consolidation for cloth-hybrid lifelong person re-identification. In *CVPR*, pages 3573–3582, 2025. 1
- [10] Yan Fan, Yu Wang, Pengfei Zhu, and Qinghua Hu. Dynamic sub-graph distillation for robust semi-supervised continual learning. In *AAAI*, pages 11927–11935, 2024. 2, 3
- [11] Enrico Fini, Pietro Astolfi, Karteek Alahari, Xavier Alameda-Pineda, Julien Mairal, Moin Nabi, and Elisa Ricci. Semi-supervised learning made simple with self-supervised clustering. In *CVPR*, pages 3187–3197, 2023. 3, 4
- [12] Lihua Fu, Yubin Du, Yu Ding, Dan Wang, Hanxu Jiang, and Haitao Zhang. Domain adaptive learning with multi-granularity features for unsupervised person re-identification. *Chinese Journal of Electronics*, 31(1):116–128, 2022. 1
- [13] Wenhong Ge, Junlong Du, Ancong Wu, Yuqiao Xian, Ke Yan, Feiyue Huang, and Wei-Shi Zheng. Lifelong person re-identification by pseudo task knowledge preservation. In *AAAI*, pages 688–696, 2022. 1
- [14] Alex Gomez-Villa, Dipam Goswami, Kai Wang, Andrew D Bagdanov, Bartłomiej Twardowski, and Joost van de Weijer. Exemplar-free continual representation learning via learnable drift compensation. In *ECCV*, pages 473–490. Springer, 2024. 2
- [15] Yunpeng Gong, Zhun Zhong, Yansong Qu, Zhiming Luo, Rongrong Ji, and Min Jiang. Cross-modality perturbation synergy attack for person re-identification. 1
- [16] Yunpeng Gong, Liqing Huang, and Lifei Chen. Person re-identification method based on color attack and joint defense. In *CVPR*, pages 4313–4322, 2022. 1
- [17] Douglas Gray and Hai Tao. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*, pages 262–275. Springer, 2008. 5
- [18] Jianyang Gu, Hao Luo, Kai Wang, Wei Jiang, Yang You, and Jian Zhao. Color prompting for data-free continual unsupervised domain adaptive person re-identification. *arXiv preprint arXiv:2308.10716*, 2023. 5
- [19] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *ICCV*, pages 14993–15002. IEEE, 2021. 1
- [20] John R Hershey and Peder A Olsen. Approximating the kullback leibler divergence between gaussian mixture models. In *ICASSP*, pages IV–317. IEEE, 2007. 4
- [21] Martin Hirzer, Csaba Belezna, Peter M Roth, and Horst Bischof. Person re-identification by descriptive and discriminative classification. In *Image Analysis*, pages 91–102. Springer, 2011. 5
- [22] Zhiqi Kang, Enrico Fini, Moin Nabi, Elisa Ricci, and Karteek Alahari. A soft nearest-neighbor framework for continual semi-supervised learning. In *ICCV*, pages 11868–11877, 2023. 2, 3
- [23] Hyuck Lee and Heeyoung Kim. Cdmad: Class-distribution-mismatch-aware debiasing for class-imbalanced semi-supervised learning. In *CVPR*, pages 23891–23900. IEEE, 2024. 6
- [24] Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive graph regularization. In *ICCV*, pages 9475–9484, 2021. 3
- [25] Qiwei Li, Kunlun Xu, Yuxin Peng, and Jiahuan Zhou. Exemplar-free lifelong person re-identification via prompt-guided adaptive knowledge consolidation. *IJCV*, pages 1–16, 2024. 1
- [26] Wei Li and Xiaogang Wang. Locally aligned feature transforms across views. In *CVPR*, pages 3594–3601. IEEE, 2013. 5
- [27] Wei Li, Rui Zhao, and Xiaogang Wang. Human reidentification with transferred metric learning. In *ACCV*, pages 31–44. Springer, 2012. 5
- [28] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, pages 152–159. IEEE, 2014. 5
- [29] Yutian Lin, Xuanyi Dong, Liang Zheng, Yan Yan, and Yi Yang. A bottom-up clustering approach to unsupervised per-

- son re-identification. In *AAAI*, pages 8738–8745, 2019. 2, 3
- [30] Lu Liu and Robby T Tan. Certainty driven consistency loss on multi-teacher networks for semi-supervised learning. *Pattern Recognition*, 120:108140, 2021. 3
- [31] Chen Change Loy, Tao Xiang, and Shaogang Gong. Time-delayed correlation analysis for multi-camera activity understanding. *IJCV*, 90(1):106–129, 2010. 5
- [32] Yan Luo, Yongkang Wong, Mohan Kankanhalli, and Qi Zhao. Learning to predict gradients for semi-supervised continual learning. *TNNLS*, 2024. 3
- [33] Nan Pu, Wei Chen, Yu Liu, Erwin M Bakker, and Michael S Lew. Lifelong person re-identification via adaptive knowledge accumulation. In *CVPR*, pages 7897–7906. IEEE, 2021. 1, 2, 3, 5
- [34] Nan Pu, Yu Liu, Wei Chen, Erwin M Bakker, and Michael S Lew. Meta reconciliation normalization for lifelong person re-identification. In *ACM MM*, pages 541–549, 2022. 3
- [35] Nan Pu, Zhun Zhong, Nicu Sebe, and Michael S Lew. A memorizing and generalizing framework for lifelong person re-identification. *IEEE TPAMI*, 45(11):13567–13585, 2023. 3
- [36] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *NeurIPS*, 29, 2016. 3
- [37] Jiangming Shi, Yachao Zhang, Xiangbo Yin, Yuan Xie, Zhizhong Zhang, Jianping Fan, Zhongchao Shi, and Yanyun Qu. Dual pseudo-labels interactive self-training for semi-supervised visible-infrared person re-identification. In *ICCV*, pages 11218–11228. IEEE, 2023. 2, 3, 5, 6
- [38] Jiangming Shi, Xiangbo Yin, Yachao Zhang, Yuan Xie, Yanyun Qu, et al. Learning commonality, divergence and variety for unsupervised visible-infrared person re-identification. *NeurIPS*, 37:99715–99734, 2024. 1
- [39] Jiangming Shi, Xiangbo Yin, Yeyun Chen, Yachao Zhang, Zhizhong Zhang, Yuan Xie, and Yanyun Qu. Multi-memory matching for unsupervised visible-infrared person re-identification. In *ECCV*, pages 456–474. Springer, 2025. 2
- [40] James Smith, Jonathan Balloch, Yen-Chang Hsu, and Zsolt Kira. Memory-efficient semi-supervised continual learning: The world is its own replay buffer. In *IJCNN*, pages 1–8. IEEE, 2021. 3
- [41] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *NeurIPS*, 33:596–608, 2020. 2, 3
- [42] Guanglu Song, Biao Leng, Yu Liu, Congrui Hetang, and Shaofan Cai. Region-based quality estimation network for large-scale person re-identification. In *AAAI*, 2018. 5
- [43] Tiberiu Sosea and Cornelia Caragea. Marginmatch: Improving semi-supervised learning with pseudo-margins. In *CVPR*, pages 15773–15782, 2023. 3
- [44] Zhicheng Sun and Yadong Mu. Patch-based knowledge distillation for lifelong person re-identification. In *ACM MM*, pages 696–707, 2022. 5, 6
- [45] Liyuan Wang, Kuo Yang, Chongxuan Li, Lanqing Hong, Zhenguo Li, and Jun Zhu. Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In *CVPR*, pages 5383–5392, 2021. 3
- [46] Yu Wang, Pengfei Zhu, Dongyue Chen, and Qinghua Hu. Persistence homology distillation for semi-supervised continual learning. *NeurIPS*, 37:76332–76355, 2024. 2
- [47] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*, pages 79–88. IEEE, 2018. 5
- [48] Ziyu Wei, Xi Yang, Nannan Wang, and Xinbo Gao. Semi-supervised learning with heterogeneous distribution consistency for visible infrared person re-identification. *IEEE TIP*, 2024. 3, 6
- [49] Ancong Wu, Wenhong Ge, and Wei-Shi Zheng. Rewarded semi-supervised re-identification on identities rarely crossing camera views. *IEEE TPAMI*, 2023. 2, 3
- [50] Guile Wu and Shaogang Gong. Generalising without forgetting for lifelong person re-identification. In *AAAI*, pages 2889–2897, 2021. 1, 2
- [51] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. End-to-end deep learning for person search. *arXiv preprint arXiv:1604.01850*, 2(2):4, 2016. 5
- [52] Kunlun Xu, Haozhuo Zhang, Yu Li, Yuxin Peng, and Jiahuan Zhou. Mitigate catastrophic remembering via continual knowledge purification for noisy lifelong person re-identification. In *ACM MM*, pages 1–9, 2024. 1, 2, 3, 5
- [53] Kunlun Xu, Xu Zou, Yuxin Peng, and Jiahuan Zhou. Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification. In *CVPR*, pages 16604–16613. IEEE, 2024. 1, 2, 3, 5, 6
- [54] Kunlun Xu, Xu Zou, and Jiahuan Zhou. Lstkc: Long short-term knowledge consolidation for lifelong person re-identification. In *AAAI*, pages 16202–16210, 2024. 1, 3, 4, 5, 6
- [55] Kunlun Xu, Chenghao Jiang, Peixi Xiong, Yuxin Peng, and Jiahuan Zhou. Dask: Distribution rehearsing via adaptive style kernel learning for exemplar-free lifelong person re-identification. In *AAAI*, pages 8915–8923, 2025. 2
- [56] Kunlun Xu, Zichen Liu, Xu Zou, Yuxin Peng, and Jiahuan Zhou. Long short-term knowledge decomposition and consolidation for lifelong person re-identification. *IEEE TPAMI*, 2025. 1
- [57] Kunlun Xu, Xu Zou, Gang Hua, and Jiahuan Zhou. Componential prompt-knowledge alignment for domain incremental learning. In *ICML*, 2025. 1
- [58] Lihe Yang, Zhen Zhao, Lei Qi, Yu Qiao, Yinghuan Shi, and Hengshuang Zhao. Shrinking class space for enhanced certainty in semi-supervised learning. In *ICCV*, pages 16187–16196, 2023. 6, 7
- [59] Zexian Yang, Dayan Wu, Wanqian Zhang, Bo Li, and Weipeng Wang. Handling label uncertainty for camera incremental person re-identification. In *ACM MM*, page 6253–6263, 2023. 1, 3
- [60] Zhaoda Ye, Xiangteng He, and Yuxin Peng. Unsupervised cross-media hashing learning via knowledge graph. *Chinese Journal of Electronics*, 31(6):1081–1091, 2022. 1

- [61] Xiangbo Yin, Jiangming Shi, Yachao Zhang, Yang Lu, Zhizhong Zhang, Yuan Xie, and Yanyun Qu. Robust pseudo-label learning with neighbor relation for unsupervised visible-infrared person re-identification. In *ACM MM*, pages 2242–2251, 2024. [1](#)
- [62] Chunlin Yu, Ye Shi, Zimo Liu, Shenghua Gao, and Jingya Wang. Lifelong person re-identification via knowledge refreshing and consolidation. In *AAAI*, pages 3295–3303, 2023. [3](#)
- [63] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *ICCV*, pages 1476–1485, 2019. [3](#)
- [64] Chenyu Zhang, Kunlun Xu, Zichen Liu, Yuxin Peng, and Jiahuan Zhou. Scap: Transductive test-time adaptation via supportive clique-based attribute prompting. In *CVPR*, pages 30032–30041, 2025. [1](#)
- [65] Haiyu Zhao, Maoqing Tian, Shuyang Sun, Jing Shao, Junjie Yan, Shuai Yi, Xiaogang Wang, and Xiaoou Tang. Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In *CVPR*, pages 907–915. IEEE, 2017. [5](#)
- [66] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pages 1116–1124. IEEE, 2015. [5](#)
- [67] Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. Simmatch: Semi-supervised learning with similarity matching. In *CVPR*, pages 14471–14481, 2022. [2](#), [5](#)
- [68] Mingkai Zheng, Shan You, Lang Huang, Chen Luo, Fei Wang, Chen Qian, and Chang Xu. Simmatchv2: Semi-supervised learning with graph consistency. In *CVPR*, pages 16432–16442. IEEE, 2023. [6](#)
- [69] Zhun Zhong, Liang Zheng, Zhedong Zheng, Shaozi Li, and Yi Yang. Camera style adaptation for person re-identification. In *CVPR*, pages 5157–5166. IEEE, 2018. [1](#)