

Lifelong Person Re-identification via Knowledge Refreshing and Consolidation

Chunlin Yu¹, Ye Shi^{1,3}, Zimo Liu², Shenghua Gao^{1,3}, Jingya Wang^{1,3*}

¹ShanghaiTech University

²Peng Cheng Laboratory

³Shanghai Engineering Research Center of Intelligent Vision and Imaging
{yuchl, shiye, gaoshh, wangjingya}@shanghaitech.edu.cn; liuzm@pcl.ac.cn

Abstract

Lifelong person re-identification (LReID) is in significant demand for real-world development as a large amount of ReID data is captured from diverse locations over time and cannot be accessed at once inherently. However, a key challenge for LReID is how to incrementally preserve old knowledge and gradually add new capabilities to the system. Unlike most existing LReID methods, which mainly focus on dealing with catastrophic forgetting, our focus is on a more challenging problem, which is, not only trying to reduce the forgetting on old tasks but also aiming to improve the model performance on both new and old tasks during the lifelong learning process. Inspired by the biological process of human cognition where the somatosensory neocortex and the hippocampus work together in memory consolidation, we formulated a model called Knowledge Refreshing and Consolidation (KRC) that achieves both positive forward and backward transfer. More specifically, a knowledge refreshing scheme is incorporated with the knowledge rehearsal mechanism to enable bi-directional knowledge transfer by introducing a dynamic memory model and an adaptive working model. Moreover, a knowledge consolidation scheme operating on the dual space further improves model stability over the long term. Extensive evaluations show KRC's superiority over the state-of-the-art LReID methods on challenging pedestrian benchmarks. Code is available at <https://github.com/cly234/LReID-KRKC>.

Introduction

Person re-identification (ReID) aims to recognize the same person reappearing in different camera views and locations. Person ReID has long been investigated in surveillance systems as a person retrieval problem, where myriad efforts have been made to tackle various challenges such as occlusions and pose variations. However, the current paradigm often sidesteps the desiderata of real-world deployments to successively learn timely-updated new tasks and retain old knowledge in an ever-changing environment.

In particular, when ReID systems are incrementally learned on a sequence of data collected from diverse locations and domains over time, they will often fit in the current data distribution to the detriment of previously gained

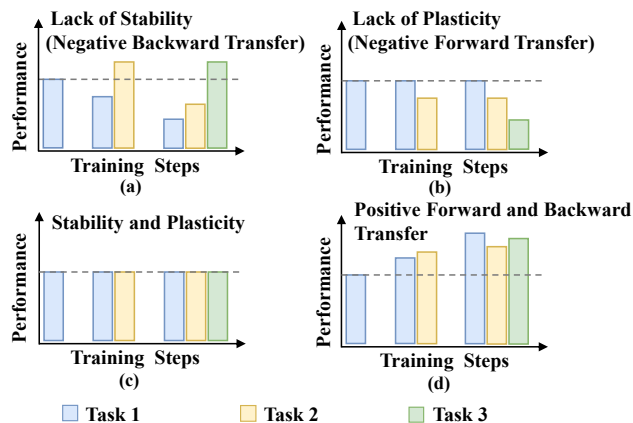


Figure 1: Four scenarios of forward and backward transfer for lifelong learning. The dashed line denotes when a single task has been trained separately at random initialization.

knowledge. As a result, *catastrophic forgetting* may occur. This suboptimal result has inspired the community to seek sustainable approaches to lifelong person ReID (LReID). Compared to conventional lifelong learning, which focuses on closed-set recognition problems, LReID is inherently a fine-grained open-set problem. The identities in the training and testing sets for LReID are non-overlapped, and one of the goals of LReID is to improve the model's generalization ability among various distributions.

To date, LReID has been the topic of several studies. For example, (Pu et al. 2021) studied knowledge accumulation, while (Wu and Gong 2021) and (Ge et al. 2022) both examined ways to replay selected exemplars with ad-hoc techniques. While these methods leverage knowledge distillation or knowledge accumulation as a countermeasure to prevent catastrophic forgetting, dealing with this issue leaves us far from solving the whole problem. Here, we target a more challenging LReID problem, which is not only trying to minimize catastrophic forgetting but also tackling the idea that model performance can be improved on both new and old tasks during the lifelong learning process. Generally, lifelong learning models can be classified into four categories according to the stability and plasticity model performance: Fig. 1(a) shows that models that lack stability are

*Corresponding author

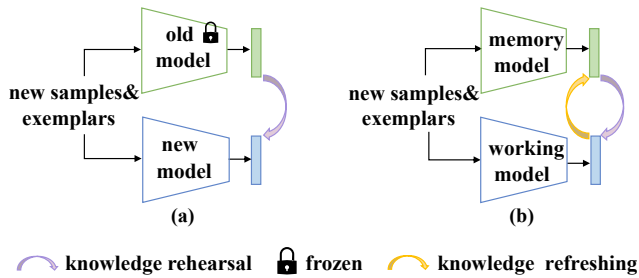


Figure 2: Illustration of our method compared with other rehearsal-based methods for LReID. (a) represents a common paradigm of other rehearsal-based approaches that utilize knowledge distillation and rehearsal. (b) represents the centerpiece of our method, which comprises knowledge refreshing with rehearsal and knowledge consolidation.

prone to forgetting old tasks; Fig. 1(b) shows a model that has overcome the catastrophic forgetting problem but lacks plasticity to handle new tasks; Fig.1(c) shows a model that maintains good performance in terms of both stability and plasticity, however, it lacks positive knowledge transfer in both directions; Lastly, Fig. 1(d) represents our target model with both positive forward and backward knowledge transfer. The model in Fig. 1(d) not only allows previous knowledge to benefit the current task, but also improves the performance of old tasks when learning the current task. Clearly, the model in Fig. 1(d) is the most desirable for LReID.

However, pursuing a lifelong model with both positive forward and backward transfer is very challenging. The first challenge is *distribution mismatch*. The natural distribution gaps are prevalent between incrementally added task data. Such sharp distribution gaps will inevitably hinder the efficacy of knowledge transfer between sequential tasks. Unfortunately, previous methods ignore the importance of a smooth transfer, leading to inferior forward transfer on the current task. The second challenge is *capacity mismatch*. Apart from distribution shifts, the incoming datasets may also have non-stationary representation capacities. If so, distillation from a past model will give biased supervision with a weaker generalization ability. This is a hindrance towards both forward and backward transfer.

In light of these challenges, we propose a novel knowledge refreshing and consolidation framework built on top of a knowledge rehearsal baseline. Our design draws inspiration from the human cognitive process, where the somatosensory neocortex and hippocampus coordinate with each other for various functions (Sirota et al. 2003; Mizuseki and Miyawaki 2017). Specifically, our model is equipped with a dynamic memory model (i.e., a hippocampus) and an adaptive working model (i.e., somatosensory neocortex). Then a knowledge refreshing scheme based on knowledge rehearsal operates in an iterative manner to alleviate the negative effects of the distribution shift and capacity mismatch. The interplay between the knowledge rehearsal and the refreshment is based on two considerations. First, the knowledge rehearsal flexibly explores new knowledge without severe forgetting; meanwhile, the knowledge refreshing

dynamically calibrates the memory model to avoid harming the performance of the current task, achieving *positive forward transfer*. Second, the knowledge refreshment gradually transfers new insights to the memory model to alleviate the response bias of the memory model, yielding *positive backward transfer*. Further, as the task sequences become longer, the memory model gradually becomes forgetful of past experiences. To this end, we further pave the way for knowledge consolidation from model and feature level to mutually enhance stability.

To further shed light on our main idea, Fig. 2 illustrates the differences between our method and other rehearsal-based methods. Standard rehearsal-based methods, coupled with distillation approaches, are highlighted as an effective solution to lifelong problems, interleaving training samples from the current task with previous exemplars (Fig. 2(a)). Unfortunately, there are several drawbacks in dealing with LReID. As training progresses, the old model is kept static with frozen weights. However, we argue that this strategy is deficient in dynamic retrospection, and therefore prevents the old model from becoming a better teacher. By contrast, our method frees the old model from being unchanged and leaves more space for mutual transfer (Fig. 2(b)).

To summarize, our contributions are as follows:

- We propose a novel algorithm named Knowledge Refreshing and Consolidation (KRC) for lifelong person ReID. In addition to dealing with catastrophic forgetting, the proposed KRC also improves model performance on both old and new tasks during the lifelong learning process.
- Beyond rehearsal, we introduce a knowledge refreshing scheme where the dynamic memory model and the adaptive working model iteratively interact with bi-directional knowledge transfers. Moreover, a knowledge consolidation scheme operating on the dual space is introduced to further improve model stability over the long term.
- Extensive experiments demonstrate that KRC outperforms the current state-of-the-art LReID methods by a large margin. Plus, it offers both positive forward transfer and positive backward transfer abilities.

Related Work

Lifelong Learning. Lifelong learning, also known as continual learning, strives to exploit past and current knowledge to adapt a model to new tasks, while maintaining stable performance with old tasks. In general, the research lines of lifelong learning can be grouped into four categories: regularization-based methods (Li and Hoiem 2017; Kirkpatrick et al. 2017; Zhao et al. 2020), structure-based methods (Serra et al. 2018; Yan, Xie, and He 2021), algorithm-based methods (Lopez-Paz and Ranzato 2017; Guo et al. 2022), and rehearsal-based methods (Rebuffi et al. 2017; Hou et al. 2019; Wu et al. 2019). As our work falls into the rehearsal-based methods, we have concentrated on this stream in this literature review. Generally, rehearsal-based methods allow one to access a finite number of previous samples, which greatly alleviates the catastrophic forgetting

problem. For example, Hou *et al.* (Hou et al. 2019) exploits cosine normalization and a less-forget constraint to learn a unified feature representation, while Wu *et al.* (Wu et al. 2019) involves a bias correction module to estimate feature drifts in a fully connected layer. Zhao *et al.* (Zhao et al. 2020) proposes a weight-aligning method, and Yan *et al.* (Yan, Xie, and He 2021) dynamically adds a new feature extractor per task to expand the representation capability. Interestingly, another method Lopez-Paz *et al.* reveals that *forward transfer* and *backward transfer* both matter for continual learning. To this end, only a few methods (Lopez-Paz and Ranzato 2017; Yan, Xie, and He 2021) make some progress in their tasks but are not applicable in LReID. With this as motivation, the proposed method attempts to pursue a strategy that simultaneously encourages positive forward and backward transfers, all within a limited memory budget.

Person Re-Identification. Person ReID has been explored from different perspectives in different settings. Fully supervised methods (FS) exploit different architectures and fine-grained features to learn a robust feature representation (Luo et al. 2019; He et al. 2021; Chen et al. 2021; Gao et al. 2020). These methods learn a stationary distribution at their best, but often require repeatedly observing the entire prepared dataset. Later approaches investigate unsupervised domain adaptive transfer (UDA) (Wang et al. 2018; Fu et al. 2019; Zheng et al. 2021) or purely unsupervised setups (PU) (Yu et al. 2019; Lin et al. 2019; Wang and Zhang 2020). These both seek to use a large corpus of unlabeled images to guide the learning process, either for adaptation purposes or from an unsupervised perspective. While these methods are more readily deployed or flexibly adapted, they usually deliver inferior performance compared with supervised methods. Apart from unsupervised learning, active learning (Wang et al. 2016; Liu et al. 2019) provides a human-in-the-loop manner to alleviate the cost of annotation. In addition to the method described above, other approaches attempt to leverage different modalities (Wu et al. 2017; Zhang et al. 2021, 2023; Han et al. 2022) or attributes (Wang et al. 2017; Tan et al. 2019; Cao, Wang, and Tao 2020) to enhance the representation of features.

Recently, Lifelong ReID has been proposed by Pu *et al.* (Pu et al. 2021) to meet the demand of long-term visual search. This paradigm is imperative for real-world applications, as it requires the ReID system to consolidate knowledge from ever-changing distributions (i.e., newly appearing classes or domains) in a task-incremental fashion. AKA (Pu et al. 2021) constructs learnable knowledge graphs which adaptively accumulate knowledge and preserve topologies. GwReID (Wu and Gong 2021), for example, includes comprehensive learning objectives so as to maintain several coherences during training. PTKP (Ge et al. 2022) casts lifelong ReID as a source-free domain adaptation problem, where old tasks are treated as a source domain. This way, the feature space for a new task can be mapped to that of an old task for domain consistency learning. However, when training on a new task, these methods typically focus on retaining old knowledge from a static teacher which may limit the upper bound of performance with the old task (Lopez-Paz and Ranzato 2017). By contrast, our work not only asks

the model to forget less of previous tasks, but it also leaves space for positive knowledge transfer on previous tasks.

Method

Preliminary

Problem Formulation This section begins by introducing the problem formulation of Lifelong person ReID. Formally, given N sequentially incoming ReID datasets $\mathcal{D} = \{\mathcal{D}^i\}_{i=1}^N$, we aim to give our ReID model the ability to consolidate sequentially-gained multi-domain knowledge and further encourage positive backward transfer. Specifically, at each step t , $\mathcal{D}^t = \{\mathcal{D}_{train}^t, \mathcal{D}_{test}^t\}$ contains the training set and testing set. As in previous works, our method builds up an exemplar memory \mathcal{M}^t , where a limited number of samples from each previous task are saved.

At training step t , we first sample a batch of images $\{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{N_b}$ from \mathcal{D}_{train}^t and a batch of exemplars $\{(\mathbf{x}_i^r, y_i^r)\}_{i=1}^{N_b}$ from \mathcal{M}^t for replaying, where $\mathbf{x}_i^t, \mathbf{x}_i^r$ represent the images, and y_i^t, y_i^r represent the corresponding labels. Plus, a working model Θ_t^w and a memory model Θ_t^m are both leveraged for representation learning. In the testing stage of step t , the classifier is discarded and the trained feature extractor is used to evaluate on the currently seen domains, which are $\{\mathcal{D}_{test}^i\}_{i=1}^t$, since the training classes and testing classes are non-overlapping.

Method Overview Our method decomposes our learning process into three stages: knowledge rehearsal, knowledge refreshing, and knowledge consolidation. Initially, the first two stages iteratively work together to enable bi-directional knowledge transfer. The knowledge rehearsal allows adaptation of the new tasks and prevents catastrophic forgetting, while the knowledge refreshing scheme mimics the behavior of the knowledge rehearsal mechanism, enabling calibration and memorization of the memory model. When the task is complete, we blend the knowledge of the working model and the memory model in the stage of knowledge consolidation to further enhance stability. In this stage, we exploit the model space consolidation for more stable training, and we explore the feature space consolidation for the generalizable ability.

Knowledge Rehearsal

In this section, we exploit knowledge rehearsal to allow the model to adapt flexibly to new tasks and to prevent old tasks from being forgotten.

To enable knowledge transfer from the memory model to the working model, we use knowledge distillation to minimize the discrepancy between the predictions. Specifically, $\{\mathbf{p}_i(\cdot)\}_{i=1}^{N_b}, \{\mathbf{q}_i(\cdot)\}_{i=1}^{N_b}$ denote the softmax predictions from the working model and the memory model, taking (\cdot) as the input. Here, new samples are exploited for anti-forgetting, which is formulated as:

$$\mathcal{L}_{anti}^w = \frac{1}{N_b} T^2 \sum_{i=1}^{N_b} \text{JS}(\mathbf{p}_i(\mathbf{x}_i^t) || \text{SD}(\mathbf{q}_i(\mathbf{x}_i^t))), \quad (1)$$

where T represents the temperature, $\text{SD}(\cdot)$ represents the stop gradient operation, $\text{JS}(\cdot || \cdot)$ refers to a Jensen-Shannon divergence (Fuglede and Topsøe 2004).

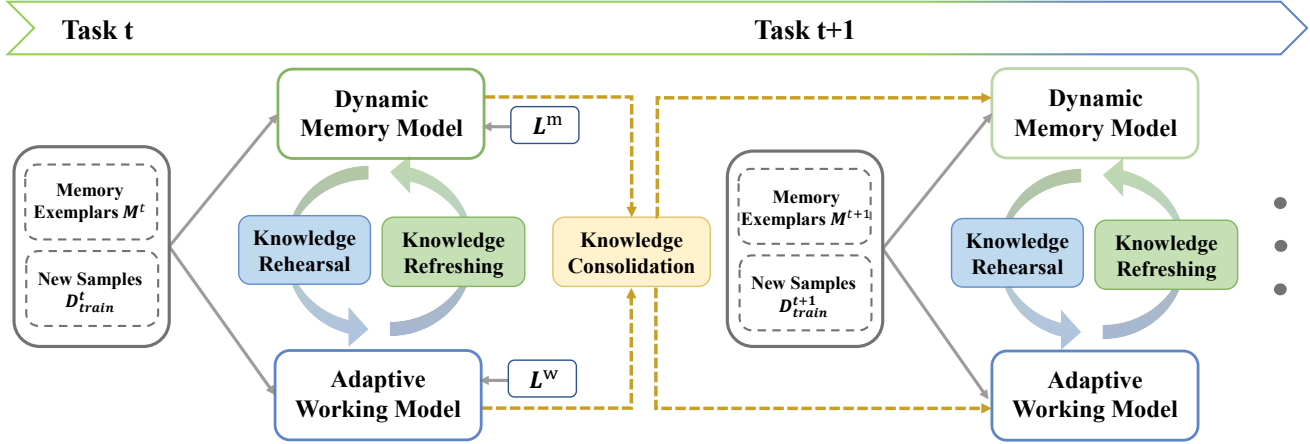


Figure 3: An overview of the proposed KRC method for Lifelong person ReID. We first equip our framework with a dynamic memory model and an adaptive working model. Then the memory exemplars and new samples are fed into these models iteratively for knowledge rehearsal and refreshing in bi-directional transfer. After training, the acquired knowledge of these two models is consolidated in preparation for future tasks. The solid lines represent the data flow of both exemplars and new samples, and the dashed lines represent the consolidation process. \mathcal{L}^w and \mathcal{L}^m denote the overall losses for knowledge rehearsal and knowledge refreshing, respectively.

Another goal of the knowledge rehearsal is to make the adaptation flexible. Therefore, a loss for adaptation is needed, which contains a cross-entropy loss and a triplet loss. Here, the cross-entropy loss is used on new samples, so we have:

$$\mathcal{L}_{ce}^w = -\frac{1}{N_b} \sum_{i=1}^{N_b} \log p_i(\mathbf{x}_i^t)^{(y_i^t)}, \quad (2)$$

where $p_i(\mathbf{x}_i^t)^{(y_i^t)}$ is the y_i^t -th digit of the softmax prediction $\mathbf{p}_i(\mathbf{x}_i^t)$. Meanwhile, we optimize the working model on both current samples and exemplars for metric learning, where the triplet loss is defined as:

$$\begin{aligned} \mathcal{L}_{trip}^w = & \frac{1}{N_b} \sum_{i=1}^{N_b} \max(d(\mathbf{a}_i^{w,t}, \mathbf{p}_i^{w,t}) - d(\mathbf{a}_i^{w,t}, \mathbf{n}_i^{w,t}) + m, 0) \\ & + \frac{1}{N_b} \sum_{i=1}^{N_b} \max(d(\mathbf{a}_i^{w,r}, \mathbf{p}_i^{w,r}) - d(\mathbf{a}_i^{w,r}, \mathbf{n}_i^{w,r}) + m, 0), \end{aligned} \quad (3)$$

where $d(\cdot, \cdot)$ represents the euclidean distance and m is the margin; $(\mathbf{a}_i^{w,t}, \mathbf{p}_i^{w,t}, \mathbf{n}_i^{w,t})$ and $(\mathbf{a}_i^{w,r}, \mathbf{p}_i^{w,r}, \mathbf{n}_i^{w,r})$ are two triplets of anchor sample, positive sample and negative sample of \mathbf{x}_i^t and \mathbf{x}_i^r , respectively. Here, the triplets are extracted from the working model.

Now the working model is equipped with two learning objectives and we get the overall knowledge rehearsal loss \mathcal{L}^w to update the working model:

$$\mathcal{L}^w = \mathcal{L}_{anti}^w + \mathcal{L}_{ce}^w + \mathcal{L}_{trip}^w \quad (4)$$

It is worth noting that, to enable positive forward transfer, exemplars for the cross-entropy loss and the distillation loss are excluded, and only exploited for the triplet loss. This

works to avoid over-confident predictions and supervision of exemplars given by the softmax functions. As such, the model can more flexibly adapt to new domains while preserving old knowledge at the same time.

Knowledge Refreshing

Whereas distillation-based knowledge rehearsal methods explore learned experiences from the memory model, such experiences could be biased. For this reason, distilling two sharply different distributions may not be very effective if the working model veers towards the current task. By contrast, our approach leverages the new identities in the current task as negative classes to update the memory model in a slow fashion, inspiring novel insights and refreshing the old knowledge. More specifically, we simply propose a knowledge refreshing scheme that mirrors the knowledge rehearsal on a similar note, however, with a smaller learning rate, which turns the memory model into a dynamic teacher.

Our first goal is to calibrate the responses of the memory model. Specifically, we use a distillation loss that is similar to \mathcal{L}_{anti}^w but we update the memory model in a smooth manner, which is different from previous practice. This way, the responses of the memory model can be better calibrated. Thus, \mathcal{L}_{cali}^m for memory calibration is formulated as:

$$\mathcal{L}_{cali}^m = \frac{1}{N_b} T^2 \sum_{i=1}^{N_b} \text{JS}(q_i(\mathbf{x}_i^t) || \text{SD}(\mathbf{p}_i(\mathbf{x}_i^t))), \quad (5)$$

In addition, we also enhance the ability of the memory model to memorize various patterns of person identities. Specifically, we include a cross-entropy loss on new sam-

ples to update the memory model:

$$\mathcal{L}_{ce}^m = -\frac{1}{N_b} \sum_{i=1}^{N_b} \log q_i(\mathbf{x}_i^t)^{(y_i^t)}, \quad (6)$$

Then we add a triplet loss on the memory model. We define $(\mathbf{a}_i^{m,t}, \mathbf{p}_i^{m,t}, \mathbf{n}_i^{m,t})$ and $(\mathbf{a}_i^{m,r}, \mathbf{p}_i^{m,r}, \mathbf{n}_i^{m,r})$ as the triplets of anchor sample, positive sample and negative sample of \mathbf{x}_i^t and \mathbf{x}_i^r . Both of these triplets are extracted from the memory model. Thus, \mathcal{L}_{trip}^m is formulated as:

$$\begin{aligned} \mathcal{L}_{trip}^m &= \frac{1}{N_b} \sum_{i=1}^{N_b} \max(d(\mathbf{a}_i^{m,t}, \mathbf{p}_i^{m,t}) - d(\mathbf{a}_i^{m,t}, \mathbf{n}_i^{m,t}) + m, 0) \\ &+ \frac{1}{N_b} \sum_{i=1}^{N_b} \max(d(\mathbf{a}_i^{m,r}, \mathbf{p}_i^{m,r}) - d(\mathbf{a}_i^{m,r}, \mathbf{n}_i^{m,r}) + m, 0), \end{aligned} \quad (7)$$

Through calibration and memorization, we cast the memory model as a dynamic teacher to offer more meaningful guidance. Therefore, the memory model is optimized through the overall knowledge refreshing loss \mathcal{L}^m , which is formulated as:

$$\mathcal{L}^m = \mathcal{L}_{cali}^m + \mathcal{L}_{ce}^m + \mathcal{L}_{trip}^m, \quad (8)$$

We remark that: (i) Similar to the knowledge rehearsal scheme, the exemplar memory only participates in metric learning, since we empirically find the cross-entropy loss or logit-level distillation on \mathcal{M}^t may possibly put the ReID model at risk of overfitting on the exemplars. This would, of course, harm the performance to generalize on previous ReID tasks. (ii) While the learning objectives are similar to those of the knowledge rehearsal scheme, there is a different motivation for this stage. With these learning objectives, the memory model was better able to communicate with the working model and could then refresh itself through gradual adaptation. (iii) The knowledge rehearsal stage and the refreshing stage iteratively coordinate with each other, mimicking the information flow between the hippocampal–entorhinal loop during the sleep/wake cycle (Mizuseki and Miyawaki 2017). The structure of our algorithm is given in Alg. 1.

Knowledge Consolidation

To further combat catastrophic forgetting and improve the model’s generalization ability for long task sequences, we carefully consolidated the old and the new knowledge in the model space and the feature space, with the purpose of retaining a minimum for all tasks.

Model Space Consolidation Normally, memory models play a key role in distilling old knowledge into working models. However, the memory model is discarded immediately after training the current task, and it is not involved in subsequent tasks. To account for this disparity, at the end of each step, our working model is reorganized with the parameters of the memory model to linearly combine their merits. Intuitively, greater importance is placed on the memory model as the task number increases since the memory model

Algorithm 1: KRKC

Input: Incoming datasets \mathcal{D}_{train}^t , Buffer \mathcal{M}^t , Learning rate γ, η .

Parameter: Parameters of working model: Θ_t^w ; Parameters of memory model: Θ_t^m .

```

1: for epoch in  $1 \rightarrow e_{max}$  do
2:   for  $\{\mathbf{x}_i^t, y_i^t\}_{i=1}^{N_b}$  in  $\mathcal{D}_{train}^t$  do
3:     Sample a mini-batch  $\{\mathbf{x}_i^r, y_i^r\}_{i=1}^{N_b}$  from  $\mathcal{M}^t$ .
4:     Knowledge rehearsal for anti-forgetting by Eq. (1).
5:     Knowledge rehearsal for adaptation by Eq. (2), Eq. (3).
6:     Calculate overall loss  $\mathcal{L}^w$  by Eq. (4).
7:     Update  $\Theta_t^w$  by gradient descent:
8:      $\vartheta \leftarrow \vartheta - \gamma \nabla_{\vartheta} \mathcal{L}^w$  for  $\vartheta$  in  $\Theta_t^w$ .
9:     Knowledge refreshing for calibration by Eq.(5).
10:    Knowledge refreshing for memorization by Eq.(6), Eq. (7).
11:    Calculate overall loss  $\mathcal{L}^m$  by Eq. (8).
12:    Update  $\Theta_t^m$  by gradient descent:
13:     $\vartheta \leftarrow \vartheta - \eta \nabla_{\vartheta} \mathcal{L}^m$  for  $\vartheta$  in  $\Theta_t^m$ .
14:  end for
15: end for
16: Model space consolidation by Eq. (9).
17: Feature space consolidation by Eq. (10).
18: Update exemplar memory:  $\mathcal{M}^t \rightarrow \mathcal{M}^{t+1}$ .
19: return updated model parameters  $\Theta_{t+1}^w, \Theta_{t+1}^m$ .

```

comprises richer knowledge of previous tasks in the long run. Formally, the update rule is defined as:

$$\Theta_{t+1}^w = \frac{1}{t+1} \Theta_t^w + \frac{t}{t+1} \Theta_t^m, \quad (9)$$

Next, a copy of Θ_{t+1}^w is retained to initialize Θ_{t+1}^m for the next task.

Feature Space Consolidation Thanks to the knowledge rehearsal and refreshment mechanism, we now have two models that perform comparably well on the current task and previous tasks, but with different focuses. To take full advantage of the working model and the memory model, we devise a simple but effective way to consolidate the feature space, which is also important for generalization ability. Specifically, once the task is complete, we have two features $\mathbf{f}_t^w, \mathbf{f}_t^m$ from the trained feature extractors Θ_t^w and Θ_t^m . These two features are concatenated and the fused feature is used for person retrieval, which is tantamount to:

$$\mathbf{f}_t^{fuse} = [\mathbf{f}_t^w; \mathbf{f}_t^m]. \quad (10)$$

Experiment

Datasets and Evaluation Protocol

To evaluate our strategy’s performance with LReID tasks, we assess our model on a challenging benchmark with four sequential datasets: VIPeR (Gray and Tao 2008), Market (Zheng et al. 2015), CUHK-SYSU (Xiao et al. 2016), and

Method	VIPeR		Market		CUHK-SYSU		MSMT17		Average	
	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	\bar{s}_{R-1}	\bar{s}_{mAP}
Joint	78.5	85.7	93.7	85.4	89.8	88.1	74.7	49.4	84.2	77.2
LwF (Li and Hoiem 2017)	45.3	54.6	74.6	53.9	83.7	81.1	32.3	14.6	59.0	51.1
iCarL (Rebuffi et al. 2017)	56.6	67.0	78.7	56.9	82.9	80.0	24.1	10.2	60.6	53.6
UCIR (Hou et al. 2019)	56.3	66.8	65.9	45.5	70.5	68.8	29.3	12.8	55.5	48.5
BiC (Wu et al. 2019)	50.4	61.2	68.9	47.4	72.5	71.6	42.1	22.4	58.5	50.7
WA (Zhao et al. 2020)	48.1	58.0	70.3	50.9	71.9	70.4	38.8	18.6	57.2	49.5
AKA (Pu et al. 2021)	50.6	61.7	50.7	28.3	79.6	76.9	28.0	13.4	52.2	45.1
PTKP (Ge et al. 2022)	56.0	66.3	77.4	58.3	78.8	77.1	48.0	25.2	65.1	56.7
Ours	67.7	76.3	82.5	64.4	90.7	88.9	67.1	43.3	77.0	68.2

Table 1: Performance comparison of the state-of-the-art methods on seen domains. All the reported results are implemented with the released code on our baseline. The training order is VIPeR \rightarrow Market \rightarrow CUHK-SYSU \rightarrow MSMT17.

Method	CUHK01	CUHK02	CUHK03	PRID
Joint	76.6	68.8	31.9	36.6
LwF	60.1	52.0	25.0	27.0
iCarL	60.4	55.1	20.8	34.9
UCIR	40.7	37.6	21.7	22.4
BiC	50.4	44.6	31.2	24.1
WA	50.2	48.5	29.6	21.9
AKA	44.3	42.1	21.2	22.7
PTKP	57.4	52.3	30.6	28.4
Ours	79.2	69.8	39.5	52.7

Table 2: Performance comparison of the state-of-the-art methods over unseen domains. The results (mAP) are given by trained models at the last step.

MSMT17 (Wei et al. 2018). These datasets mimic real-world scenarios. With CUHK-SYSU, we follow the same data pre-processing for training and testing as in PTKP (Ge et al. 2022). Then, to further evaluate our method’s generalization ability on unseen datasets, we compare a range of methods on the CUHK01 (Li, Zhao, and Wang 2012), CUHK02 (Li and Wang 2013), CUHK03 (Li et al. 2014), and PRID (Hirzer et al. 2011) benchmarks.

The evaluation metrics for each task include standard mean Average Precision (mAP) and Rank-1 accuracy (R-1). For the lifelong evaluations on the seen tasks, we report the “Average” accuracy which is the average incremental accuracy of mAP/Rank-1 after the last step (Ge et al. 2022).

Implementation Details

We use a ResNet50 model pre-trained on ImageNet as our backbone. Global average pooling is replaced with generalized mean pooling. All person images are resized to 256×128 . The batch size for the current task and the exemplar task is set to 128. For the memory update, we follow PTKP (Ge et al. 2022), which saves the farthest two samples for each identity, with 250 identities per task. We use Adam for optimization and train each task for 60 epochs.

The learning rate is initialized at 3.5×10^{-4} , which is then decreased by 0.1 after the 40th epoch for the first task. The learning rate in the knowledge refreshing scheme is initialized as 3.5×10^{-5} for subsequent tasks and 3.5×10^{-6} for the last task, all decreased by 0.1 after the 30th epoch. The initial learning rate of the knowledge rehearsal mechanism is set to 3.5×10^{-4} , with the other settings kept the same as the knowledge refreshing scheme. The weights for all losses are set to 1 and the temperature T for distillation is set to 2.

Comparison with the State-of-the-Arts

Results on Seen Domains Table 1 summarizes the results of LReID benchmark on seen domains. As shown, our method consistently outperforms the other rehearsal-based methods and LReID methods by a clear margin, with an average gain of +11.9% w.r.t. the average incremental R-1 accuracy (\bar{s}_{R-1}) and +11.5% w.r.t. the average incremental mAP accuracy (\bar{s}_{mAP}), respectively. Additionally, our method adds more plasticity to new tasks, achieving 67.1% Rank-1 and 43.3% mAP on MSMT with at least +19.1% performance boost. We argue that this owes to our knowledge refreshing scheme which unleashes sufficient flexibility for further exploration of new tasks.

Results on Unseen Domains Table 2 summarizes the results of generalization on unseen domains. Our method reaches a more stable result over four unseen datasets. Compared with other works, our method even outperforms joint training on CUHK01, CUHK03, and PRID. Specifically, our method surpasses joint training by +16.1% on PRID, which demonstrates that our method does enable a smooth transfer and robust generalization ability over various unseen datasets.

Ablation Study

Individual Component Analysis To better understand the contribution of each module and how these modules corroborate each other, we reported the results w.r.t. the accuracy of four tasks at the last training step. From the results in Table 3, we made the following interpretations.

KRH	KRF	MSC	FSC	VIPeR		Market		CUHK-SYSU		MSMT17		Average	
				R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	\bar{s}_{R-1}	\bar{s}_{mAP}
✓				49.4	59.8	62.0	37.2	81.8	78.5	30.1	12.2	55.8	46.9
✓	✓			57.0	67.8	72.1	47.6	87.3	85.4	66.4	42.9	70.7	60.9
✓		✓	✓	60.1	70.0	72.7	51.3	84.6	82.5	34.4	14.8	63.0	54.7
✓	✓	✓		58.5	69.9	72.3	47.9	86.4	84.4	69.9	46.1	71.8	62.1
✓	✓		✓	61.7	71.9	81.7	61.6	91.4	89.6	59.0	36.0	73.5	64.8
✓	✓	✓	✓	67.7	76.3	82.5	64.4	90.7	88.9	67.1	43.3	77.0	68.2

Table 3: Ablation study of individual component analysis. We reported the accuracy of four tasks at the last training step.

First, on top of knowledge rehearsal, the knowledge refreshing scheme (KRF) consistently and gradually adds new insights to the memory model, which in turn helps to learn the working model. This is an indispensable part of maximally encouraging positive backward transfer for previous tasks. Second, knowledge consolidation (KC) comprehensively equips our model with stability and resilience against ever-increasing task sequences. Specifically, model space consolidation (MSC) and feature space consolidation (FSC) both mitigate catastrophic forgetting. More importantly, our model attains significant performance gains when these two modules work together. Further, knowledge consolidation (KC) and knowledge refreshing (KR) can complement each other as their collaboration considers both short-term plasticity and long-term stability.

Stability and Plasticity Analysis

Positive Backward Transfer with Stability To assess the ability to use new data to improve performance on previous tasks, we introduce a memory refreshing scheme to gradually absorb new knowledge and encourage a smooth transfer, which helps preclude the response bias of the memory model. Formally, we define $R_{i,j}$ as the test accuracy on task j after observing task i and T as the total number of tasks. At each step i , we can observe the performance change of task j from $R_{j,j}$ to $R_{i,j}$. By averaging these local accuracy changes we compute the overall backward transfer value (BWT) of all previous tasks as:

$$\text{BWT} = \frac{1}{T-1} \sum_i \frac{1}{i} \sum_{j=1}^i (R_{i,j} - R_{j,j}) \quad (11)$$

As seen in Table 4, our method gives +5.3% positive backward transfer.

Positive Forward Transfer with Plasticity To measure the influence of previous knowledge on the performance of new tasks, we also evaluate the forward transfer ability (FWT), which is defined by:

$$\text{FWT} = \frac{1}{T-1} \sum_{i=2}^T (R_{i,i} - \bar{R}_{i,i}) \quad (12)$$

where $\bar{R}_{i,i}$ represents the accuracy given by the model trained at random initialization. As shown in Table 4, most

Method	LwF	iCaRL	WA	AKA	PTKP	Ours
BWT	-0.9	3.7	-0.5	-1.4	3.4	5.3
FWT	-14.6	-19.1	-17.6	-22.4	-13.2	3.9

Table 4: Backward transfer (BWT) and forward transfer (FWT) measured in R-1 accuracy.

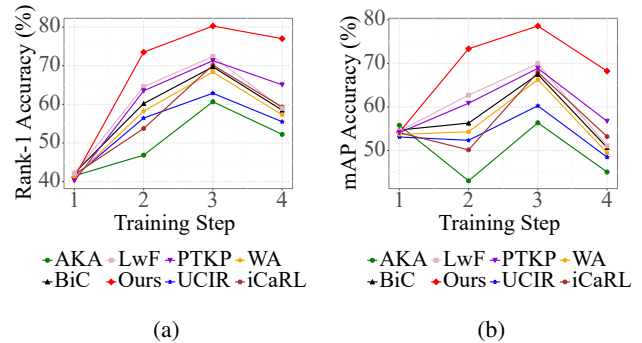


Figure 4: (a) Average incremental accuracy w.r.t. \bar{s}_{R-1} . (b) Average incremental accuracy w.r.t. \bar{s}_{mAP} .

methods sacrifice plasticity to maintain stability, achieving negative forward transfer. Nevertheless, our method achieves +3.9% positive forward transfer in R-1 accuracy, which benefits from our KRC module that forms a closed-loop positive feedback system for subsequent tasks.

Conclusion

In this paper, we propose a novel rehearsal-based framework that not only serves the purpose of anti-forgetting, but also enables both positive forward and backward transfer. To meet this goal, we introduce a knowledge refreshing scheme to iteratively transfer knowledge between the memory and the working modules. Moreover, the knowledge consolidation scheme further improves model stability in the long term. Extensive experiments are conducted on seen and unseen domains. The performance gains over the state-of-the-art methods demonstrate that our method achieves both positive forward and backward transfer for LReID.

Acknowledgements

This work was supported by Shanghai Local College Capacity Building Program (23010503100), Shanghai Sailing Program (21YF1429400, 22YF1428800), Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), National Key R&D Program of China (2018AAA0100704), NSFC (61932020, 62172279), Science and Technology Commission of Shanghai Municipality (20ZR1436000), Program of Shanghai Academic Research Leader, and Shuguang Program supported by Shanghai Education Development Foundation.

References

- Cao, Y.-T.; Wang, J.; and Tao, D. 2020. Symbiotic adversarial learning for attribute-based person search. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Chen, X.; Xu, J.; Xu, J.; and Gao, S. 2021. OH-Former: Omni-Relational High-Order Transformer for Person Re-Identification. *arXiv preprint arXiv:2109.11159*.
- Fu, Y.; Wei, Y.; Wang, G.; Zhou, Y.; Shi, H.; and Huang, T. S. 2019. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Fuglede, B.; and Topsoe, F. 2004. Jensen-Shannon divergence and Hilbert space embedding. In *International Symposium on Information Theory*.
- Gao, S.; Wang, J.; Lu, H.; and Liu, Z. 2020. Pose-Guided Visible Part Matching for Occluded Person ReID. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ge, W.; Du, J.; Wu, A.; Xian, Y.; Yan, K.; Huang, F.; and Zheng, W.-S. 2022. Lifelong Person Re-identification by Pseudo Task Knowledge Preservation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Gray, D.; and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Guo, Y.; Hu, W.; Zhao, D.; and Liu, B. 2022. Adaptive Orthogonal Projection for Batch and Online Continual Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Han, X.; Cong, P.; Xu, L.; Wang, J.; Yu, J.; and Ma, Y. 2022. LiCamGait: Gait Recognition in the Wild by Using LiDAR and Camera Multi-modal Visual Sensors. *arXiv preprint arXiv:2211.12371*.
- He, S.; Luo, H.; Wang, P.; Wang, F.; Li, H.; and Jiang, W. 2021. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Hirzer, M.; Belezni, C.; Roth, P. M.; and Bischof, H. 2011. Person re-identification by descriptive and discriminative classification. In *Scandinavian Conference on Image Analysis*.
- Hou, S.; Pan, X.; Loy, C. C.; Wang, Z.; and Lin, D. 2019. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences (PNAS)*.
- Li, W.; and Wang, X. 2013. Locally aligned feature transforms across views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, W.; Zhao, R.; and Wang, X. 2012. Human reidentification with transferred metric learning. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*.
- Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Lin, Y.; Dong, X.; Zheng, L.; Yan, Y.; and Yang, Y. 2019. A bottom-up clustering approach to unsupervised person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Liu, Z.; Wang, J.; Gong, S.; Lu, H.; and Tao, D. 2019. Deep reinforcement active learning for human-in-the-loop person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Luo, H.; Jiang, W.; Gu, Y.; Liu, F.; Liao, X.; Lai, S.; and Gu, J. 2019. A strong baseline and batch normalization neck for deep person re-identification. *IEEE Transactions on Multimedia (TMM)*.
- Mizuseki, K.; and Miyawaki, H. 2017. Hippocampal information processing across sleep/wake cycles. *Neuroscience research*, 118: 30–47.
- Pu, N.; Chen, W.; Liu, Y.; Bakker, E. M.; and Lew, M. S. 2021. Lifelong person re-identification via adaptive knowledge accumulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. Incremental classifier and representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Serra, J.; Suris, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning (ICML)*.

- Sirota, A.; Csicsvari, J.; Buhl, D.; and Buzsáki, G. 2003. Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences (PNAS)*.
- Tan, Z.; Yang, Y.; Wan, J.; Hang, H.; Guo, G.; and Li, S. Z. 2019. Attention-based pedestrian attribute analysis. *IEEE transactions on image processing*, 28(12): 6126–6140.
- Wang, D.; and Zhang, S. 2020. Unsupervised person re-identification via multi-label classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, H.; Gong, S.; Zhu, X.; and Xiang, T. 2016. Human-in-the-loop person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Wang, J.; Zhu, X.; Gong, S.; and Li, W. 2017. Attribute recognition by joint recurrent learning of context and correlation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Wang, J.; Zhu, X.; Gong, S.; and Li, W. 2018. Transferable Joint Attribute-Identity Deep Learning for Unsupervised Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wu, A.; Zheng, W.-S.; Yu, H.-X.; Gong, S.; and Lai, J. 2017. RGB-infrared cross-modality person re-identification. In *Proceedings of the IEEE international conference on computer vision (ICCV)*.
- Wu, G.; and Gong, S. 2021. Generalising without forgetting for lifelong person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Wu, Y.; Chen, Y.; Wang, L.; Ye, Y.; Liu, Z.; Guo, Y.; and Fu, Y. 2019. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2016. End-to-end deep learning for person search. *arXiv preprint arXiv:1604.01850*.
- Yan, S.; Xie, J.; and He, X. 2021. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yu, H.-X.; Zheng, W.-S.; Wu, A.; Guo, X.; Gong, S.; and Lai, J.-H. 2019. Unsupervised person re-identification by soft multilabel learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zhang, J.; Shi, Y.; Xu, L.; Yu, J.; and Wang, J. 2023. IKOL: Inverse kinematics optimization layer for 3D human pose and shape estimation via Gauss-Newton differentiation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Zhang, R.; Xu, L.; Yu, Z.; Shi, Y.; Mu, C.; and Xu, M. 2021. Deep-irtarget: An automatic target detector in infrared imagery using dual-domain feature extraction and allocation. *IEEE Transactions on Multimedia*, 24: 1735–1749.
- Zhao, B.; Xiao, X.; Gan, G.; Zhang, B.; and Xia, S.-T. 2020. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zheng, K.; Lan, C.; Zeng, W.; Zhang, Z.; and Zha, Z.-J. 2021. Exploiting sample uncertainty for domain adaptive person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*.
- Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.