

Lifelong Unsupervised Domain Adaptive Person Re-identification with Coordinated Anti-forgetting and Adaptation

Zhipeng Huang^{1†*} Zhizheng Zhang^{2†‡} Cuiling Lan^{2‡} Wenjun Zeng³ Peng Chu²
Quanzeng You² Jiang Wang² Zicheng Liu² Zheng-jun Zha¹

¹University of Science and Technology of China

²Microsoft ³EIT Institute for Advanced Study

hzp1104@mail.ustc.edu.cn zengw2011@hotmail.com zhazj@ustc.edu.cn

{zhizhzhang, culan, pengchu, quyou, jiangwang, zliu}@microsoft.com

Abstract

Unsupervised domain adaptive person re-identification (ReID) has been extensively investigated to mitigate the adverse effects of domain gaps. Those works assume the target domain data can be accessible all at once. However, for the real-world streaming data, this hinders the timely adaptation to changing data statistics and sufficient exploitation of increasing samples. In this paper, to address more practical scenarios, we propose a new task, Lifelong Unsupervised Domain Adaptive (LUDA) person ReID. This is challenging because it requires the model to continuously adapt to unlabeled data in the target environments while alleviating catastrophic forgetting for such a fine-grained person retrieval task. We design an effective scheme for this task, dubbed CLUDA-ReID, where the anti-forgetting is harmoniously coordinated with the adaptation. Specifically, a meta-based Coordinated Data Replay strategy is proposed to replay old data and update the network with a coordinated optimization direction for both adaptation and memorization. Moreover, we propose Relational Consistency Learning for old knowledge distillation/inheritance in line with the objective of retrieval-based tasks. We set up two evaluation settings to simulate the practical application scenarios. Extensive experiments demonstrate the effectiveness of our CLUDA-ReID for both scenarios with stationary target streams and scenarios with dynamic target streams.

1. Introduction

Person re-identification (ReID) aims to identify the same person across different locations, time instances and cameras. It is of high value in a wide range of applications [36]

*This work was done when Zhipeng Huang was an intern at MSRA.

†Equal contribution.

‡Corresponding authors.




Figure 1. Illustration of the regular UDA (a), the lifelong UDA in the stationary target scenario (b) and dynamic target scenario (c).

in diverse environments, such as customer activity analysis, missing children finding. In practice, a trained ReID model usually suffers from severe performance drops when deployed in new environments due to the domain gaps between the training data and the data in new environments. Unsupervised domain adaptive (UDA) person ReID has been widely studied to help deployed person ReID models better adapt to new environments, by transferring learned knowledge from labeled source domain to unlabeled target domain [1, 9, 10, 31, 45, 46].

Existing UDA person ReID methods [1, 9, 10, 31, 45, 46] assume that the data in the target domain can be accessible all at once, which does not accord well with practical scenarios. In deployment environments, new people are often showing up constantly and the data statistics may vary over time. This imposes strong demands in developing a practical person ReID model/system that can support 1) flexible adaptation to data statistics on the fly so that we can timely obtain adapted performance without waiting for a long time to collect massive data and 2) continuous knowledge acquisition from increasing data. Recent works begin to look into lifelong (supervised) person ReID [27, 39]. They greatly advance practical applications by supporting continuous su-

ervised domain adaptation, but still leave much exploration room in the following aspects. First, these works require costly annotations in the target domain. Besides human labour efforts, data labelling is time consuming, which is not conducive to quick adaptation and privacy protection. Second, they all aim to balance anti-forgetting and adaptation, lacking explicit consideration for their coordination. In fact, anti-forgetting can facilitate flexible/quick adaptation when harnessed in the right manner. In this paper, we propose and study a more practical task, *i.e.*, Lifelong Unsupervised Domain Adaptive (LUDA) person ReID, which does not require annotations in the target domain.

LUDA person ReID enables deployed models to achieve continuous domain adaptation using unlabeled target stream. To the best of our knowledge, we are the first to define this task. Different from conventional life-long learning [15, 20, 34], our proposed task is challenging attributable to two reasons. First, person ReID is a fine-grained retrieval task. Critical to this task is to capture the inter-instance relations for accurately ranking instances. Second, in the context of lifelong learning, person ReID requires higher generalization ability on unseen identities/categories since there are constantly new people showing up, beyond the requirements for remembering how to classify seen categories as in classification tasks. Thus, for *LUDA person ReID*, we expect that anti-forgetting can promote adaptation by avoiding over-fitting new data and facilitating learning more generalizable features so that we can achieve fast adaptation especially when the target domain changes or reappears.

In this paper, we design an effective scheme for *LUDA person ReID* in which anti-forgetting and adaptation are coordinated, dubbed *CLUDA-ReID*. In previous studies [2, 5, 25, 40], anti-forgetting is commonly placed into a joint learning framework together with the training for adaptation, lacking explicit considerations on their coordination. We propose Coordinated Data Replay (CDR) to solve this problem. Particularly, we store old data with a modified ID-wise reservoir sampling algorithm, and design a meta-based optimization strategy to align the objective of anti-forgetting with that of adaptation when replaying stored old data. Moreover, we maintain a model that is timely updated and a momentum-updated model that accumulates learned knowledge over time. We further introduce Relational Consistency Learning (RCL) to promote knowledge distillation/inheritance from the historical model to the current model when performing adaptation. Different from other knowledge distillation strategies [39] in conventional life-long learning, RCL is carefully designed to match the task characteristics of person ReID. We introduce two scenarios with stationary or dynamic target streams as illustrated in Fig. 1, and evaluate our *CLUDA-ReID* on these settings. Besides, we build a new dataset named MMP-Retrieval upon the MMPTRACK dataset for unseen-domain gener-

alization evaluation, released at <https://iccv2021-mmp.github.io/subpage/dataset.html>.

Our contributions can be summarized in three aspects:

- We propose a novel yet realistic task, Lifelong Unsupervised Domain Adaptive (LUDA) person re-identification, to endow person ReID models with automatic adaptation using continuously collected unlabeled data.
- We design an effective *LUDA person ReID* scheme, *i.e.*, *CLUDA-ReID*, wherein Coordinated Data Replay (CDR) and Relational Consistency Learning (RCL) are proposed to explicitly coordinate anti-forgetting and adaptation.
- We set up two practical evaluation schemes to simulate real application scenarios. Extensive experiments demonstrate that our proposed method is effective in achieving LUDA for scenarios with stationary or dynamic target streams and enhancing unseen-domain generalization.

2. Related Works

2.1. Domain Adaptive Person Re-identification

Person re-identification (ReID) has been widely investigated and applied in many real-world scenarios. Unsupervised domain adaptive (UDA) person ReID is of high practical values in transferring the knowledge from labeled source domain to unlabeled target domain for improving the accuracy on the target domain after supervised fine-tuning. Approaches in this field can be grouped into three main categories: (1) Style translation based methods [12, 22]; (2) Pseudo labelling based methods [10, 43, 44, 49]; (3) Domain-invariant feature learning based methods [16, 21, 23]. Albeit their effectiveness, they all assume that the target domain can be entirely accessible but this may not hold in practice. In the real-world, target data usually comes in a stream and its statistics may vary over time.

Recently, a few works [27, 39] study the lifelong (supervised) person ReID problem wherein supervised domain adaptation is continually performed. In [39], consistencies from different perspectives (*i.e.*, classification, distribution and representation) are incorporated into a comprehensive training objective for distilling knowledge from the old model. But it does not explicitly consider the optimization in line with the objective of learning to rank. Pu *et al.* [27] propose an adaptive knowledge accumulation method for this task. However, both of them consider the supervised setting and lack explicit consideration on coordinating anti-forgetting with adaptation, in the sense that enabling anti-forgetting to enhance adaptation. In this work, we propose and address a new task, *i.e.*, *LUDA person ReID*.

2.2. Lifelong Learning

Lifelong learning aims to learn from ever-expanding data, which is challenging in capturing new knowledge as well as maintaining old knowledge. Approaches in this field

can be summarized into three categories. (1) *Knowledge distillation based methods* [17, 20, 39] transfer already acquired knowledge from a fixed old model maintained at a certain moment to the new model currently being trained via adding consistency regularization between them. Existing works in this view ignore the design of a consistency regularization in line with the task characteristics of person ReID, *i.e.*, learning to rank. This problem is addressed in this paper via our proposed RCL. (2) *Sample storing/generation based methods* recall old knowledge on the stored [26, 29, 34] or synthetic [3, 38] data, interleaved with the training on new collected data. They aim to strike a balance between the knowledge learned on old and new data respectively, lacking explicit consideration on the coordination between different optimization objectives at the task level. (3) *Dynamic model based methods* [4, 30, 42] handle expanding data/knowledge with progressively modified model architecture. Our framework is based on combining the first two categories.

3. Proposed Problem: LUDA Person ReID

3.1. Problem Definition

Consider a realistic scenario: a person ReID model or system is offline trained on labeled source data, then deployed in new environments wherein new data arrives continuously. Usually, the new data has domain gaps with the source data and is hard to be annotated in time. *Can we utilize such unlabeled new data to perform continual domain adaptation so that the deployed model can quickly adapt to new environments?* This motivates the problem of LUDA person ReID.

For *LUDA person ReID* task, we assume that a labeled source dataset $\mathcal{X}^s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$ is available at first, which comprises N_s samples $\mathbf{x}_i^s \in \mathcal{X}^s$ with corresponding ID labels $y_i^s \in \mathcal{Y}_s$. The target stream is divided into T splits over time, denoted by $\mathcal{X}^t = \{\mathbf{x}_i^t\}_{i=1}^{N_t^j}$. For an interval of time indexed by j (called “stage” in the following), N_t^j new unlabeled samples \mathbf{x}_i^t are collected in this stage and compose a set \mathcal{X}_j^t . After each interval/stage, \mathcal{X}_j^t is used to automatically update the deployed person ReID model for continual adaptation, in the sense that we need not wait to collect the entire target dataset or train the model from scratch on the combination of newly collected and historical data many times.

3.2. Learning Setups

In the real world, the statistics of target streams may vary over time in different degrees, corresponding to different scenarios as follows: (i) In the *stationary target scenario*, such as fixed cameras in retail, supermarket, *etc.*, only identities increase as new data comes while the domain/style characteristics maintain a relatively stable state. We can

suppose that there are no distribution shifts between the \mathcal{X}_j^t in different stages. (ii) In the *dynamic target scenario*, such as cameras mounted in cars, drones, *etc.*, the number of seen identities keeps increasing and their statistic characteristics also vary over time, leaving this setup being not only class-incremental but also domain-incremental.

4. Proposed Solution: A Coordinated Scheme

4.1. Overall Pipeline

In our proposed scheme, we first pre-train the person ReID model on labeled source data $\mathcal{X}^s = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$. Then, we deploy the pre-trained model in the target environment and update the model with target stream data $\mathcal{X}^t = \{\mathbf{x}_i^t\}_{i=1}^{N_t^j}$ for achieving lifelong UDA. For each sample, from \mathcal{X}^s or \mathcal{X}^t , we first use a backbone network to extract its corresponding feature tensor, then adopt a global spatial GeM pooling [28] on this feature tensor to obtain a feature vector as its final ID representation. As common practices in this field [14, 24, 48], we employ a triplet loss and a cross-entropy based classification loss (each ID is taken as one category) as person ReID losses on the extracted feature vector for training. In the j -th stage, to acquire new knowledge for unsupervised domain adaptation, we train a timely updated model (called “current model” for brevity) on $\mathcal{X}_j^t = \{\mathbf{x}_i^t\}_{i=1}^{N_t^j}$ with person ReID losses. Here, we generate or update the pseudo ID labels for \mathcal{X}_j^t via a clustering algorithm every few training epochs.

Meanwhile, we recall old knowledge via data replay and knowledge distillation for anti-forgetting. To this end, we set up a memory buffer with a limited size for data replay and maintain a historical model for knowledge distillation. Specifically, we sample historical data stored in the memory to update the current model together with new data in each iteration, and store new samples into the memory at the end of each stage. Besides, we perform knowledge distillation from the maintained historical model to the current model by adding consistency constraints on their outputs.

4.2. Coordinated Data Replay

Recalling old knowledge (for anti-forgetting) and capturing new knowledge (for adaptation) correspond to two different objectives, thus can be viewed as two “tasks”. However, as we discussed in our introduction, these two tasks do not have to conflict with each other. In fact, they can be coordinated to enhance the model’s generalization capability for unseen identities and new domains. Thus, we propose Coordinated Data Replay (CDR) to coordinate the optimization of anti-forgetting and adaptation.

The trainable network parameters of a *LUDA person ReID* model comprise a backbone network θ (such as the ResNet-50 in Fig.2) for ReID feature extraction, and a classifier ϕ for adding the classification loss. In the learning




Figure 2. The overall architecture of the proposed scheme *CLUDA-ReID* for *LUDA person ReID*. In *CLUDA-ReID*, we set up a memory buffer to store samples that have been seen/trained as the old data, and preserve a momentum-updated model as the historical model to accumulate acquired knowledge over time. We fed both the old data (dashed lines) and new data (solid lines) into the current model and the historical model for capturing new knowledge for timely adaptation while enhancing the memorization of old knowledge. We achieve coordinated anti-forgetting and adaptation learning via our proposed Coordinated Data Replay and Relational Consistency Learning.

stage of lifelong unsupervised domain adaptation, the optimization objective of **adaptation** can be formulated as:

$$\min_{\theta, \phi} \mathcal{L}_{Adap}(\theta, \phi) = \min_{\theta, \phi} \mathcal{L}_{Tri}(\theta) + \mathcal{L}_{Cls}(\theta, \phi), \quad (1)$$

where $\mathcal{L}_{Tri}(\theta)$ is the triplet loss applied to the extracted feature vectors of new data, $\mathcal{L}_{Cls}(\theta, \phi)$ denotes the Cross-Entropy based classification loss applied to the logits of new data. Here, the new data refers to the training batch sampled from \mathcal{X}_j^t in the j -th stage.

To recall previously learned knowledge, we set up a limited-size memory buffer for storing data to replay these samples. This memory is updated at the end of each stage with a modified reservoir sampling algorithm. The classical reservoir sampling algorithm [35] is designed to randomly keep N samples in memory from a sequence, with equal probability for each one. Here, to preserve as diverse ID information as possible, we modify the reservoir algorithm from an instance-wise one to an ID-wise one. (See more details in the supplementary.) With both the stored old data and the new data together, the optimization objective of **anti-forgetting** can be formulated as:

$$\min_{\theta, \phi} \mathcal{L}_{AntiF}(\theta, \phi) = \min_{\theta, \phi} \mathcal{L}_{KD}(\theta, \phi) + \mathcal{L}_{Tri}(\theta) + \mathcal{L}_{Cls}(\theta, \phi), \quad (2)$$

where $\mathcal{L}_{KD}(\theta)$ represents the knowledge distillation loss for transferring already acquired knowledge from the maintained historical model to the current model, applied to the

inferred results of both old and new data. We design a novel knowledge distillation loss in line with person ReID task and leave its detailed introduction in the subsequent section. $\mathcal{L}_{Tri}(\theta)$ and $\mathcal{L}_{Cls}(\theta, \phi)$ denote the triplet loss and the classification loss respectively, similar to Eq.(1). Note that they are both applied for replaying the old data stored in the memory buffer. The general optimization objective of lifelong learning with the regular data replay is:

$$\min_{\theta, \phi} \mathcal{L}_{Adap}(\theta, \phi) + \mathcal{L}_{AntiF}(\theta, \phi). \quad (3)$$

In our proposed CDR, we aim to align two different optimization objectives as described above via a meta-optimization strategy. We draw inspiration from Model-Agnostic Meta-learning (MAML) [8] which splits the optimization into meta-train and meta-test processes and then involves a meta-update with a gradient-through-gradient mechanism such that the model is updated along an aligned direction between the meta-train and the meta-test. To coordinate **adaptation** and **anti-forgetting**, we propose to treat these two tasks as meta-train and meta-test in each iteration of parameter updating rather than splitting samples into meta-train and meta-test as in [19, 29]. With the task of adaptation as meta-train while the task of anti-forgetting as meta-test, the overall meta-optimization objective is:

$$\begin{aligned} \min_{\theta, \phi} \mathcal{L}_{Adap}(\theta, \phi) + \mathcal{L}_{AntiF}(\theta - \Delta\theta, \phi - \Delta\phi), \\ \Delta\theta = \alpha \nabla_{\theta} \mathcal{L}_{Adap}(\theta, \phi), \quad \Delta\phi = \alpha \nabla_{\phi} \mathcal{L}_{Adap}(\theta, \phi), \end{aligned} \quad (4)$$

where α denotes the meta-learning rate. With the meta-optimization objective, the optimization of anti-forgetting is considered through the gradient calculated upon the pre-updated parameters after the meta-train process via one gradient descent step: $\theta' \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_{Adap}(\theta, \phi)$ and $\phi' \leftarrow \phi - \alpha \nabla_{\phi} \mathcal{L}_{Adap}(\theta, \phi)$. As common practices in [7, 8] for the efficient implementation of meta-learning, we omit higher-order items during performing gradient back-propagation.

One natural question may raise: *Why can the meta-optimization objective presented in Eq.(4) coordinate anti-forgetting and adaptation for lifelong learning?* Inspired by the analysis in [19], we recall the first-order Taylor expansion for a general function $f(x, y)$ of two variables at $x = x_0, y = y_0$ as:

$$f(x, y) = f(x_0, y_0) + f'_x(x_0, y_0)(x - x_0) + f'_y(x_0, y_0)(y - y_0), \quad (5)$$

where x_0 is any point close to x . We instantiate the function f as \mathcal{L}_{AntiF} , and set x and y to be $\theta - \alpha \nabla_{\theta} \mathcal{L}_{Adap}(\theta, \phi)$ and $\phi - \alpha \nabla_{\phi} \mathcal{L}_{Adap}(\theta, \phi)$, respectively. Further, we choose x_0 and y_0 to be θ and ϕ , respectively. Thus, we can replace the second term $\mathcal{L}_{AntiF}(\theta - \Delta\theta, \phi - \Delta\phi)$ in Eq.(4) with its corresponding first-order Taylor expansion, and reformulate the overall meta-optimization objective of Eq.(4), as:

$$\begin{aligned} \min_{\theta, \phi} \quad & \mathcal{L}_{Adap}(\theta, \phi) + \mathcal{L}_{AntiF}(\theta, \phi) \\ & - \alpha \nabla_{\theta} \mathcal{L}_{Adap}(\theta, \phi) \nabla_{\theta} \mathcal{L}_{AntiF}(\theta, \phi) \\ & - \alpha \nabla_{\phi} \mathcal{L}_{Adap}(\theta, \phi) \nabla_{\phi} \mathcal{L}_{AntiF}(\theta, \phi). \end{aligned} \quad (6)$$

Compared to the general optimization objective in Eq.(3), the above meta-optimization objective, *i.e.*, Eq.(6), additionally maximizes the dot product of gradients of $\mathcal{L}_{Adap}(\theta, \phi)$ and $\mathcal{L}_{AntiF}(\theta, \phi)$ with respect to θ and ϕ . Therefore it encourages both the feature extractor and the classifier to be optimized along a consistent/aligned direction between anti-forgetting and adaptation so that these two tasks are coordinated for more effective life-long learning.

4.3. Relational Consistency Learning

As aforementioned, a classical method to alleviate the catastrophic forgetting problem is to transfer the knowledge from a model preserved at a certain historical time instance to the current model via knowledge distillation [17, 20, 39]. These methods are indeed effective but leave two problems for *LUDA person ReID*: 1) The task characteristics and objectives of person ReID, *i.e.*, learning to rank, have not been sufficiently considered in designing the loss functions for knowledge distillation; 2) The old model preserved at a certain historical moment lacks sufficient robustness for the potential training perturbation over time. We propose Relational Consistency Learning (RCL) to address these problem towards more effective knowledge distillation in line with the task characteristics of *LUDA person ReID*.

We first introduce the idea of *temporal ensembling* from semi-supervised learning [33] into *LUDA person ReID* to build the historical model. Specifically, we initialize the historical model with the weights of source pre-trained model, then update it via a weighted summation where a momentum scalar β is applied to historical model while $1 - \beta$ is applied to the current model.

The goal of person ReID is learning to rank. In line with this goal, we design a regular term to encourage the historical model and the current model to output consistent ranking results for input samples. In each meta-optimization step, one batch of historically stored samples (abbreviated as ‘‘old batch’’, denoted by \mathcal{B}^o) is sampled from the memory while another batch of newly collected samples (abbreviated as ‘‘new batch’’, denoted by \mathcal{B}^n) is sampled from the current stream. Both \mathcal{B}^o and \mathcal{B}^n are fed into the historical model and the current model to get their corresponding ID feature vectors inferred by these two models.

We denote the feature vectors of the i -th and the j -th samples inferred by the current model as \mathbf{f}_i^c and \mathbf{f}_j^c , respectively. Likewise, \mathbf{f}_i^h and \mathbf{f}_j^h represent their corresponding feature vectors extracted by the historical model. The relation/affinity between the i -th and the j -th samples inferred by the current model can be calculated by $r_{i,j}^c = \mathbf{f}_i^{c\top} \mathbf{f}_j^c / (\|\mathbf{f}_i^c\|_2 \|\mathbf{f}_j^c\|_2)$. The relation/affinity inferred by the historical model $r_{i,j}^h$ is computed in the same way. As the core of our RCL, the regular term on inter-instance relations can be formulated as:

$$\mathcal{L}_{Rel.} = \sum_{i,j \in \mathcal{B}^n} (r_{i,j}^c - r_{i,j}^h)^2 + \sum_{i',j' \in \mathcal{B}^o} (r_{i',j'}^c - r_{i',j'}^h)^2. \quad (7)$$

We empirically find that merging the old batch \mathcal{B}^o and the new batch \mathcal{B}^n into one to calculate such regular term delivers very close performance to the version in Eq.(7). To be computation-efficient, we propose to calculate the affinity matrix for \mathcal{B}^n and \mathcal{B}^o individually. This regular term plays the regularization role in transferring the knowledge from the historical model to the current model on learning relative relations, which is required by ranking. Complementary to this, we additionally add another regular term to transfer the knowledge on learning absolute ID-related features, which is formulated as:

$$\mathcal{L}_{KL} = \sum_{i \in \mathcal{B}^n \cap \mathcal{B}^o} KL(\mathbf{g}_i^c \parallel \mathbf{g}_i^h), \quad (8)$$

where $KL(\cdot \parallel \cdot)$ denotes the Kullback Leibler (KL) divergence distance between two distributions. $\mathbf{g}_i^c = \sigma(\phi_c(\mathbf{f}_i^c))$ and $\mathbf{g}_i^h = \sigma(\phi_h(\mathbf{f}_i^h))$ denotes the classification logits inferred by the current model and the historical model, respectively. Here, $\phi_c(\cdot)$ and $\phi_h(\cdot)$ are the corresponding classifiers while $\sigma(\cdot)$ is the Softmax function. As a result, in our proposed RCL, we implement the knowledge distillation loss $\mathcal{L}_{KD}(\theta, \phi)$ in Eq.(2) as: $\mathcal{L}_{KD} = \mathcal{L}_{Rel.} + \mathcal{L}_{KL}$.




Figure 3. Effectiveness analysis of *CLUDA-ReID* on the source domain (PersonX) and on the target domain (Market1501) in the stationary target scenario. In “*Stage-wise UDA (Base.)*”, we fine-tune the pre-trained model using the data collected in the current stage sequentially, without explicit considerations on anti-forgetting. In “*Base.+Data Replay*”, the data is stored into a memory buffer for replaying, wherein \mathcal{L}_{Tri} and \mathcal{L}_{Cls} are applied to both old data and new data for joint training, without \mathcal{L}_{KD} used. “*Base.+CDR*” denotes the model using our proposed CDR as in Eq.(4) (wherein \mathcal{L}_{KD} is not adopted). In “*Base.+CDR+KL Reg.*”, we apply \mathcal{L}_{KL} to both old and new data to implement \mathcal{L}_{AntiF} in Eq.(4). In “*Base.+CDR+RCL*”, we further apply the proposed \mathcal{L}_{Rel} to both old and new data to achieve our proposed RCL. “*All-in-one UDA*” refers to the model combining all available target domain data for regular UDA training.

5. Experiments

5.1. Datasets and Evaluation Metrics

We employ a synthetic dataset PersonX (PX) [32] for source pre-training and three public real-world image datasets Market1501 (MA) [51], CUHK-SYSU (SY) [41], MSMT17 (MS) [37] for unsupervised fine-tuning. This shows a privacy-friendly practice since no ID annotations of real persons are required. Besides, for unseen-domain generalization evaluation, we build a new dataset called MMP-Retrieval upon the training and validation splits of MMPTRACK dataset released in ICCV 2021 multi-camera multiple people tracking workshop. MMP-Retrieval comprises 21 people in 5 simulated environments. We have released it as aforementioned. More details of the used public datasets can be found in the supplementary. We use the cumulative matching characteristics (CMC) at Rank-1 (R-1) and mean average precision (mAP) for evaluation.

5.2. Implementation Details

Following common practices [14,24,47] in person ReID, we adopt ResNet50 [13] pretrained on Imagenet [6] as our backbone. Similar to [24], the last spatial down-sampling in the “conv5_x” block is removed. We resize images to 256×128 , and set the batch size of both \mathcal{B}^o and \mathcal{B}^n to 64, including 16 identities and 4 images per identity for each batch. We adopt the commonly used data augmentation strategies of horizontal flipping, random cropping. Following [11,50], we use the clustering algorithm of DBSCAN to generate pseudo labels. Unless otherwise specified, we train the model on the source domain for 60 epochs, and train the model on the stationary/dynamic target streams for 40/60 epochs per stage respectively. We adopt Adam [18] optimizer with its learning rate initialized as 3.5×10^{-4} . The

size of memory buffer is set to 512. We provide the ablation study for the memory buffer size and more implementation details in the supplementary.

5.3. Evaluation in the Stationary Target Scenario

Experiment configuration. We evaluate our proposed *CLUDA-ReID* in the first scenario introduced in Sec.3.2. We employ the PersonX [32] as the source for supervised pre-training while employing the Market1501 [51] for performing LUDA. We sample 750 identities of the training data of Market1501, and uniformly split them into 5 subsets for a 5-stage LUDA training. At the end of each stage, we compute the R-1 and mAP accuracy on the entire test set of the source data and the target data to respectively assess the model capacities on anti-forgetting and adaptation.

Effectiveness of *CLUDA-ReID*. We start with a basic model “*Stage-wise UDA*” without anti-forgetting as the base model, then incrementally add different components upon it to assess their effectiveness. As shown in Fig.3, the performances on the source domain (PersonX) and the target domain (Market1501) reflect the *anti-forgetting* ability and the timely *adaptation* ability of models, respectively. Except for *All-in-one UDA*, all other models are trained stage by stage as new data arrives over time. *Stage-wise UDA* does not explicitly consider the optimization for anti-forgetting. It suffers from a dramatic performance drop on the source domain data and is inferior to adapting to the target domain. As the results of *Base.+Data Replay* show, data replaying boosts the performance on both source and target data by revisiting old data. Relative to *Stage-wise UDA*, the model *Base.+CDR* achieves significant improvements. This demonstrates our proposed CDR effectively promotes both the old knowledge memorization and the new knowledge capturing through their coordination. Upon our proposed




Figure 4. The ablation study on the meta-optimization strategies. \mathcal{L}_{KD} denotes the loss function for knowledge distillation. \mathcal{B}^o and \mathcal{B}^n denotes the batches of old data and new data, respectively. ‘‘M-train’’ is short for meta-train, ‘‘M-test’’ is short for meta-test. In all shown models, the person ReID losses (\mathcal{L}_{Tri} and \mathcal{L}_{Cls}) applied on the new data are adopted for the meta-train while the person ReID losses applied to the old data are adopted for the meta-test.

coordinated optimization scheme, we further enhance the anti-forgetting ability by adding KL consistency regularization \mathcal{L}_{KL} (KL Reg.) and relational consistency regularization $\mathcal{L}_{Rel.}$. The comparison of the model *Base.+CDR+RCL* to the model *Base.+CDR* shows striking improvements of our proposed RCL, thanks to the loss function designs on old knowledge distillation/inheritance in line with the task objective of person ReID.

Ablation study for meta-optimization strategies. In this part, we empirically study the practices of using different data \mathcal{B}^o and \mathcal{B}^n for \mathcal{L}_{KD} in meta-optimization. As shown in Fig.4, by comparing our proposed scheme (\mathcal{L}_{KD} on $\mathcal{B}^o \cup \mathcal{B}^n$ for meta-test) to the one applying \mathcal{L}_{KD} to \mathcal{B}^o for meta-test and the one applying \mathcal{L}_{KD} to \mathcal{B}^n for meta-test, we find that both \mathcal{B}^o and \mathcal{B}^n can contribute to the knowledge distillation for anti-forgetting. Besides, we further compare our proposed task-oriented meta-train/meta-test splitting in CDR with the data-based meta-train/meta-test splitting (*i.e.*, \mathcal{L}_{KD} on \mathcal{B}^n for meta-train while \mathcal{L}_{KD} on \mathcal{B}^o for meta-test) as in [19, 29]. We find our task-oriented splitting of meta-train and meta-test is more effective in coordinating the optimization objectives of anti-forgetting and adaptation.

Ablation study for different design choices. We conduct an ablation study on different design choices in *CLUDA-ReID* by replacing the components in the *Full Model* of *CLUDA-ReID* (the same as *Base.+CDR+RCL* in Fig.3) with other design choices. The results are shown in Fig.5. Compared to the First-In-First-Out (FIFO) rule and the regular (instance-wise) reservoir sampling algorithm, our modified ID-wise reservoir sampling adopted in *Full Model* is superior since it enables stored samples to be diverse and ID-balanced. Compared to the historical model preserved at a certain historical moment, our design in *Full Model* is more effective thanks to the temporal averaging via the momentum updating strategy. In terms of the loss function




Figure 5. The ablation study on different design choices for data replay, historical model updating and consistency regularization designs. This is conducted in the stationary target scenario. ‘‘Data Replay (FIFO)’’ denotes the scheme storing seen samples into the memory buffer with the rule of First-In-First-Out. ‘‘Instance-wise RS’’ refers to the regular version of Reservoir Sampling as introduced in [35]. ‘‘Old Model (wo. momentum)’’ denotes we preserve a model at a certain old moment as the historical model. ‘‘only KL’’ and ‘‘only Rel.’’ means using \mathcal{L}_{KL} and using $\mathcal{L}_{Rel.}$ as the knowledge distillation loss \mathcal{L}_{KD} , respectively.

design for consistency learning, we find \mathcal{L}_{KL} and $\mathcal{L}_{Rel.}$ are complementary for the old knowledge distillation from the historical model to the current model, since \mathcal{L}_{KL} aims to learn absolute ID-related features while $\mathcal{L}_{Rel.}$ captures the relative information about inter-instance relations.

5.4. Evaluation in the Dynamic Target Scenario

Experiment configuration. We evaluate *CLUDA-ReID* in the scenario with dynamic target stream introduced in Sec.3.2. This scenario is not only ID-incremental but also domain-incremental. One synthesis dataset PersonX (PX) is employed to simulate the source domain while three real-world datasets Market1501 (MA), CUHK-SYSU (SY), and MSMT17 (MS) are employed to simulate the dynamic target domain. We perform LUDA training on MA, SY, and MS domains in order, and repeat this process again to simulate the domain recurring cases in the real-world dynamic scenario. There are 6 stages in all for LUDA training in our experiment. Each stage includes a sub-set of 350 identities selected from the corresponding dataset via a random sampling without replacement. Furthermore, we also evaluate the model at the end of the final stage on MMP-Retrieval to assess its generalization ability on unseen domains.

Adaptation performance. We evaluate the models at the end of each stage. As shown in Tab.1, compared to the scheme *Stage-wise UDA (Base.)*, the scheme *Base.+CDR* achieves 0.89%/1.79%, 4.64%/3.78%, 3.76%/2.09% improvements in R-1/mAP on MA, SY and MS respectively when these domains first appear, and achieves 4.01%/3.00%, 9.53%/9.00%, 5.67%/3.25% improvements in R-1/mAP on MA, SY and MS respectively when these domains show up for the second time. The scheme

Methods	MA (t=1)		SY (t=2)		MS (t=3)		MA (t=4)		SY (t=5)		MS (t=6)	
	R-1	mAP										
Stage-wise UDA (Base.)	72.23	49.32	64.59	60.76	19.80	7.24	72.92	50.46	63.93	60.29	20.23	7.18
Base. + Data Replay	72.51	50.82	66.22	61.75	20.98	7.86	74.12	51.03	69.43	64.44	21.73	8.48
Base. + CDR	73.12	51.11	69.23	64.54	23.56	9.33	76.93	53.46	73.46	69.29	25.90	10.43
Base. + CDR + KL Reg.	74.24	51.76	71.55	67.88	26.16	9.94	79.26	59.48	76.76	74.45	29.79	12.44
Base. + CDR + RCL	75.21	52.96	75.38	71.35	28.67	11.04	82.79	64.87	81.86	78.39	33.91	14.64
All-in-one UDA	85.12	68.59	86.76	84.70	36.37	17.84	85.12	68.59	86.76	84.70	36.37	17.84

Table 1. Adaptation performance (%) evaluation in the dynamic target scenario. We test the model instantly at the end of each stage, on its corresponding test set. The training order is PX→MA→SY→MS→MA→SY→MS. Our *CLUDA-ReID* is marked in gray shading.

Methods	t=3						t=6					
	PX		MA		SY		PX		MA		SY	
	R-1	mAP										
Stage-wise UDA (Base.)	52.71	30.76	49.76	24.59	52.97	55.99	46.92	24.59	50.63	19.64	51.17	51.65
Base. + Data Replay	56.33	38.21	57.46	31.73	54.46	56.33	48.78	27.33	60.34	32.37	58.93	57.11
Base. + CDR	63.93	44.43	61.43	33.46	60.52	60.43	58.46	39.64	65.34	39.46	65.33	62.77
Base. + CDR + KL Reg.	73.75	54.33	65.83	40.30	68.55	62.11	65.45	47.43	71.43	44.10	72.33	69.48
Base. + CDR + RCL	80.66	69.48	69.97	45.60	73.24	69.44	76.47	58.23	76.25	49.12	78.28	75.43
Pre-trained Model (t=0)	94.86	86.34	54.04	27.67	46.52	48.62	94.86	86.34	54.04	27.67	46.52	48.62
All-in-one UDA	40.43	16.44	85.12	68.59	86.76	84.70	40.43	16.44	85.12	68.59	86.76	84.70

Table 2. Anti-forgetting performance (%) evaluation in the dynamic target scenario. We test the model on the test-sets of all domains at the end of the 3rd stage (t=3) and the 6th stage (t=6). The training order is PX→MA→SY→MS→MA→SY→MS. Note that we omit the MS dataset for brevity because it appears in the final, whose results can not reflect the anti-forgetting ability. *CLUDA-ReID* is marked in gray.

Base.+CDR+RCL is superior to *Stage-wise UDA* by 2.98%/3.64%, 10.79%/10.59%, 8.87%/3.80% respectively in R-1/mAP on MA, SY and MS when these domains first appear, by 9.87%/14.41%, 17.93%/18.10%, 13.68%/7.46% respectively in R-1/mAP on MA, SY and MS when they appear for the second time. These demonstrate the effectiveness of our proposed CDR and RCL in timely adapting to new environments. The improvements are especially significant when the similar or same domain re-appears thanks to the coordinated old knowledge memorization.

Anti-forgetting performance. We measure the performance at the end of the 3rd stage (*i.e.*, t=3, all domains are traversed once) and at the end of the 6th stage (*i.e.*, t=6, all domains are seen for the second time), for evaluating the anti-forgetting ability of our proposed method. As shown in Tab.2, relative to the model pre-trained on PX, the scheme *Base.+CDR+RCL* ranks the first with the lowest performance degradation on PX at the end of the 3rd and the 6th stages. Besides, comparing the results in Tab.2 to the timely measured results in Tab.1, we find the scheme *Base.+CDR+RCL* is of the lowest performance degradation over all domains. It demonstrates the effectiveness of our *CLUDA-ReID* in alleviating catastrophic forgetting.

Generalization performance on unseen domains. We evaluate the trained model directly on MMP-Retrieval. As shown in Fig.6, the Rank-1 and mAP of *Base.+CDR* and *Base.+CDR+RCL* on MMP-Retrieval increases constantly as the LUDA training goes on. This demonstrates the effectiveness of our proposed CDR and RCL of *CLUDA-ReID* in improving the generalization ability for unseen domains through coordinated anti-forgetting and adaptation.




Figure 6. Performance (%) of unseen-domain generalization evaluation on our built MMP-Retrieval dataset.

6. Conclusion

In this paper, we propose a new and practical task, *LUDA person ReID*, which aims to enable person ReID models to achieve continuous domain adaptation using unlabeled streaming data. This facilitates exploiting increased samples and achieves timely adaptation for real-world streaming data. Besides, it helps privacy protection of customer data as the practices in our experiments. For this task, we design an effective scheme *CLUDA-ReID* in which the Coordinated Data Replay (CDR) and Relational Consistency Learning (RCL) are proposed to explicitly coordinate anti-forgetting and adaptation. We set up two practical evaluation schemes to simulate the real applications for method evaluation. Extensive experiment results demonstrate the effectiveness of *CLUDA-ReID* in achieving LUDA for scenarios with stationary or dynamic target streams and enhancing the generalization capacity on unseen domains.

References

- [1] Zechen Bai, Zhigang Wang, Jian Wang, Di Hu, and Errui Ding. Unsupervised multi-source domain adaptation for person re-identification. In *CVPR*, 2021. 1
- [2] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *ECCV*, pages 532–547, 2018. 2
- [3] Ali Cheraghian, Shafin Rahman, Sameera Ramasinghe, Pengfei Fang, Christian Simon, Lars Petersson, and Mehrtash Harandi. Synthesized feature based few-shot class-incremental learning on a mixture of subspaces. In *ICCV*, pages 8661–8670, 2021. 3
- [4] Corinna Cortes, Xavier Gonzalvo, Vitaly Kuznetsov, Mehryar Mohri, and Scott Yang. Adanet: Adaptive structural learning of artificial neural networks. In *ICML*, pages 874–883. PMLR, 2017. 3
- [5] Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *TPAMI*, 2021. 2
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009. 6
- [7] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *NeurIPS*, volume 32, pages 6450–6461, 2019. 5
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135. PMLR, 2017. 4, 5
- [9] Yang Fu, Yunchao Wei, Guanshuo Wang, Yuqian Zhou, Honghui Shi, and Thomas S Huang. Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In *ICCV*, pages 6112–6121, 2019. 1
- [10] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020. 1, 2
- [11] Yixiao Ge, Feng Zhu, Dapeng Chen, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *NeurIPS*, 2020. 6
- [12] Yixiao Ge, Feng Zhu, Rui Zhao, and Hongsheng Li. Structured domain adaptation for unsupervised person re-identification. *arXiv preprint arXiv:2003.06650*, 2020. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [14] Lingxiao He, Xingyu Liao, Wu Liu, Xinchen Liu, Peng Cheng, and Tao Mei. Fastreid: A pytorch toolbox for general instance re-identification. *arXiv preprint arXiv:2006.02631*, 2020. 3, 6
- [15] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *CVPR*, pages 3957–3966, 2021. 2
- [16] Yangru Huang, Peixi Peng, Yi Jin, Junliang Xing, Congyan Lang, and Songhe Feng. Domain adaptive attention model for unsupervised cross-domain person re-identification. In *AAAI*, 2020. 2
- [17] Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental learning through feature adaptation. In *ECCV*, pages 699–715. Springer, 2020. 3, 5
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [19] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 4, 5, 7
- [20] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. 2, 3, 5
- [21] Shan Lin, Haoliang Li, Chang-Tsun Li, and Alex Chichung Kot. Multi-task mid-level feature alignment network for unsupervised cross-dataset person re-identification. In *BMVC*, 2018. 2
- [22] Jiawei Liu, Zheng-Jun Zha, Di Chen, Richang Hong, and Meng Wang. Adaptive transfer network for cross-domain person re-identification. In *CVPR*, pages 7202–7211, 2019. 2
- [23] Xiaobin Liu and Shiliang Zhang. Domain adaptive person re-identification via coupling optimization. In *ACMMM*, pages 547–555, 2020. 2
- [24] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *CVPR Workshops*, pages 0–0, 2019. 3, 6
- [25] Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 2021. 2
- [26] Andrea Maracani, Umberto Michieli, Marco Toldo, and Pietro Zanuttigh. Recall: Replay-based continual learning in semantic segmentation. In *ICCV*, pages 7026–7035, 2021. 3
- [27] Nan Pu, Wei Chen, Yu Liu, Erwin M Bakker, and Michael S Lew. Lifelong person re-identification via adaptive knowledge accumulation. In *CVPR*, pages 7901–7910, 2021. 1, 2
- [28] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *TPAMI*, 2018. 3
- [29] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *ICLR*, 2019. 3, 4, 7
- [30] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 3
- [31] Liangchen Song, Cheng Wang, Lefei Zhang, Bo Du, Qian Zhang, Chang Huang, and Xinggang Wang. Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognit*, 102:107173, 2020. 1

- [32] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *CVPR*, pages 608–617, 2019. [6](#)
- [33] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. [5](#)
- [34] Eli Verwimp, Matthias De Lange, and Tinne Tuytelaars. Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In *ICCV*, 2021. [2](#), [3](#)
- [35] Jeffrey S Vitter. Random sampling with a reservoir. *TOMS*, 11(1):37–57, 1985. [4](#), [7](#)
- [36] Xiaogang Wang. Intelligent multi-camera video surveillance: A review. *Pattern Recognit. Lett.*, 34(1):3–19, 2013. [1](#)
- [37] 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, 2018. [6](#)
- [38] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to generate new categories without forgetting. In *NeurIPS*, volume 31, pages 5962–5972, 2018. [3](#)
- [39] Guile Wu and Shaogang Gong. Generalising without forgetting for lifelong person re-identification. In *AAAI*, volume 35, pages 2889–2897, 2021. [1](#), [2](#), [3](#), [5](#)
- [40] Guile Wu, Shaogang Gong, and Pan Li. Striking a balance between stability and plasticity for class-incremental learning. In *ICCV*, pages 1124–1133, 2021. [2](#)
- [41] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. Joint detection and identification feature learning for person search. In *CVPR*, pages 3415–3424, 2017. [6](#)
- [42] Tianjun Xiao, Jiaying Zhang, Kuiyuan Yang, Yuxin Peng, and Zheng Zhang. Error-driven incremental learning in deep convolutional neural network for large-scale image classification. In *ACMMM*, pages 177–186, 2014. [3](#)
- [43] Hong-Xing Yu, Wei-Shi Zheng, Ancong Wu, Xiaowei Guo, Shaogang Gong, and Jian-Huang Lai. Unsupervised person re-identification by soft multilabel learning. In *CVPR*, pages 2148–2157, 2019. [2](#)
- [44] Yunpeng Zhai, Qixiang Ye, Shijian Lu, Mengxi Jia, Ronrong Ji, and Yonghong Tian. Multiple expert brainstorming for domain adaptive person re-identification. In *ECCV*, 2020. [2](#)
- [45] Minying Zhang, Kai Liu, Yidong Li, Shihui Guo, Hongtao Duan, Yimin Long, and Yi Jin. Unsupervised domain adaptation for person re-identification via heterogeneous graph alignment. In *AAAI*, 2021. [1](#)
- [46] Xinyu Zhang, Jiewei Cao, Chunhua Shen, and Mingyu You. Self-training with progressive augmentation for unsupervised cross-domain person re-identification. In *ICCV*, pages 8222–8231, 2019. [1](#)
- [47] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Densely semantically aligned person re-identification. In *CVPR*, pages 667–676, 2019. [6](#)
- [48] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Xin Jin, and Zhibo Chen. Relation-aware global attention for person re-identification. In *CVPR*, pages 3186–3195, 2020. [3](#)
- [49] Fang Zhao, Shengcai Liao, Guo-Sen Xie, Jian Zhao, Kaihao Zhang, and Ling Shao. Unsupervised domain adaptation with noise resistible mutual-training for person re-identification. In *ECCV*, pages 1–18, 2020. [2](#)
- [50] Kecheng Zheng, Cuiling Lan, Wenjun Zeng, Zhizheng Zhang, and Zheng-Jun Zha. Exploiting sample uncertainty for domain adaptive person re-identification. In *AAAI*, volume 35, pages 3538–3546, 2021. [6](#)
- [51] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, pages 1116–1124, 2015. [6](#)