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

Zhipeng Huang[1†*] Zhizheng Zhang[2†‡] Cuiling Lan[2‡] Wenjun Zeng[3] Peng Chu[2]
Quanzeng You[2] Jiang Wang[2] Zicheng Liu[2] Zheng-jun Zha[1]
[1]University of Science and Technology of China
[2]Microsoft [3]EIT Institute for Advanced Study
hzp1104@mail.ustc.edu.cn zengw2011@hotmail.com zhazj@ustc.edu.cn
{zhizzhang, culan, pengchu, quyou, jiangwang, zliu}@microsoft.com

In this supplementary material, we first detail our modified ID-wise reservoir sampling algorithm which is originally an instance-wise one. Then we provide detailed introduction for datasets used in this work and give more implementation details. Finally, we present more experiment results and discuss the ethical impact of this work.

1. ID-wise Reservoir Sampling Algorithm

As introduced in our main body, we set up a memory buffer with limited size for data replay, which is updated using ID-wise Reservoir Sampling algorithm at the end of each stage. Here, we detail this algorithm which is modified from the regular one that is instance-wise.

The regular reservoir sampling algorithm [8] is designed to choose a subset of k individuals at random, without replacement, from a population of size N in a sequence. Here, N is allowed to be unknown and typically large. In this algorithm, each sample is chosen with an equal probability over the part of the population seen so far. However, for the data replaying adopted in our targeted *LUDA person ReID* problem, we expect the stored samples are of diverse data statistics to avoid over-fitting them. Thus, we propose to modify the regular reservoir sampling algorithm to be an ID-wise one, in the sense that each ID will be randomly chosen for storing with an equal probability over all identities that have been seen so far.

We describe the ID-wise reservoir sampling algorithm in Alg. 1. The $record$ function summarizes all samples into a dictionary structure according to their identities, where the keys of the dictionary are the recorded identities while its values are the corresponding sample indexes. The $select$ function randomly selects K samples for a given ID. Here, if the number of samples of the given ID is less than K, we will adopt data augmentation (*i.e.*, random horizontal flipping, random cropping and random erasing) to augment the current samples and store augmented samples. The $random$ function randomly produces an integer between the provided minimum and maximum values inclusively. We compare it to the regular instance-wise reservoir sampling algorithm in the Fig.5 of our main body. The experiment results demonstrate the superiority of our modified ID-wise reservoir sampling algorithm relative to its original version when adopted to our proposed scheme *CLUDA-ReID* for *LUDA person ReID*.

Algorithm 1 ID-wise Reservoir Sampling Algorithm

$\mathcal{M}$: the memory buffer with the size of $\|\mathcal{M}\|$;
N_{id}: the total number of IDs observed so far;
K: the number of samples stored for each identity;
$\mathcal{X}_j^t$: the data set in the j-th stage of the target stream;
$\mathcal{Y}_j^t$: the pseudo labels of $\mathcal{X}_j^t$, obtained via clustering.

procedure ID-RESERVOIR($\mathcal{M}, \|\mathcal{M}\|, N_{id}, K, \mathcal{X}_j^t, \mathcal{Y}_j^t$)
  $id\_dicts = record(\mathcal{X}_j^t, \mathcal{Y}_j^t)$
  for $i = 0$ to $len(id\_dicts.keys()) - 1$ **do**
    $indices = id\_dicts[i]$
    $(\mathbf{x}_k^t, y_k^t)_{k=1}^K = select(\mathcal{X}_j^t, \mathcal{Y}_j^t, indices, K)$
    if $\|\mathcal{M}\| // K > N_{id}$ **then**
      $\mathcal{M}[N_{id} * K, (N_{id}+1) * K] \leftarrow (\mathbf{x}_k^t, y_k^t)_{k=1}^K$
    else
      $r = random(min = 0, max = N_{id})$
      if $r < \|\mathcal{M}\| // K$ **then**
        $\mathcal{M}[r * K, (r+1) * K] \leftarrow (\mathbf{x}_k^t, y_k^t)_{k=1}^K$
      end if
    end if
    $N_{id} \leftarrow N_{id} + 1$
  end for
  return $\mathcal{M}$
end procedure

*This work was done when Zhipeng Huang was an intern at MSRA.
†Equal contribution.
‡Corresponding authors.

Algorithm 2 Coordinated Data Replay (CDR) Algorithm

$\mathcal{M}$: the memory buffer for storing old samples;
θ: the parameters of the backbone network;
ϕ: the parameters of the classifier;
α: the learning rate of the meta-train optimization;
η: the learning rate of the meta-optimization;
$\mathcal{L}_{Adap}$: the optimization objective for ***adaptation***;
$\mathcal{L}_{AntiF}$: the optimization objective for ***anti-forgetting***.

Init: parameters $\Psi = \{\theta, \phi\}$, learning rates η, α.
for t **in** iterations **do**
  Meta-train:
  Sample a batch $\mathcal{B}^n$ from the new data.
  Compute $\mathcal{L}_{Adap}(\Psi_t)$ with $\mathcal{B}^n$. ▷ Eq. (1)
  Update Ψ w.r.t. $\mathcal{L}_{Adap}$:
  $$\Psi_t' \leftarrow \Psi_t - \alpha \nabla_{\Psi_t} \mathcal{L}_{Adap}(\Psi_t)$$
  Meta-test:
  Sample a batch $\mathcal{B}^o$ from the old data stored in $\mathcal{M}$.
  Compute $\mathcal{L}_{AntiF}(\Psi_t')$ with $\mathcal{B}^n$ and $\mathcal{B}^o$. ▷ Eq. (2)
  Meta optimization:
  $\mathcal{L}_{CDR} = \mathcal{L}_{Adap}(\Psi_t) + \mathcal{L}_{AntiF}(\Psi_t')$. ▷ Eq. (4)
  Update model parameters:
  $$\Psi_{t+1} \leftarrow \Psi_t - \eta \bigtriangledown_{\Psi_t} \mathcal{L}_{CDR}.$$
end for
Output: θ, ϕ

2. Optimization of Coordinated Data Replay

We describe the training procedure of *CLUDA-ReID* using the proposed coordinated data replay in Alg. 2. For simplicity, we show the case where we take the "task" of adaptation as meta-train while taking the "task" of anti-forgetting as meta-test. Theoretically, as indicated in the Eq.(6) of our manuscript, the orders of these two tasks are exchangeable for the meta optimization. Our empirical study also demonstrates that iteratively choosing one of these two tasks (*i.e.*, adaptation and anti-forgetting) as meta-train while taking the other one as meta-test delivers very close performance with our reported ones.

3. Detailed Introduction for Datasets

In this work, we use four public datasets PersonX (PX) [7], Market1501 (MA) [11], CUHK-SYSU (SY) [10], MSMT17 (MS) [9] for pre-training and lifelong unsupervised domain adaptation, and build a new dataset MMP-Retrieval for unseen domain generalization evaluation.

PersonX [7] is a synthetic dataset generated based on Unity [6], containing 45,792 images, where 410 identities are used for training and 856 identities are used for testing.

Market1501 [11] has 12,936 images of 751 identities for training and 19,732 images of 750 identities for testing. Its test split has 3,368 query images and 16,364 gallery images.

CUHK-SYSU [10] is originally a large-scale dataset for person search task, containing 18,184 images and 8,432 identities. Following the previous work [5], we employ a subset of CUHK-SYSU using the bounding box annotations, in which each identity includes at least 4 person crops. As a result, the training set of this subset includes 942 identities while its test set includes 2,900 identities.

MSMT17 [9] is a large-scale person ReID dataset which is captured from 15 cameras (including 12 outdoor cameras and 3 indoor cameras). It contains 126,441 images of 4,101 identities, where 1,041 identities and 3,060 identities are used for training and testing respectively.

MMP-Retrieval, the new dataset we propose in this paper, is built upon the Multi-camera Multiple People Tracking dataset (*i.e.*, MMPTRACK) which is released in ICCV 2021 Multi-camera Multiple People Tracking Workshop*. It is available at https://iccv2021-mmp.github.io/subpage/dataset.html. In MMP-Retrieval, all videos are collected in 5 simulated environments: retail, lobby, industry, cafe and office, including 28 people participating in recording (14 in training, 7 in validation and 7 in testing). All people were paid and signed an agreement to release their data to the public for research usage. So there is no privacy issue. MMP-Retrieval is built using the combination of the training and validation splits of MMPTRACK dataset, comprising 21 identities in all. To make the samples in MMP-Retrieval diverse, we uniformly downsample the original video sequences of MMPTRACK with a ratio of 128, then divide each downsampled sequence into two halves. We use the cropped persons in the first half as the query set and those in the second half as the gallery set. Although this dataset contains a limited number of identities, the diverse camera angles and cluttered backgrounds (as illustrated in Fig. 1) still make MMP-Retrieval a challenging test dataset for unseen-domain generalization evaluation. To avoid cloth-changing cases crossing different environments, we report the Rank-1 and mAP scores averaged over all environments in our experiments.

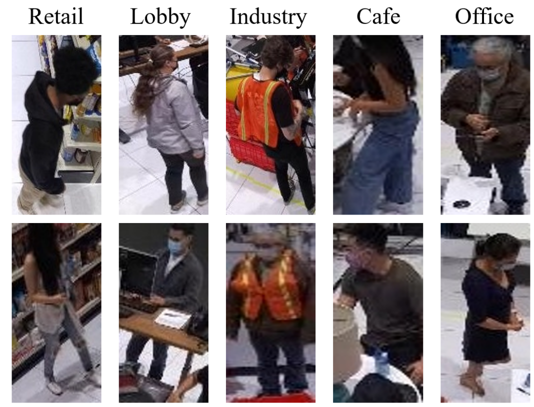

Figure 1. Examples of different environments in MMP-Retrieval.

*https://iccv2021-mmp.github.io

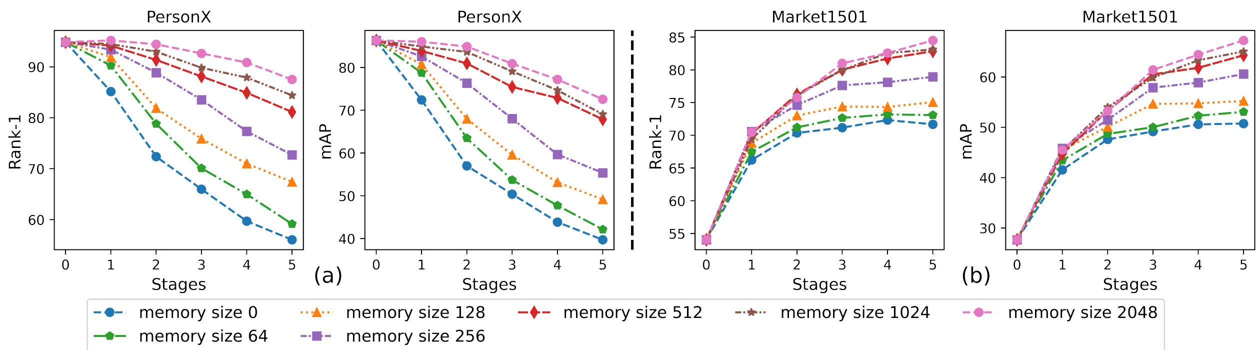

Figure 2. The empirical study results for the size of the memory buffer. This experiment is conducted in the stationary target scenario. The performances on the source domain (PersonX) are in (a) while those on the target domain (Market1501) are in (b).

4. More Implementation Details

Following the prior works for UDA person ReID [2, 3], we use the clustering algorithm of DBSCAN [1] to generate pesudo labels, where the maximum distance between neighbors is set to 0.6 and the minimal number of neighbors for a dense point is set to 4. We balance different terms in $\mathcal{L}_{Adap}$ and $\mathcal{L}_{AntiF}$ to make their corresponding gradients lie in similar ranges, and we empirically find that setting all of their weights to be 1 works well. We train the model on the source domain for 60 epochs, and train the model on each domain of the stationary/dynamic scenarios for 40/60 epochs, respectively. We adopt Adam [4] optimizer with a weight decay of 5×10^{-4}. In the pre-training stage on the source domain or each fine-tuning stage on the target domain, as in Alg. 2, the learning rate η for the meta optimization is initialized to 3.5×10^{-4} and the learning rate α is initialized to 3.5×10^{-3}. Both η and α is decayed with a factor of 0.1 in the 40^{th} epoch. Unless otherwise specified, the size of the memory buffer is set to 512, and 8 randomly sampled images are stored for each sampled identity using the Alg. 1.

5. More Experiment Results

5.1. Empirical Study for the Memory Buffer Size

In this section, we conduct an empirical study for the size of the memory buffer. As shown in Fig.2, the performances on the source and target domains both become higher as the size of the memory buffer increases. This is because more historical data will deliver more benefits for capturing new knowledge while recalling old knowledge, thanks to the coordination between timely adaptation and anti-forgetting in our proposed scheme. In practice, we need to consider the trade-off between the storage burdens and performance. Thus, we recommend using the memory size of 512.

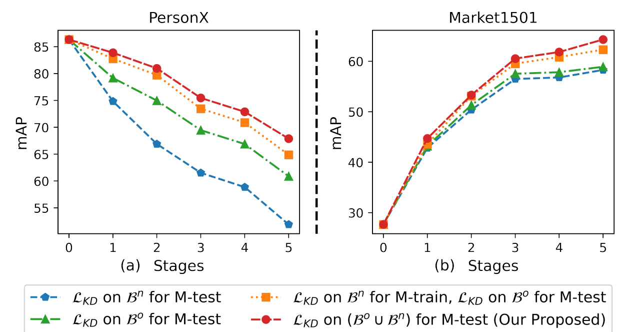

Figure 3. The mAP results of the ablation study on the meta-optimization strategies.

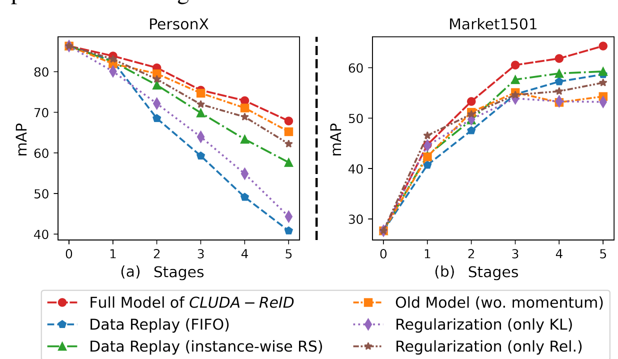

Figure 4. The mAP results of the ablation study on different design choices of the technical components in our *CLUDA-ReID*.

5.2. More Results with the mAP Metric

For simplicity, in our manuscript, we only report the experiment results of our ablation study using the Rank-1 metric. In this section, we provide more experiment results using mAP as the evaluation metric in Fig.3 and Fig.4. Comparing the Fig.3 and Fig.4 in this supplementary to the Fig.4 and Fig.5 in our manuscript, we can find the experiment results in mAP has very consistent trend with those in Rank-1. This further demonstrates the effectiveness and superiority of different components in our proposed *CLUDA-ReID*.

Methods	MS (t=1)		MA (t=2)		SY (t=3)		MS (t=4)		MA (t=5)		SY (t=6)	
	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP
Stage-wise UDA (Base.)	24.16	9.13	66.69	40.99	57.83	54.42	23.43	8.90	68.92	43.43	56.33	52.83
Base. + Data Replay	26.08	9.45	68.08	43.26	59.14	59.03	27.38	9.83	70.34	47.34	61.43	59.24
Base. + CDR	28.03	11.03	70.32	46.42	63.12	64.34	29.34	11.63	72.43	52.31	67.21	67.82
Base. + CDR + KL Reg.	29.64	11.82	72.80	48.63	69.11	68.48	32.43	12.88	78.36	56.73	74.43	73.33
Base. + CDR + RCL	**31.66**	**12.15**	**74.44**	**51.55**	**75.90**	**74.56**	**35.48**	**14.98**	**81.32**	**61.33**	**84.23**	**82.13**
pretrain (t=0)	18.02	5.02	54.04	27.67	46.52	48.62	18.02	5.02	54.04	27.67	46.52	48.62
All-in-one UDA	36.37	17.84	85.12	68.59	86.76	84.70	36.37	17.84	85.12	68.59	86.76	84.70

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→MS→MA→SY→MS→MA→SY. Our *CLUDA-ReID* is marked in gray shading.

Methods	t=3						t=6					
	PX		MS		MA		PX		MS		MA	
	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP	R-1	mAP
Stage-wise UDA (Base.)	48.95	27.13	10.83	3.01	50.98	25.22	40.83	16.45	8.83	2.79	48.42	21.42
Base. + Data Replay	58.26	39.29	18.60	6.01	56.15	29.41	52.39	24.28	20.11	7.83	58.32	31.21
Base. + CDR	63.93	44.33	20.49	8.21	59.74	32.70	58.46	39.64	22.21	9.03	63.74	35.71
Base. + CDR + KL Reg.	72.75	54.24	22.09	8.94	63.34	36.77	67.45	48.43	25.38	10.95	69.24	42.43
Base. + CDR + RCL	**85.57**	**71.16**	**26.16**	**11.07**	**69.75**	**43.70**	**70.42**	**51.36**	**31.46**	**13.91**	**77.34**	**50.11**
Pre-trained Model (t=0)	94.86	86.34	18.02	5.02	54.04	27.67	94.86	86.34	18.02	5.02	54.04	27.67
All-in-one UDA	40.43	16.44	36.37	17.84	85.12	68.59	40.43	16.44	36.37	17.84	85.12	68.59

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→MS→MA→SY→MS→MA→SY. Note that we omit the SY dataset for brevity because it appears in the final, whose results can not reflect the anti-forgetting ability. *CLUDA-ReID* is marked in gray.

5.3. Experimental Analysis for the Training Orders

In the manuscript, we demonstrate the effectiveness of *CLUDA-ReID* in the dynamic target scenario with the training order of PX→MA→SY→MS→MA→SY→MS. In practice, the order of target domains is agnostic. To verify this, we conduct an experiment in the dynamic target scenario with another domain order: PX→MS→MA→SY →MS→MA→SY. Other configurations are kept the same as that in the Sec.5.4 of the manuscript.

In Tab.1, we measure the performance at the end of each stage for evaluating the timely adaptation performance. The scheme *Base.+CDR* achieves 3.87%/1.90%, 3.63%/5.43%, 5.29%/9.92% improvements in R-1/mAP on MS, MA and SY respectively when these domains first appear, and achieves 5.91%/2.73%, 3.51%/8.88%, 10.88%/14.99% improvements in R-1/mAP on MS, MA and SY respectively when these domains show up for the second time. The scheme *Base.+CDR+RCL* is superior to the baseline *Stage-wise UDA* by 7.50%/3.02%, 7.75%/10.56%, 18.07%/20.14% improvements in R-1/mAP on MS, MA and SY when these domains first appear, and by 12.05%/6.08%, 12.40%/17.90%, 27.90%/29.30% improvements in R-1/mAP on MS, MA and SY when they appear for the second time. The trends of experiment results in this domain order is consistent with those in the Sec.5.4 of our manuscript, further demonstrating the effectiveness of *CLUDA-ReID* in timely adapting to new environments.

In Tab.2, 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 capacity of our proposed method. Relative to the model pre-trained on PX, the scheme *Base.+CDR+RCL* ranks the first with performance degradation of 9.29%/15.18% in R-1/mAP on PX at the end of the 3rd stage, and with performance degradation of 24.44%/34.98% in R-1/mAP on PX at the end of the 6th stage. Besides, comparing the results in Tab.2 to the timely measured results in Tab.1, we find that the scheme *Base.+CDR+RCL* is of the lowest performance degradation over all domains. It also shows consistent trend with that in the Sec.5.2 of our manuscript, demonstrating the effectiveness of our proposed *CLUDA-ReID* in anti-forgetting. The above experiment results present that our proposed *CLUDA-ReID* performs consistently with different domain orders in the dynamic target scenario. This shows the practicability of *CLUDA-ReID*.

6. Ethical Impact

Our method is proposed to address more practical scenarios by achieving continuous domain adaptation using unlabeled streaming data in deployment environments. Our proposed task *LUDA person ReID* and its corresponding scheme *CLUDA-ReID* take an important next step in privacy protection since we do not need any identity annotations of real persons. (Note that we pre-train the person ReID models on synthetic data.) Despite this, it may still cause a violation of human privacy. Therefore, governments and officials need to carefully formulate strict regulations and laws to ensure the legal use of person ReID related technologies and strictly protect the data.

References

[1] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996. 3

[2] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *ICLR*, 2020. 3

[3] 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. 3

[4] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 3

[5] 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. 2

[6] John Riccitiello. John riccitiello sets out to identify the engine of growth for unity technologies (interview). *VentureBeat. Interview with Dean Takahashi. Retrieved January*, 18(3), 2015. 2

[7] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *CVPR*, pages 608–617, 2019. 2

[8] Jeffrey S Vitter. Random sampling with a reservoir. *TOMS*, 11(1):37–57, 1985. 1

[9] 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. 2

[10] 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. 2

[11] 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. 2