

# Camera Alignment and Weighted Contrastive Learning for Domain Adaptation in Video Person ReID

Djebril Mekhazni\*, Maximilien Dufau\*, Christian Desrosiers, Marco Pedersoli, Eric Granger  
LIVIA, Dept. of Systems Engineering, ETS Montreal, Canada  
{djebril.mekhazni, maximilien-francois.dufau.1}@ens.etsmtl.ca,  
{christian.desrosiers, marco.pedersoli, eric.granger}@etsmtl.ca

## Abstract

*Systems for person re-identification (ReID) can achieve a high accuracy when trained on large fully-labeled image datasets. However, the domain shift typically associated with diverse operational capture conditions (e.g., camera viewpoints and lighting) may translate to a significant decline in performance. This paper focuses on unsupervised domain adaptation (UDA) for video-based ReID – a relevant scenario that is less explored in the literature. In this scenario, the ReID model must adapt to a complex target domain defined by a network of diverse video cameras based on tracklet information. State-of-art methods cluster unlabeled target data, yet domain shifts across target cameras (sub-domains) can lead to poor initialization of clustering methods that propagates noise across epochs, thus preventing the ReID model to accurately associate samples of same identity. In this paper, an UDA method is introduced for video person ReID that leverages knowledge on video tracklets, and on the distribution of frames captured over target cameras to improve the performance of CNN backbones trained using pseudo-labels. Our method relies on an adversarial approach, where a camera-discriminator network is introduced to extract discriminant camera-independent representations, facilitating the subsequent clustering. In addition, a weighted contrastive loss is proposed to leverage the confidence of clusters, and mitigate the risk of incorrect identity associations. Experimental results obtained on three challenging video-based person ReID datasets – PRID2011, iLIDS-VID, and MARS – indicate that our proposed method can outperform related state-of-the-art methods. Our code is available at: <https://github.com/dmekhazni/CAWCL-ReID>*

## 1. Introduction

Person re-identification (ReID) is an important computer vision function aiming to match individuals captured over distributed network of non-overlapping video cameras. It has a growing interest in real-world applications, such as video surveillance, search and retrieval, and pedestrian tracking for autonomous driving. The non-rigid structure of human bodies, variations in capture conditions in the wild, due to changes in illumination, motion blur, resolution, pose, view point, and occlusion and background clutter leads to high intra-person and low inter-person variability. Despite these considerable challenges, recent progress in deep learning (DL) has allowed to develop state-of-art person ReID models that can achieve an impressive level of performances, especially then they are trained end-to-end in a supervised fashion on large fully-annotated image datasets [12].

Given the increasing demand in surveillance and camera networks, video-based ReID has received much attention in recent years. Video-based ReID systems rely on tracklets – a consecutive set of bounding box images captured from the same individual – for person matching, rather than single image. This approach may provide a richer representations by leveraging the temporal coherence of tracklets, allowing for temporal modeling to, e.g., attention-based aggregation of the more important features over time. This reduces the volume of data that must be stored for person ReID, and the complexity for similarity matching [28].

The distribution of image data captured with different cameras and operational conditions may differ considerably, a problem known as domain shift [24]. Given this problem, state-of-the-art DL models that undergo supervised training with a labeled image dataset (from the source domain) often generalize poorly for images captured in a target operational domain, leading to a decline in ReID accuracy. Moreover, labeled data are

---

\*These authors contributed equally to this work

costly to obtain especially in ReID due to the nature of the task where unknown and new identities are captured very often. Unsupervised Domain Adaptation (UDA) focuses on transferring the knowledge from a labeled source ( $S$ ) dataset to an unlabeled target ( $T$ ) dataset. It seeks to resolve the domain shift problem by leveraging unlabeled data from the target domain, in conjunction with labeled source domain data, to bridge the gap between the different domains. Most UDA methods for pair-wise similarity matching have been proposed for image-based person ReID, and adopt approaches for domain-invariant feature learning [30, 16, 38, 24, 25, 27], adversarial training [44, 43, 34, 7, 2, 31], and clustering [8]. Recently work on adversarial multi-source and multi-target UDA [29, 26, 25] has spurred a growing interest in leveraging the diverse data from individual domains or sub-domains (e.g., video cameras). For instance, the typically apply extended multi-domain adversarial classifiers with information-theoretic weighting strategies allow to regulate training.

Given the high cost of annotation, video-based ReID has mostly focused on adaptation with an unsupervised setup, where the source dataset is not available. Clustering and noisy pseudo-labels are used to train the ReID network with no target ground-truth labels [3, 17, 35, 19]. To the best of our knowledge, the UDA problem has not been addressed for video-based person ReID. However, given the complexity of this problem, leverage knowledge provided by video tracklets, or by a small quantity of labeled data may be extremely valuable. We observed that the clustering strategy may consequently require tuning of its hyper-parameters [19] due to a complex feature space at the initialization. Indeed, a misalignment between the data distribution of target domains cameras at initialization can prevent convergence of the training process. Inspired by the clustering and contrastive loss in [8], we propose to improve the training process on pseudo-labels by improving the alignment of camera data before initialization of pseudo-labels with an adversarial approach. In addition, to mitigate the impact of noisy pseudo-labels during training, we propose a leveraging information from video tracklets, and novel soft contrastive loss. An aspect with these new techniques, and that is not exploited in the literature, is the use of meta-information, such as the camera label and the atomicity of tracklets, i.e, the fact that all the images in a tracklet have one same identity. That were not used before in the literature and is an observation essential for video person ReID.

Our contributions can be summarized as follows. (1) An adversarial camera-level confusion loss is introduced to address the shift in data distributions observed across individual cameras in the target domain. (2) An ap-

proach for association of pseudo-labels with the positive pairs (video clips), build from sub-parts of tracklets. These strong positive pairs provide a meaningful heuristic to the clustering algorithm. (3) A soft contrastive loss with dynamic weighting is proposed that relies on a  $k$ -Nearest Neighbors weighing strategy and video clip-based augmentation to learn discriminant tracklet representations. (4) An extensive set of experiments conducted on three challenging datasets for video-based ReID datasets, PRID2011 [11], iLIDS-VID [33] and MARS [41], that show our proposed approach outperforming related state-of-the-art models.

## 2. Related Work

### 2.1. Video-Based ReID

In video based person ReID, input samples are video-sequences of the same identity (tracklets). The main objective is to find a good representation of the whole tracklet in order to use it for ReID. This representation is richer than the one extracted from a single image as the temporal information is exploitable, as well as more features can be extracted than in the single image case. The exploitation of this new information required the development of spatio-temporal aggregator networks.

**RNNs:** First attempts of aggregation were mostly via Recurrent Neural Networks (RNNs) [23, 4, 1, 39, 45]. It consists in aggregating frame-level features for each individual. Temporal context is stored from previous frames in the hidden state of the RNN. McLaughlin et al. [23] proposed a RNN architecture optimizing both final recurrent layer and temporal pooling layer. A weighted version of it was proposed by Chung et al. [4].

**Attention Networks:** Another approach for aggregation is to use Attention networks, using temporal information to select the most discriminating frames. While Liu et al. [21] consider only temporal attention, Li et al. [18] and Fu et al. [6] combine spatial and temporal attention. Hou et al. [12] differ from previous approaches avoiding redundancy on features of different frames by designing a temporal complementary learning mechanism. They proposes an integral characteristic of each identity by extracting complementary features for consecutive frames of a video. These works leverage the temporal information brought by the video to build richer representations. However these techniques do not leverage video-specific meta-information such as the fact that all the persons appearing in a single tracklet belong to the same identity. Moreover, those methods do not address the problem of domain shift.

## 2.2. Unsupervised Domain Adaptation for ReID

Unsupervised Domain Adaptation rely on transferring knowledge from a labeled source domain to an unlabeled target domain. The goal is to learn discriminant and domain-invariant representations from both domains. Despite the great results achieved by DL methods on supervised person ReID, a domain-shift between source and target data is often present due to the difference across domains (pose, viewpoint, illumination, ...). This leads models trained on source to a drastic performance drop when deployed on target. Different approaches are proposed to address this problem, separated into: domain-invariant feature learning and clustering based.

**Domain-invariant feature learning:** Discrepancy-based approaches measure the shift between source and target distribution. Their objective is to minimize the loss characterising the domain-shift gap, trying to align the confused target distribution on the well separated source distribution [30, 16, 38, 24, 25]. Methods that generate synthetic labeled data aim to produce similar images on source and target data [44, 43, 34]. However results can vary depending on the quality of the image generation. Adversarial-based methods are designed to train a domain discriminator whose goal is to fool the feature extractor using a gradient reversal layer [7, 2, 31]. This allows the feature extractor to make abstraction of the origin of the domain. A supervised loss is present during training not to forgot the initial task and avoid collapse of the representation.

**Clustering:** Clustering-based approaches work by clustering unlabeled target data and assign them pseudo-labels. Then those labels are used to train the network with a self-supervised method such as the minimization of a contrastive loss defined with pseudo-positives and negatives pairs [19, 3, 17, 8]. The main issue with this family of approaches is about the quality of the pseudo-labels. Indeed, noise from false positive or false negative pairs could lead to a diverging training. In this context, some criteria are designed to evaluate the coherence of clusters [8]. Contrastive Learning typically requires large batch size to reach satisfying performances. However, the batch size is limited by the size of the GPU memory available. This is particularly important in video-based ReID where each representation is created from multiples images, thus requiring a great amount of memory during training. [37, 9] allow to contrast samples with the whole dataset by storing in memory all the representation calculated at the beginning of each epoch and then update the network batch-wise. He et al. [9] extend this idea with a momentum update after each batch to follow the update by the back-propagation of selected samples.

## Adversarial Approaches for Multi-Source UDA:

Shen et al. [29] proposed a source-free multi-source free DA approach. They estimate the importance of each pretrained model on source to conduct the adaptation onto the target domain incrementally. Park and Lee [26] proposed another adversarial approach in the context of multi-source domain adaptation based on classifiers. They dealt with the computational scalability, and the very sensitive gradient variations presents in this kind of approach through a information-theoretic method.

## 2.3. Self-paced learning

Self-paced learning [15] refers to add progressively more difficult samples across the epochs of the training. It requires to be able to estimate a confidence score and it can help the convergence of the training. In case of semi-supervised or unsupervised learning in a closed-set environment [32, 29] it is possible to base the evaluation on the output of a classifier to work with its probabilities. Also we need to differentiate the weighting of domains in case of multiple source or target [29] which is different and does not apply in our case. However, it is not possible to obtain this value in the context of the unsupervised open-set problem. In this case, identities between source and target are totally different and we do not know the number of classes also. As mentioned above, clustering is very noise-sensitive. In person ReID, Ge et al. [8] proposed two criteria to manage the clustering based on tightness and compactness. In contrast, we propose to assimilate a weight to pairs in the contrastive loss directly, allowing less negative impact in case of bad associations.

## 2.4. Critical Analysis

Although unsupervised ReID has been widely studied recently, most of the works were designed in an image-based scenario, not taking advantage of temporal features [8, 44, 34] yet full of information inside a tracklet specific to its identity. Moreover, in this work we are the first at our best knowledge who consider the camera alignment issue for a given domain, which can lead to a poor representation of the tracklets. As clustering-approaches are very noise-sensitive, the camera alignment is really important for good results and we propose an adversarial loss to fulfill it. Finally, few existing works in unsupervised video-based ReID leveraged the information of a source dataset in a video context. In a more general context, video-based approaches are missing a more comprehensive study regarding the importance of the tracklet [19, 3, 17]. Although the development of pertinent features is an important point to be approached, managing in the first place the raw data across tracklets requires still some

efforts to fully benefit the hidden information available from the raw data.

### 3. Proposed Approach

In this section, the overall approach shown in Figure 1 is presented module by module. Most methods for video person ReID learn to associate tracklets to the correct identity through a clustering approach based on pseudo-labels. However, those method strongly rely on the quality of the pseudo-labels in the early training iterations. Poor initialization will prevent iterative convergence and the whole training process will fail. We propose a novel training process based on domain adaptation that starts by consolidating the performances of the encoder on the Target dataset with an Adversarial Camera confusion technique. After improving the cluster quality thanks to previous step, the network can start the self-paced pseudo-label based training. Self-paced is essential to estimate bad clusters and avoid encouraging wrong directions during the training. From our best knowledge, we are the first using the camera ID information to decrease the encoder discrimination on the Target set.

#### 3.1. Tracklet Representation

Given a set  $\mathcal{D}$  of  $N$  tracklets, each tracklet is composed of a sequence of  $N_\tau$  images  $\{x_1, \dots, x_{|N_\tau|}\}$ . We consider that the first  $N_s < N$  tracklets are from a source domain  $S$  and come with a personID label  $y$ , while the remaining  $N_t = N - N_s$  ones are from a target domain for which there is no person ID. The goal is to learn a function  $f$  mapping images  $x$  to a representation vector  $f(x)$  which is able to discriminate identities, i.e. find separable distances to consider two pairs as same individual or no, while being robust to differences between the source and target domains.

Since input data are tracklets, a features aggregator is employed so each video-sequence has its own unique representation. The aggregation method TCLNet proposed by Hou [13] is adopted. By exploiting spatio-temporal information, TCLNet is the association of Temporal Saliency Erasing (TSE) and Temporal Saliency Boosting (TSB) modules. We denote the backbone  $\mathcal{M}$  as the layer just before the classification layer. It produces the representation of a tracklet  $g(\tau)$ .

**Clips:** In the video-based context, it happens that tracklets contain different mode. Between the starting frame and the final frame, events could have arrived (occlusion, change in trajectory, ...). From this observation, we proposed to split each tracklet  $\tau$  into  $N_\Omega$  equal-sized clips  $\Omega$ . The tracklet becomes then  $\tau = \{\Omega_0, \dots, \Omega_{N_\Omega}\}$

#### 3.2. Training Strategy

An association of a supervised loss, a camera confusion loss and a self-paced contrastive learning loss for domain adaptation is proposed as follow:

$$\mathcal{L} = \delta_1 \mathcal{L}_{\text{CE}} + \delta_2 \mathcal{L}_{\text{cam}} + \delta_3 \mathcal{L}_{\text{contr}}. \quad (1)$$

where  $\mathcal{L}_{\text{CE}}$  refers to the supervised training on source = S domain,  $\mathcal{L}_{\text{cam}}$  refers to the camera classifier to avoid learning camera-identity related, and  $\mathcal{L}_{\text{contr}}$  refers to the contrastive loss on clustered samples.  $\delta_i$  are weighting factors used to balance the magnitude of the losses.

#### 3.3. Supervised training

The first loss, which uses only source data, ensures that the tracklet representation  $g(\tau)$  can be used for identity classification through an identity classifier  $\theta_I$ . It converts  $g(\tau)$  to a class probability vector with a linear layer followed by a softmax, and uses cross-entropy as classification loss. It can be expressed as:

$$\mathcal{L}^{\text{CE}} = \frac{1}{K} \sum_i^K -\log \left( \frac{\exp(W_{y_i}^T g(\tau_i) + b_{y_i})}{\sum_{j=1}^{N_I} \exp(W_j^T g(\tau_i) + b_j)} \right) \quad (2)$$

where  $K$  is the batch size. Class label  $y_i \in \{1, 2, \dots, N_I\}$  is associated with training tracklet  $g(\tau_i)$ , the  $i^{\text{th}}$  training tracklet. Weight vectors  $W_{y_i}$  and bias  $b_{y_i}$  of last fully connected (FC) layer corresponds to class  $y$  of  $i^{\text{th}}$  tracklet.  $W_j$  and  $b_j$  are weights and bias of last FC corresponding of the  $j^{\text{th}}$  class ( $j \in [1, N_I]$ ).  $W_j$  and  $b_j$  are respectively the  $j^{\text{th}}$  and  $y_i^{\text{th}}$  column of  $W = \{w_{ij} : i = 1, 2, \dots, F; j = 1, 2, \dots, N_I\}$ , where  $F$  is the size of the last FC layer of the backbone.

#### 3.4. Camera-level confusion loss

We empirically observed a mis-alignment of target = T inter camera distributions after the pre-training on the Source dataset. This lack of alignment yield poor initial pseudo-labels. Inspired by [7] we use an adversarial distribution alignment approach. We propose to align the per-camera distribution rather than the domains. A camera classifier  $\theta_c$  is trained to discriminate the target = T cameras with a cross entropy loss, the gradient of this loss is inverted with a Gradient Reversal Layer (GRL) and back-propagated through the encoder. This reversed objective aims to remove the information related to the camera scene from the representation and helps to make the source and target domains indistinguishable for the classifier. Using the cross-entropy loss to train the camera classifier incentives the encoder to minimize the classifier accuracy, which can be unstable.

$$\mathcal{L}_{\text{cam}}^{\text{CE}} = \frac{1}{K} \sum_i^K -\log \left( \frac{\exp(W_{c_i}^T \theta_c(g(\tau_i)))}{\sum_{j=1}^{N_c} \exp(W_j^T \theta_c(g(\tau_i)))} \right) \quad (3)$$

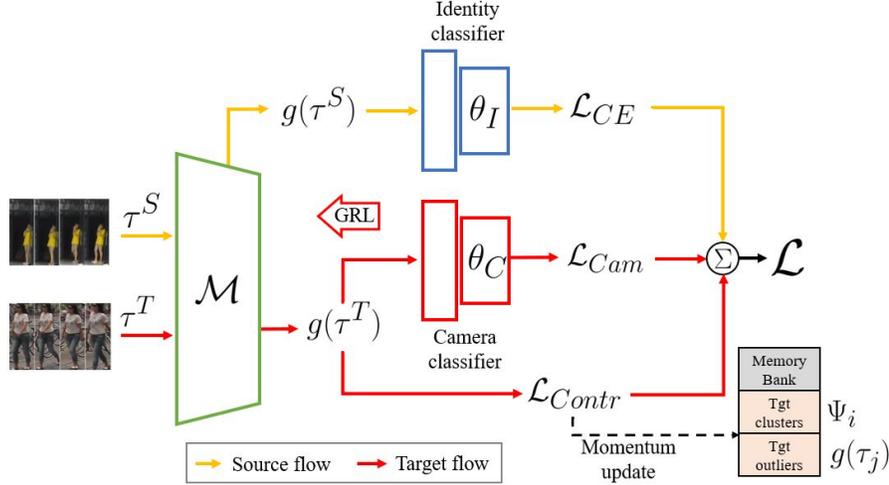


Figure 1: Training architecture of the proposed model for video-based UDA in ReID. Source and target tracklets are processed by the backbone  $\mathcal{M}$  and a representation  $g(Tk)$  is computed. The training loss  $\mathcal{L}$  is composed of three part: the supervised cross-entropy loss on the source tracklets ( $\mathcal{L}_{CE}$ ), camera confusion loss on source and target tracklets ( $\mathcal{L}_{Cam}$ ) and cotrastive loss ( $\mathcal{L}_{Contr}$ ), which uses a memory bank of clustered and unclustered tracklets.

where  $N_c$  is the number of cameras  $T$  in the target domain. We therefore propose a novel Camera Level Confusion loss  $\mathcal{L}_{cam}^{CLC}$  to train the camera classifier as:

$$\mathcal{L}_{cam}^{CLC} = \frac{1}{K^2} \left[ \sum_{i=1}^K \log \frac{\exp c_{y_i, i}}{\sum_{j=1}^{N_c} \exp(c_{j, i})} \right] \left[ \sum_{i=1}^K \log \frac{\exp(1 - c_{y_i, i})}{\sum_{j=1}^{N_c} \exp(1 - c_{j, i})} \right] \quad (4)$$

where  $c_{j, i}$  is the output of the classifier for camera  $j$  and sample  $i$ . This loss, encourages the classifier to be sure about his predictions, may they be right or wrong. The encoder is then encouraged to maximise the classifier confusion, i.e maximise the conditional entropy  $H(C|g(\tau))$  where  $C$  is the distribution of cameras and the distribution of tracklets. A similar confusion based loss for adversarial approaches was proposed in [26].

### 3.5. Contrastive loss

The contrastive loss is based on pseudo-positive/negatives pairs of samples assimilated to the same or different clusters from  $T$  data. It is defined as:

$$\mathcal{L}_{contr} = \frac{1}{N_t} \sum_{i=1}^{N_t} \sum_{j \neq i} \log \frac{\exp(\tilde{g}(\tau_i)^\top \tilde{g}(\tau_j))}{\sum_{j' \neq i} \exp(\tilde{g}(\tau_i)^\top \tilde{g}(\tau_{j'}))} \quad (5)$$

where  $\tilde{g}(\tau) = g(\tau) / \|g(\tau)\|$ . This loss helps to adapt the tracklets representations such that different identities are easily distinguished on both source and target data.

**1) Generation of Pseudo-Labels:** To generate pseudo-labels, DBSCAN clustering algorithm is used. A Jaccard distance matrix, calculated based on the  $k$ -reciprocal nearest neighbours sets, is used for this clustering phase. The choice of such a metric enforce more robust clusters. Indeed, samples that share reciprocal neighbors tend to form positive pairs [8, 42].

**2) Self-Paced Strategy:** Let  $w_{ij} \in [0, 1]$  be a self-paced weight representing the likeliness that tracklets  $\tau_i$  and  $\tau_j$  correspond to the same person [14].

For target examples, we need to optimize both embedding function  $f$  and self-paced weights  $w_{ji}$ . The update of self-paced weights is controlled by a self-paced regularization function  $R_\gamma$ , parameterized by the learning pace  $\gamma > 0$ . This function has the property of being monotone decreasing with respect to the loss (i.e., harder examples are given less importance) and monotone increasing with respect to the learning (i.e., a larger  $\gamma$  increases the weights). Here, we consider for this function a linear self-paced learning:

$$R_\gamma(w_{ij}) = -\gamma w_{ij}. \quad (6)$$

Given the self-paced contrastive loss (5):

$$\mathcal{L}_{contr} = \frac{1}{N_t} \sum_{i=1}^{N_t} \sum_{j \neq i} w_{ij} l_{ij} + R_\gamma(w_{ij}), \text{ where} \quad (7)$$

$$l_{ij} = -\log \frac{\exp(\tilde{g}(\tau_i)^\top \tilde{g}(\tau_j))}{\sum_{j' \neq i} \exp(\tilde{g}(\tau_i)^\top \tilde{g}(\tau_{j'}))} \quad (8)$$

Since  $l_{ij} \geq 0$  and  $w_{ij} \in [0, 1]$ ,  $w$  can be optimized in closed-form for a fixed embedding function  $g$  as:

$$w_{ij} = \max\left(1 - \frac{1}{\gamma}l_{ij}, 0\right). \quad (9)$$

The overall optimization process starts with a small learning rate  $\gamma_0 \approx 0$ , at each iteration, we update the embedding function  $g$  gradient descent and updating the weights  $w_{ij}$  as in (9). At the end of each iteration, the learning pace can be increased according to a fixed policy, for instance  $\gamma := (1 + \alpha)\gamma$ , with  $\alpha > 0$ . In our method, we set  $\gamma$  based on the maximum distance between an example and its  $k$ -nearest neighbor. This policy treats  $k$ -nearest neighbors as partially positive pairs and thus mitigates the harm of false-negative close pairs. Indeed, close pseudo-negative pairs are more likely to actually be positive pairs.

**3) Memory bank:** To alleviate memory issues and allow the full use of all data for the contrastive loss, allowing a more reliable comparison pairs to pairs at each batch step, a memory bank is employed to store the features of all samples at the beginning of each epoch. Clustered samples have the mean of the features belonging to the same cluster  $\Psi_i$ , with  $i$  the  $i$ -th cluster, saved in the memory. For unclustered samples we save their features  $g(\tau_j)$  directly, with  $j$  the  $j$ -th sample unclustered. It consist into considering a query  $q$  and a set of encoded samples  $\{k_0, k_1, \dots\}$  that are keys of a dictionary. The contrastive loss is then calculated considering samples in the batch as  $q$  and all other samples store in the dict as keys. The backpropagation and weight update for the encoder are made only with samples in the batch, not including the keys being detached from the graph. To avoid memory issues in storing all the graph for them, we update the keys differently, according to the following rule [9]:

$$\theta_k \leftarrow m\theta_k + (1 - m)\theta_q. \quad (10)$$

In this context,  $m \in [0, 1]$  and  $\theta_k$  (resp.  $\theta_q$ ) refers to the representation of the keys (resp. query).

## 4. Results and Discussion

### 4.1. Experimental Methodology

The proposed method is evaluated on the three most challenging public video re-identification benchmarks.

PRID2011 [11] is the least challenging of the three benchmarks, it features 356 videos from 178 identities captured by two cameras monitoring crosswalks and sidewalks. iLIDS-VID [33] is captured in an airport hall. It is a two camera benchmark with 600 videos from 300 identities. This benchmark is challenging due

to pervasive background clutter, mutual occlusions, and lighting variations. MARS [41] is the current largest open video-based pedestrian retrieval benchmark. It features six cameras on a university campus, for a total of 20,478 videos from 1,261 identities.

Property	MARS	PRID2011	iLIDS
# Identities	1251	178	300
ID Train/Query/Gallery	625/626/622	89/89/89	150/150/150
Train/Query/Gallery	8298/1980/9330	178/89/89	300/150/150
# Tracklets	19608	356	600
# Image per Tracklet	2~920 (59.5)	28~675 (108.1)	22~192 (70.8)
# Cameras	6	2	2
# Images	1,191,003	24,541	42,495
Annotation Method	DPM+GMMCP	Hand	Hand
Crop Size	256x128	128x64	Vary

Table 1: Description of the different datasets used in experiments – MARS, PRID2011 and iLIDS-VID.

The proposed method is implemented using PyTorch. We use ResNet50 [10] pretrained on ImageNet as the backbone of the spatio-temporal aggregator TCLNet [13]. We separate each video in  $K = 2$  clips, from which we randomly select four images to form our input tracklet. We train the model for 150 epochs with batch size 24. Adam optimizer is used to perform gradient descent with an initial learning rate of 0.0003, a scheduled decay at epochs 40, 80 and 120 and weight decay parameter  $\gamma$  equal 0.1.  $\delta_1$  is set at 1,  $\delta_2$  is set at 0.2, and  $\delta_3$  at 1. At evaluation time the features are  $\ell_2$  normalized. We use the  $\ell_2$  distance between the gallery and query set for the final matching. For the clips approach, we proposed to split the tracklet into  $N$  parts by taking  $N$  times successive frames from the beginning to the end. In our case  $N$  is set at 2 for best results as stated Table 4, meaning the first 50% frame of each tracklet are the first clip and the rest the second clip. Then each clips is separated by a fixed number of equal chunks (4 in our case) and randomly take a frame from each chunks. We believe that it can be used for video-based person ReID and even be extended for Video-based tasks such as action recognition as an example.

### 4.2. Ablation Study

In this section, we assess the impact on performance of different losses for the camera classifier and contrastive learning, and the use of clips for the video tracklets.

**1) Adversarial Adaptation for Camera-Wise Domain Shift:** Table 2 compare the two proposed losses for the camera classifier  $\mathcal{L}_{cam}^{CE}$  and  $\mathcal{L}_{cam}^{CLC}$  in the first section. We observe that the impact on results is significant for all scenario compared to the baseline (deployment of the source trained model on target directly). We achieve 61.8% (PRID2011), 35.7% (iLIDS-VID) and 30.8% (MARS) Rank-1 for  $\mathcal{L}_{cam}^{CE}$ . Even better is our proposed loss  $\mathcal{L}_{cam}^{CLC}$  that improve results by 4.5% (PRID2011), 3.3% (iLIDS-VID) and 0.7% (MARS).

Loss Functions	iLIDS → PRID	PRID → iLIDS	iLIDS → MARS
	(2 cameras)	(2 cameras)	(6 cameras)
Lower Bound (sup. $S$ only)	49.0	12.7	19.8
$\mathcal{L}_{\text{cam}}^{\text{CE}}$	61.8	32.0	30.8
$\mathcal{L}_{\text{cam}}^{\text{CLC}}$	66.3	35.3	31.5
$\mathcal{L}_{\text{cam+clips}}^{\text{CLC}}$	75.3	36.0	32.3
$\mathcal{L}_{\text{contr}}$	77.6	41.9	37.6
$\mathcal{L}_{\text{contr}}^{\text{kNN}}$	78.2	44.0	38.5
$\mathcal{L}_{\text{contr+clips}}^{\text{kNN}}$	79.8	44.7	38.9
$\mathcal{L}_{\text{contr}} + \mathcal{L}_{\text{cam}}^{\text{CLC}}$	83.9	54.0	59.6
$\mathcal{L}_{\text{contr}}^{\text{kNN}} + \mathcal{L}_{\text{cam}}^{\text{CLC}}$	84.2	58.0	60.2
$\mathcal{L}_{\text{contr+clips}}^{\text{kNN}} + \mathcal{L}_{\text{cam+clips}}^{\text{CLC}}$	<b>86.5</b>	<b>58.3</b>	<b>62.2</b>

Table 2: Rank-1 accuracy with different losses functions for the camera classifier and contrastive learning, and when subdividing tracklets into 2 clips (best result obtained in Table 4).

Method	Setting	iLIDS → PRID		PRID → iLIDS		iLIDS → MARS	
		(2 cameras)		(2 cameras)		(6 cameras)	
		Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
Lower Bound (sup. $S$ only)	–	49.0	60.0	12.7	20.9	19.8	10.1
DGM+IDE [40], ICCV’19	OneEx	–	–	–	–	36.8	16.8
Stepwise [22], ICCV’17	OneEx	–	–	–	–	41.2	19.6
EUG [36], CVPR’18	OneEx	–	–	–	–	<b>62.2</b>	42.5
TAUDL [3], BMVC’18	Unsup	85.3	–	56.9	–	46.8	21.4
UTAL [17], TPAMF’19	Unsup	54.7	–	35.1	–	49.9	35.2
UGA [35], ICCV’19	Unsup	80.9	–	57.3	–	58.1	39.3
BUC [19], AAAI’19	Unsup	–	–	–	–	61.1	38.0
Soft Sim. [20], CVPR’20	Unsup	–	–	–	–	61.9	43.6
SPCL* [8], NeurIPS’20	UDA	77.6	82.1	41.9	47.6	37.6	20.4
Ours ( $\mathcal{L}_{\text{cam}}^{\text{CLC}}$ )	UDA	70.8	77.3	32.0	42.6	31.5	16.3
Ours ( $\mathcal{L}_{\text{cam}}^{\text{CLC}} + \mathcal{L}_{\text{contr}}^{\text{kNN}}$ )	UDA	<b>86.5</b>	<b>89.9</b>	<b>58.3</b>	<b>66.7</b>	<b>62.2</b>	<b>44.8</b>
Upper Bound (sup. $S \cup T$ )	Tuning	92.1	94.5	76.0	84.0	86.9	81.8

Table 3: Accuracy of the proposed and state-of-art methods when alternating iLIDS, PRID and MARS datasets as sources target domains ( $S \rightarrow T$ ). Lower bound performance is obtained through supervised training on the labeled source dataset only, while the upper bound after supervised fine-tuning on both source and target data. “OneEx” (One Example) settings indicates approaches that labeled one tracklet per identity. “Unsup” (Unsupervised) is for end-to-end setting without using source data, while UDA involves using source data. \*SPCL was designed for image-based person ReID, while unsupervised methods were designed considering temporal information in tracklets.

Figure 2 illustrates a before and after the distribution shift across cameras of the target domain. We observe that our method successfully shorten the distribution shift between the different cameras on PRID2011. We conducted experiments using source cameras also, but since the source already associated pairs from different cameras close to perfectly, results were similar or a little bit under because of potential degradation from the camera classifier. Moreover, the introduction of the confusion loss yields a significant improvement in performances. This is a consequence of the more consistent target for the encoder network that aims at

maximising the conditional entropy introduced in Section 3.4 rather than minimizing the accuracy of the classifier, which can lead to unstable training performances.

**2) Soft Contrastive Loss:** In the second part of Table 2, we compare the performance of the initial technique from SPCL [8] applied for Video-based ReID scenario, that we named  $\mathcal{L}_{\text{contr}}$ , and the proposed amelioration  $\mathcal{L}_{\text{contr}}^{\text{kNN}}$  where we weight the contrastive loss based on the criterion using k- nearest neighbors. We observe that the kNN based weighting improve performances by 0.6% (PRID2011), 2.1%

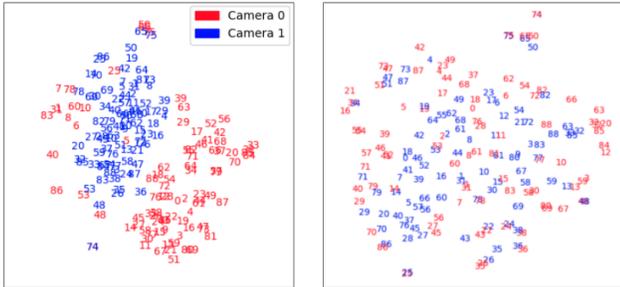


Figure 2: T-SNE plot showing the shift in data distribution among PRID2011 cameras (T) before (left) and after (right) adversarial adaptation. Each number represents an identity.

(Iuids-Vid) and 0.9% (MARS) for Rank-1.

**3) Combination of Losses:** On one side, the contrastive loss [8] train the network based on a self-paced clustering. On the other side, the camera classifier eliminates camera shift observed in Video-based and allow an improvement on results. These two components both aim to improve performance by mitigating the harm caused by incorrect pseudo labels in different manners and can be used in combination. The combination of the losses allows the method to improve results by a significant margin as presented in the last section of 2. Especially, the harder the dataset is and the larger the impact of the combination. Indeed, for PRID we only have 6.0% of improvement compared to Section 2 while we not that for iLIDS we have 14.0% improvement and 21.7% for MARS. This indicates the need for the camera classifier, especially for complex domains where the generation of pairs is hard due to domain shift across cameras.

**4) Video ReID Using Clips:** Since Video-based person ReID has access to each frame of the whole tracklet, it is possible to split it in clips. We observe in Table 2 that dividing each tracklet in clips, as explained in Section 3 improves the performance of our method for all loss combinations. It allows the model in its best configuration to reach 86.5% Rank-1 accuracy for PRID (+ 2.3%), 58.3% for iLIDS-VID (+ 0.3%) and 62.2% for MARS (+ 2.0%). This can be explained by multiple facts. First, the number of samples in the dataset increases, so it can be seen as data augmentation but without transformation. For smallest dataset it is pertinent because it allow to avoid overfitting. Second, it generates strong reliable positives pairs, which provides a supervised signal even with unlabeled data. From our experiments we observed that splitting a tracklet into two clips is the most promising as stated in Table 4.

### 4.3. Comparison with the State-of-Art

Table 3 compare results to the existing literature for Unsupervised Person ReID [3, 17, 35, 19, 5, 20, 8] have been conducted. For more reference we included few shot techniques [40, 22, 36] where they allow one tracklet per identity to be labeled, giving them a considerable advantage,

#clips	PRID	iLIDS	MARS
1	84.2	58.0	60.2
2	86.5	58.3	62.2
3	86.1	58.1	61.9
4	85.4	58.0	60.9
8	83.1	56.6	59.5

Table 4: Rank-1 accuracy of the proposed method as a function of the number of clips per tracklet.

especially because it informs about the total number of classes. In our case, the Domain Adaptation part is present only for the camera classifier where we need to maintain a supervised training to adjust properly the confusion on the backbone, so the backbone is not fully updated via the camera classifier. The clusters and the contrastive loss are performed without any source or supervised information. For end-to-end approaches, TAUDL [3] and UTAL [17] have trouble to perform well on the three proposed dataset. Indeed as example TAUDL perform well on PRID (85.3% Rank-1 compared to 86.5 % for ours) but fails when dealing with complex dataset (46.8% Rank-1 compared to 62.2% for us). BUC [19] perform well but is clustering based approach and requires high amount of resources in term of hyper-parameters search. We have a margin of 1.1% Rank-1 and 6.8% mAP compared to them. In the same idea is DBC [5] but we did not include them because it is hard to reproduce their results. Compared to methods with one example setup, we outperform DGM+IDE and Stepwise [40, 22] by a significant margin. We obtain similar results compared to EUG [36] for Rank-1 but reach 2.3% more mAP in our case without any labeled tracklets.

## 5. Conclusion

This work proposes a novel training approach for the domain adaptation video-based person ReID task, that can be applied to most existing encoders architectures. To solve this challenging task we propose to use the meta-information of the problem to reduce our dependence to the noisy pseudo-labels. Our method achieves competitive results with state-of-the-art. Moreover, apart from the camera classifier, our ideas aren't specific to the video-based person ReID task and could be use in others time-series representation learning tasks. Future direction of the project include improving the inter-Clips diversity with ReID specific augmentation and sampling, as well as improving the clustering with the prior obtained with the usage of clips.

## Acknowledgments:

This research was supported in part by the Ericsson Global AI Accelerator, Compute Canada and MITACS.

## References

- [1] Jean-Baptiste Boin, Andre F. de Araújo, and Bernd Girod. Recurrent neural networks for person re-identification revisited. *CVPR*, 2016.
- [2] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. *Advances in neural information processing systems*, 2016.
- [3] Yanbei Chen, Xiatian Zhu, and Shaogang Gong. Deep association learning for unsupervised video person re-identification. *arXiv preprint arXiv:1808.07301*, 2018.
- [4] Dahjung Chung, Khalid Tahboub, and Edward J Delp. A two stream siamese convolutional neural network for person re-identification. *article of the IEEE international conference on computer vision*, 2017.
- [5] Guodong Ding, Salman H Khan, Zhenmin Tang, Jian Zhang, and Fatih Porikli. Dispersion based clustering for unsupervised person re-identification. *BMVC*, 2019.
- [6] Yang Fu, Xiaoyang Wang, Yunchao Wei, and Thomas S. Huang. STA: spatial-temporal attention for large-scale video-based person re-identification. *AAAI*, 2018.
- [7] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation, 2015.
- [8] Yixiao Ge, Dapeng Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. *Neurips*, 2020.
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. *CVPR*, 2020.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, 2015.
- [11] Martin Hirzer, Csaba Beleznaï, Peter M. Roth, and Horst Bischof. Person Re-Identification by Descriptive and Discriminative Classification. *Proc. Scandinavian Conference on Image Analysis (SCIA)*, 2011.
- [12] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Temporal complementary learning for video person re-identification. *ECCV*, 2020.
- [13] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Temporal complementary learning for video person re-identification. *European conference on computer vision*, 2020.
- [14] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- [15] M. Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010.
- [16] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. MMD GAN: towards deeper understanding of moment matching network. *CoRR*, 2017.
- [17] Minxian Li, Xiatian Zhu, and Shaogang Gong. Unsupervised tracklet person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [18] Shuang Li, Slawomir Bak, Peter Carr, and Xiaogang Wang. Diversity regularized spatiotemporal attention for video-based person re-identification. *CVPR*, 2018.
- [19] Yutian Lin, Xuanyi Dong, Liang Zheng, Yan Yan, and Yi Yang. A bottom-up clustering approach to unsupervised person re-identification. *article of the AAAI Conference on Artificial Intelligence*, 2019.
- [20] Yutian Lin, Lingxi Xie, Yu Wu, Chenggang Yan, and Qi Tian. Unsupervised person re-identification via softened similarity learning. *CVPR*, 2020.
- [21] Yu Liu, Junjie Yan, and Wanli Ouyang. Quality aware network for set to set recognition. *CVPR*, 2017.
- [22] Zimo Liu, Dong Wang, and Huchuan Lu. Stepwise metric promotion for unsupervised video person re-identification. *ICCV*, 2017.
- [23] Niall McLaughlin, Jesus Martinez Del Rincon, and Paul Miller. Recurrent convolutional network for video-based person re-identification. *CVPR*, 2016.
- [24] Djebri Mekhazni, Amran Bhuiyan, George S. Eskander Ekladios, and Eric Granger. Unsupervised domain adaptation in the dissimilarity space for person re-identification. *ECCV*, 2020.
- [25] Le Thanh Nguyen-Meidine, Atif Belal, Madhu Kiran, Jose Dolz, Louis-Antoine Blais-Morin, and Eric Granger. Unsupervised multi-target domain knowledge distillation. *WACV*, 2021.
- [26] Geon Yeong Park and Sang Wan Lee. Information-theoretic regularization for multi-source domain adaptation, 2021.
- [27] Félix Remigereau, Djebri Mekhazni, Sajjad Abdoli, Le Thanh Nguyen-Meidine, Rafael MO Cruz, and Eric Granger. Knowledge distillation for multi-target domain adaptation in real-time person re-identification. *ICIP*, 2022.
- [28] Arif Ahmed Sekh, Debi Prosad Dogra, Heeseung Choi, Seungho Chae, and Ig-Jae Kim. Person re-identification in videos by analyzing spatio-temporal tubes. *Multi-media Tools and Applications*, 2020.
- [29] Maohao Shen, Yuheng Bu, and Gregory Wornell. On the benefits of selectivity in pseudo-labeling for unsupervised multi-source-free domain adaptation. *arXiv preprint arXiv:2202.00796*, 2022.
- [30] Baochen Sun and Kate Saenko. Deep CORAL: correlation alignment for deep domain adaptation. *ECCV*, 2016.
- [31] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. *CVPR*, 2017.
- [32] Ping Wang, Jizong Peng, Marco Pedersoli, Yuanfeng Zhou, Caiming Zhang, and Christian Desrosiers.

- Self-paced and self-consistent co-training for semi-supervised image segmentation. *Medical Image Analysis*, 2021.
- [33] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. Person re-identification by video ranking. *ECCV*, 2014.
- [34] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer GAN to bridge domain gap for person re-identification. *CVPR*, 2017.
- [35] Jinlin Wu, Yang Yang, Hao Liu, Shengcai Liao, Zhen Lei, and Stan Z Li. Unsupervised graph association for person re-identification. *ICCV*, 2019.
- [36] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. *CVPR*, 2018.
- [37] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. *CVPR*, 2018.
- [38] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, 2017.
- [39] Yichao Yan, Bingbing Ni, Zhichao Song, Chao Ma, Yan Yan, and Xiaokang Yang. Person re-identification via recurrent feature aggregation. *ECCV*, 2016.
- [40] Mang Ye, Andy Jinhua Ma, Liang Zheng, Jiawei Li, and Pong C. Yuen. Dynamic label graph matching for unsupervised video re-identification. *ICCV*, 2017.
- [41] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. *ECCV*, 2016.
- [42] Zhun Zhong, Liang Zheng, Donglin Cao, and Shaozi Li. Re-ranking person re-identification with k-reciprocal encoding. *CoRR*, abs/1701.08398, 2017.
- [43] Zhun Zhong, Liang Zheng, Shaozi Li, and Yi Yang. Generalizing a person retrieval model hetero-and homogeneously. In *ECCV*, 2018.
- [44] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. *CVPR*, 2019.
- [45] Zhen Zhou, Yan Huang, Wei Wang, Liang Wang, and Tieniu Tan. See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification. *CVPR*, 2017.