

Single Camera Training for Person Re-Identification

Tianyu Zhang,¹ Lingxi Xie,² Longhui Wei,³ Yongfei Zhang,^{*,1,4} Bo Li,^{1,4} Qi Tian⁵

¹Beijing Key Laboratory of Digital Media, School of Computer Science and Engineering, Beihang University

²Johns Hopkins University, ³Peking University

⁴State Key Laboratory of Virtual Reality Technology and Systems, Beihang University

⁵School of Electronic Engineering, Xidian University

{zhangtianyu, yfzhang, boli}@buaa.edu.cn, {198808xc, wywqtian}@gmail.com, longhuiwei@pku.edu.cn

Abstract

Person re-identification (ReID) aims at finding the same person in different cameras. Training such systems usually requires a large amount of cross-camera pedestrians to be annotated from surveillance videos, which is labor-consuming especially when the number of cameras is large. Differently, this paper investigates ReID in an unexplored single-camera-training (SCT) setting, where each person in the training set appears in only one camera. To the best of our knowledge, this setting was never studied before. SCT enjoys the advantage of low-cost data collection and annotation, and thus eases ReID systems to be trained in a brand new environment. However, it raises major challenges due to the lack of cross-camera person occurrences, which conventional approaches heavily rely on to extract discriminative features. The key to dealing with the challenges in the SCT setting lies in designing an effective mechanism to complement cross-camera annotation. We start with a regular deep network for feature extraction, upon which we propose a novel loss function named multi-camera negative loss (MCNL). This is a metric learning loss motivated by probability, suggesting that in a multi-camera system, one image is more likely to be closer to the most similar negative sample in other cameras than to the most similar negative sample in the same camera. In experiments, MCNL significantly boosts ReID accuracy in the SCT setting, which paves the way of fast deployment of ReID systems with good performance on new target scenes.

1 Introduction

Person re-identification (ReID) aims to retrieve a certain person appearing in a camera network. With increasing concerns on public security, ReID has attracted more and more research attention from both academia and industry. In the past years, many algorithms (Sun et al. 2018; Wang et al. 2019; Suh et al. 2018; Zheng, Zheng, and Yang 2018) and datasets (Zheng et al. 2015; Zheng, Zheng, and Yang 2017b; Wei et al. 2018; Zheng, Karanam, and Radke 2018) have been proposed, which significantly boosted the progress of this research field. Despite the higher and higher accuracy obtained by specifically designed approaches on standard

*Corresponding author.

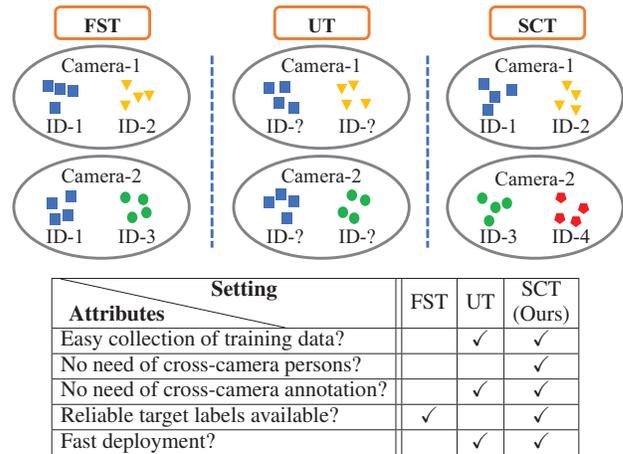


Figure 1: The comparison between SCT and previous settings in person ReID. *Fully-supervised-training* (FST) data are composed of annotated pedestrians appearing in multiple cameras. *Unsupervised-training* (UT) data have no identity annotations. Under our *single-camera-training* (SCT) setting, each pedestrian appears in only one camera and identity labels are easy to obtain.

ReID benchmarks, many issues of this task remain unsolved. As discussed in (Wei et al. 2018; Deng et al. 2018), a ReID model trained on one dataset performs poorly on other datasets due to dataset bias. Thus, to deploy a ReID system to a new environment, labelers have to annotate a training dataset from the target scene, which is often time-consuming and even impractical in large-scale application scenarios. To tackle this issue, researchers make a common assumption is that there exist a number of available unlabeled images in the target scene, based on which they design some unsupervised learning (Lin et al. 2019) or domain adaptation approaches (Wang et al. 2018b) to improve ReID performance. However, these methods depend on extra modules to predict the pseudo label of each image or generate fake images. Therefore, it is not always reliable and reports less satisfying performance compared to supervised learning methods.

Different from previous work, this paper investigates

ReID under a novel single-camera-training (SCT) setting, where each pedestrian appears in only one camera. We compare our setting to previous ones in Fig. 1. Without the heavy burden of annotating cross-camera pedestrians, training data with labels are easy to obtain under SCT. For example, using off-the-shelf tracking techniques (Keuper et al. 2018; Luo et al. 2019), researchers can quickly collect a large number of tracklets under each camera at different time periods, and thus it is very likely that each of them corresponds to a unique ID. Therefore, compared to the fully-supervised-training (FST) setting, *i.e.*, learning knowledge from cross-camera annotations, SCT requires much fewer efforts in preparing for training data. Compared to the unsupervised-training (UT) setting, which requires frequent cross-camera person occurrences, SCT makes a mild assumption on camera-independence, so as to provide weak but reliable supervision signals for learning. Therefore, SCT has the potential of being deployed to a wider range of application scenarios.

It remains an issue of how to make use of the camera-independence assumption to learn discriminative features for ReID. The most important one lies in camera isolation, which implies that there are no cross-camera pedestrians in the entire training set. Conventional methods heavily rely on cross-camera annotations because this is the key supervision that a model receives for metric learning. That is to say, by pulling the images of the same person appearing in different cameras close, conventional methods can learn camera-unrelated features so that they perform well on the testing set. With camera isolation in SCT, we must turn to other types of supervision to achieve the goal of metric learning.

To this end, we propose a novel loss term named Multi-Camera Negative Loss (MCNL). The design of MCNL is inspired by a simple hypothesis, that given an arbitrary person in a multi-camera network, it is more likely that the most similar person is found in another camera, rather than in the same camera, because there are simply more candidates in other cameras. To verify this, we perform statistical analysis on several public datasets, and the results indeed support our assumption (please see Fig. 2). Based on the above observation, our MCNL adjusts feature distributions and alleviates camera isolation problem by ranking the distances of cross-camera negative pairs and within-camera negative pairs. Extensive experiments show that MCNL can force the backbone network to learn more person-related features but ignore camera-unrelated representations, and then achieves good performance under the SCT setting.

Our major contributions can be summarized as follows:

- To the best of our knowledge, this paper is the first to present the SCT setting. Moreover, this paper analyzes the advantages and challenges under the SCT setting compared to existing settings in person ReID.
- To solve the issue of camera isolation under the SCT setting, this paper proposes a simple yet effective loss term named MCNL. Extensive experiments show that MCNL significantly boosts the ReID performance under SCT, and it is not sensitive to wrong annotations.
- Last but not least, by solving SCT, this paper sheds light

on fast deployment of ReID systems in new environments, implying a wide range of real-world applications.

2 Related Work

Our work is proposed under the new single-camera-training setting, which is relative to previous FST setting and UT setting. In this section, we mainly summarize the existing methods of these settings and then elaborate the differences between these settings and our SCT.

2.1 Fully-Supervised-Training Setting

The FST setting implies that there are a large number of annotated cross-camera pedestrian images for training. In previous works under the FST setting, most of them formulated person ReID as a classification task and trained a classification model with the labeled training data (Zheng, Yang, and Hauptmann 2016; Zheng, Zheng, and Yang 2017a; Sun et al. 2017). With the advantages of large-scale training data and deep neural networks, these methods achieve good results. In addition, some researchers designed complex network architectures to extract more robust and discriminative features (Wei et al. 2017; Zhang et al. 2017; Liu et al. 2018). Differently, other researchers argue that the surrogate loss for classification may not be suitable when the number of identities increases (Hermans, Beyer, and Leibe 2017). Therefore, end-to-end deep metric learning methods were proposed and widely used under the FST setting (Wen et al. 2016; Hermans, Beyer, and Leibe 2017; Chen et al. 2017). For example, Hermans *et al.* (Hermans, Beyer, and Leibe 2017) demonstrated that the triplet loss is more effective for person ReID task. Chen *et al.* (Chen et al. 2017) proposed a deep quadruplet network to improve the ReID performance further. Although the performance has been boosted significantly, the demand for annotating large-scale training data hinders their real-world applications, *e.g.*, the fast deployment of ReID systems in new target scenes is almost impossible. This is because it is rather expensive to collect this kind of training data for the FST setting. Different from the FST setting, the SCT setting requires much less time in the training data collection process, since there is no need to collect and annotate cross-camera pedestrian images. Therefore, our SCT setting is more suitable for fast deployment of ReID systems in new target scenes.

2.2 Unsupervised-Training Setting

Different from FST, the UT setting means there are no labeled training data. Although hand-crafted features like LOMO (Liao et al. 2015), BOW (Zheng et al. 2015) and ELF (Gray and Tao 2008) can be used directly, the ReID performance is relatively low. Therefore, some researchers designed novel unsupervised learning methods to improve ReID performance under the UT setting. Liang *et al.* (Liang et al. 2015) proposed a saliency weighted model. Lin *et al.* (Lin et al. 2019) adopted a bottom-up clustering approach for purely unsupervised ReID. Without the supervision of identity labels, the performance of their methods is still not satisfactory. To further boost ReID accuracy, many unsupervised domain adaptation methods have been

proposed. They conducted supervised learning on the source domain and transferred to the target domain, thus can benefit from FST and produce better results. The ways of transferring domain knowledge include image-image translation (Wei et al. 2018; Deng et al. 2018), attribute consistency scheme (Wang et al. 2018b), and so on (Zhong et al. 2018; Peng et al. 2016; Zhong et al. 2019b). These methods perform well when the target domain and the source domain are very similar, but may not be suitable when the domain gap is large (Li, Zhu, and Gong 2018). Differently, the problem does not exist in our proposed SCT setting because under SCT, ReID models are only trained with training data from the target scene. One-view learning (Zhong et al. 2019a) is another direction for reducing annotation labor, in which identities in only one specific camera are annotated. More recently, Li *et al.* (Li, Zhu, and Gong 2018) built the cross-camera tracklet association to learn a robust ReID model from automatically generated person tracklets. This method (Li, Zhu, and Gong 2018) assumes that cross-camera pedestrians are common, and thus camera relations can be learned by matching person tracklets. However, in a large-scale camera network, the average number of cameras pedestrians pass through is quite small, *e.g.*, one person appears in only five cameras from thousands of cameras. Moreover, the tracklet association method (Li, Zhu, and Gong 2018) is not so reliable to make sure each matched tracklets belonging to the same person. The wrong matched tracklets may cause the learned ReID model to perform poorly. Inspired by the above discussion, we propose a more reliable setting, *i.e.*, single-camera-training, and further design the multi-camera negative loss to improve the ReID performance under this setting.

3 Problem: Single-Camera-Training

Researchers report major difficulty in collecting and annotating data for person ReID, and such difficulty is positively related to the number of cameras in the network. We take MSMT17 (Wei et al. 2018), a large-scale ReID dataset, as an example. To construct it, researchers collected high-resolution videos covering 180 hours from 15 cameras, after which three labelers worked on the data for two months for cross-camera annotation. In another dataset named RPIfield (Zheng, Karanam, and Radke 2018), there are two types of pedestrians known as actors and distractors, respectively. A small number of actors followed pre-defined paths to walk in the camera network, and thus it is easy to associate the images captured by different cameras. However, a large number of distractors, without being controlled, are walking randomly, so that it is rather expensive to annotate these pedestrians among cameras. This annotation process, from a side view, verifies the difficulty of collecting and annotating cross-camera pedestrians.

On the other hand, cross-camera information plays the central role in person ReID, because for the conventional approaches, this is the main source of supervision that the same person appears differently in the camera network – this is exactly what we hope to learn. We quantify how existing datasets provide cross-camera information by computing the average number of occurrences of each person in the

Table 1: The camera-per-person (CP) value of a few ReID datasets. N_{cam} denotes the number of cameras.

Dataset	N_{cam}	CP	CP/N_{cam}
MSMT17	15	3.81	0.254
DukeMTMC-reID	8	3.13	0.391
Market-1501	6	4.34	0.724
RPIfield (distractors)	12	1.25	0.104
RPIfield (actors)	12	6.99	0.583
RPIfield (total)	12	1.40	0.117

camera network, *i.e.*, if a person appears in three cameras, his/her number of occurrences is 3. We name it the camera-per-person (CP) value and list a few examples in Tab. 1. We desire a perfect dataset in which all persons are annotated in all cameras, *i.e.*, CP equals to the number of cameras, but for a large camera network, this is often impossible, *e.g.*, in MSMT17, the CP value is 3.81, far smaller than 15, the number of cameras.

To alleviate the burden of data annotation, we propose to consider the scenario that no cross-camera annotations are available, *i.e.*, CP equals to 1 regardless of the number of cameras in the network. We name this setting to be **single-camera-training** (SCT). This requirement can be achieved by collecting data from different cameras in different time periods (*e.g.*, recording camera A from 8 am to 9 am while camera B from 10 am to 11 am) – although this cannot guarantee our assumption, as we shall see in experiments, our approach is robust to a small fraction of ‘outliers’, *i.e.*, two or more occurrences of the same person in different cameras are assumed to be different identities.

This setting greatly eases the fast deployment of a ReID system. With off-the-shelf person tracking algorithms (Keuper et al. 2018; Luo et al. 2019), we can easily extract a large number of *tracklets* in videos, each of which forms an identity in the training set. However, such a training dataset is less powerful than those specifically designed for the ReID task, as it lacks supervision of how a person can appear differently in different cameras. We call this challenge **camera isolation**, and will elaborate on this point carefully in the next section.

4 Our Approach

4.1 Baseline and Motivation

Existing ReID approaches often start with a backbone which extracts a feature vector \mathbf{x}_k from an input image \mathbf{I}_k . On top of these features, there are mainly two types of loss functions, and sometimes they are used together towards higher accuracy. The first type is named the *cross entropy* (CE) loss, which requires the model to perform a classification task in which the same person in different cameras are categorized into one class. The second type is named the *triplet margin* (TM) loss, which assumes that the largest distance between two appearances of the same person should be smaller than the smallest distance between this person and another. When built upon a ResNet-50 (He et al. 2016) backbone, CE and TM achieve 78.9% and 79.0% rank-1 accuracy on the

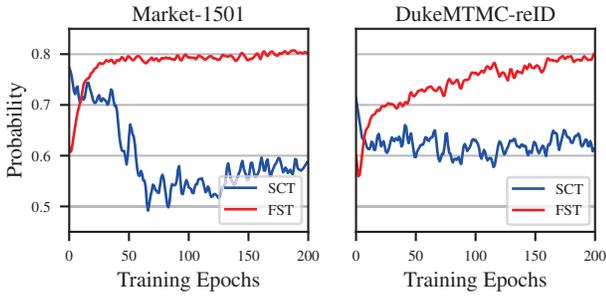


Figure 2: Curves of the probability produced by the triplet margin loss, with respect to the number of elapsed training epochs, of finding the most similar person (of a different ID) in another camera. The figures on the left and right show results on Market-1501 and DukeMTMC-reID, respectively.

DukeMTMC-reID dataset (Zheng, Zheng, and Yang 2017b), respectively, without any bells and whistles.

However, in the SCT setting, both of them fail dramatically due to camera isolation. From the DukeMTMC-reID dataset, we sample 5,993 images from the training set which satisfy the SCT setting, and the corresponding models, with CE and TM losses, report 40.2% and 21.2%, respectively. In comparison, we sample another training subset with the same number of images but equipped with cross-camera annotation, and these numbers become 69.3% and 75.8%, which verifies our hypothesis.

To explain this dramatic accuracy drop, we first point out that a ReID system needs to learn feature embedding which is independent to cameras (*i.e.*, camera-unrelated features), which is to say, the learned feature distribution is approximately the same under different cameras. However, we point out that both CE and TM losses cannot achieve this goal by themselves – they heavily rely on cross-camera annotations. Without such annotations, existing ReID systems often learn camera-related features.

Here, we provide another metric to quantify the impact of camera-related/unrelated features, which is the core observation that motivates our algorithm design. Intuitively, for a set of camera-unrelated features, the feature distribution over the entire camera network should be approximately the same as the distribution over any single camera. In other words, the expectation of similarity between two different persons in the same camera should **not** be higher than that between two different persons from two cameras. Therefore, given an anchor image, the probability that the most similar person appears in the same camera is only $1/N_{\text{cam}}$, *i.e.*, in a multi-camera system, the most similar person mostly appears in another camera. Thus, we perform statistics during the training process with the TM loss, under both SCT and FST, and show results in Fig. 2. We can see that, under the FST setting, this probability is mostly increasing during the training process, and eventually reaches a plateau at around 0.8; while under the SCT setting, the curve is less stable and the stabilized probability is much lower.

Thus, our motivation is to facilitate the learned features to satisfy that the most similar person appears in another cam-

era. **This is considered as the extra, weakly-supervised cue to be explored in the SCT setting.** This leads to a novel loss function, the Multi-Camera Negative Loss (MCNL), which is detailed in the next subsection.

4.2 Multi-Camera Negative Loss

Inspired by the analyses above, we design the Multi-Camera Negative Loss (MCNL) to ensure that, given any anchor image in one camera, the most similar negative image is more likely to be found from other cameras, and the negative image should be less similar to the anchor image, compared to the most dissimilar positive image.

In a mini-batch with C cameras, P identities from each camera and K images of each identity (*i.e.*, the batch size is $C \times P \times K$), given an anchor image $\mathbf{I}_k^{c,p}$, let $f_\theta(\mathbf{I}_k^{c,p})$ denote the feature mapping function learned by our network, and $\|f_1 - f_2\|$ represent the Euclidean distance between two feature vectors. The hardest positive distance of $\mathbf{I}_k^{c,p}$ is defined as:

$$\text{dist}_+^{c,p,k} = \max_{l=1 \dots K, l \neq k} \|f_\theta(\mathbf{I}_k^{c,p}) - f_\theta(\mathbf{I}_l^{c,p})\|. \quad (1)$$

Then, we have the hardest negative distance in the same camera:

$$\text{dist}_{-, \text{same}}^{c,p,k} = \min_{\substack{l=1 \dots K, \\ q=1 \dots P, q \neq p}} \|f_\theta(\mathbf{I}_k^{c,p}) - f_\theta(\mathbf{I}_l^{c,q})\|, \quad (2)$$

and the hardest negative distance in other cameras:

$$\text{dist}_{-, \text{other}}^{c,p,k} = \min_{\substack{l=1 \dots K, \\ q=1 \dots P, \\ o=1 \dots C, o \neq c}} \|(f_\theta(\mathbf{I}_k^{c,p}) - f_\theta(\mathbf{I}_l^{o,q}))\|. \quad (3)$$

With these terms, MCNL is formulated as follows:

$$\begin{aligned} \mathcal{L}_{\text{MCNL}} = & \sum_{c=1}^C \sum_{p=1}^P \sum_{k=1}^K [m_1 + \text{dist}_+^{c,p,k} - \text{dist}_{-, \text{other}}^{c,p,k}]_+ \\ & + [m_2 + \text{dist}_{-, \text{other}}^{c,p,k} - \text{dist}_{-, \text{same}}^{c,p,k}]_+, \end{aligned} \quad (4)$$

where $[z]_+ = \max(z, 0)$, and both of m_1 and m_2 denote the values of margins.

As shown in Eq. (4), the second loss term is to ensure the most similar negative image is found from other cameras, and the first loss term is to force this negative image to be less similar than the most dissimilar positive image. These two parts together provide boundaries to restrict $\text{dist}_{-, \text{other}}^{c,p,k}$ between $\text{dist}_+^{c,p,k}$ and $\text{dist}_{-, \text{same}}^{c,p,k}$, which meets the motivation described in the previous section.

Moreover, the proposed MCNL also ensures that the learned feature is discriminative and camera-unrelated. Given the most similar cross-camera negative image differs in camera factors, it is more likely that the similarity lies in person appearance. By pulling the most similar cross-camera negative pairs a little closer, MCNL encourages the model to focus more on person appearance. For the most similar within-camera negative pair, as camera factors are shared with the anchors, pushing them away further reduces the impact of cameras. In addition, MCNL also ensures positive pairs closer than cross-camera negative pairs, which meets the basic correctness of metric learning.

Table 2: Details of datasets used in our experiments.

Dataset	#Train IDs	#Train Images	#Test IDs	#Test Images	With cross-camera persons?
Market	751	12,936	750	15,913	True
Market-SCT	751	3,561	750	15,913	False
Duke	702	16,522	1,110	17,661	True
Duke-SCT	702	5,993	1,110	17,661	False

Differences from prior work. Previously, researchers proposed many triplet-based or quadruplet-based loss functions to improve ReID performance (Hermans, Beyer, and Leibe 2017; Schroff, Kalenichenko, and Philbin 2015; Shi et al. 2016). The largest difference between our approach and theirs lies in that they pushed away the hardest negative images from other cameras without constraints, while we do not. In a dataset constructed under the SCT setting, these methods tend to learn camera-related cues to separate negative images from another camera, which further aggravates the camera isolation problem. Moreover, we evaluate several state-of-the-art methods related to metric learning and ReID under SCT. The experiment results demonstrate that existing methods are not suitable for this new setting.

Advantages. Based on the above discussions, the advantages of our proposed MCNL can be summarized as two folds. (i) MCNL can alleviate the camera isolation problem. Through pulling the cross-camera negative pairs closer and pushing the within-camera negative pairs away, MCNL forces the feature extraction model to ignore the camera clues. (ii) Same with previous metric learning approaches, MCNL can force the feature extraction model to learn a more discriminative representation by adding the constraint that, the hardest positive image should be closer to the anchor image, compared with the negative images (both cross-camera and within-camera negative images).

5 Experiments

5.1 Datasets

To evaluate the effectiveness of our proposed method, we mainly conduct experiments on two large-scale person ReID datasets, *i.e.*, Market-1501 (Zheng et al. 2015) and DukeMTMC-reID (Zheng, Zheng, and Yang 2017b). For short, we refer to Market-1501 and DukeMTMC-reID as Market and Duke, respectively.

Both Market and Duke are widely used person ReID datasets. For each person in the training sets, there are multiple images from different cameras. To better evaluate our method, we reconstruct these training sets for the SCT setting. More specifically, we randomly choose one camera for each person and take those images of the person under the selected camera as training images. Finally, we sample 5,993 images from the training set of Duke and 3,561 images from the training set of Market. In this paper, we denote these sampled datasets as Duke-SCT and Market-SCT, respectively. Note that, we still keep the original testing data and strictly follow the standard testing protocols. The detailed statistics of the datasets are shown in Tab. 2.

Table 3: ReID accuracy (%) produced by different loss terms, among which MCNL reports the best results.

Methods	Duke-SCT		Market-SCT	
	Rank-1	mAP	Rank-1	mAP
Triplet	21.2	11.3	39.7	18.2
Triplet-other	9.9	3.6	25.2	8.8
Triplet-same	54.6	35.9	51.3	28.0
MCNL	66.4	45.3	66.2	40.6

5.2 Implementation Details

We adopt ResNet-50 (He et al. 2016) which is pre-trained on ImageNet (Deng et al. 2009) as our network backbone. The final fully connected layers are removed, and we conduct global averaging pooling (GAP) to the output of the fourth block of ResNet-50. The GAP feature is used for metric learning. In each batch, we randomly select 8 cameras, and sample 4 identities for each selected camera. Then, we randomly sample 8 images for each identity, leading to the batch size of 256 for Duke-SCT. For Market-SCT, there are only 6 cameras in the training set. Hence, we sample 6 cameras, 5 identities for each camera, and 8 images for each identity, thus the batch size is 240 for Market-SCT. We empirically set m_1 and m_2 as 0.1, respectively. For baseline, we implement the batch hard triplet loss (Hermans, Beyer, and Leibe 2017), which is one of the most effective implementations of the TM loss. For short, we use Triplet to denote the batch hard triplet loss in the following sections. The margin of Triplet is set to be 0.3, as it achieves excellent performance under the FST setting. The input images are resized as 256×128 , and Adam (Kingma and Ba 2014) optimizer is adopted. Weight decay is set as 5×10^{-4} . The learning rate ϵ is initialized as 2×10^{-4} and exponentially decays following the Eq. (5) proposed in (Hermans, Beyer, and Leibe 2017):

$$\epsilon(t) = \begin{cases} \epsilon_0, & t \leq t_0 \\ \epsilon_0 \times 0.001^{\frac{t-t_0}{t_1-t_0}}, & t_0 \leq t \leq t_1. \end{cases} \quad (5)$$

For all datasets, we update the learning rate every epoch after 100 epochs and stop training when reaching 200 epochs, *i.e.*, $t_0 = 100$ and $t_1 = 200$, respectively. All experiments are conducted on two NVIDIA GTX 1080Ti GPUs.

5.3 Diagnostic Studies

The effectiveness of MCNL. MCNL is designed based on Triplet (Hermans, Beyer, and Leibe 2017). To better demonstrate the effectiveness of MCNL, we evaluate the performance of Triplet and its two variations, Triplet-same and Triplet-other. Triplet-same represents the hardest negative image is selected from the same camera as the anchor image, and Triplet-other means the hardest negative image is found from other cameras. The performance of these methods is summarized in Tab. 3.

As shown in Tab. 3, MCNL achieves huge improvements compared to Triplet and its two variations. For example, MCNL outperforms Triplet with 45.2% performance gains in Rank-1 accuracy on Duke and boosts the ReID performance with 11.8% compared to Triplet-same. It is worth

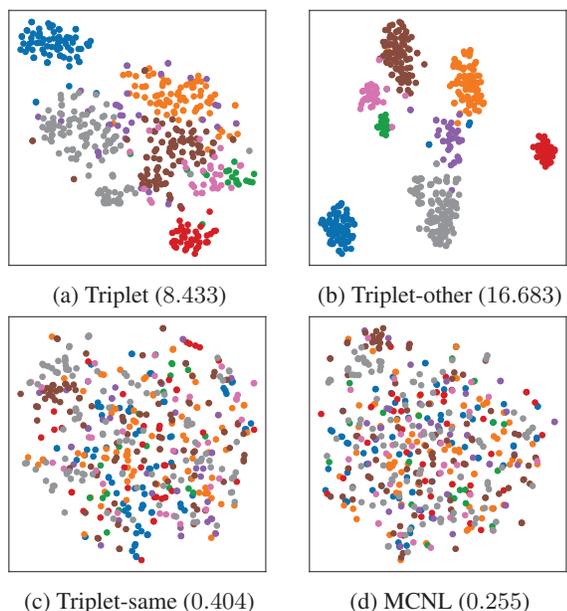


Figure 3: Visualization of feature distributions. Pseudo F statistics are shown in parentheses. Each color indicates features from a camera. This figure is best viewed in color.

noticing that, compared with Triplet, Triplet-same also improves the ReID performance under the SCT setting. That is because Triplet-same aims to maximize the distance of the negative pair, of which the two images come from the same camera. To achieve the above goal, Triplet-same forces the feature extraction model to focus on foreground area and extract more camera-unrelated features because camera-related clues are very similar. It is obvious that Triplet-other aims to push cross-camera negative pairs away. Therefore, the model will focus more on background and get worse performance. Similar to Triplet-same, MCNL also aims to push within-camera negative pairs as far as possible. Moreover, by restricting the distance of the hardest cross-camera negative pair smaller than the distance of the hardest within-camera negative pair, the model further solves the camera isolation problem and ignores camera-related features.

To better evaluate the above discussions, we utilize t-SNE (Van Der Maaten 2014) to visualize the feature distributions extracted by different methods. To achieve this, we randomly sample 500 images from the testing set of Duke, and then extract the features on these images through four models trained with Triplet, Triplet-same, Triplet-other, and MCNL, respectively. Moreover, we use the pseudo F statistics (Caliński and Harabasz 1974) to evaluate the relations of feature distributions of different cameras quantitatively. A larger value of pseudo F indicates more distinct clusters, which means the extracted features are more related to cameras. In other words, a smaller value of pseudo F implies that features are better learned.

As shown in Fig. 3, features extracted by Triplet and Triplet-other are separable according to cameras, which is bad for ReID systems. Differently, Triplet-same and MCNL

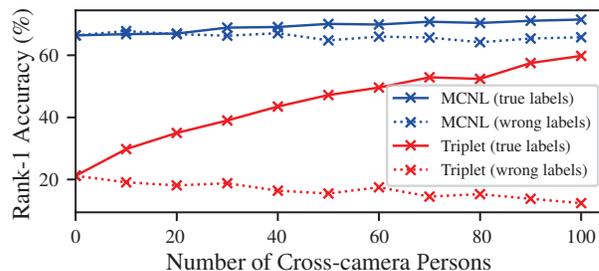


Figure 4: Rank-1 accuracy (%) on Duke-SCT with randomly selected cross-camera persons. MCNL shows great robustness against outliers. Solid or dashed line: whether the model receives accurate annotations.

both map images to a camera-unrelated feature space.

Stability analysis. Although the data collection process is restricted, there are inevitably some persons appearing in not only one camera. To evaluate the robustness of our MCNL under this setting, we conduct experiments on Duke to show how the accuracy changes with respect to the percentage of people showing up in multiple cameras. As shown in Fig. 4, when these people are annotated truly according to their identities, Triplet loss benefits largely from ground-truth cross-camera annotations. Nevertheless, with a considerable portion (14%, 100 out of 702) of outliers, MCNL still holds an advantage. On the other hand, when they are annotated under the SCT setting, *i.e.*, the images of the same person but different cameras are assigned with different labels, MCNL is quite robust with a small accuracy drop.

This result further demonstrates that our proposed MCNL improves the ReID accuracy under the SCT setting with great robustness against outliers. In real-world applications, it is easy to control the portion of outliers under a low ratio.

5.4 Comparison to Previous Work

Comparisons to FST methods. We evaluate a few popular FST methods under the SCT setting and compare our method with other advanced metric learning algorithms. As shown in Tab. 4, previous state-of-the-art methods for the FST setting fail dramatically under the SCT setting while MCNL shows great advantages. This is because, without cross-camera annotations, these methods are unable to extract camera-unrelated features.

Comparisons to UT methods. Our motivation of the SCT setting is for fast deployment of ReID systems on new target scenes, which is the same as the motivation of the UT setting. Thus, as shown in Tab. 5, we also compare MCNL with previous unsupervised training methods, including purely unsupervised methods, tracklet association learning method, and domain adaptation methods.

Compared to the state-of-the-art purely unsupervised methods (labels denoted as *None*), our proposed MCNL significantly outperforms BUC (Lin et al. 2019) with 19.0% performance gains in Rank-1 accuracy on Duke. In one-view learning (Zhong et al. 2019a), labeled data are available in only one camera. Compared to one-view learning, SCT makes full use of camera information towards better

Table 4: Comparisons of ReID accuracy (%) when training with SCT datasets. MCNL reports the best performance on SCT datasets while other methods undergo dramatic accuracy drop.

Methods	Ref.	Duke-SCT		Market-SCT	
		Rank-1	mAP	Rank-1	mAP
Center Loss (Wen et al. 2016)	ECCV'16	38.7	23.2	40.3	18.5
A-Softmax (Liu et al. 2017)	CVPR'17	34.8	22.9	41.9	23.2
ArcFace (Deng et al. 2019)	CVPR'19	35.8	22.8	39.4	19.8
PCB (Sun et al. 2018)	ECCV'18	32.7	22.2	43.5	23.5
Suh's method (Suh et al. 2018)	ECCV'18	38.5	25.4	48.0	27.3
MGN (Wang et al. 2018a)	ACMMM'18	27.1	18.7	38.1	24.7
MCNL	This paper	66.4	45.3	66.2	40.6

Table 5: ReID accuracy (%) comparisons to UT methods. *None* denotes purely unsupervised training without any labels. *One-view* denotes identities in only one camera are labeled. *Tracklet* denotes using tracklet labels. *Transfer* denotes utilizing other labeled source datasets and unlabeled target datasets.

Methods	Ref.	Labels	Duke		Market	
			Rank-1	mAP	Rank-1	mAP
BOW (Zheng et al. 2015)	ICCV'15	<i>None</i>	17.1	8.3	35.8	14.8
DECAMEL (Yu, Wu, and Zheng 2018)	TPAMI'18	<i>None</i>	-	-	60.2	32.4
BUC (Lin et al. 2019)	AAAI'19	<i>None</i>	47.4	27.5	66.2	38.3
CamStyle (Zhong et al. 2019a)	TIP'19	<i>One-view</i>	48.7	25.7	57.6	29.6
TAUDL (Li, Zhu, and Gong 2018)	ECCV'18	<i>Tracklet</i>	61.7	43.5	63.7	41.2
TJ-AIDL (Wang et al. 2018b)	CVPR'18	<i>Transfer</i>	44.3	23.0	58.2	26.5
SPGAN (Deng et al. 2018)	CVPR'18	<i>Transfer</i>	46.9	26.4	58.1	26.9
HHL (Zhong et al. 2018)	ECCV'18	<i>Transfer</i>	46.9	27.2	62.2	31.4
MAR (Yu et al. 2019)	CVPR'19	<i>Transfer</i>	67.1	48.0	67.7	40.0
ECN (Zhong et al. 2019b)	CVPR'19	<i>Transfer</i>	63.3	40.4	75.1	43.0
MCNL	This paper	SCT	66.4	45.3	66.2	40.6
MCNL+MAR (Yu et al. 2019)	This paper	<i>Transfer</i> +SCT	71.4	53.3	72.3	48.0
MCNL+ECN (Zhong et al. 2019b)	This paper	<i>Transfer</i> +SCT	67.3	45.5	76.3	51.2

performance, e.g., a 17.7% advantage in Rank-1 accuracy on Duke. As for TAUDL (Li, Zhu, and Gong 2018) that uses *Tracklet* labels, the entire training sets are used to train the models. Our method constructs SCT datasets for training, and thus only a small portion of training data are used, but still surpasses TAUDL in Rank-1 accuracy.

Recently, many domain adaptation methods that use other labeled datasets for extra supervision obtain good ReID accuracy. Our MCNL alone achieves competitive results compared to them. Moreover, our method is also complementary to current domain adaptation methods and can be easily combined by replacing the target datasets with SCT datasets. Such combination instantly brings significant improvement. We take MAR (Yu et al. 2019) and ECN (Zhong et al. 2019b), for example. After using SCT data and MCNL in MAR, we boost 4.3% in Rank-1 accuracy and 5.3% in mAP on Duke dataset; the combination of MCNL and ECN improves mAP on Market by 8.2% compared to ECN only. Taking the advantages of reliable target domain annotations and extra transferred information, we achieve the best ReID performance on Duke and Market, respectively.

Note that, our method achieves good performance with much fewer training data. Because of giving up collecting cross-camera pedestrian images under the SCT setting, the

above training data can be easily collected and annotated. Therefore, compared with prior work, our method and proposed SCT setting are more suitable for fast deployment of ReID systems with good performance on new target scenes.

6 Conclusions

In this paper, we explore a new setting named single-camera-training (SCT) for person ReID. With the advantage of low costs in data collection and annotation, SCT lays the foundation of fast deployment of ReID systems in new environments. To work under SCT, we propose a novel loss term named multi-camera negative loss (MCNL). Experiments demonstrate that under SCT, the proposed approach boosts ReID performance of existing approaches by a large margin.

Our approach reveals the possibility of learning cross-camera correspondence without cross-camera annotations. In the future, we will explore more cues to leverage under the SCT setting and consider the mixture of single-camera and cross-camera annotations to improve ReID accuracy.

Acknowledgement. This work was partially supported by the National Natural Science Foundation of China (No. 61772054), the NSFC Key Project (No. 61632001) and the Fundamental Research Funds for the Central Universities.

References

- Caliński, T., and Harabasz, J. 1974. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*.
- Chen, W.; Chen, X.; Zhang, J.; and Huang, K. 2017. Beyond triplet loss: a deep quadruplet network for person re-identification. In *CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Deng, W.; Zheng, L.; Kang, G.; Yang, Y.; Ye, Q.; and Jiao, J. 2018. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person reidentification. In *CVPR*.
- Deng, J.; Guo, J.; Niannan, X.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*.
- Gray, D., and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Keuper, M.; Tang, S.; Andres, B.; Brox, T.; and Schiele, B. 2018. Motion segmentation & multiple object tracking by correlation co-clustering. *IEEE transactions on pattern analysis and machine intelligence*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, M.; Zhu, X.; and Gong, S. 2018. Unsupervised person re-identification by deep learning tracklet association. In *ECCV*.
- Liang, C.; Huang, B.; Hu, R.; Zhang, C.; Jing, X.; and Xiao, J. 2015. A unsupervised person re-identification method using model based representation and ranking. In *ACM MM*.
- Liao, S.; Hu, Y.; Zhu, X.; and Li, S. Z. 2015. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*.
- Lin, Y.; Dong, X.; Zheng, L.; Yan, Y.; and Yang, Y. 2019. A bottom-up clustering approach to unsupervised person re-identification. In *AAAI*.
- Liu, W.; Wen, Y.; Yu, Z.; Li, M.; Raj, B.; and Song, L. 2017. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*.
- Liu, J.; Ni, B.; Yan, Y.; Zhou, P.; Cheng, S.; and Hu, J. 2018. Pose transferrable person re-identification. In *CVPR*.
- Luo, W.; Stenger, B.; Zhao, X.; and Kim, T.-K. 2019. Trajectories as topics: Multi-object tracking by topic discovery. *IEEE Transactions on Image Processing* 28(1):240–252.
- Peng, P.; Xiang, T.; Wang, Y.; Pontil, M.; Gong, S.; Huang, T.; and Tian, Y. 2016. Unsupervised cross-dataset transfer learning for person re-identification. In *CVPR*.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *CVPR*.
- Shi, H.; Yang, Y.; Zhu, X.; Liao, S.; Lei, Z.; Zheng, W.; and Li, S. Z. 2016. Embedding deep metric for person re-identification: A study against large variations. In *ECCV*.
- Suh, Y.; Wang, J.; Tang, S.; Mei, T.; and Lee, K. M. 2018. Part-aligned bilinear representations for person re-identification. In *ECCV*.
- Sun, Y.; Zheng, L.; Deng, W.; and Wang, S. 2017. Svdnet for pedestrian retrieval. In *ICCV*.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*.
- Van Der Maaten, L. 2014. Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research* 15(1):3221–3245.
- Wang, G.; Yuan, Y.; Chen, X.; Li, J.; and Zhou, X. 2018a. Learning discriminative features with multiple granularities for person re-identification. In *ACM MM*.
- Wang, J.; Zhu, X.; Gong, S.; and Li, W. 2018b. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*.
- Wang, G.; Lai, J.; Huang, P.; and Xie, X. 2019. Spatial-temporal person re-identification. In *AAAI*.
- Wei, L.; Zhang, S.; Yao, H.; Gao, W.; and Tian, Q. 2017. Glad: global-local-alignment descriptor for pedestrian retrieval. In *ACM MM*.
- Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*.
- Wen, Y.; Zhang, K.; Li, Z.; and Qiao, Y. 2016. A discriminative feature learning approach for deep face recognition. In *ECCV*.
- Yu, H.-X.; Zheng, W.-S.; Wu, A.; Guo, X.; Gong, S.; and Lai, J.-H. 2019. Unsupervised person re-identification by soft multilabel learning. In *CVPR*.
- Yu, H.; Wu, A.; and Zheng, W. 2018. Unsupervised person re-identification by deep asymmetric metric embedding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, X.; Luo, H.; Fan, X.; Xiang, W.; Sun, Y.; Xiao, Q.; Jiang, W.; Zhang, C.; and Sun, J. 2017. Alignedreid: Surpassing human-level performance in person re-identification. *arXiv preprint arXiv:1711.08184*.
- Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. In *ICCV*.
- Zheng, M.; Karanam, S.; and Radke, R. J. 2018. Rpfifield: A new dataset for temporally evaluating person re-identification. In *CVPR Workshops*.
- Zheng, L.; Yang, Y.; and Hauptmann, A. G. 2016. Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984*.
- Zheng, Z.; Zheng, L.; and Yang, Y. 2017a. A discriminatively learned cnn embedding for person reidentification. *ACM Transactions on Multimedia Computing, Communications, and Applications* 14(1):13.
- Zheng, Z.; Zheng, L.; and Yang, Y. 2017b. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*.
- Zheng, Z.; Zheng, L.; and Yang, Y. 2018. Pedestrian alignment network for large-scale person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Zhong, Z.; Zheng, L.; Li, S.; and Yang, Y. 2018. Generalizing a person retrieval model hetero-and homogeneously. In *ECCV*.
- Zhong, Z.; Zheng, L.; Zheng, Z.; Li, S.; and Yang, Y. 2019a. Camstyle: A novel data augmentation method for person re-identification. *IEEE Transactions on Image Processing* 28(3):1176–1190.
- Zhong, Z.; Zheng, L.; Luo, Z.; Li, S.; and Yang, Y. 2019b. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *CVPR*.