HAL
open science

Unsupervised domain adaptation for person re-identification through source-guided pseudo-labeling

Fabian Dubourvieux, Romaric Audigier, Angelique Loesch, Samia Ainouz, Stephane Canu

To cite this version:

Fabian Dubourvieux, Romaric Audigier, Angelique Loesch, Samia Ainouz, Stephane Canu. Unsupervised domain adaptation for person re-identification through source-guided pseudo-labeling. ICPR 2020 - 25th International Conference on Pattern Recognition, Jan 2021, Milano (Virtual conference), Italy. 10.1109/ICPR48806.2021.9412964 . cea-03251813

HAL Id: cea-03251813

https://cea.hal.science/cea-03251813v1

Submitted on 7 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unsupervised Domain Adaptation for Person Re-Identification through Source-Guided Pseudo-Labeling

Fabian Dubourvieux[⋆,†], Romaric Audigier[⋆], Angelique Loesch[⋆], Samia Ainouz[†], Stephane Canu[†]
[⋆] *CEA, LIST, Vision and Learning Lab for Scene Analysis*
PC 184, F-91191 Gif-sur-Yvette, France
{firstname.lastname}@cea.fr
[†] *Normandie Univ, INSA Rouen, LITIS*
Av. de l'Université le Madrillet 76801 Saint Etienne du Rouvray, France
{firstname.lastname}@insa-rouen.fr

***Abstract*—Person Re-Identification (re-ID) aims at retrieving images of the same person taken by different cameras. A challenge for re-ID is the performance preservation when a model is used on data of interest (target data) which belong to a different domain from the training data domain (source data). Unsupervised Domain Adaptation (UDA) is an interesting research direction for this challenge as it avoids a costly annotation of the target data. Pseudo-labeling methods achieve the best results in UDA-based re-ID. They incrementally learn with identity pseudo-labels which are initialized by clustering features in the source re-ID encoder space. Surprisingly, labeled source data are discarded after this initialization step. However, we believe that pseudo-labeling could further leverage the labeled source data in order to improve the post-initialization training steps. In order to improve robustness against erroneous pseudo-labels, we advocate the exploitation of both labeled source data and pseudo-labeled target data during all training iterations. To support our guideline, we introduce a framework which relies on a two-branch architecture optimizing classification and triplet loss based metric learning in source and target domains, respectively, in order to allow *adaptability to the target domain* while ensuring *robustness to noisy pseudo-labels*. Indeed, shared low and mid-level parameters benefit from the source classification and triplet loss signal while high-level parameters of the target branch learn domain-specific features. Our method is simple enough to be easily combined with existing pseudo-labeling UDA approaches. We show experimentally that it is efficient and improves performance when the base method has no mechanism to deal with pseudo-label noise. And it maintains performance when combined with base method that already manages pseudo-label noise. Our approach reaches state-of-the-art performance when evaluated on commonly used datasets, Market-1501 and DukeMTMC-reID, and outperforms the state of the art when targeting the bigger and more challenging dataset MSMT.**

I. Introduction

Person re-identification (re-ID) aims at retrieving images of a person of interest captured by different cameras. It is known as an *open-set problem* because identities (i.e., classes) seen at test time are different from those at training time. Effective representations must be learned in order to discriminate people (i.e., classes) never seen during training. Convolutional Neural Networks (CNNs) have achieved excellent performance on various re-ID datasets, due to their capacity to learn camera invariant and identity discriminative representations robust to appearance changes [1].

In practice, one may want to deploy a re-ID system on a different set of cameras than those used for training data, i.e. on a test set from a different distribution (brightness, colorimetry and angle of the cameras, background...). In this context of *cross-dataset* testing, we observe in practice a sharp drop in re-ID performance due to the *domain gap* [2]. A solution may be the annotation of larger scale re-ID datasets or the target dataset itself so that a new better performing model can be trained on the target domain. To avoid these costly manual annotations, the research community focused on designing Unsupervised Domain Adaptation (UDA) algorithms. UDA aims at learning an efficient model on the target domain. It only requires labeled training samples from one or multiple domains (source domains) and unlabeled samples from the domain of interest (target domain) on which the model is tested. While extensive research is conducted for classification UDA, the open-set nature of the re-ID task makes it impossible or inefficient to directly apply closed-set approaches designed for classification [3]. Therefore, a part of computer vision researchers focuses on the peculiarities of re-ID being an open-set retrieval problem.

In this setting, pseudo-labeling approaches have proven to be the best UDA methods to learn ID-discriminative features for the target domain [4]. These ID pseudo-labels are initially predicted by clustering the target data embedded in an ID-discriminative feature space learned on source data. To do so, pseudo-labeling UDA algorithms alternate between training phases and uour baselinetes of pseudo-labels with the lastly trained model to refine them to a certain extent. Through this iterative process, the source data is not reused beyond the initialization step of the model. It can be assumed that this choice not to integrate the labeled source data aims at not biasing the learning of ID-discriminative features on target data. One could think that using source data to constrain the training would decrease adaptability to target domain and

would degrade the re-ID performance on target domain. We believe that it is possible to subtly exploit the source data and their ground-truth labels to improve the representation on the target domain, while reducing the potential undesirable effects of the domain gap between the source and target data.

Our contribution can be summarized as follows.

- We propose a *source-guided pseudo-labeling* framework to solve *cross-dataset re-ID*. It leverages the labeled source data during pseudo-labeling scheme of all training phases.
- The key element of this framework is a two-branch architecture that simultaneously optimizes classification and triplet-based metric learning in both source and target domains, in order to allow *adaptability to the target domain* (high-level parameters of target branch optimized with pseudo-labels) while ensuring *robustness to noisy pseudo-labels* (shared low and mid-level parameters constrained by supervised learning with source labels).
- The proposed framework is simple enough to be easily combined with existing pseudo-labeling UDA approaches. Experiments show that it is particularly effective when the base method has no special mechanism to deal with pseudo-label noise. It also improves the stability of re-ID performance with relation to the tricky choice of clustering parameters.
- Our framework combined with MMT method outperforms the state-of-the-art on the challenging cross-domain scenario of MSMT dataset.

In Section II, we review the related work. Then, we detail our proposed framework in Section III. The experiment settings are given in Section IV and results are presented in Section V.

II. Related Work

The state of the art contributions for UDA re-ID can be divided into 3 families which leverage the source data differently.

A. Image-to-Image translation

Image-to-Image translation methods [5] [2] [6] [7] are based on learning how to transform images from one domain to another, preserving class information. The majority of these approaches use a CycleGAN [8] model to transfer the style of the target images to the source images, while constraining the person's appearance preservation after transfer. The target-style source images are then used to learn in a supervised way a re-ID model for the target domain. Although these methods fully exploit the source images (in the target style) and their labels, they depend on the quality of image generation and the preservation of identity information after transfer. In practice, the performance obtained by these methods shows that the source images in the target style are not sufficiently representative of the target domain.

B. Domain invariant feature learning

Domain invariant feature learning methods look for a domain invariant discriminative feature space [9] [10] [11] [12]. In addition to the supervised ID discriminative loss on source samples, these approaches seek to align the source and target domain feature distributions by penalizing an unsupervised domain discrepancy loss term [9] [10] or learn domain invariant space by domain feature disentanglement [11] [12] sometimes with auxiliary information for supervision (semantic attribute labels, pose labels...) [13]. As with methods that operate at the pixel level, domain invariant feature learning makes full use of source data and their labels. While they can outperform Image-to-image translation approaches [9] [9], these methods cannot learn efficient target specific features because of domain invariance constraints and the presence of discriminative information (labels) only for the source data.

C. Pseudo-labeling

Pseudo-labeling methods exploit a source-trained model to initialize pseudo-identity labels for target data by clustering their feature representation by this model [14] [15]. Most of pseudo-labeling UDA methods are based on an iterative paradigm which alternates between optimizing and pseudo-label refinery by uour baselinetes with the lastly optimized model on target images [16] [17] [18] [19]. However, the source biased representation and the clustering introduce some errors that can persist through iterations, be over-fitted by the model and thus degrade the pseudo-label refinery process. To avoid over-fitting the noise of these generated labels, recent work focused on improving Pseudo-labeling methods by improving their robustness to noisy labels using asymmetric co-teaching [20] or mutual mean teaching with soft target labels [4]. Even though recent Pseudo-labeling methods achieved near supervised training performance on some re-ID adaptation tasks, they still suffer from sensitivity to clustering or additional hyperparameters which can't be easily estimated in the UDA setting where no labeled target validation set is available. [4] Moreover, we believe that they under-exploit the labeled source dataset since it is discarded after the initialization phase. While UDAP method [16] uses an additional weight ratio term between source and target to measure feature similarity, it is only used during the pseudo-label predictions by clustering. Moreover, their ablation studies show negligible performance gain and the source samples are still discarded from the optimization iterations.

Our contribution focuses on these pseudo-label iterative methods. In contrast to the existing pseudo-labeling approaches mentioned above, we propose to better leverage the labeled source samples, beyond the initialization phase, in order to further improve their target discriminative features. Contrary to UDAP [16], we leverage the source samples as well as their labels directly during the optimization phases of the iterative training.

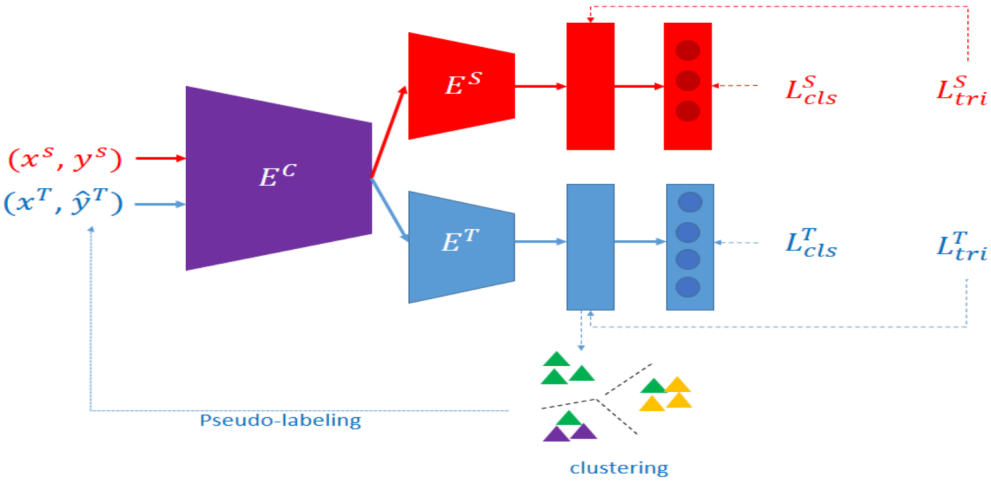

Fig. 1. Unlike existing pseudo-labeling strategies (1), our framework is a source-guided (3) as it leverages the labeled source dataset during pseudo-label training. It is composed of a shared encoder E^C and two domain-specific branches E^S and E^T. The numbers of shared layers and specific layers can vary. Configuration (2) represents the extreme case when all layers are shared and there is no domain-specific branch.

III. Proposed Framework

In this section, we describe our source-guided pseudo label framework. As other pseudo-labeling UDA frameworks, it is composed of two distinct parts:
- the supervised training of the initialization model on source: we follow the supervised training guideline from used [1] by most pseudo-labeling UDA re-ID [4] [20]
- our source guided pseudo-label iterative training to efficiently leverage the source samples

A. Supervised training on source domain for initialization

Let $D^S = \{(x^S_k, y^S_k)_{1 \leq k \leq N^S}\}$ be the labeled source training set of N^S samples from M_S different people, where x_k represents the k-th person image and y_k its one-hot identity label. Similarly let $D^T = \{(x^T_k)_{1 \leq k \leq N^T}\}$ the unlabaled target training set of N^T samples. E denotes a learnable feature encoder.

1) Triplet Loss: To learn re-ID discriminative features, we use the batch hard Tiplet Loss [21] to pull the same ID sample features closer to each other than other ID ones. In a batch of N samples $\{(x_k, y_k)_{1 \leq k \leq N^S}\}$, it is given by:

$$L_{tri} = \sum_{i=1}^{N} \max(0, \left\|E(x_i) - E(x_{p(i)})\right\|_2 - \left\|E(x_i) - E(x_{n(i)})\right\|_2 + m)$$

where m is a fixed margin, $x_{p(i)}$ and $x_{n(i)}$ are resp. the hardest positive (the farthest batch feature sample to x_i for the L2 distance) to and negative (the closest) samples in the batch for x_i.

2) Cross-Entropy Loss: To further improve the ID discriminativeness of features, we use the classification of ID labels with the Cross-Entropy Loss. For this, we consider a classification layer given by the parameters $W = [W_1, ..., W_M]$ where M is the number of ID in the training set. We use the Softmax classification loss given by:

$$L_{cls} = -\sum_{i=1}^{N} y_i \log(p_i) \tag{1}$$

where $p_i = \frac{exp(W_{y_i} E(x_i))}{\sum_{k=1}^{M} exp(W_k E(x_i))}$.

3) Initialization phase loss: By considering a source domain classification layer $W^S = [W^S_1, ..., W^S_M]$, we can define the softmax cross-entropy loss and the triplet for the source labeled training dataset as described in Eq. III-A1 and 1. Then, we train E by optimizing the total loss function on source domain L^S given by:

$$L^S = L^S_{cls} + L^S_{tri} \tag{2}$$

B. Source-guided pseudo-labeling UDA

1) Noisy label regularization: In order to further exploit the source data and their labels, we propose the learning of source identity discriminative features by the feature encoder. This auxiliary task is added to the main task of learning target identity discriminative features with pseudo-labels. Since the source labels are ground-truth, this source-based auxiliary task is expected to guide the training in order to reduce the negative impact of over-fitting on erroneous target pseudo-labels, by reducting the amount of noisy data in this total (source + target) training set compared to the commonly-used target only set.

2) Two-branch architecture: To avoid biasing the model and thus the discriminative power of the target re-ID features with the source data, we choose a two-branch neural network architecture for the feature encoders as illustrated on Fig 1. It is composed of a domain-shared encoder E^C for low and mid-level features and two domain-specific encoders E^S and E_T resp. for source and target high level features. Our choice of modeling is supported by work that shows that features specialize for tasks in the top layers of the network. Our

feature encoder learns two separate domain-specific feature spaces given by $E^S \circ E^C$ for the source and $E_T \circ E^C$ for the target.

3) Domain-specific batch-normalization: Common neural network architectures for re-ID contain batch normalization to improve the training convergence. Experiments from the paper [22] shows that domain shift in data can reduce performance if the statistics of batch normalization layers are not computed separately for each domain. Since we learn with data from two different domains, we follow the paper suggestion in our pipeline and compute statistics separately for source and target data.

4) Our source-guided optimization criterion: Similarly to the previous section, we define for the target pseudo-labeled data, a target domain classifier W^T and then analogously L^T 2 for the target samples by considering the pseudo-labels as ground-truth labels. The optimization criterion of our Source-Guided Pseudo-labeling domain adaptation framework is given by:

$$L = L^S + L^T \tag{3}$$

We can note that we do not introduce any additional hyperparameter pondering the source and target term. Besides the difficulty of hyperparameter estimation in the UDA setting as mentioned earlier, our framework gives positive experimental results as it is on various datasets IV. The complete optimization procedure of our framework is detailed in the general Algorithm 1, where we highlight in **bold** how our source-guided baseline differentiates from the classical target-only pseudo-labeling UDA.

Algorithm 1 Source-guided Pseudo-label Domain Adaptation

Require: Labeled source data $D^S = (X^S, Y^S)$, unlabeled target data $D^T = X^T$, clustering algorithm C, a source re-ID loss function L^S, a target re-ID loss function L^T, number of training epochs N_{epoch}, number of pseudo-labeling iterations N_{iter}, an initial encoder $E^{(0)}$, a **two-branch** encoder E

Train the initial encoder $E^{(0)}$ on D^S by optimizing $L^S(E(X^S), Y^S)$
Initialize **two-branch** E such that $\mathbf{E^S} \circ \mathbf{E^C} = \mathbf{E^{(0)}}$ and $E^T \circ E^C = E^{(0)}$
for $t = 1$ to N_{iter} **do**
  Compute target features: $F^T \leftarrow \mathrm{E}(\mathrm{X}^T)$
  Compute pairwise target feature distances: $d(F^T) \leftarrow d(F_i^T, F_j^T)_{1 \leq i,j \leq N_T}$
  Pseudo-label some/all target samples by clustering: $(X^T, \hat{Y^T}) \leftarrow C(d(F^T), D^T)$
  Train E during N_{epoch} by optimizing $\mathbf{L^S}((\mathbf{E^S} \circ \mathbf{E^C})(\mathbf{X^S}), \mathbf{Y^S}) + L^T((E^T \circ E^C)(X^T), \hat{Y}^T)$
end for
Return $E^T \circ E^C$

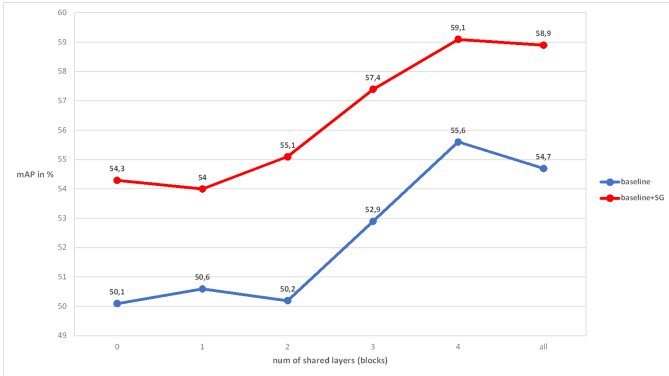

Fig. 2. Impact on mAP (in %) of the number of shared layers used in the shared encoder E^C of baseline+SG, on Duke-to-Market and Market-to-Duke.

IV. Experiments

A. Datasets

We evaluate our framework on three commonly-used re-ID datasets: Market-1501 (Market) [23], DukeMTMC-re-ID (Duke) [24] and MSMT17 (MSMT) [5]. Market-1501 and DukeMTMC-re-ID are the most commonly-used datasets for re-ID evaluation. Each of them defines a domain.

Market-1501 is composed of 32,668 labeled images from 1501 people captured by 6 outdoor cameras. It is divided into a training set of 12,936 images of 751 identities and a test set with 19,732 images of 750 identities different from the training ones.

DukeMTMC-re-ID contains 36,411 labeled images of 702 IDs taken by 8 outdoor cameras. It is split into a training set with 6,522 images of 702 identities and 19,889 images of 702 other identities for the test set.

MSMT17 is a larger dataset, with 126,441 labeled images of 4,101 identities collected by 15 indoor and outdoor cameras. The training set contains 32,621 images of 1,041 identities and the testing set 93,820 images of 3,060 other identities. It is worth noticing that MSMT17 is a much more challenging dataset than the other two: due to the size of its test set, its number of identities and cameras, MSMT17 is the closest dataset to the conditions of a large-scale re-ID system deployment. Therefore, re-ID UDA from Market or Duke to a larger-scale dataset such as MSMT is more difficult but also more interesting since in practice we often have few labeled data and a lot of unlaballed data.

We evaluate our framework on the commonly tested Duke-to-Market (Duke being the labeled source and Market the unlabeled target dataset) and Market-to-Duke UDA tasks. Besides, we test our approach on the more challenging adaptation tasks Market-to-MSMT and Duke-to-MSMT UDA tasks. Mean average precision (mAP) and CMC top-1 accuracy are reported to measure our framework's performance.

B. Experimental settings

1) Tested models: To show that our guideline can easily be added and contribute to various pseudo-labeling UDA

approaches, we choose to integrate it into two target-only frameworks:

- Our baseline: The classical pseudo-labeling UDA algorithm based on the UDAP approach [16] which is not designed to be robust to overfiting pseudo-labels'errors. It corresponds to the MMT [4] framework without Mutual Mean Teaching and k-means replaced by DBSCAN clustering algorithm on the pairwise matrix of k-reciprocal encoding distances [25] of the target training features.
- Our baseline+MMT [4] framework: actually the best state-of-the-art pseudo-labeling UDA method which mitigates for bad effects due to pseudo-label errors. It uses k-means as the clusterer in the feature space to predict target pseudo-labels based on the parirwise L2 distance matrix of the target training set features, as well as the mutual-mean teaching strategy described in their paper for pseudo-label error robustness.

2) Implementation details: **Initialization phase.** We follow the guidelines for supervised training from the paper [1] adopted by MMT [4].For fair comparison, we choose the ResNet-50 [26] initialized on the pretrained ImageNet weights [27] as our feature extractor. We use batch of 64 images composed of 16 identities and 4 shots per identity. Images are randomly flipped and resized to 256x128. Random Erasing Data augmentation [28] is not used during the initialization phase since it may reduce transferability of the source model features to the target domain thus generating more errors in pseudo-labels after UDA initialization [1]. We use ADAM as the optimizer and a weight decay of $5 \cdot 10^{-4}$. The initial learning rate is set to 0.00035 and is decreased to 1/10 of its previous value on the 40th and 70th epoch in a total of 80 epochs.

Pseudo-labeling phase. Unless otherwise specified, we use as a common encoder all but the layers from the last convolutional block and after (4 first blocs of layers), as motivated by our parameter analysis 2. We use the same initialization phase preprocessing with two batches of 64 images, adding Random Erasing Data augmentation [28]: one for source images and another one for target. We feed separately the network with the source and target batches to ensure domain-specific batch normalization statistics as explained in Section III. For our baseline, after each uour baselinete phase of the pseudo-labels, the pseudo-ID and their number can change. Therefore we initialize randomly a new classification layer for target after each pseudo-labels our baseline. Other hyperparameters (clusterer parameter, triplet loss margin, number of iterations for pseudo-labeling,...) used after the initialization phase are kept the same as the UDA paper's ones (resp MMT's ones): they correspond to the best hyperparameters found after validation on the target test set in their papers.

Source-Guided (SG) versions. For the baseline+SG framework, it is the direct application of our source-guided Algorithm 1. For our baseline+MMT+SG, the two-branch architecture is adopted for the two models (mean teacher and student) that train mutually in their baseline. It implies that we do not modify the soft and hard label loss weighting and we directly add the source term (on hard labels) as described in Eq. 3.

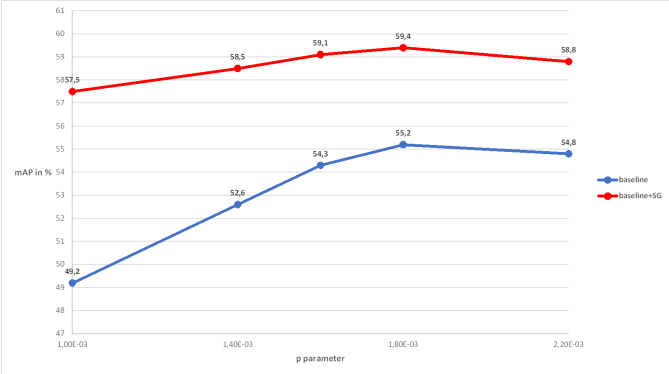

Fig. 3. Robustness of our baseline+SG to p parameter's changes (p controls DBSCAN neighborhood distance parameter) compared to the target-only framework our baseline on Duke-to-Market.

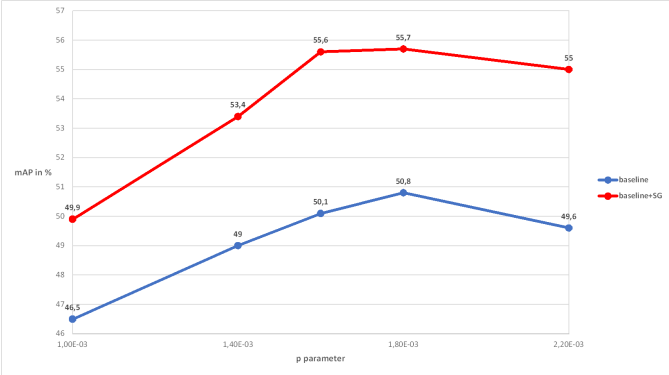

Fig. 4. Robustness of our baseline+SG to p parameter's changes (p controls DBSCAN neighborhood distance parameter) compared to the target-only framework our baseline on Market-to-Duke.

V. Results

In this part, we conduct parameter analysis and comparison with existing state-of-the-art re-ID UDA methods.

A. Parameter analysis.

1) Do the source domain help the model to learn better target features ?: As explained and motivated in Section III, we propose a two-branch architecture to learn domain-specific high level ID discriminative features based on low and mid level domain-shared features learned with labeled source data and pseudo-labeled target data. We can wonder if our two-branch encoder manages to leverage the source samples to improve the target features. Furthermore, we would like to know how many layers we should share to take advantage from the labeled source data without negatively biasing the target features.

To answer these two questions, we vary the number of ResNet-50 layers shared between source and target domain through

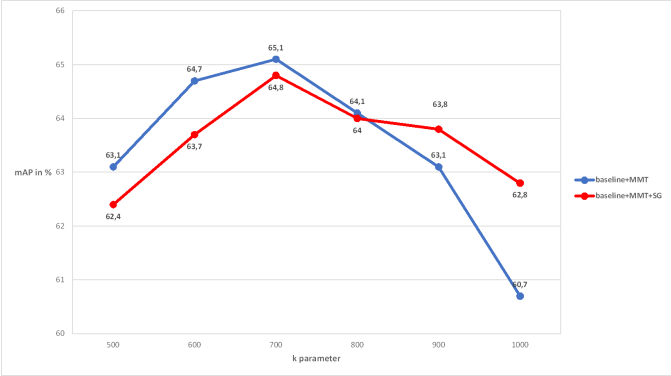

Fig. 5. Robustness of our baseline+MMT+SG to k parameter's changes (k controls k-means number of clusters) compared to the target-only framework MMT on Market-to-Duke.

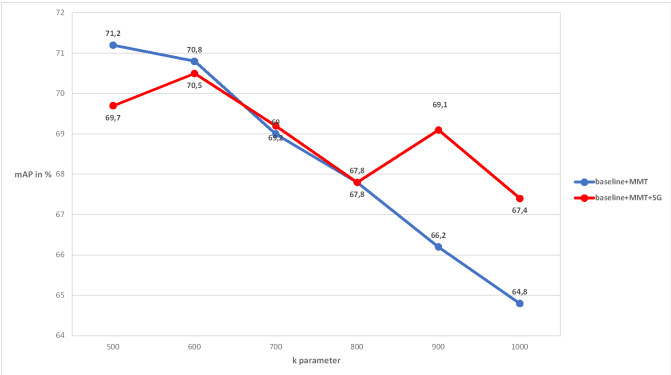

Fig. 6. Robustness of our our baseline+MMT+SG to k parameter's changes (k controls k-means number of clusters) compared to the target-only framework MMT on Duke-to-Market.

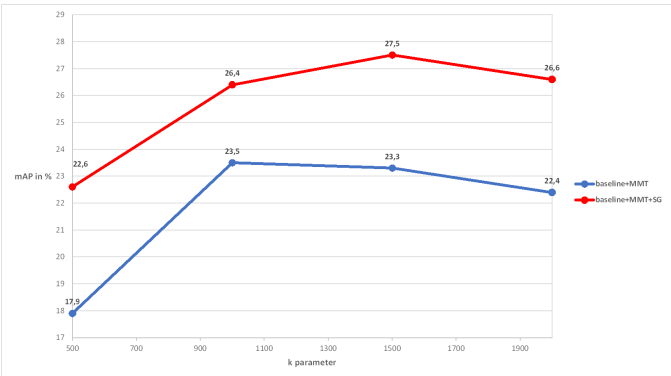

Fig. 7. Robustness of our baseline+MMT+SG to k parameter's changes (k controls k-means number of clusters) compared to the target-only framework MMT on Duke-to-MSMT.

the E^C encoder of our Source-Guided baseline. The ResNet-50 architecture can be divided into 5 convolutional blocks of layers defined in the ResNet paper [26] to which we refer to vary the number of shared layers. Case "0 shared

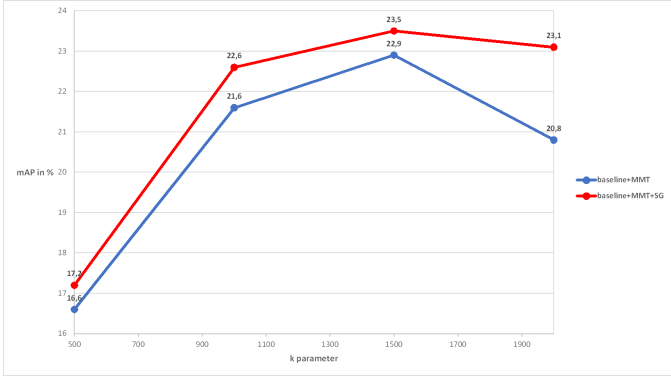

Fig. 8. Robustness of our baseline+MMT+SG to k parameter's changes (k controls k-means number of clusters) compared to the target-only framework MMT on Market-to-MSMT.

TABLE I
IMPACT OF DOMAIN-SPECIFIC BATCH-NORMALIZATION ON DOMAIN ADAPTATION PERFORMANCE (MAP IN %) .

Methods	Market-to-Duke	Duke-to-Market
	mAP	mAP
baseline	50.1	54.3
Domain-Shared BatchNorm baseline+SG	36.7	43.1
Domain-Specific BatchNorm baseline+SG	**55.6**	**59.1**

layer" corresponds to the classical target-only pseudo-labeling methods (Figure 1(1)) which corresponds to our baseline, where case "5 shared layers" to sharing the whole ResNet50 between source and target domain. Experiments show on Fig.2 increasing performances when the number shared layers increase. More precisely, the best mAP is reached for 4 shared block of layers: 59.1% mAP for Duke-to-Market and 55.6% mAP for Market-to-Duke, increasing resp. the performances by $4.8p.p.$ and $5.5p.p.$ compared to the target only model. Our source-guided baseline outperforms the classical target-only pseudo-labeling baseline and our partially-shared strategy gives the best results for Duke-to-Market and Market-to-Duke. Moreover, we also notice that the slight drop of performances for a fully-shared model may highlight a trade-off between:

- Sharing the highest level (specifically the last one) layers which can benefit the most from the auxiliary source re-ID task regularization
- Biasing the highest level feature space which is directly used for re-ID on target domain.

2) Efficiency of Specific Batch Normalization: We study the influence of the use of domain-specific batch nornmalization as motivated in Section III. To study its effectiveness, we compare our baseline+SG framework to a version that shares the batch normalization between domains (Domain Shared BatchNorm baseline+SG). From Table I, we notice that sharing the batch normalization deteriorates the performance on both couples of Market-to-Duke and Duke-to-Market adaptation datasets. mAP are dropping more than -10 p.p. below the model using only the target data (our baseline). Only the addition of domain-specific batch normalization increases the

performances of our framework above the our baseline model. These experiments therefore show that the use of domain-specific batch normalization is an essential key of our framework in order not to deteriorate the learning of discriminative target features by biasing the batch normalization statistics.

3) Is our strategy of using source samples robust to clustering parameters changes ?: In the UDA setting, choosing or tuning hyperparameter is a tricky task due to the absence of a labeled validation set for the target domain. It is therefore important in practice to design UDA methods robust to hyperparameter changes. In particular, pseudo-labeling UDA methods [16] [4] give experimental evidences that performance can be very sensible to clustering parameters changes. That's why we would like to focus in this part on the performance of our source-guided frameworks when these clustering parameters change. For the baseline+SG baseline, we study the p parameter of DBSCAN as defined and studied in UDAP paper [16] and for our baseline+MMT+SG we focus on the k paremeter of k-means as in the MMT paper [4]. We compare our two frameworks to their target-only versions. The p parameter directly controls the DBSCAN neighborhood distance parameter in our baseline and baseline+SG: it determines the percentage of the smallest distances to be used to define clusters in the feature space. We vary p in the same interval of the UDAP paper [16]. By doing so, we notice that for Duke-to-Market in Figure 3 and Market-to-Duke in Figure 4, in baseline+SG the source benefits the domain adaptation by pseudo-labeling: we go from 50.6% to 55.9% on Market-to-Duke and from 54.6% to 59.4% by comparison to the target only version. For every parameter value p, there is an improvement in performance for baseline+SG ranging from at least 3.4 p.p. for Market-to-Duke and 4 p.p. for Duke-to-Market in comparison to the target only baseline our baseline: *** our framework seems to be robust to the change of parameters p.

The k parameter determines the number of clusters in the MMT frameworks. We choose the same interval of values as in the MMT paper [4] for varying the k parameter. For Maket-to-Duke and Duke-to-Market in Figures 5 and 6, the addition of the source term with our baseline+MMT+SG does not seem to increase the maximum performance, which is reached by the classical target-only MMT model. Nevertheless, there are quite different performance curve trends between MMT and our baseline+MMT+SG. our baseline+MMT+SG seems to be more robust for k values above 800 (Market and Duke containing 751 and 702 identities respectively), i.e. when a number of clusters is chosen above the actual number of identitiesWhile MMT already proposes a strategy of resistance to pseudo-label noise, which can explain the non improvement of the best mAP, the addition of the source-guidance in our baseline+MMT+SG seems to contribute to errors in pseudo-labels. Numerically, this stability can also be observed by calculating the standard deviation (std) of the mAPs over k: using the source, we go from 2.5 p.p. to 1.2 p.p. on Duke-to-Market and from 1.6 p.p. to 0.9 p.p. on Market-to-Duke, with on average equivalent performance for the two pairs of data sets. This stability conferred by the source is interesting given that we not know the number of identities of the training set target, which can only be estimated at best.

TABLE II
COMPARISON WITH STATE-OF-THE-ART METHODS.

Methods	Market-to-Duke		Duke-to-Market	
	mAP	top-1	mAP	top-1
SPGAN [2]	22.3	41.1	22.8	51.5
TJ-AIDL [13]	23.0	44.3	26.5	58.2
MMFA [10]	24.7	45.3	38.3	66.2
HHL [7]	27.2	46.9	31.4	62.2
CFSM [9]	27.3	49.8	28.3	61.2
UCDA-CCE [6]	31.0	47.7	30.9	60.4
ARN [11]	33.4	60.2	39.4	70.3
ECN [15]	40.4	63.3	43.0	75.1
PoseDA-Net [12]	45.1	63.2	47.6	75.2
UDAP [16]	49.0	68.4	53.7	75.8
SSG [18]	53.4	73.0	58.3	80.0
ISSDA-re-ID [19]	54.1	72.8	63.1	81.3
PCB-PAST [17]	54.3	72.4	54.6	78.4
ACT [20]	54.5	72.4	60.6	80.5
MMT [4]	**65.1**	78.0	**71.2**	87.7
Our (target-only) baseline	50.1	70.1	54.3	73.5
Our baseline+SG	55.6	73.2	59.1	80.8
Our baseline+MMT+SG	64.8	**78.5**	70.5	**88.1**

Methods	Market-to-MSMT		Duke-to-MSMT	
	mAP	top-1	mAP	top-1
PTGAN [5]	2.9	10.2	3.3	11.8
ECN [6]	8.5	25.3	10.2	30.2
UDAP [16]	12.0	30.5	16.0	39.2
SSG [18]	13.2	49.6	13.3	32.2
MMT [4]	22.9	49.2	23.5	50.1
Our (target-only) baseline	11.6	29.8	14.8	36.1
Our baseline+SG	14.9	35.4	19.3	45.6
Our baseline+MMT+SG	**23.5**	**50.2**	**27.5**	**56.1**

In the more challenging cases where MSMT is the target dataset, there is a clear contribution from the source. We can see in Figures 8 and 7 that it is stable to the change in k and allowed to increase the maximum performance: from 23.5% to 27.5% for Duke-to-MSMT and from 22.9% to 23.5% for Market-to-MSMT. There is also a higher source contribution at high k values. It can be assumed that our baseline+MMT+SG works better in this more challenging case of adaptation because of the presence of more noisy labels during the transfer of the source model for initialization of pseudo-labeling: adding our strategy of exploiting source data therefore presents less redundancy with the one already implemented in the MMT framework, and even more if we "over-estimate" the number of clusters.

B. Comparison with state-of-the-art methods

We compare in Table II our two source-guided frameworks with the state of the art on Duke-to-Market, Market-to-Duke, Market-to-MSMT and Duke-to-MSMT.

On Market-to-Duke and Duke-to-Market, our two frameworks far exceed those that do not use pseudo-labels at all. Our frameworks outperform all state of the art methods except MMT on Market-to-Duke and Duke-to-Market. Specifically, baseline+SG outperforms similar Pseudo-labeling methods that do not exploit the labeled source data after pseudo-label initialization: UDAP and PCB-PAST, but also

those that integrate noise resistance or error filtering strategies in pseudo-labels such as ISSDA-re-ID, SSG and ACT on Maket-to-Duke. On Duke-to-Market, ISSDA-re-ID surpasses baseline+SG, even if baseline+SG has the advantage of simple integration in any framework and needs few additional parameters (no need to train a CycleGAN or a second model for asymmetric co-teaching).

Even if our baseline+MMT+SG, has a slightly lower maximum mAP than MMT, it offers the advantage of a better stability on the parameter of the number of clusters (estimated) as seen in the parameter analysis part, which is a major asset in domain adaptation.

On Market-to-MSMT and Duke-to-MSMT, our baseline+SG framework exceeds all methods except MMT. In particular, it exceeds SSG by $+1.7$ p.p. mAP on Market-to-MSMT and UDAP by $+3.3$ p.p. on Duke-to-MSMT. The our baseline+MMT+SG framework exceeds the state of the art on these two pairs of data sets: on Market-to-MSMT and Duke-to-MSMT, our baseline+MMT+SG resp. increases performance by $+0.6$ p.p. and by $+4$ p.p. on Duke-to-MSMT.

VI. Conclusion

In this paper, we propose a guideline to leverage the commonly under-used ground-truth labeled source samples during pseudo-labeling domain adaptation. It consists of guiding the target pseudo-label training stage with an auxiliary ID-discriminative source feature learning task, while preventing the source samples from biasing the training by using domain-specific batch normalization and an architecture with two domain-specific branches. Experiments on combining our framework with two state-of-the-art methods are carried out on different datasets. They show that leveraging the source samples brings more stability with relation to the choice of clustering parameters and improves performance on particularly challenging adaptation settings.

References

[1] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, "Bag of tricks and a strong baseline for deep person re-identification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.

[2] W. Deng, L. Zheng, Q. Ye, G. Kang, Y. Yang, and J. Jiao, "Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 994–1003.

[3] P. Panareda Busto and J. Gall, "Open set domain adaptation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 754–763.

[4] Y. Ge, D. Chen, and H. Li, "Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification," *arXiv preprint arXiv:2001.01526*, 2020.

[5] L. Wei, S. Zhang, W. Gao, and Q. Tian, "Person transfer gan to bridge domain gap for person re-identification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 79–88.

[6] L. Qi, L. Wang, J. Huo, L. Zhou, Y. Shi, and Y. Gao, "A novel unsupervised camera-aware domain adaptation framework for person re-identification," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 8080–8089.

[7] Z. Zhong, L. Zheng, S. Li, and Y. Yang, "Generalizing a person retrieval model hetero-and homogeneously," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 172–188.

[8] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232.

[9] X. Chang, Y. Yang, T. Xiang, and T. M. Hospedales, "Disjoint label space transfer learning with common factorised space," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3288–3295.

[10] S. Lin, H. Li, C.-T. Li, and A. C. Kot, "Multi-task mid-level feature alignment network for unsupervised cross-dataset person re-identification," *arXiv preprint arXiv:1807.01440*, 2018.

[11] Y.-J. Li, F.-E. Yang, Y.-C. Liu, Y.-Y. Yeh, X. Du, and Y.-C. Frank Wang, "Adaptation and re-identification network: An unsupervised deep transfer learning approach to person re-identification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018, pp. 172–178.

[12] Y.-J. Li, C.-S. Lin, Y.-B. Lin, and Y.-C. F. Wang, "Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 7919–7929.

[13] J. Wang, X. Zhu, S. Gong, and W. Li, "Transferable joint attribute-identity deep learning for unsupervised person re-identification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2275–2284.

[14] H.-X. Yu, W.-S. Zheng, A. Wu, X. Guo, S. Gong, and J.-H. Lai, "Unsupervised person re-identification by soft multilabel learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2148–2157.

[15] Z. Zhong, L. Zheng, Z. Luo, S. Li, and Y. Yang, "Invariance matters: Exemplar memory for domain adaptive person re-identification," *CoRR*, vol. abs/1904.01990, 2019. [Online]. Available: http://arxiv.org/abs/1904.01990

[16] L. Song, C. Wang, L. Zhang, B. Du, Q. Zhang, C. Huang, and X. Wang, "Unsupervised domain adaptive re-identification: Theory and practice," *Pattern Recognition*, p. 107173, 2020.

[17] X. Zhang, J. Cao, C. Shen, and M. You, "Self-training with progressive augmentation for unsupervised cross-domain person re-identification," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 8222–8231.

[18] Y. Fu, Y. Wei, G. Wang, Y. Zhou, H. Shi, and T. S. Huang, "Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 6112–6121.

[19] H. Tang, Y. Zhao, and H. Lu, "Unsupervised person re-identification with iterative self-supervised domain adaptation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.

[20] F. Yang, K. Li, Z. Zhong, Z. Luo, X. Sun, H. Cheng, X. Guo, F. Huang, R. Ji, and S. Li, "Asymmetric co-teaching for unsupervised cross domain person re-identification," *arXiv preprint arXiv:1912.01349*, 2019.

[21] A. Hermans, L. Beyer, and B. Leibe, "In defense of the triplet loss for person re-identification," *arXiv preprint arXiv:1703.07737*, 2017.

[22] M. Zajac, K. Zolna, and S. Jastrzebski, "Split batch normalization: Improving semi-supervised learning under domain shift," 2019.

[23] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian, "Scalable person re-identification: A benchmark," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1116–1124.

[24] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *European Conference on Computer Vision*. Springer, 2016, pp. 17–35.

[25] Z. Zhong, L. Zheng, D. Cao, and S. Li, "Re-ranking person re-identification with k-reciprocal encoding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1318–1327.

[26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[27] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.

[28] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, "Random erasing data augmentation," *arXiv preprint arXiv:1708.04896*, 2017.