

Large-scale Training Data Search for Object Re-identification

Yue Yao¹ Tom Gedeon² Liang Zheng¹
¹Australian National University ²Curtin University
 {first name.second name}@{anu | curtin}.edu.au

Abstract

We consider a scenario where we have access to the target domain, but cannot afford on-the-fly training data annotation, and instead would like to construct an alternative training set from a large-scale data pool such that a competitive model can be obtained. We propose a search and pruning (SnP) solution to this training data search problem, tailored to object re-identification (re-ID), an application aiming to match the same object captured by different cameras. Specifically, the search stage identifies and merges clusters of source identities which exhibit similar distributions with the target domain. The second stage, subject to a budget, then selects identities and their images from the Stage I output, to control the size of the resulting training set for efficient training. The two steps provide us with training sets 80% smaller than the source pool while achieving a similar or even higher re-ID accuracy. These training sets are also shown to be superior to a few existing search methods such as random sampling and greedy sampling under the same budget on training data size. If we release the budget, training sets resulting from the first stage alone allow even higher re-ID accuracy. We provide interesting discussions on the specificity of our method to the re-ID problem and particularly its role in bridging the re-ID domain gap. The code is available at <https://github.com/yorkeyao/SnP>

1. Introduction

The success of a deep learning-based object re-ID relies on one of its critical prerequisites: the labeled training data. To achieve high accuracy, typically a massive amount of data needs to be used to train deep learning models. However, creating large-scale object re-ID training data with manual labels is expensive. Furthermore, collecting training data that contributes to the high test accuracy of the trained model is even more challenging. Recent years have seen a large amount of datasets proposed and a significant increase in the data size of a single dataset. For example, the RandPerson [37] dataset has 8000 identities, which is more than

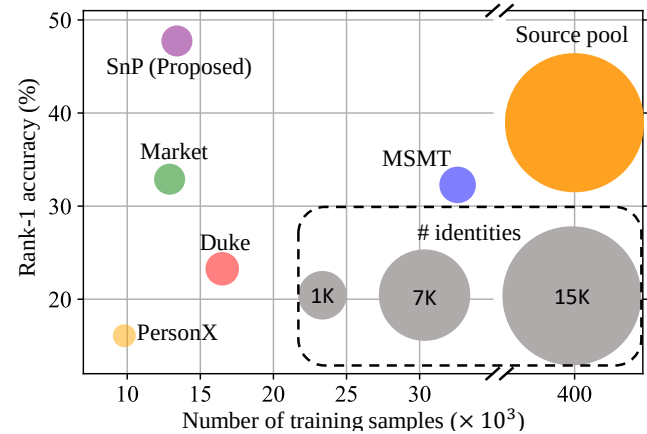


Figure 1. We present a search and pruning (SnP) solution to the training data search problem in object re-ID. The source data pool is 1 order of magnitude larger than existing re-ID training sets in terms of the number of images and the number of identities. When the target is AlicePerson [1], from the source pool, our method (SnP) results in a training set 80% smaller than the source pool while achieving a similar or even higher re-ID accuracy. The searched training set is also superior to existing individual training sets such as Market-1501 [54], Duke [55], and MSMT [45].

6 \times larger than the previous PersonX [37] dataset.

However, these datasets generally have their own dataset bias, making the model trained on one dataset unable to generalize well to another. For example, depending on the filming scenario, different person re-ID datasets generally have biases on camera positions, human races, and clothing style. Such dataset biases usually lead to model bias, which results in the model's difficulty performing well in an unseen filming scenario. To address this, many try to improve learning algorithms, including domain adaptation and domain generalization methods [2, 10, 27, 33, 35, 56]. Whereas these algorithms are well-studied and have proven successful in many re-ID applications, deciding what kind of data to use for training the re-ID model is still an open research problem, and has received relatively little attention in the community. We argue that this is a crucial problem

to be answered in light of the ever-increasing scale of the available datasets.

In this paper, we introduce SnP, a search and pruning solution for sampling an efficient re-ID training set to a target domain. SnP is designed for the scenario in that we have a target dataset that does not have labeled training data. Our collected source pool, in replace, provides suitable labeled data to train a competitive model. Specifically, given a user-specified budget (*e.g.*, maximum desired data size), we sample a subset of the source pool, which satisfies the budget requirement and desirably has high-quality data to train deep learning models. This scenario is especially helpful for deploying a re-ID system for unknown test environments, as it is difficult to manually label a training set for these new environments. We note that due to the absence of an in-distribution training set, the searched data are directly used for training the re-ID model rather than pre-training.

In particular, we combine several popular re-ID datasets into a source pool, and represent each image in the pool with features. Those features are extracted from an Imagenet-pretrained model [39]. The images with features are stored on the dataserver to serve as a gallery. When there is a query from the target, we extract the feature of the query image, and search in the gallery for similar images on a feature level. Specifically, in the search stage, we calculate feature-level distance, *i.e.*, Fréchet Inception Distance (FID) [16]. Given the constraint of a budget, we select the most representative samples in the pruning stage, based on the outputs from the search stage. This limits the size of the constructed training set and enables efficient training.

Combining search and pruning, we construct a training dataset that empirically shows significant accuracy improvements on several object re-ID targets, compared to the baselines. Without budget constraints, our searched training sets allow higher re-ID accuracy than the complete source pool, due to its target-specificity. With budget constraints, the pruned training set still achieves comparable or better performance than the source pool. The proposed SnP is demonstrated to be superior to random or greedy sampling. We show in Fig. 1 that the training set constructed by SnP leads to the best performance on the target compared to the others. We provide discussions on the specificity of our method and its role in bridging the re-ID domain gap.

2. Related Work

Active learning gradually searches over unlabeled data to find samples to be labeled by the oracle [32]. It is an iterative training process, where data search is performed at each iteration to train a task model. Our task is different from active learning. First, active learning is designed to acquire a training set and train models gradually. This requires multiple real training processes which are computationally expensive. In comparison, we directly search the

whole training set all at once. Then the searched training set is used for model training. Second, many active learning methods require access to the target task model to provide selection metrics, *e.g.*, uncertainty-based metrics [9, 12, 14, 49]. Our task does not require task model information, in other words, a real training process, during the process of the training set search. Thus enabling a fast training data retrieval.

Neural data server is the closest inspiring work [32, 48]. They also aim to search training data all at once from a large database. However, compared with us, firstly, [48] and [32] are designed for searching pretraining data rather than direct training data. Such a design is understandable as they are mainly for the classification task. Searching direct training data require careful class alignment. In comparison, We are targeting the re-ID task, where its training set can have a different class from the target, and then be directly used for training models. Furthermore, [48] and [32] require unsupervised pretrained experts to measure the domain gap. While we do not require so, which saves extraction time and simplifies the solution.

Transfer learning is a long-standing problem for re-ID tasks, and many attempts on learning algorithms have been made to reduce the effect of domain gap [7, 26, 29, 30, 42]. Common strategies contain feature-level [24] and pseudo-label based [10, 35, 56] domain adaptation, and domain generalization [2, 27, 33]. In this paper, we focus on training data that is orthogonal to existing training algorithms. To be shown in experiments, together with some domain adaptation (*i.e.*, pseudo-label) methods, SnP can achieve higher re-ID accuracy.

Learning to generate synthetic training data. Data simulation is also an inexpensive alternative to increasing a training set scale while providing accurate image labels [23, 37, 50–52]. These methods aim to lower the domain gap between synthetic data and real data by searching a set of parameters that control the 3D rendering process [50, 51]. In comparison, our search is not conducted on predefined parameters but on data directly. This enables more direct research on how to form a good training set.

Object re-ID has received increasing attention in the past few years, and many effective systems have been proposed [17, 40, 44, 57]. In this paper, we study object re-ID datasets rather than algorithms. Depending on the camera condition, location and environment, existing object re-ID datasets usually have their own distinct characteristics or bias [36, 37, 51]. We show details in §3.

3. Motivation: Tackling Target Bias

Data bias commonly exists in the re-ID datasets (for examples see below) [25, 43, 47, 52], and it becomes problematic when the training and testing have different biases. Given a target domain with a certain bias, we aim to find



Figure 2. Dataset bias in existing object re-ID datasets. (A) Viewpoint distribution of different re-ID datasets. Each blue dot indicates a sample filmed from a specific viewpoint. Compared with the bi-modal viewpoint distribution in Market camera 1, Duke camera 4 has more diverse viewpoints. (B) Different human races and clothing styles of different person re-ID datasets. (C) Different vehicle makes in different datasets. For example, taxis in VeRi are unlikely to be found in CityFlow. (D) Different backgrounds in different datasets. We show a city background in AliceVehicle versus an urban background in CityFlow. To tackle such bias in the target domain, we design an automatic way to generate a training set with similar distribution or bias.

a target-specific training set that has similar distribution or bias. Depending on the filming scenario, there are four major types of data biases in existing re-ID datasets. We show examples of each type in Fig. 2.

Viewpoint. Viewpoint bias applies to generic objects, including persons and vehicles. We visualize the viewpoint distributions of two representative person re-ID datasets in Fig. 2(A), *i.e.*, Market-1501 [54] (denoted as Market) and Duke-reID [55] (denoted as Duke)¹.

Human races and clothing style. Subject to the places where the data is collected, identities in the person re-ID datasets can have distinctive patterns. We show in Fig. 2(B) that the humans in Market and Duke datasets evince different races and clothing styles.

Vehicle model. Similar to the identity bias (*i.e.*, human races and clothing style) in person re-ID, the vehicle identities in vehicle re-ID also hold distinct patterns across different datasets. We show examples in Fig. 2(C) using the VeRi [22] (denoted as VeRi) and the CityFlow [41].

Background. Background bias exists in both person and vehicle re-ID datasets, which is similar to viewpoint bias. Fig. 2(D) compares the background difference of images from two different vehicle re-ID datasets.

4. Method

Given a target unlabeled dataset, we aim to construct a source labeled dataset that has minimal data bias inconsistencies with the target, under certain budget constraints. It induces the model trained on the source to show good performance on the target. To achieve the construction of the source (training) dataset, we propose the search and pruning (SnP) framework.

¹ We understand that it is no longer encouraged to use the Duke dataset. In fact, we are not using it individually for algorithm design, but moving forward to find solutions to replace such individual use.

4.1. Overview

We denote the target set as $\mathcal{D}_T = \{(\mathbf{x}_i, y_i)\}_{i \in [m_t]}$ where m_t indicates the number of image-label pairs in the target and $[m_t] = \{1, 2, \dots, m_t\}$. It follows the distribution p_T , *i.e.*, $\mathcal{D}_T \sim p_T$. Let \mathcal{D}_S be the source set to be constructed under a budget b . The budget is specified by the number of identities n and the number of images m allowed in \mathcal{D}_S , denoted as $b = (n, m)$. A high budget can lead to an unwanted increase in the training cost, in terms of either training time or model size.

To construct the training set \mathcal{D}_S , we create a source pool \mathcal{S} , which is a collection of multiple object re-ID datasets. It is represented as $\mathcal{S} = \mathcal{D}_S^1 \cup \mathcal{D}_S^2 \cdots \cup \mathcal{D}_S^K$. Here each \mathcal{D}_S^k , $k \in [K]$, indicates the k -th source re-ID dataset. Given the source pool, we firstly build a subset \mathbf{S}^* of \mathcal{S} regardless of the budget constraint. Let h_S be a model h trained on an arbitrary dataset \mathbf{S} . The prediction risk of h_S on the test sample \mathbf{x} with ground truth label y is computed as $\ell(h_S(\mathbf{x}), y)$. We build \mathbf{S}^* by ensuring that the model $h_{\mathbf{S}^*}$ has minimized risk on \mathcal{D}_T , *i.e.*,

$$\mathbf{S}^* = \arg \min_{\mathbf{S} \in 2^{\mathcal{S}}} \mathbb{E}_{\mathbf{x}, y \sim p_T} [\ell(h_S(\mathbf{x}), y)]. \quad (1)$$

We apply target-specific search in §4.3 to construct \mathbf{S}^* .

It can be seen that the construction of \mathbf{S}^* does not take the budget constraint $b = (n, m)$ into consideration, which is otherwise important in reality. Therefore, we build the training set \mathcal{D}_S by pruning \mathbf{S}^* to comprise no more than n identities and no more than m images. Details of the budgeted pruning process are introduced in §4.4.

4.2. Correlation Study

In order to know how to obtain \mathbf{S}^* , we conduct the correlation study to learn the relationships between the dataset bias differences (*i.e.*, domain gap measured in measured in FID [16]), the number of IDs and the training set quality

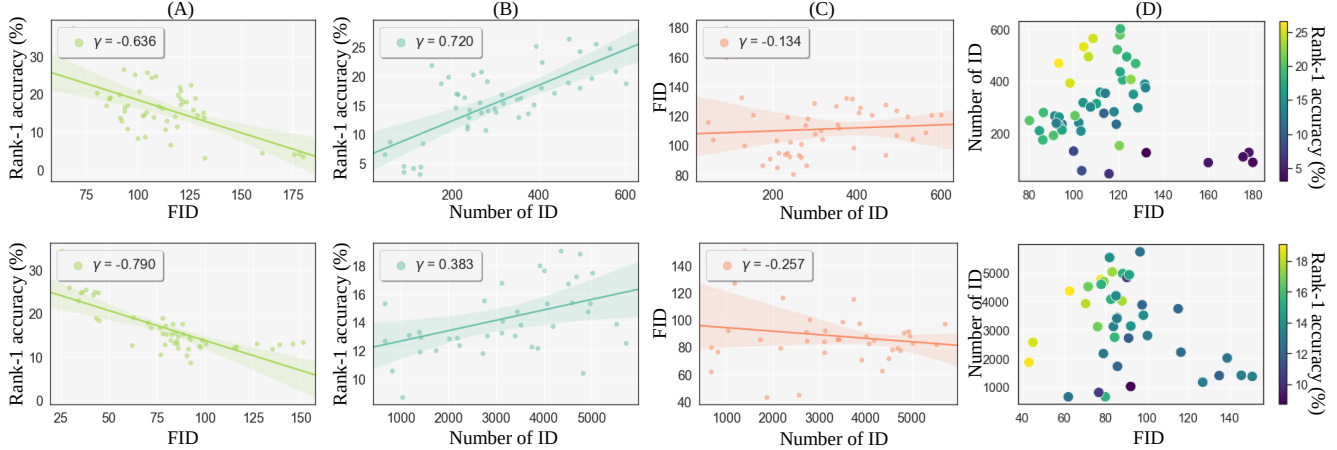


Figure 3. Relationships between the domain gap (measured using FID), the number of ID and the rank-1 accuracy on the target. **Top:** correlations when AlicePerson [1] is targeted. **Bottom:** correlations when AliceVehicle [1] is targeted. The Pearson correlation is used to measure the relationship between them. (A) FID vs. the rank-1 accuracy. They have a relatively strong negative correlation (≤ -0.636). (B) The number of ID vs. the rank-1 accuracy. There exists a positive correlation but such a correlation is not stable. (C) FID vs. number of ID. The correlation is weak between them. (D) The joint influence of FID and the number of ID on rank-1 accuracy. The top left corner, *i.e.*, training sets that have low FID scores and a large number of IDs have high rank-1 accuracy on the target.

(*i.e.*, rank-1 test accuracy with the model trained on the such training set). Shown in Fig. 3, in each subfigure, each point represents a training set, which is clustered from the source pool. For each training set, we calculate its domain gap to the target domain, count its number of IDs, and evaluate the rank-1 accuracy. We use the Pearson correlation coefficient (γ) [5] to measure the correlation between them. The Pearson correlation coefficients range from $[-1, 1]$. A value closer to -1 or 1 indicates a stronger negative or positive correlation, respectively, and 0 implies no correlation.

From this Fig. 3(A), we observe a relatively strong negative correlation between domain gap and rank-1 accuracy. This indicates minimizing the domain gap between the source and target is highly likely to improve training set quality, *i.e.*, test set rank-1 accuracy. From this Fig. 3(B), we observe a positive but unstable correlation between the number of ID and the rank-1 accuracy. For example, when AlicePerson is targeted, they have a relatively strong positive correlation (0.720). However, such correlation is only 0.383 when AliceVehicle is targeted. Fig. 3(C) shows the correlation between FID and the number of ID is weak. They are independent factors that influence training set quality. Fig. 3(D) shows the joint influence of FID and the number of ID on the training set quality. The top left corner, *i.e.*, training sets that have a low domain gap to the target and a large number of IDs are of high quality, thereby can train a model that has higher rank-1 accuracy on the target.

4.3. Target-specific Subset Search

The theory of domain adaptation [3] states that:

$$\varepsilon_T(h) < \varepsilon_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}, \mathcal{D}_T). \quad (2)$$

Here $h \in \mathcal{H}$ represents the hypothesis function (*i.e.*, the model). $\varepsilon_T(h)$ is the risk of model h on the target set \mathcal{D}_T , while $\varepsilon_S(h)$ is its risk on the source set \mathbf{S} . $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}, \mathcal{D}_T)$ is the unlabelled $\mathcal{H}\Delta\mathcal{H}$ divergence [3] between \mathbf{S} and \mathcal{D}_T . Equation 2 shows that the target risk $\varepsilon_T(h)$ is upper bounded by $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}, \mathcal{D}_T)$.

In the common practice of feature-level domain adaptation, the source training set \mathbf{S} is fixed while the joint feature extraction model is used to minimize $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}, \mathcal{D}_T)$ [24]. In our design, the feature extraction model is fixed instead. To minimize $\varepsilon_T(h)$, the problem is reformulated as

$$\mathbf{S}^* = \arg \min_{\mathbf{S} \in 2^{\mathcal{S}}} d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}, \mathcal{D}_T). \quad (3)$$

Generally, $d_{\mathcal{H}\Delta\mathcal{H}}(\mathbf{S}^*, \mathcal{D}_T)$ is difficult to compute, but many alternatives exist in the literature [24]. We use Fréchet Inception Distance (FID) [16], which is defined as:

$$\text{FID}(\mathbf{S}^*, \mathcal{D}_T) = \|\boldsymbol{\mu}_s - \boldsymbol{\mu}_t\|_2^2 + \text{Tr}(\boldsymbol{\Sigma}_s + \boldsymbol{\Sigma}_t - 2(\boldsymbol{\Sigma}_s \boldsymbol{\Sigma}_t)^{\frac{1}{2}}). \quad (4)$$

In Eq. 4, $\boldsymbol{\mu}_s \in \mathbb{R}^d$ and $\boldsymbol{\Sigma}_s \in \mathbb{R}^{d \times d}$ are the mean and covariance matrix of the image descriptors of \mathbf{S}^* , respectively. $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ are those of \mathcal{D}_T . $\text{Tr}(\cdot)$ represents the trace of a square matrix. d is the dimension of the image descriptors. Consequently, the objective function is reduced as,

$$\mathbf{S}^* = \arg \min_{\mathbf{S} \in 2^{\mathcal{S}}} \text{FID}(\mathbf{S}, \mathcal{D}_T). \quad (5)$$

We build \mathbf{S}^* with the greedy algorithm below.

Firstly, we divide the entire dataset \mathcal{S} into J clusters $\{\mathbf{S}^1, \dots, \mathbf{S}^J\}$, as shown in Fig. 4 (A). Specifically, we average all image descriptors that belong to the same identity,

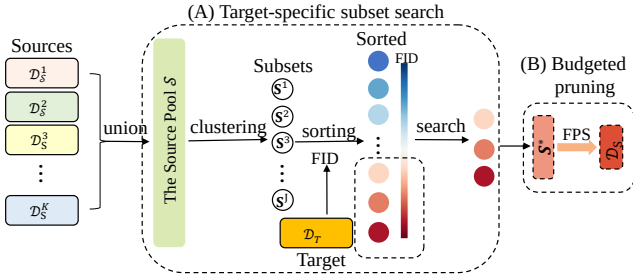


Figure 4. Workflow of the proposed SnP method. We are given sources composed of K existing datasets. From these sources, we aim to construct a training set, which satisfies the budget of no more than n IDs and m images. To achieve this, we perform (A) target-specific subset search to obtain a subset \mathbf{S}^* with similar distributions to the target, then perform (B) budgeted pruning to select n IDs and m representative images, forming the final training set.

and use this ID-averaged descriptor to represent all corresponding images. Afterwards, we cluster the ID-averaged features into J groups using the k-means method [20]. Each subset \mathbf{S}^j , $j \in [J]$, is composed of all images with the corresponding IDs in that group. Secondly, we calculate the FID between each subset \mathbf{S}^j and the target \mathcal{D}_T , and sort $\{\text{FID}(\mathbf{S}^j, \mathcal{D}_T)\}_{j \in [J]}$ in ascending order. To build \mathbf{S}^* , we keep adding the subsets with lower FID to \mathbf{S}^* until $\text{FID}(\mathbf{S}^*, \mathcal{D}_T)$ stops to decrease. This indicates that the constructed \mathbf{S}^* holds the minimum FID to the target set. Algorithm 1 summarizes the above procedures. We empirically demonstrate in Fig. 5 that, the subset \mathbf{S}^* with the minimum $\text{FID}(\mathbf{S}^*, \mathcal{D}_T)$ results in the model to produce highest accuracy on the target dataset.

4.4. Budget-constrained Pruning

Using the target-specific subset search, we construct a candidate training set \mathbf{S}^* . Yet, it can violate the budget constraint of $b = (n, m)$, which prefers the training set to include no more than n identities and no more than m images. Let $\mathcal{Y}(\mathbf{S}^*) = \{y^1, y^2, \dots, y^a\}$, be the set of unique identities in \mathbf{S}^* , and $\mathbf{s}(y^i)$, $i \in [a]$ be the set of all images with identity y^i . Generally, we have $a \geq n$. To get the training set \mathcal{D}_S under budget b , we randomly sample a subset \mathbf{y} of n identities from \mathcal{Y} . All images with identities in \mathbf{y} are combined to form a subset $\hat{\mathbf{s}} = \{\mathbf{s}(y^i) | y^i \in \mathbf{y}\}$. We also initialize \mathcal{D}_S by random sampling a seed image from each $\mathbf{s}(y^i)$, $y^i \in \mathbf{y}$. This guarantees \mathcal{D}_S to cover all the identities in the subset $\hat{\mathbf{s}}$.

Suppose the number of images in $\hat{\mathbf{s}}$ is larger than m . We construct \mathcal{D}_S by sampling images iteratively from $\hat{\mathbf{s}}$ until $|\mathcal{D}_S| = m$. To ensure similar performance between the model trained on \mathcal{D}_S and the model trained on $\hat{\mathbf{s}}$, we minimize the risk differences between them,

$$\mathcal{D}_S = \arg \min_{\mathcal{D}_S \in 2^{\hat{\mathbf{s}}}} |L(h_{\hat{\mathbf{s}}}(\mathbf{x}), y) - L(h_{\mathcal{D}_S}(\mathbf{x}), y)|. \quad (6)$$

$L(h_{\hat{\mathbf{s}}}(\mathbf{x}), y)$ and $L(h_{\mathcal{D}_S}(\mathbf{x}), y)$ are the respective risks of

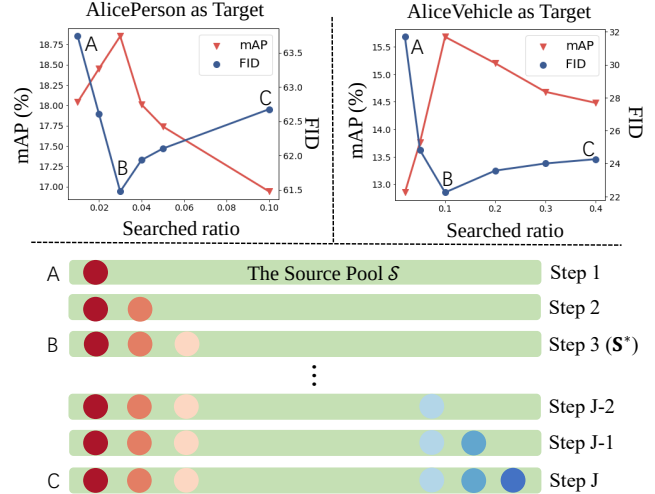


Figure 5. Performance of constructed training set at different iterations on the target set. We greedily add clustered subset to form \mathbf{S}^* which has the smallest domain gap to the target \mathcal{D}_T . During this process, we have relatively high FID and low accuracy at (A). Gradually, we reach situation (B) of the smallest domain gap and highest accuracy. After (B), adding subsets leads to an increase in the domain gap, and the accuracy drops, *e.g.*, (C). To study the correlation between FID and mAP during the search process, we hope to eliminate the impact of dataset size. Thus we use the same number of IDs from the searched set (*i.e.*, 2% ID).

model h on the dataset $\hat{\mathbf{s}}$ and \mathcal{D}_S . For explicitly, we define the risk of model h on an arbitrary dataset \mathbf{S} as

$$L(h_{\mathbf{S}}(\mathbf{x}), y) = \frac{1}{|\mathbf{S}|} \sum_{(\mathbf{x}_i, y_i) \in \mathbf{S}} \ell(h_{\mathbf{S}}(\mathbf{x}_i), y_i). \quad (7)$$

As in Eq. 1, $\ell(h_{\mathbf{S}}(\mathbf{x}_i), y_i)$ is the risk on individual samples.

From the theory of core set [31], if \mathcal{D}_S is the δ_s cover of the set $\hat{\mathbf{s}}$ and shares the same number of classes with $\hat{\mathbf{s}}$, the risk difference between model $h_{\hat{\mathbf{s}}}$ and $h_{\mathcal{D}_S}$ is bounded by

$$|L(h_{\hat{\mathbf{s}}}(\mathbf{x}), y) - L(h_{\mathcal{D}_S}(\mathbf{x}), y)| \leq \mathcal{O}(\delta_s) + \mathcal{O}\left(\frac{1}{\sqrt{|\mathcal{D}_S|}}\right). \quad (8)$$

δ_s is the radius of the cover, and $\mathcal{O}(\delta_s)$ is a polynomial function over δ_s . The problem can be reduced as a K-center problem [11] by optimizing $\mathcal{O}(\delta_s)$. We apply a 2-approximation algorithm [46] to iteratively find optimal samples in $\hat{\mathbf{s}}$ and add to \mathcal{D}_S . Specifically, each optimal sample \mathbf{z}^* is computed as

$$\mathbf{z}^* = \arg \max_{\mathbf{z}_i \in \hat{\mathbf{s}} \setminus \mathcal{D}_S} \min_{\mathbf{z}_j \in \mathcal{D}_S} \|\mathbf{f}(\mathbf{x}_i) - \mathbf{f}(\mathbf{x}_j)\|_2, \quad (9)$$

where $\mathbf{z} = (\mathbf{x}, y)$, and $\mathbf{f}(\mathbf{x})$ represents the descriptor of an image \mathbf{x} . Equation 9 relates to the furthest point sampling method [8], which enables the most representative samples from a dataset to be selected. We summarize the pruning process in Algorithm 2.

Algorithm 1 Target-specific Subset Search

```
1: Input: Source pool  $\mathcal{S}$ , target set  $\mathcal{D}_T$ , and the number  
   of clusters  $J$ .  
2: Begin:  
3: Cluster  $(\mathcal{S}, J) \rightarrow \{\mathbf{S}^1, \dots, \mathbf{S}^J\}$   
4:  $\mathbf{s} = \emptyset, \epsilon = \infty, \mathbf{S}^* = \emptyset$   
5:  $\mathcal{J} = \text{argsort}(\{\text{FID}(\mathbf{S}^j, \mathcal{D}_T)\}_{j \in [J]})$   $\triangleright$  Ascending  
6: for  $j$  in  $\mathcal{J}$  do  
7:    $\mathbf{s} = \mathbf{s} \cup \mathbf{S}^j$   
8:   if  $\text{FID}(\mathbf{s}, T) < \epsilon$  then  
9:      $\epsilon = \text{FID}(\mathbf{s}, \mathcal{D}_T)$   
10:     $\mathbf{S}^* = \mathbf{s}$   
11: return  $\mathbf{S}^*$ 
```

4.5. Discussion

What to prune first, ID or image? In our design, we select the IDs first and then images, aiming to build a dataset that has a small domain gap to target and is meanwhile in small-scale. Selecting the images first and then IDs is possible. However, it leads to a significant increase in the time complexity. As shown in Algorithm 2, the most time-consuming part of our algorithm is the FPS for finding image core sets, which takes $\mathcal{O}(m|\hat{\mathbf{s}}|)$. If we select images first, the time complexity becomes $\mathcal{O}(m|\mathbf{S}^*|)$, which is significantly higher than $\mathcal{O}(m|\hat{\mathbf{s}}|)$ as $|\mathbf{S}^*| \gg |\hat{\mathbf{s}}|$.

Is SnP applicable to tasks beyond object re-ID? It is possible to apply our search and pruning (SnP) framework to other tasks, *e.g.*, classification. Yet, achieving this requires the SnP algorithm to be partially redesigned. The reason is that, in object re-ID, the major domain gap comes from the class difference, *i.e.*, IDs. However, such a gap does not exist in the classification tasks as the train/val/test datasets always share the same classes. Therefore, our current design of selecting similar IDs for similar distributions is not directly applicable, and has to be adapted.

5. Experiment

We evaluate the effectiveness of SnP on person re-ID and vehicle re-ID. In both tasks, given target data, we use the SnP pipeline to find a training set that has a similar distribution to the target and simultaneously meets a budget.

5.1. Source and Target Datasets

Person re-ID. We create the source pool for person re-ID using 10 public datasets, including Market [54], Duke [55], MSMT17 [45] (denoted as MSMT), CUHK03 [19], RAiD [6], PersonX [37], UnrealPerson [52], RandPerson [43], PKU-Reid [28] and VIPeR [4]. Those datasets cover both synthetic and real-world data. Ten re-ID datasets constitute a source pool that contains in total of 15,060 IDs and 399,715 images.

Algorithm 2 Budgeted Pruning

```
1: Input: Initial source set  $\mathbf{S}^*$ , budget  $b = (n, m)$ .  
2: Begin:  
3:  $\mathbf{y} = \mathcal{U}(\mathcal{Y}(\mathbf{S}^*), n)$   $\triangleright$  Sample  $n$  IDs  
4:  $\hat{\mathbf{s}} = \{\mathbf{s}(y_i) | y_i \in \mathbf{y}\}$   
5:  $\mathcal{D}_S = \{\mathcal{U}(\mathbf{s}(y^i), 1) | y_i \in \mathbf{y}\}$   
6: if  $|\hat{\mathbf{s}}| > m$  then  $\triangleright$  Sample  $m$  images  
7:   repeat  
8:      $\mathbf{z}^* = \arg \max_{\mathbf{z}_i \in \hat{\mathbf{s}} \setminus \mathcal{D}_S} \min_{\mathbf{z}_j \in \mathcal{D}_S} \|\mathbf{x}_i - \mathbf{x}_j\|_2$   
9:      $\mathcal{D}_S = \mathcal{D}_S \cup \{\mathbf{z}^*\}$   
10:    until  $|\mathcal{D}_S| = m$   
11:  else  
12:     $\mathcal{D}_S = \hat{\mathbf{s}}$   
13: return  $\mathcal{D}_S$ 
```

We use two real-world datasets as targets: AlicePerson [1] and Market [54]. Specifically, AlicePerson is specially designed for domain adaptation as it contains unlabeled training images. Note that in Market, the label for its training set is not be used. When Market is the target, Market training set is excluded from the source pool.

Vehicle re-ID. For vehicle re-ID, we create the source pool by combining 8 datasets, which are VeRi [22], CityFlow [41], VehicleID [21], VeRi-Wild [25], VehicleX [50], Stanford Cars [18], PKU-vd1 [47] and PKU-vd2 [47]. It totally has 156,512 IDs and 1,284,272 images.

AliceVehicle [1] and VeRi [22] are separately used as the target domains. Similar to AlicePerson, AliceVehicle is also designed for domain adaptation. When VeRi is the target, the VeRi training set is excluded from the source pool.

Evaluation protocol. For object re-ID, we use mean average precision (mAP) and cumulative match curve (CMC) scores to measure system accuracy, *e.g.*, “rank-1” and “rank-5”. “rank-1” denotes the success rate of finding the true match in the first rank, and “rank-5” means the success rate of ranking at least one true match within top-5 matches.

5.2. Results

SnP framework vs. random sampling and greedy sampling. Given a target domain, the SnP framework allows us to construct a budgeted dataset with a similar distribution. We show the superiority of SnP framework in Table 1. After sampling target-specific data with SnP, we train the subsequent re-ID model with the sampled target-specific data only. In Table 1, we compare the sampled dataset with those created by greedy sampling and random sampling. Random sampling means we randomly select IDs according to the uniform distribution. For greedy sampling, we reproduce [48]. Specifically, We assign each ID a score, which is calculated using FID. IDs will then be selected greedily from the lowest FID value to the highest FID value.

Table 1. Comparing different methods in training data search: SnP, random sampling, and greedy sampling. We set the budget as 2%, 5%, and 20% of the total source IDs. We use four targets: AlicePerson, Market, AliceVehicle and VeRi. The task model is IDE [53]. FID, rank-1 accuracy (%), and mAP (%) are reported.

Training data		Person re-ID targets						Vehicle re-ID targets						
		AlicePerson			Market			AliceVehicle			VeRi			
		FID↓	R1↑	mAP↑	FID↓	R1↑	mAP↑	FID↓	R1↑	mAP↑	FID↓	R1↑	mAP↑	
Source pool		81.67	38.96	17.62	37.53	55.55	30.62	43.95	30.47	14.64	24.39	55.90	25.03	
Searched		60.95	48.19	25.51	30.42	61.49	34.40	23.44	46.78	25.46	17.92	74.13	41.71	
Pruned	2% IDs	Random	80.06	23.67	9.80	39.03	40.77	19.54	45.97	31.35	12.87	24.68	67.16	26.48
		Greedy [48]	61.78	33.22	15.10	34.43	42.19	19.35	38.51	31.82	12.91	29.53	66.57	26.24
		SnP	60.42	38.23	18.17	31.29	44.80	21.65	23.48	38.48	17.79	17.70	69.96	31.10
	5% IDs	Random	81.41	33.16	14.49	39.65	47.39	23.97	44.52	36.36	14.17	25.27	70.38	30.44
		Greedy [48]	61.01	44.63	22.81	31.63	49.17	24.77	32.48	41.32	17.77	26.06	71.23	32.53
		SnP	60.64	47.26	25.45	30.37	51.96	26.56	23.92	44.58	21.79	18.09	72.05	36.01
	20% IDs	Random	79.33	38.10	17.79	38.63	53.15	28.39	43.90	40.89	18.13	24.43	68.71	34.10
		Greedy [48]	63.15	46.74	22.65	32.42	53.53	28.19	24.15	44.58	22.82	18.74	71.04	38.07
		SnP	61.87	47.20	25.36	30.58	57.14	33.09	23.47	46.07	25.24	17.93	73.48	40.75

Table 2. The effectiveness of the target-specific subset search. mAP (%) and CMC scores are reported. “R1” and “R5” denote rank-1 accuracy (%) and rank-5 accuracy (%), respectively.

Training data	AlicePerson			Training data	AliceVehicle		
	R1	R5	mAP		R1	R5	mAP
Market [54]	32.89	52.54	16.06	VeRi [22]	30.69	43.66	11.05
Duke [55]	23.27	41.13	8.59	CityFlow [41]	23.95	36.00	7.45
MSMT [45]	31.29	52.79	15.08	VehicleID [21]	16.3	29.13	4.73
PersonX [37]	16.08	27.42	6.41	VehicleX [50]	18.85	32.18	8.89
Source pool	38.96	57.48	17.62	Source pool	30.47	48.33	14.64
Searched	48.19	67.57	25.51	Searched	46.78	64.41	25.46

From the results shown in Table 1 and Table 3, we observe that training sets selected using SnP achieve consistently lower domain gap to the target and higher re-ID accuracy than those found by random sampling and greedy sampling. For example, when creating a training set with only 2% source IDs for AlicePerson as the target, SnP results in -19.64 and -1.36 improvement in FID value, and +14.56% and +5.01% improvement in rank-1 accuracy over using random sampling and greedy sampling, respectively. When sampling training set with 2% of the source IDs for AliceVehicle as the target, the rank-1 accuracy improvement is +7.13% and +6.66%, respectively.

Effectiveness of the search step. We analyze this step in Table 1 and Table 2. In this search process that poses no limit on the number of IDs and images, we aim to identify and merge clusters of source identities that exhibit similar distributions with the target. We show the searched datasets contribute to improved accuracy over the entire data pool. For example, for AlicePerson, we observe a +9.23% improvement in rank-1 accuracy over the entire data pool. We further show the searched data is superior to individual training sets in Table 2. For example, when AliceVehicle is our target, we show our searched dataset results in +16.09%, +22.83%, +30.48%, +27.93% higher rank-1 accuracy than VeRi, CityFlow, VehicleID, and VehicleX, respectively.

Of note, accuracy under this application scenario is usu-

Table 3. The superiority of SnP over random sampling and greedy sampling, when different direct transfer and pseudo-label re-ID models are used. We report accuracy when 2% IDs are selected for target AlicePerson. Notations and evaluation metrics are the same as those in the previous table.

Type	Model	Training data	R1	R5	mAP
Direct Transfer	IDE [53]	Random	23.67	42.32	9.80
		Greedy [48]	33.22	54.45	15.10
		SnP	38.23	58.40	18.17
	PCB [38]	Random	24.79	41.60	9.91
		Greedy [48]	29.07	46.01	12.59
		SnP	32.43	49.90	15.18
TransReid [15]	Random	52.04	69.73	28.47	
	Greedy [48]	63.73	80.24	41.88	
	SnP	64.31	80.46	42.74	
Pseudo-labeling	UDA [34]	Random	32.17	54.67	15.32
		Greedy [48]	36.47	52.06	17.34
		SnP	41.41	55.74	20.47
	MMT [13]	Random	35.94	51.91	17.25
		Greedy [48]	38.64	56.48	21.18
		SnP	43.36	60.38	23.34

ally lower than that produced by in-distribution training sets. This difference is understandable, because searched data have a relatively lower resemblance to the target data compared with in-distribution training sets. That said, annotating in-distribution training sets is usually expensive, especially considering the complex, specific and ever-changing target environments, where creating a training set on-the-fly with good performance is of practical value.

Effectiveness of the pruning step is analyzed in Table 1 and Fig. 7. Pruning aims to find a subset that has no more than n IDs and m images. From Table 1 and Fig. 7, admittedly, the pruning of both IDs and images will lead to an accuracy decrease. For example, when the target is AlicePerson, if we select 2% IDs, there is a -9.23% decrease in rank-1 accuracy. From Fig. 7, when the target is AlicePerson, if we further select 40% images, there is a -9.96% decrease in rank-1 accuracy. However, even though only 2% IDs are used, rank-1 accuracy obtained from the pruned

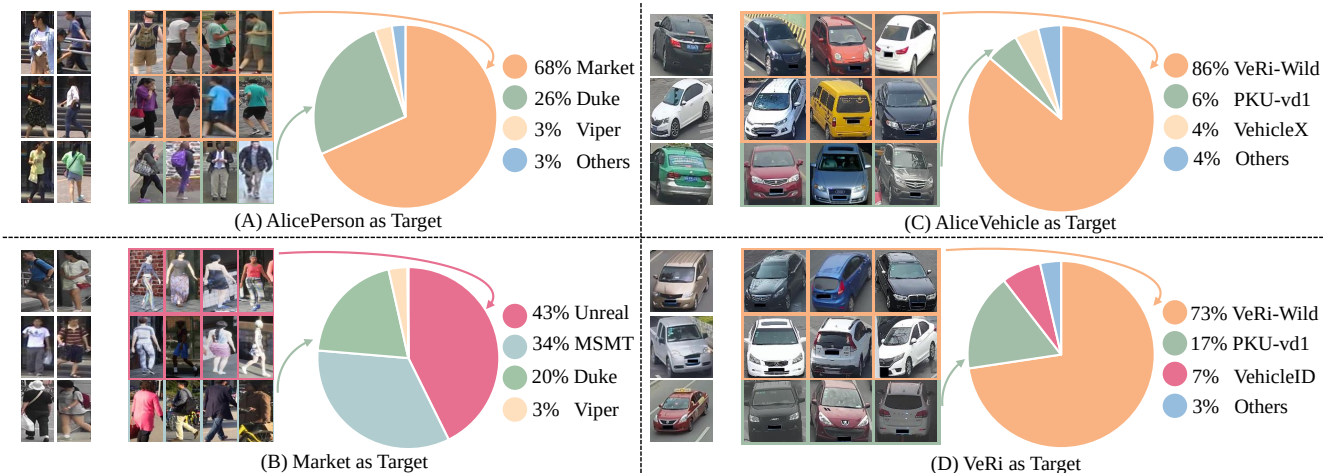


Figure 6. Composition statistics and images samples of searched training sets. For each subfigure (A), (B), (C) and (D), the columns on the left presents unlabeled target samples; columns the middle provide samples of the searched training set; the pie chart Community-verified icon Verified on the right shows composition statistics of the searched training set.

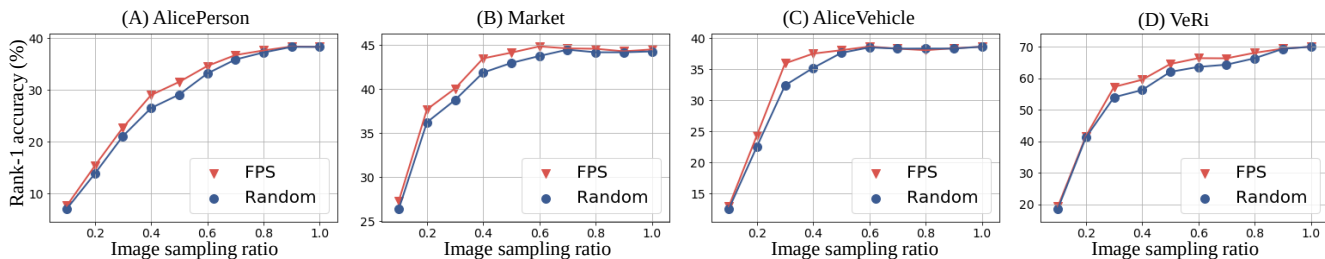


Figure 7. Comparing FPS with random sampling. The search step provides us with 2% of the source IDs. We further select various ratios of images from these IDs. Four different targets are used, and rank-1 accuracy is the evaluation metric. Score is averaged over three runs.

training data is still competitive over the source pool, which is just -0.73% lower than the source pool. It shows the pruning method significantly reduces the training set scale while being able to train a model of reasonable accuracy.

Composition of searched training sets. We visualize four examples in in Fig. 6. It is clear that searched training sets have different compositions under different targets. If we use AlicePerson as the target, in the searched training set, images from Market and Duke take up 68% and 26%, respectively. In comparison, when using Market as the target, the resulting training set contains, 43% images from UnrealPerson, and 34% from MSMT. An interesting observation is that synthetic data (UnrealPerson) has a major role under Market as the target. It demonstrates the potential use of synthetic data for real-world target domains.

Comparison between FPS and random sampling. Both can be used for sampling images resulting from the search step. In Fig. 7, we sample different ratios of the images resulting from the search step. We observe FPS is consistently superior to random sampling, under different selection ratios and targets. Improvement of FPS over random sampling first increase and then decreases, with peak

improvement happening at the 30-60% ratio.

6. Conclusion

This paper studies the training data set search problem, for object re-ID applications. Under a certain budget, we aim to find a target-specific training set that gives a competitive re-ID model. We show our method is overall superior to existing strategies such as random sampling and greedy sampling in terms of accuracy on the target domain. We analyze various components in the SnP system and find them to be stable under various source pools and targets. We also point out the correlation between domain gap, dataset size, and training set quality, and would like to further study the data-centric problems in the community.

7. Acknowledgement

This work was supported in part by the ARC Discovery Early Career Researcher Award (DE200101283), the ARC Discovery Project (DP210102801), Oracle Cloud credits, and related resources provided by Oracle for Research. We thank Dr. Lei Huan for her help in general in paper editing.

References

- [1] Alice benchmark. https://competitions.codalab.org/competitions/27739?secret_key=289a61c4-4566-4431-89e2-be09ab16b848. 1, 4, 6
- [2] Yan Bai, Jile Jiao, Wang Ce, Jun Liu, Yihang Lou, Xuetao Feng, and Ling-Yu Duan. Person30k: A dual-meta generalization network for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2123–2132, 2021. 1, 2
- [3] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010. 4
- [4] Yeong-Jun Cho and Kuk-Jin Yoon. Pamm: Pose-aware multi-shot matching for improving person re-identification. *IEEE Transactions on Image Processing*, 27(8):3739–3752, 2018. 6
- [5] Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4, 2009. 4
- [6] Abir Das, Anirban Chakraborty, and Amit K Roy-Chowdhury. Consistent re-identification in a camera network. In *European conference on computer vision*, pages 330–345. Springer, 2014. 6
- [7] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [8] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y Zeevi. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9):1305–1315, 1997. 5
- [9] Ehsan Elhamifar, Guillermo Sapiro, Allen Yang, and S Shankar Sasrty. A convex optimization framework for active learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 209–216, 2013. 2
- [10] Hehe Fan, Liang Zheng, Chenggang Yan, and Yi Yang. Unsupervised person re-identification: Clustering and fine-tuning. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 14(4):1–18, 2018. 1, 2
- [11] Reza Zanjirani Farahani and Masoud Hekmatfar. *Facility location: concepts, models, algorithms and case studies*. Springer Science & Business Media, 2009. 5
- [12] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017. 2
- [13] Yixiao Ge, Dapeng Chen, and Hongsheng Li. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. In *Proceedings of the International Conference on Learning Representations*, 2020. 7
- [14] Yuhong Guo. Active instance sampling via matrix partition. *Advances in Neural Information Processing Systems*, 23, 2010. 2
- [15] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 7
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, 2017. 2, 3, 4
- [17] Pirazh Khorramshahi, Amit Kumar, Neehar Peri, Sai Saketh Rambhatla, Jun-Cheng Chen, and Rama Chellappa. A dual path model with adaptive attention for vehicle re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [18] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. 6
- [19] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014. 6
- [20] Aristidis Likas, Nikos Vlassis, and Jakob J Verbeek. The global k-means clustering algorithm. *Pattern recognition*, 36(2):451–461, 2003. 5
- [21] Hongye Liu, Yonghong Tian, Yaowei Yang, Lu Pang, and Tiejun Huang. Deep relative distance learning: Tell the difference between similar vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 6, 7
- [22] Xinchen Liu, Wu Liu, Huadong Ma, and Huiyuan Fu. Large-scale vehicle re-identification in urban surveillance videos. In *The IEEE International Conference on Multimedia and Expo*, 2016. 3, 6, 7
- [23] Yuchi Liu, Zhongdao Wang, Tom Gedeon, and Liang Zheng. How to synthesize a large-scale and trainable micro-expression dataset? In *Proceedings of the European Conference on Computer Vision*, 2022. 2
- [24] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of International Conference on Machine Learning*, pages 97–105, 2015. 2, 4
- [25] Yihang Lou, Yan Bai, Jun Liu, Shiqi Wang, and Lingyu Duan. Veri-wild: A large dataset and a new method for vehicle re-identification in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3235–3243, 2019. 2, 6
- [26] Yihang Lou, Yan Bai, Jun Liu, Shiqi Wang, and Ling-Yu Duan. Embedding adversarial learning for vehicle re-identification. *IEEE Transactions on Image Processing*, 28(8):3794–3807, 2019. 2
- [27] Chuanchen Luo, Chunfeng Song, and Zhaoxiang Zhang. Generalizing person re-identification by camera-aware invariance learning and cross-domain mixup. In *European*

- Conference on Computer Vision*, pages 224–241. Springer, 2020. 1, 2
- [28] Liqian Ma, Hong Liu, Liang Hu, Can Wang, and Qianru Sun. Orientation driven bag of appearances for person re-identification. *arXiv preprint arXiv:1605.02464*, 2016. 6
- [29] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. Improving the fisher kernel for large-scale image classification. In *Proceedings of the European Conference on Computer Vision*, 2010. 2
- [30] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proceedings of the European Conference on Computer Vision*, 2010. 2
- [31] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. 5
- [32] Burr Settles. Active learning literature survey. 2009. 2
- [33] Jifei Song, Yongxin Yang, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Generalizable person re-identification by domain-invariant mapping network. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 719–728, 2019. 1, 2
- [34] Liangchen Song, Cheng Wang, Lefei Zhang, Bo Du, Qian Zhang, Chang Huang, and Xinggang Wang. Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognition*, 102:107173, 2020. 7
- [35] Liangchen Song, Yonghao Xu, Lefei Zhang, Bo Du, Qian Zhang, and Xinggang Wang. Learning from synthetic images via active pseudo-labeling. *IEEE Transactions on Image Processing*, 2020. 1, 2
- [36] Xiaoxiao Sun, Yunzhong Hou, Weijian Deng, Hongdong Li, and Liang Zheng. Ranking models in unlabeled new environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11761–11771, 2021. 2
- [37] Xiaoxiao Sun and Liang Zheng. Dissecting person re-identification from the viewpoint of viewpoint. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 6, 7
- [38] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *European Conference on Computer Vision*, 2018. 7
- [39] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [40] Zheng Tang, Milind Naphade, Stan Birchfield, Jonathan Tremblay, William Hodge, Ratnesh Kumar, Shuo Wang, and Xiaodong Yang. Pamtri: Pose-aware multi-task learning for vehicle re-identification using highly randomized synthetic data. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. 2
- [41] Zheng Tang, Milind Naphade, Ming-Yu Liu, Xiaodong Yang, Stan Birchfield, Shuo Wang, Ratnesh Kumar, David Anastasiu, and Jenq-Neng Hwang. Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 3, 6, 7
- [42] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Proceedings of the European Conference on Computer Vision*, 2011. 2
- [43] Yanan Wang, Shengcai Liao, and Ling Shao. Surpassing real-world source training data: Random 3d characters for generalizable person re-identification. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3422–3430, 2020. 2, 6
- [44] Zhongdao Wang, Luming Tang, Xihui Liu, Zhuliang Yao, Shuai Yi, Jing Shao, Junjie Yan, Shengjin Wang, Hongsheng Li, and Xiaogang Wang. Orientation invariant feature embedding and spatial temporal regularization for vehicle re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017. 2
- [45] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 79–88, 2018. 1, 6, 7
- [46] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011. 5
- [47] Ke Yan, Yonghong Tian, Yaowei Wang, Wei Zeng, and Tiejun Huang. Exploiting multi-grain ranking constraints for precisely searching visually-similar vehicles. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 562–570, 2017. 2, 6
- [48] Xi Yan, David Acuna, and Sanja Fidler. Neural data server: A large-scale search engine for transfer learning data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3893–3902, 2020. 2, 6, 7
- [49] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127, 2015. 2
- [50] Yue Yao, Liang Zheng, Xiaodong Yang, Milind Naphade, and Tom Gedeon. Simulating content consistent vehicle datasets with attribute descent. In *Proceedings of the European Conference on Computer Vision*, 2020. 2, 6, 7
- [51] Yue Yao, Liang Zheng, Xiaodong Yang, Milind Naphade, and Tom Gedeon. Attribute descent: Simulating object-centric datasets on the content level and beyond. *arXiv preprint arXiv:2202.14034*, 2022. 2
- [52] Tianyu Zhang, Lingxi Xie, Longhui Wei, Zijie Zhuang, Yongfei Zhang, Bo Li, and Qi Tian. Unrealperson: An adaptive pipeline towards costless person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11506–11515, 2021. 2, 6
- [53] 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. In *European Conference on Computer Vision*, 2016. 7

- [54] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. [1](#), [3](#), [6](#), [7](#)
- [55] Zhedong Zheng, Liang Zheng, and Yi Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3754–3762, 2017. [1](#), [3](#), [6](#), [7](#)
- [56] Zhun Zhong, Liang Zheng, Zhiming Luo, Shaozi Li, and Yi Yang. Invariance matters: Exemplar memory for domain adaptive person re-identification. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [1](#), [2](#)
- [57] Yi Zhou and Ling Shao. Aware attentive multi-view inference for vehicle re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. [2](#)