

# Domain Adaptive Person Search

Junjie Li<sup>1,2\*</sup>, Yichao Yan<sup>1†</sup>, Guanshuo Wang<sup>2</sup>, Fufu Yu<sup>2</sup>,  
Qiong Jia<sup>2</sup>, Shouhong Ding<sup>2</sup>

<sup>1</sup>MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

<sup>2</sup>Tencent Youtu Lab

serenitycapo@gmail.com, yanyichao@sjtu.edu.cn,  
{mediswang, fufuyu, boajia, ericshding}@tencent.com

**Abstract.** Person search is a challenging task which aims to achieve joint pedestrian detection and person re-identification (ReID). Previous works have made significant advances under fully and weakly supervised settings. However, existing methods ignore the generalization ability of the person search models. In this paper, we take a further step and present Domain Adaptive Person Search (DAPS), which aims to generalize the model from a labeled source domain to the unlabeled target domain. Two major challenges arise under this new setting: one is how to simultaneously solve the domain misalignment issue for both detection and ReID tasks, and the other is how to train the ReID subtask without reliable detection results on the target domain. To address these challenges, we propose a strong baseline framework with two dedicated designs. 1) We design a domain alignment module including image-level and task-sensitive instance-level alignments, to minimize the domain discrepancy. 2) We take full advantage of the unlabeled data with a dynamic clustering strategy, and employ pseudo bounding boxes to support ReID and detection training on the target domain. With the above designs, our framework achieves 34.7% in mAP and 80.6% in top-1 on PRW dataset, surpassing the direct transferring baseline by a large margin. Surprisingly, the performance of our unsupervised DAPS model even surpasses some of the fully and weakly supervised methods. The code is available at <https://github.com/caposerenity/DAPS>.

**Keywords:** Person Search, Domain Adaptation

## 1 Introduction

Person search [44,39] aims to detect and identify the query person from natural images. The mainstream approaches to tackling this task is to simultaneously address both tasks in an end-to-end manner, where supervised learning [44,34,6,26] that rely on both pedestrian bounding boxes annotation and identity labels have been actively investigated. However, these supervised methods may suffer from significant performance degradation on unseen domains due to domain gaps.

---

\*This work was done during Li’s internship at Tencent Youtu Lab.

†Corresponding author: Yichao Yan.



Fig. 1: Comparison of three person search settings. (a) Fully supervised setting: bounding boxes and identity annotations are available. (b) Weakly supervised setting: only bounding boxes annotations are available. (c) Domain adaptive setting: neither bounding boxes nor identity annotations on the target domain is accessible, and there exists obvious domain gaps between different domains, *e.g.*, the size of human crops. The network is trained with both the labeled source domain and the unlabeled target domain images.

To address this problem, several recent works [40,20] propose the weakly supervised person search (WSPS) setting without accessible ID annotations, shown in Fig. 1. Nevertheless, several limitations are still waiting to be addressed. First, these works still require manual annotation of the ground-truth bounding boxes for the detection task, which obviously is not an economical option for real-world applications. Second, there exist several large-scale annotated person search datasets, *e.g.*, CUHK-SYSU [39] and PRW [44], which can serve as supervised source domains and help improve the performance on the unlabeled target data. Unfortunately, the weakly supervised setting does not fully unleash the potential of the available training data. Third, these methods adopt an inconsistent training strategy with supervised detection and unsupervised ReID, which ignores the essential correlation between the two sub-tasks.

Inspired by the unsupervised domain adaptation (UDA) [16,23,36], as shown in Fig. 1, we present the Domain Adaptive Person Search (DAPS) framework, where person search models trained on labeled source domain are transferred to unlabeled target domains. Compared to weakly supervised person search, neither the identity labels nor the bounding boxes are accessible in DAPS. Our framework faces two major challenges: (1) Both the detection and the ReID sub-tasks suffer from domain gap. However, detection focuses on the commonness of people regardless of the identities, while ReID needs to learn the uniqueness of different persons. This conflict can be more serious in domain adaptation. (2) Since the ground-truth detection boxes are not available, it will be extremely challenging to accurately localize the pedestrians in the target domain, which further increases the difficulty for the ReID sub-task. Therefore, directly extending WSPS methods to take advantage of target domain data is infeasible.

To address the first challenge, we explore domain alignment for robust domain invariant feature learning. In the context of pedestrian detection, this is typically achieved by domain adversarial training [8] on both image-level and instance-level features. Following this line of research, we design a domain align-

ment module (DAM) to alleviate the discrepancy between different domains. Specifically, on the one hand, we introduce domain discriminators at intermediate backbone layers. On the other hand, we perform a task-sensitive instance-level alignment to mitigate the conflicts between two sub-tasks. We observe that such a domain alignment operation is beneficial for both branches.

To tackle the second challenge, we generate pseudo bounding boxes on the target domain images iteratively, and perform the training process with GT and pseudo boxes for domain adaptation. Furthermore, we present a dynamic clustering strategy to generate pseudo identity labels on the target domain. To fully release the potential of the target domain training data, the proposed framework refines the detection task with selected proposals, and enhances the interaction between the two sub-tasks with hybrid hard case mining. Experimental results demonstrate that this design surprisingly achieves comparative performance with directly adopting ground-truth bounding boxes.

Our contributions are summarized as three-fold:

- We introduce a novel unsupervised domain adaptation paradigm for person search. This setting requires neither bounding boxes nor identity annotations on the target domain, making it more practical for real-world applications.
- We present the DAPS framework to overcome the challenges caused by cross-domain discrepancy and cross-task dependency. We propose domain alignment for person search to enhance domain-invariant feature learning. Meanwhile, a dynamic clustering and a hybrid hard case mining strategy are introduced to facilitate unsupervised target domain learning.
- Without any auxiliary label in the target domain, our framework achieves promising performance on two target person search benchmarks, surprisingly outperforming several weakly and fully supervised models.

## 2 Related Work

### 2.1 Person Search

With the development of deep learning and large scale benchmarks [39,44], person search [4] has recently become a popular research topic. Existing fully supervised person search models can be divided into two-step and one-step frameworks. Two-step frameworks typically consist of separately trained detection and ReID models [34,21]. Zheng *et al.* [44] make a systematic evaluation on different combination of detection and ReID models. Wang *et al.* [34] solve the inconsistency between detection and person ReID tasks. One-step frameworks [6,26,41] design a unified model to jointly solve detection and ReID tasks in an end-to-end manner, making the pipeline more efficient. Yan *et al.* [43] introduce a graph model to explore the impact of contextual information for identity matching. Chen *et al.* [6] disentangle the person representation into norm and angle to eliminate the cross-task conflict. Li *et al.* [26] develop a sequential structure to reduce the low-quality proposals. Several recently studies [40,20] adopt the weakly supervised setting without no accessible person ID labels. In this work,

we explore a novel person search setting to generalize labeled source to unlabeled target domain without any bounding boxes and ID labels annotation.

## 2.2 Domain Adaptation for Person ReID

Unsupervised domain adaptation (UDA) ReID [7,10,28,15,17,32,42] typically trains a model with labelled source domain and transfers to the target domain under the unsupervised setting. Mainstream UDA ReID methods can be divided into two categories. The first category employs generative adversarial networks [19] to mitigate the style discrepancy and translate the labelled source domain data into the target domain [7,10,28]. For the second category, they generate pseudo labels by clustering [15,17,32] or assigning soft labels [35] on target domain, and use these pseudo labels to further supervise target domain training. Recently, pseudo label-based methods raise more attention due to their superior performance. However, UDA ReID requires the cropped images, which cannot be directly extended to adaptive person search due to the lack of bounding boxes on target domain. To address this, we propose a dynamic clustering strategy to generate high-quality pseudo boxes to facilitate target domain training.

## 2.3 Domain Adaptive Object Detection

Existing Domain Adaptive approaches object detection can be categorized into three main branches, including adversarial-based methods [33,8,45,37,31], discrepancy-based methods [24,2,3] and reconstruction-based methods [27,1,11]. Adversarial-based methods utilize a domain discriminator to distinguish the domain of input data, the adversarial training is performed to encourage domain confusion between the source domain and the target domain. The discrepancy-based strategy utilizes the unlabeled target domain images to fine-tune the detector, and further followed by mean-teacher learning [2] or auto-annotation [3]. The reconstruction-based approaches bridge the domain gap by reconstructing the source or target samples, which is usually realized by image-to-image translation [1,27]. In this work, we consider the conflicts between sub-tasks of person search, and develop a task-sensitive alignment module to alleviate such conflicts.

# 3 Methodology

## 3.1 Framework Overview

The general pipeline of the proposed DAPS framework is illustrated in Fig. 2. Given the input images from both the source and the target domain, the image-level feature maps are extracted with a backbone network. Then, these features are input into the Region Proposal Network (RPN) to generate candidate bounding boxes, which are subsequently fed into the ROI-Align layer to represent instance-level feature maps. To close the domain gaps for the downstream detection and ReID tasks, we design a domain alignment module (DAM) to align both image-level and instance-level features from different domains.

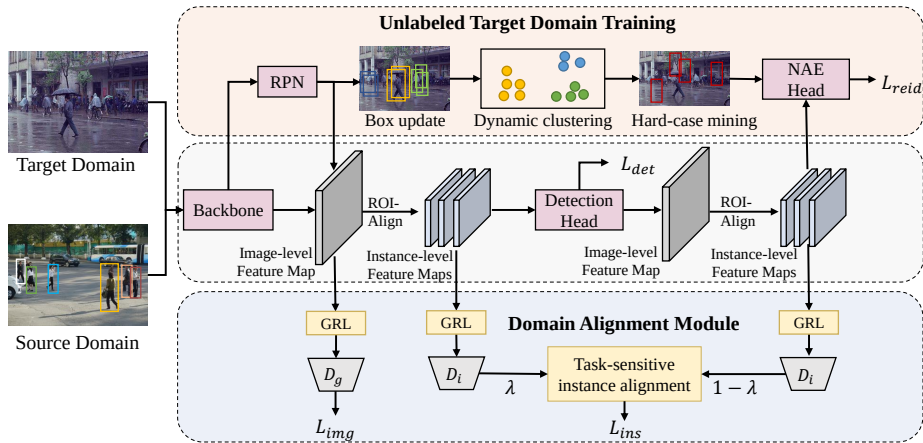


Fig. 2: Architecture of the DAPS framework. “GRL” denotes the gradient reverse layer [16]. The backbone follows SeqNet [26], and we employ a domain alignment module to minimize domain discrepancy on both image-level and instance-level. We further impose dynamic clustering, hybrid hard case mining and target detection training to take full advantage of the unlabeled target domain data.

Subsequently, the domain-aligned instance-level feature maps are input into both the detection and the ReID branch. Since the ground-truth bounding boxes are not available in the target domain, the model will generate different pedestrian detection results for each training epoch. Therefore, it is infeasible to follow the traditional UDA ReID methods, which generally perform clustering on a fixed size of instances to generate pseudo labels. To address this issue, we design a novel dynamic clustering strategy, which continuously associates the bounding boxes generated from consecutive epochs, to guarantee the stability of instance-level ReID features. Based on dynamic clustering strategy, we further introduce the hybrid hard case mining and the target domain detection refinement to sufficiently take advantage of the unlabeled training data.

### 3.2 Domain Alignment Module

**Image-level Alignment.** As discussed in [36,8,7,10], minimizing domain discrepancy is beneficial for both sub-tasks of person search, and an effective way is to guide the model to learn domain-invariant representation. Motivated by the recent progress in domain adaptive detectors [8,45,37,31], where intermediate features are imposed with image-level alignment constraints, we introduce a domain alignment module into our DAPS framework. As shown in Fig. 2, DAM employs a patch-based domain classifier to predict the domain where the input feature comes from. A min-max formulation is adopted to misdirect the domain classifier and encourage domain-invariant representation learning.




Fig. 3: Details of the two heads and the task-sensitive instance-level alignment.

Suppose we have  $N$  training images  $\{I_1, \dots, I_N\}$  with corresponding domain labels  $\{d_1, \dots, d_N\}$ . Particularly,  $d_i = 0$  indicates that image  $I_i$  comes from the source domain, while  $d_i = 1$  denotes the target domain. We denote the backbone of DAPS as  $\Phi$  and the image-level domain classifier as  $D_g$ , and further represent the domain prediction result of input  $I_i$  as  $p_i$ . We apply a cross entropy loss to perform domain alignment in an adversarial training manner:

$$\mathcal{L}_{img} = - \sum_i [d_i \log p_i + (1 - d_i) \log (1 - p_i)]. \quad (1)$$

We have tried to conduct image-level alignment on different intermediate features and multi-scale alignment, but achieve no better results.

**Task-sensitive Instance-level Alignment.** As illustrated in Fig. 3, our framework consists of two head networks, where the detection performance mainly depends on the first standard Faster R-CNN [18] head, while the NAE [6] head is highly relevant to ReID. When the scale of the source domain is much smaller than the unlabeled target, the target pseudo bounding boxes predicted by detector trained on the source can be severely overfitted to the smaller domain, but no reliable target detection guidance can relieve this issue. When the target is much smaller, pseudo target ID labels can be easily obtained by clustering, but these might provide insufficient generalizing for the ReID sub-task.

According to the characteristics of the up- and down-stream tasks, we propose the task-sensitive instance-level alignment module by balancing the alignment weight on instance-level features for both sub-tasks. Suppose we have  $K_1$  instances in the standard head and  $K_2$  instances for the NAE head, two domain classifiers  $\{D_i^d, D_i^r\}$  are built in the same way with image-level alignment, and the domain predictions of the two local classifiers are denoted as  $\{p_{i,1}^d, \dots, p_{i,K_1}^d\}$ ,  $\{p_{i,1}^r, \dots, p_{i,K_2}^r\}$ , respectively. The instance-level loss can be formulated as:

$$\begin{aligned} \mathcal{L}_{ins} = & - \lambda \sum_{i,j} [d_i \log p_{i,j}^d + (1 - d_i) \log (1 - p_{i,j}^d)] \\ & - (1 - \lambda) \sum_{i,k} [d_i \log p_{i,k}^r + (1 - d_i) \log (1 - p_{i,k}^r)]. \end{aligned} \quad (2)$$

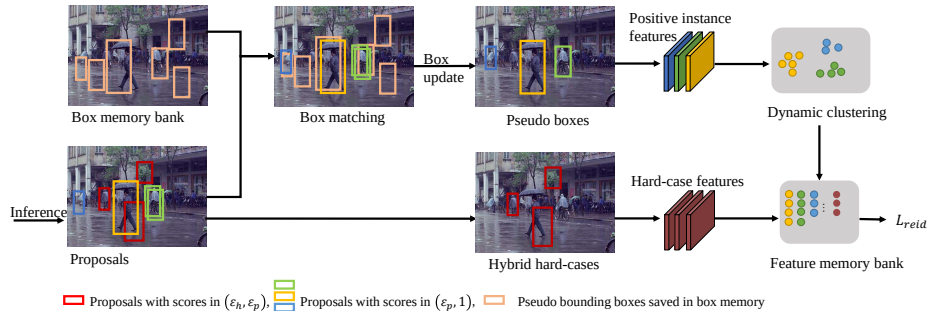


Fig. 4: Illustration of the dynamic clustering and hard case mining. At start of each epoch, we employ generated proposals, including both qualified ones and hard cases, to update the memory bank. Qualified proposals are adopted for matching pseudo boxes memory, and hard cases will directly be added.

where  $j \in \{1, \dots, K_1\}$ , and  $k \in \{1, \dots, K_2\}$ . The source and target domain contains  $N_s$  and  $N_t$  images respectively, and the balancing factor  $\lambda$  is obtained by

$$\lambda = \sigma \left( 4 \cdot \text{sign}(N_t - N_s) \left( \frac{\max(N_s, N_t)}{\min(N_s, N_t)} - 1 \right) \right). \quad (3)$$

where the  $\sigma(\cdot)$  is Sigmoid function to normalize the domain scale ratio. Moreover, we impose a L2-norm regularizer to ensure the consistency between image-level and instance-level classifiers.

### 3.3 Training on Unlabeled Target Domain

**Dynamic Clustering.** UDA ReID models typically employ the clustering strategy (e.g., DBSCAN) to generate pseudo labels for the target domain instances, and employ memory-based losses [17] for metric learning. However, without ground-truth bounding boxes on target domain, the instances can be only generated from the detection results, which varies with the training process. This makes it infeasible to directly apply typical clustering approach to DAPS. To address this issue, we propose a novel dynamic clustering strategy to make full use of the detection results for continuous ReID training.

As illustrated in Fig. 4, an asynchronous training strategy is introduced to progressively update pseudo bounding boxes with the selected proposals as ground-truth boxes on the target domain. Specifically, for the beginning  $\alpha$  epochs, DAPS is trained only on the source dataset labeled with both bounding boxes and ID labels. After that, we maintain a bounding box memory  $\mathbf{M}_{\mathbf{B}} = \{B_1, \dots, B_{N_t}\}$  and a feature vector memory  $\mathbf{M}_{\mathbf{V}} = \{V_1, \dots, V_{N_t}\}$ , corresponding to each of  $N_t$  target domain images. At the start of each subsequent epochs, DAPS filters out high-confidence candidate proposals  $\{c_1, \dots, c_m\}$  from  $x_i^t$ , and employ them to match pseudo bounding boxes in box memory  $B_i = \{b_1, \dots, b_n\}$  according to IOU scores. Every proposal is assigned to the

most relevant box in memory if their IOU score is above the threshold, and the boxes which fail to match any qualified proposal will be removed from memory  $B_i$ . The remaining boxes in the memory are continuously updated in the Exponential Moving Average (EMA) method.

For example, suppose the proposals  $c_{j1}, c_{j2}, c_{j3}$  are mapped to the box  $b_k$ , then  $b_k$  is updated by:

$$b_k \leftarrow \gamma b_k + (1 - \gamma) \text{avg}(c_{j1}, c_{j2}, c_{j3}), \quad (4)$$

where  $\gamma \in [0, 1]$  controls the update rate. Eventually, the proposals without any matched box will also be fed into the memory  $B_i$ , and further, the feature memory  $\mathbf{M}_{\mathbf{V}}$  is updated in the same way. Afterwards, we perform clustering upon  $\mathbf{M}_{\mathbf{V}}$  to obtain  $N_t^c$  clusters  $\{C_1, \dots, C_{N_t^c}\}$  with centroids  $\mathbf{W} = \{w_1, \dots, w_{N_t^c}\}$ , and  $N_t^o$  instances  $\mathbf{F} = \{f_1, \dots, f_{N_t^o}\}$  not belonging to any cluster. By extracting the identity features  $\mathbf{V}$  in the source domain, we eventually build a unified memory  $\mathbf{M} = \{\mathbf{V}, \mathbf{W}, \mathbf{F}\}$  for ReID training. The loss function can be expressed :

$$\mathcal{L} = -\log \frac{\exp(x \cdot z^+ / \tau)}{\sum_{k=1}^{N_t^c} \exp(x \cdot w_k / \tau) + \sum_{k=1}^{N_t^o} \exp(x \cdot f_k / \tau) + \sum_{k=1}^{N_s^c} \exp(x \cdot v_k / \tau)}, \quad (5)$$

where  $w$ ,  $f$ , and  $v$  represents the target domain clusters, the independent instances and the source domain classes, respectively.  $z^+$  is the corresponding class prototype of the input feature  $x$ , and  $\cdot$  denotes the inner product to measure the feature similarity. The features in the memory will be updated in a momentum way during backward stage:

$$z_t \leftarrow \gamma z_t + (1 - \gamma)x, \quad (6)$$

where  $z_t$  is the  $t$ -th prototype in the memory bank  $\mathbf{M}$ .

**Hybrid Hard Case Mining.** A significant challenge for dynamic clustering is to generate reliable bounding boxes. We treat those boxes with lower confidence than a threshold as negative samples. In order to sufficiently exploit target domain information, we explore the potential of adding these “negative” samples to the ReID training. Proposals with relatively low confidence scores can be divided into highly overlapped with high-confidence boxes, the undetected persons and the background clutters. It is undesirable to enhance the ReID sub-task by treating all these proposals as negative samples. As a result, we design a hierarchical scheme to categorize the candidate proposals, and employ both of the low-confidence person proposals and non-trivial background clutters to enhance the discrimination of the ReID branch.

Specifically, proposals with confidence score in the range of  $(\epsilon_h, \epsilon_p)$  defined by upper and lower bound thresholds are regarded as non-trivial cases. We exclude highly overlapped duplicates by further screening IOUs with positive proposals, while the hybrid of undetected persons and the negative clutters are reserved for training. The features of these hard cases will be added to  $\mathbf{M}$ , and be used for



the contrastive learning process. The memory loss in Eq. 5 is modified as:

$$\begin{aligned} \mathcal{L} &= -\log \frac{\exp(x \cdot z^+ / \tau)}{\sum_{z \in \mathbf{M}} \exp(x \cdot z / \tau)}, \\ \sum_{z \in \mathbf{M}} \exp(x \cdot z / \tau) &= \sum_{k=1}^{N_t^c} \exp(x \cdot w_k / \tau) + \sum_{k=1}^{N_t^o} \exp(x \cdot f_k / \tau) + \\ &\quad \sum_{k=1}^{N_s^c} \exp(x \cdot v_k / \tau) + \sum_{k=1}^{N_t^n} \exp(x \cdot h_k / \tau), \end{aligned} \quad (7)$$

where  $h$  denotes the hybrid hard cases. It is noteworthy that the hybrid hard cases will be involved into the dynamic clustering before the next epoch. Once a hard case is matched with new qualified proposals, it will be treated as a positive sample and updated in a momentum way.

**Target Detection Training.** Although DAM can minimize the domain discrepancy, the over-fitting towards the source domain is still likely to take place, especially when the source domain data is extremely less complex and comprehensive than the target domain images. To this end, simultaneously training detection with both of the source and the target domain data is beneficial for the generalization ability of model. DAM and dynamic clustering provide relatively reliable pseudo bounding boxes, and specifically, we employ such pseudo bounding boxes after the  $\alpha$  epoch to supervise detection on the target domain. In this way, the potential of unlabeled target domain images is released for both ReID and detection training.

## 4 Experiment

### 4.1 Datasets and Evaluation Protocols

**Datasets.** We employ two large-scale benchmark datasets, CUHK-SYSU [39] and PRW [44] in our experiments. CUHK-SYSU is one of the largest public datasets for person search, composed of 18,184 images and 96,143 bounding boxes from 8,432 different identities. It is divided into a training set of 11,206 images with 5,532 identities, and a test set with 6,978 gallery images and 2,900 query images. The widely used PRW dataset contains 11,816 images, 43,110 annotated bounding boxes from 932 identities. The training set includes 5,704 images and 482 labelled persons, while the other 6,112 images and 2,057 probe persons from 450 identities are adopted as test set.

**Evaluation Protocols.** Our experiments employ the default splits for both datasets. For domain adaptation settings, the annotations of dataset used as the source domain is accessible, while neither bounding boxes nor identity labels of datasets as the target domain are available. All evaluations are performed on the test set of target domain. We adopt the widely used mean average precision (mAP) and cumulative matching characteristic (CMC) top-1 accuracy as evaluation metrics for ReID sub-task, while average precision (AP) and recall rate are adopted as the metrics for detection.

Table 1: Comparative results when combining different components. DAM: Domain Alignment Module. DC: Dynamic Clustering. HM: Hybrid hard case Mining. DTD: Detection on Target Domain.

DAM	DC	HM	DTD	Target: PRW				Target: CUHK-SYSU			
				mAP	top-1	recall	AP	mAP	top-1	recall	AP
×	×	×	×	30.3	77.7	94.0	88.3	52.5	54.8	55.2	55.1
✓	×	×	×	30.9	79.3	96.3	90.7	62.2	63.6	70.8	63.1
×	✓	×	×	32.2	79.4	96.8	90.3	70.9	72.3	67.8	62.2
✓	✓	×	×	32.7	79.6	95.9	90.4	72.6	74.3	68.3	63.2
✓	✓	✓	×	34.5	<b>80.7</b>	97.0	91.0	73.2	74.8	70.4	64.1
✓	✓	×	✓	33.1	79.9	96.6	<b>91.2</b>	76.8	78.7	<b>79.4</b>	<b>71.1</b>
✓	✓	✓	✓	<b>34.7</b>	80.6	<b>97.2</b>	90.9	<b>77.6</b>	<b>79.6</b>	77.7	69.9

Table 2: Comparative results of task-sensitive instance-level alignment.

instance da	Target: PRW				Target: CUHK-SYSU			
	mAP	top-1	recall	AP	mAP	top-1	recall	AP
normal	21.7	76.0	<b>96.7</b>	<b>91.1</b>	58.2	60.5	66.3	56.3
task-sensitive	<b>30.9</b>	<b>79.3</b>	96.3	90.7	<b>62.2</b>	<b>63.6</b>	<b>70.8</b>	<b>63.1</b>

## 4.2 Implementation Details

We adopt ResNet50 [22] pretrained on ImageNet-1k [9] as our default backbone network. DBSCAN [14] with self-paced learning strategy [25] is employed as the basic clustering method, we set default hyper-parameters  $\epsilon_p = 0.95$ ,  $\epsilon_h = 0.8$  and  $\lambda_t = 0.1$ . During training, the input images are resized to  $1500 \times 900$ , and random horizontal flip is applied for data augmentation. Our model is optimized by Stochastic Gradient Descent (SGD) for 20 epochs. We set a mini-batch size of 4, and an initial learning rate of 0.0024, which is reduced by a factor of 0.1 at epoch 16 with warmed up in the first epoch. The momentum and weight decay are set to 0.9 and  $5 \times 10^{-4}$ , respectively. We set the momentum factor  $\gamma$  for memory updating to 0.2. The starting epoch of  $\alpha$  is set to 8 when PRW is chosen as target domain, and 0 for CUHK-SYSU. All experiments are implemented with one NVIDIA Tesla A100 GPU. We also plan to support this project with MindSpore in our future work.

## 4.3 Ablation Study

We perform analytical experiments to verify the effectiveness of each detailed component in our proposed framework. In table 1, we compare the baseline method with different combinations of proposed components, and report the results on both CUHK-SYSU and PRW datasets. For example, when we use CUHK-SYSU as the target domain dataset, the directly transferring baseline

Table 3: Comparative results when employing different strategies to handle the lack of bounding boxes. ‘GT’ refers to using the ground truth bounding boxes for all the training process of ReID, and ‘GT for init’ only employs these boxes to initialize the memory bank. ‘static’ means directly employing the qualified proposals before each epoch.

strategy	Target: PRW				Target: CUHK-SYSU			
	mAP	top-1	recall	AP	mAP	top-1	recall	AP
GT	34.9	79.9	94.9	89.5	73.6	76.0	74.6	68.2
GT for init	33.5	79.6	92.9	88.5	73.5	75.4	64.4	60.8
Static	25.3	77.3	96.6	90.8	64.0	66.1	67.6	62.5
Dynamic Update	32.7	79.6	95.9	90.4	72.6	74.3	68.3	63.2

Table 4: Comparative results of when to start asynchronous training.

starting epoch	Target: PRW				Target: CUHK-SYSU			
	mAP	top-1	recall	AP	mAP	top-1	recall	AP
0	31.5	79.7	95.8	89.4	<b>77.6</b>	<b>79.6</b>	<b>77.7</b>	<b>69.9</b>
4	31.4	79.4	95.8	89.1	73.6	75.3	76.6	67.7
8	<b>34.7</b>	<b>80.6</b>	97.2	<b>90.9</b>	73.2	74.7	76.7	69.0
10	33.4	80.6	<b>97.5</b>	90.7	71.4	73.3	74.8	65.8

model achieves 52.5% mAP and 54.8% top-1. After individually adding the domain adaptive module (DAM) and dynamic clustering (DC), the performance improves 9.3% and 18.4% in terms of mAP. When combining DAM and DC, the mAP is further promoted to 72.6%, surpassing the 52.5% of baseline by a large margin. Furthermore, to make full use of the unlabeled target data, we implement the hybrid hard case mining (HM) and detection on target domain (DTD). HM improves the ReID performance by 0.6% in mAP, and DTD prominently enhances the detection branch with a 7.0% gain for AP. Eventually, DAPS achieves 77.6% mAP and 79.6% top-1 with all designed modules, outperforming the baseline by 25.1% in mAP, 24.8% in top-1, 22.5% in recall, and 14.8% in AP.

**Effectiveness of task-sensitive instance-level alignment.** To validate the effectiveness of our task-sensitive instance-level alignment design, we compare it with normal instance-level alignment, which conducts instance alignment on both head networks without balancing between them. As observed in Table 2, the task-sensitive design successfully alleviates the inner task conflicts and outperforms normal strategy by a large margin.

**Effectiveness of dynamic clustering.** As aforementioned, the key to utilizing unlabeled target domain data is generating reliable pseudo bounding boxes. To validate the quality of the pseudo bounding boxes we use, we compare different strategies of obtaining bounding boxes, and the results are reported in Table 3. We first measure the performance achieved by using ground-truth




Fig. 5: Target domain performance with different  $\epsilon_p$  on CUHK-SYSU dataset. (a): ReID accuracy results; (b): Numbers of generated positive proposals. The ground truth instances number is 55,260.

bounding boxes for training the ReID task. Furthermore, we report the performance achieved by directly employing the qualified proposals before each epoch, which is denoted as ‘static’ in Table 3. The results reveal that our proposed dynamic clustering strategy can generate trustworthy pseudo bounding boxes to achieve comparable performance with using ground-truth boxes.

**Effectiveness of asynchronous training.** We conduct experiments for influences by the training stage hyper-parameter  $\alpha$  on final performance. As shown in Table 4, when PRW is adopted as the target domain, the best performance is achieved with  $\alpha = 8$ , while with  $\alpha = 0$  for CUHK-SYSU. The results might be counterintuitive but indeed validate our task-sensitive motivation. For smaller source dataset, even limited additional target information might be helpful for cross-domain generalization. In contrast, for larger source dataset, unreliable target proposals can be harmful for domain gap bridging.

**Analysis on hyper-parameter  $\epsilon_p$ .** We visualize the influence of hyper-parameter  $\epsilon_p$  in Fig. 5. We observe that the value of  $\epsilon_p$  influences the ReID performance to a large extent, and the best performance is achieved with  $\epsilon_p = 0.95$ . From Fig. 5b, it can be observed that the selection of  $\epsilon_p$  is a trade-off between recall rate and proposal quality. Setting it to a extremely high value leads to discarding useful proposals, while a low threshold will induct clutters to undermine the quality of clustering.

#### 4.4 Comparison with State-of-the-Art Methods

Since no existing person search methods with such domain adaptation settings can be directly compared, we further compare DAPS with fully supervised methods in Table 5, including both of the two-step methods and one-step ones. It is surprising that our framework even surpasses some supervised methods. For example, DAPS outperforms MGTS [5], OIM [39], IAN [38], NPSM [29] and

Table 5: Comparison with fully supervised person search models

Method	PRW		CUHK-SYSU	
	mAP	top-1	mAP	top-1
DPM [18]	20.5	48.3	-	-
MGTS [5]	32.6	72.1	83.0	83.7
RDLR [21]	42.9	70.2	93.0	94.2
IGPN [13]	47.2	87.0	90.3	91.4
TCTS [34]	46.8	87.5	93.9	95.1
OIM [39]	21.3	49.9	75.5	78.7
IAN [38]	23.0	61.9	76.3	80.1
NPSM [29]	24.2	53.1	77.9	81.2
CTXGraph [43]	33.4	73.6	84.1	86.5
QEEPS [30]	37.1	76.7	88.9	89.1
HOIM [4]	39.8	80.4	89.7	90.8
BINet [12]	45.3	81.7	90.0	90.7
NAE [6]	44.0	81.1	92.1	92.9
AlignPS [41]	45.9	81.9	93.1	93.4
SeqNet [26]	46.7	83.4	93.8	94.6
DAPS (ours)	<b>34.7</b>	<b>80.6</b>	<b>77.6</b>	<b>79.6</b>

CTXGraph [43] on PRW. The comparison with the state-of-the-art fully supervised methods indicate that there exists a large performance gap, and we hope our work will encourage more explorations for this setting. Moreover, to make measure the theoretical upper limit of DAPS setting, we train some state-of-the-art method with both datasets in a supervised manner, and more details are described in the supplementary material.

The comparisons with existing weakly supervised methods are shown in Table 6, and we also present the results of training R-SiamNet with both datasets in the weakly supervised manner. When evaluated on the PRW dataset, DAPS outperforms all existing weakly supervised methods by a significant margin. For the CUHK-SYSU dataset, DAPS still underperforms the state-of-the-art weakly supervised models, which is mainly caused by the limitation brought by detection capabilities. As mentioned in Sec. 4.1, the images and identities in PRW are prominently fewer than those in CUHK-SYSU, and this further leads to the poor detection performance of adopting CUHK-SYSU as target domain.

#### 4.5 Qualitative Results

To better illustrate the distributions of our hybrid hard cases, we visualize some qualitative results from both datasets in Fig. 6. As is observed, the hybrid hard cases consist of undetected persons (column a), highly overlapped human crops (column b) and background clutters(column c,d). These qualitative results

Table 6: Comparison with weakly supervised person search models. \* denotes training R-SiamNet together with both of CUHK-SYSU and PRW.

Method	PRW		CUHK-SYSU	
	mAP	top-1	mAP	top-1
CGPS [40]	16.2	68.0	80.0	82.3
R-SiamNet [20]	21.4	75.2	86.0	87.1
R-SiamNet* [20]	23.5	76.0	86.2	87.6
DAPS (ours)	<b>34.7</b>	<b>80.6</b>	<b>77.6</b>	<b>79.6</b>

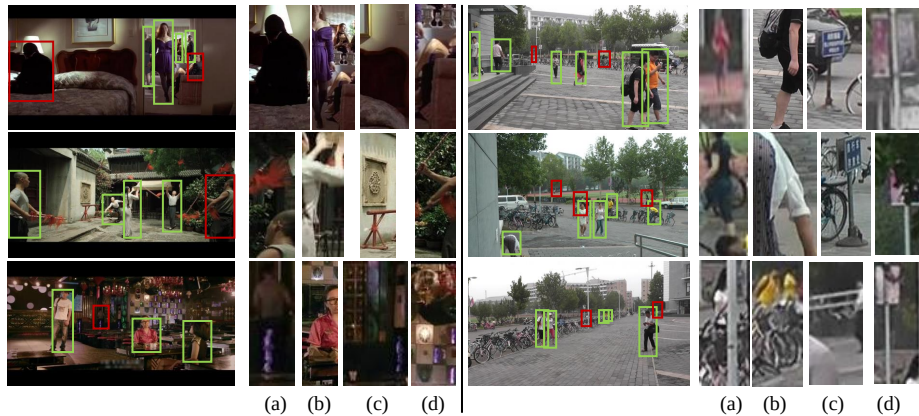


Fig. 6: Visualization of some hard cases, the green bounding boxes denote the qualified proposals, while the red ones denote the undetected persons. The crops of the hybrid hard cases are presented on the right of the images.

demonstrate the diversity of our hybrid hard cases, and validate the rationality of adding such cases to the memory bank.

## 5 Conclusions

In this paper, we introduce a novel Domain Adaptive Person Search setting, where neither bounding boxes nor identity labels for target domain are required. Based on this new setting, we propose a strong baseline framework by investigating domain alignment and taking advantage of unlabeled target domain data. Extensive results on two large-scale benchmarks demonstrate the promising performance our framework achieves and the effectiveness of designed modules. We hope this work will encourage more exploration in this direction.

**Acknowledgment** This work was supported by Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), CAAI-Huawei MindSpore Open Fund.

## References

1. Arruda, V.F., Paixão, T.M., Berriel, R.F., Souza, A.F.D., Badue, C., Sebe, N., Oliveira-Santos, T.: Cross-domain car detection using unsupervised image-to-image translation: From day to night. In: IJCNN. pp. 1–8 (2019)
2. Cai, Q., Pan, Y., Ngo, C., Tian, X., Duan, L., Yao, T.: Exploring object relation in mean teacher for cross-domain detection. In: CVPR. pp. 11457–11466 (2019)
3. Cao, Y., Guan, D., Huang, W., Yang, J., Cao, Y., Qiao, Y.: Pedestrian detection with unsupervised multispectral feature learning using deep neural networks. *Inf. Fusion* **46**, 206–217 (2019)
4. Chen, D., Zhang, S., Ouyang, W., Yang, J., Schiele, B.: Hierarchical online instance matching for person search. In: AAAI. pp. 10518–10525 (2020)
5. Chen, D., Zhang, S., Ouyang, W., Yang, J., Tai, Y.: Person search via a mask-guided two-stream CNN model. In: ECCV (7). vol. 11211, pp. 764–781 (2018)
6. Chen, D., Zhang, S., Yang, J., Schiele, B.: Norm-aware embedding for efficient person search. In: CVPR. pp. 12612–12621 (2020)
7. Chen, Y., Zhu, X., Gong, S.: Instance-guided context rendering for cross-domain person re-identification. In: ICCV. pp. 232–242 (2019)
8. Chen, Y., Li, W., Sakaridis, C., Dai, D., Gool, L.V.: Domain adaptive faster R-CNN for object detection in the wild. In: CVPR. pp. 3339–3348 (2018)
9. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255 (2009)
10. Deng, W., Zheng, L., Ye, Q., Kang, G., Yang, Y., Jiao, J.: Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In: CVPR. pp. 994–1003 (2018)
11. Devaguptapu, C., Akolekar, N., Sharma, M.M., Balasubramanian, V.N.: Borrow from anywhere: Pseudo multi-modal object detection in thermal imagery. In: CVPR Workshops. pp. 1029–1038 (2019)
12. Dong, W., Zhang, Z., Song, C., Tan, T.: Bi-directional interaction network for person search. In: CVPR. pp. 2836–2845 (2020)
13. Dong, W., Zhang, Z., Song, C., Tan, T.: Instance guided proposal network for person search. In: CVPR. pp. 2582–2591 (2020)
14. Ester, M., Kriegel, H., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: KDD. pp. 226–231 (1996)
15. Fu, Y., Wei, Y., Wang, G., Zhou, Y., Shi, H., Huang, T.S.: Self-similarity grouping: A simple unsupervised cross domain adaptation approach for person re-identification. In: ICCV. pp. 6111–6120 (2019)
16. Ganin, Y., Lempitsky, V.S.: Unsupervised domain adaptation by backpropagation. In: ICML. JMLR Workshop and Conference Proceedings, vol. 37, pp. 1180–1189 (2015)
17. Ge, Y., Zhu, F., Chen, D., Zhao, R., Li, H.: Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In: NeurIPS (2020)
18. Girshick, R.B., Iandola, F.N., Darrell, T., Malik, J.: Deformable part models are convolutional neural networks. In: CVPR. pp. 437–446 (2015)
19. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: NIPS. pp. 2672–2680 (2014)
20. Han, C., Su, K., Yu, D., Yuan, Z., Gao, C., Sang, N., Yang, Y., Wang, C.: Weakly supervised person search with region siamese networks. In: ICCV. pp. 12006–12015 (2021)

21. Han, C., Ye, J., Zhong, Y., Tan, X., Zhang, C., Gao, C., Sang, N.: Re-id driven localization refinement for person search. In: ICCV. pp. 9813–9822 (2019)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
23. Kang, G., Jiang, L., Yang, Y., Hauptmann, A.G.: Contrastive adaptation network for unsupervised domain adaptation. In: CVPR. pp. 4893–4902 (2019)
24. Khodabandeh, M., Vahdat, A., Ranjbar, M., Macready, W.G.: A robust learning approach to domain adaptive object detection. In: ICCV. pp. 480–490 (2019)
25. Kumar, M.P., Packer, B., Koller, D.: Self-paced learning for latent variable models. In: NIPS. pp. 1189–1197 (2010)
26. Li, Z., Miao, D.: Sequential end-to-end network for efficient person search. In: AAAI. pp. 2011–2019 (2021)
27. Lin, C.: Cross domain adaptation for on-road object detection using multimodal structure-consistent image-to-image translation. In: ICIP. pp. 3029–3030 (2019)
28. Liu, C., Chang, X., Shen, Y.: Unity style transfer for person re-identification. In: CVPR. pp. 6886–6895 (2020)
29. Liu, H., Feng, J., Jie, Z., Karlekar, J., Zhao, B., Qi, M., Jiang, J., Yan, S.: Neural person search machines. In: ICCV. pp. 493–501 (2017)
30. Munjal, B., Amin, S., Tombari, F., Galasso, F.: Query-guided end-to-end person search. In: CVPR. pp. 811–820 (2019)
31. Saito, K., Ushiku, Y., Harada, T., Saenko, K.: Strong-weak distribution alignment for adaptive object detection. In: CVPR. pp. 6956–6965 (2019)
32. Song, L., Wang, C., Zhang, L., Du, B., Zhang, Q., Huang, C., Wang, X.: Unsupervised domain adaptive re-identification: Theory and practice. *Pattern Recognit.* **102**, 107173 (2020)
33. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR. pp. 2962–2971 (2017)
34. Wang, C., Ma, B., Chang, H., Shan, S., Chen, X.: TCTS: A task-consistent two-stage framework for person search. In: CVPR. pp. 11949–11958 (2020)
35. Wang, D., Zhang, S.: Unsupervised person re-identification via multi-label classification. In: CVPR. pp. 10978–10987 (2020)
36. Wang, M., Deng, W.: Deep visual domain adaptation: A survey. *Neurocomputing* **312**, 135–153 (2018)
37. Wang, T., Zhang, X., Yuan, L., Feng, J.: Few-shot adaptive faster R-CNN. In: CVPR. pp. 7173–7182 (2019)
38. Xiao, J., Xie, Y., Tillo, T., Huang, K., Wei, Y., Feng, J.: IAN: the individual aggregation network for person search. *Pattern Recognit.* **87**, 332–340 (2019)
39. Xiao, T., Li, S., Wang, B., Lin, L., Wang, X.: Joint detection and identification feature learning for person search. In: CVPR. pp. 3376–3385 (2017)
40. Yan, Y., Li, J., Liao, S., Qin, J., Ni, B., Lu, K., Yang, X.: Exploring visual context for weakly supervised person search. In: AAAI. vol. 36, pp. 3027–3035 (2022)
41. Yan, Y., Li, J., Qin, J., Bai, S., Liao, S., Liu, L., Zhu, F., Shao, L.: Anchor-free person search. In: CVPR. pp. 7690–7699 (2021)
42. Yan, Y., Li, J., Liao, S., Qin, J., Ni, B., Yang, X.: TAL: two-stream adaptive learning for generalizable person re-identification. *CoRR* **abs/2111.14290** (2021)
43. Yan, Y., Zhang, Q., Ni, B., Zhang, W., Xu, M., Yang, X.: Learning context graph for person search. In: CVPR. pp. 2158–2167 (2019)
44. Zheng, L., Zhang, H., Sun, S., Chandraker, M., Yang, Y., Tian, Q.: Person re-identification in the wild. In: CVPR. pp. 3346–3355 (2017)
45. Zhu, X., Pang, J., Yang, C., Shi, J., Lin, D.: Adapting object detectors via selective cross-domain alignment. In: CVPR. pp. 687–696 (2019)