

Efficient Person Search: An Anchor-Free Approach

Yichao Yan, Jinpeng Li, Jie Qin, *Member, IEEE*, Shengcai Liao, *Senior Member, IEEE*, and Xiaokang Yang, *Fellow, IEEE*

Abstract—Person search aims to simultaneously localize and identify a query person from realistic, uncropped images. To achieve this goal, state-of-the-art models typically add a re-id branch upon two-stage detectors like Faster R-CNN. Owing to the ROI-Align operation, this pipeline yields promising accuracy as re-id features are explicitly aligned with the corresponding object regions, but in the meantime, it introduces high computational overhead due to dense object anchors. In this work, we present an anchor-free approach to efficiently tackling this challenging task, by introducing the following dedicated designs. *First*, we select an anchor-free detector (i.e., FCOS) as the prototype of our framework. Due to the lack of dense object anchors, it exhibits significantly higher efficiency compared with existing person search models. *Second*, when directly accommodating this anchor-free detector for person search, there exist several major challenges in learning robust re-id features, which we summarize as the misalignment issues in different levels (i.e., scale, region, and task). To address these issues, we propose an aligned feature aggregation module to generate more discriminative and robust feature embeddings. Accordingly, we name our model as Feature-Aligned Person Search Network (AlignPS). *Third*, by investigating the advantages of both anchor-based and anchor-free models, we further augment AlignPS with an ROI-Align head, which significantly improves the robustness of re-id features while still keeping our model highly efficient. Extensive experiments conducted on two challenging benchmarks (i.e., CUHK-SYSU and PRW) demonstrate that our framework achieves state-of-the-art or competitive performance, while displaying higher efficiency. All the source codes, data, and trained models are available at: <https://github.com/daodaofr/alignps>.

Index Terms—Person search, anchor-free model, efficient learning, feature alignment.

I. INTRODUCTION

PERSON search [1], [2] aims to localize and identify a target person from a gallery of realistic, uncropped scene images, and it has recently emerged as a practical task with real-world applications, e.g., video surveillance. Two fundamental computer vision tasks, i.e., pedestrian detection [3], [4] and person re-identification (re-id) [5], [6], need to be addressed to tackle this task. Both detection and re-id are very challenging tasks and have received tremendous attention in the past decade. In person search, we need to not only address the challenges (e.g., occlusions, pose/viewpoint variations, and background clutter) of the two individual tasks, but also pursue a unified and optimized framework to simultaneously perform detection and re-id.

Previous person search frameworks can be generally divided into two categories. The first line of works [1], [7], [8] can be

summarized as *two-step* approaches, which attempt to deal with detection and re-id separately. As shown in Fig. 1a, multiple persons are first localized with off-the-shelf detection models, and then cropped out and fed to re-id networks to extract feature representations. Although two-step models can obtain satisfactory results, the disentangled treatment of the two tasks is time- and resource-consuming. In contrast, the second category [2], [9], [10], [11], [12] provides a *one-step* solution that unifies detection and re-id in an end-to-end manner. As shown in Fig. 1b, one-step models first apply an ROI-Align layer to aggregate features in the detected bounding boxes. The features are then shared by detection and re-id; with an additional re-id loss, the simultaneous optimization of the two tasks becomes feasible. Since these models adopt two-stage detectors like Faster R-CNN [13], we refer to them as *one-step two-stage* models. However, these methods inevitably inherit the limitations of two-stage detectors, e.g., high computational complexity caused by dense anchors, and high sensitivity to the hyperparameters including the size, aspect ratio and number of anchor boxes, etc.

Compared with two-stage detectors, anchor-free models exhibit unique advantages (e.g., simpler structure and higher speed), and have been actively studied in recent years [14], [15], [16], [17]. Inspired by this, an open question is naturally thrown at us - *Is it possible to develop an anchor-free framework for person search?* Our answer is yes. However, this is a non-trivial task due to the following three misalignment issues. **1)** Many anchor-free models learn multi-scale features using feature pyramid networks (FPNs) [18] to achieve scale invariance for object detection. However, this introduces the misalignment issue for re-id (i.e., scale misalignment), as a query person needs to be compared with all the people of various scales in the gallery set, while the re-id features would be inconsistently taken from different FPN levels. **2)** In the absence of operations like ROI-Align, anchor-free models cannot align the features for re-id and detection according to a specific region. Therefore, re-id embeddings must be directly learned from feature maps without explicit region alignment, which brings additional challenges as re-id features are sensitive to the foreground regions [19], [20]. **3)** Person search can be intuitively formulated as a multi-task learning framework with detection and re-id as its sub-tasks. However, there exist conflicts between the objectives of these two tasks, i.e., pedestrian detection tries to learn features that are commonly shared by all the people, while re-id tries to extract unique features for individual persons. Hence, we need to find a better tradeoff/alignment between the two tasks.

In this work, we present an anchor-free framework for efficient person search, which we name the Feature-Aligned Person Search Network (**AlignPS**). Our model employs the

Y. Yan and X. Yang are with MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, Shanghai, China. E-mail: {yanyichao, xkyang}@sjtu.edu.cn.

J. Li, J. Qin and S. Liao are with the Inception Institute of Artificial Intelligence, Abu Dhabi, United Arab Emirates. E-mail: ljpadam@gmail.com, qinjiebuua@gmail.com, scliao@ieee.org.

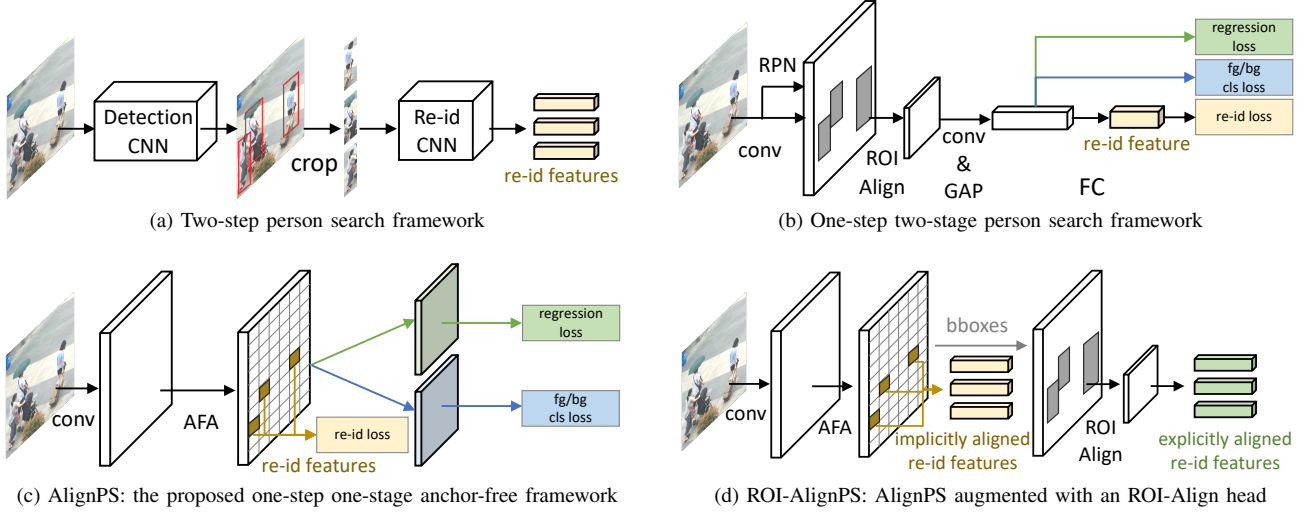


Fig. 1: Comparison of four person search frameworks. (a) The two-step framework addresses detection and re-id as two separate tasks. (b) The one-step model enables end-to-end training of detection and re-id with an ROI-Align operation based on a two-stage detector; however, re-id is considered as a secondary task after detection. (c) The proposed framework enables single-stage inference for both detection and re-id, while making re-id the primary task. (d) By appending an ROI-Align head to the bounding box output of AlignPS, the re-id features are further improved with explicit region alignment.

typical architecture of an anchor-free detection model (i.e., FCOS [21]), which allows our framework to be more efficient than prior person search models. To address the above-mentioned challenges, we design an aligned feature aggregation (AFA) module to make our model focus more on the re-id subtask. Specifically, AFA reshapes some building blocks of FPN to overcome the issues of region and scale misalignment in re-id feature learning. For example, we exploit deformable convolution to make the re-id embeddings adaptively aligned with the foreground regions. In the meantime, we design a feature fusion scheme to better aggregate features from different FPN levels, which makes the re-id features more robust to scale variations. We also optimize the training procedures of re-id and detection to place more emphasis on generating robust re-id embeddings (as shown in Fig. 1c). These simple yet effective designs successfully transform a classic anchor-free detector into a powerful and efficient person search framework, and allow the proposed model to outperform its anchor-based competitors. Moreover, we find that although the proposed AlignPS framework implicitly aligns re-id features with the corresponding regions, there inevitably exists a gap between the adapted regions and the foreground bounding boxes. As observed in previous works [19], re-id features are sensitive to the context outside the foreground regions. Inspired by the learning scheme in two-stage models (as shown in Fig. 1b), we further augment AlignPS with an ROI-Align head, to explicitly extract more robust re-id features corresponding to the foreground regions. We name this variant as **ROI-AlignPS**, as shown in Fig. 1d. These explicitly aligned re-id features can be viewed as complements to the implicitly aligned features from AlignPS, and these two kinds of features are fused to yield even better re-id representations. More importantly, the ROI-Align head only receives the output bounding boxes

from AlignPS during inference, avoiding computing dense region proposals. Therefore, ROI-AlignPS still inherits the high efficiency from AlignPS.

In summary, our main contributions include:

- We propose the first *anchor-free* framework (AlignPS) for efficient person search, which will significantly foster future research in this direction.
- We design an AFA module that simultaneously addresses the issues of scale, region, and task misalignment to successfully accommodate an anchor-free detector for the task of person search.
- We further propose a novel model variant (ROI-AlignPS) by augmenting AlignPS with an ROI-Align head, which complements the implicitly aligned re-id features. With mutual learning strategy, this variant further improves the performance while remaining highly efficient.
- As an anchor-free framework, our model surprisingly achieves state-of-the-art or competitive performance on two challenging person search benchmarks, while running at a higher speed.

Part of this work has been published in [22]. In this paper, we further make the following extensions: **1)** We augment AlignPS with an ROI-Align head, which explicitly aligns the foreground regions with their re-id features. This architecture variant notably improves the performance, while keeping the framework efficient. **2)** We investigate several mutual feature learning strategies, which allow the explicitly and implicitly aligned re-id features to promote each other during training. We find these strategies yield more robust re-id representations. **3)** We provide more thorough ablation studies and component analysis on both CUHK-SYSU and PRW, and present more qualitative results. These analyses further illustrate the effectiveness of the proposed framework.

II. RELATED WORK

Pedestrian Detection. Pedestrian or object detection can be considered as a preliminary task of person search. Current deep learning-based detectors are generally categorized into one-stage and two-stage models, according to whether they employ a region proposal network (RPN) to generate object proposals. Alternatively, object detectors can also be categorized into anchor-based and anchor-free detectors, depending on whether they utilize anchor boxes to associate objects. One of the most representative two-stage anchor-based detectors is Faster R-CNN [13], which has been extended into numerous variants [23], [24], [25], [26]. Notably, some one-stage detectors [27], [28], [29], [30] also work with anchor boxes. Compared with the above models, one-stage anchor-free detectors [14], [15], [16], [31], [32], [21], [33] have been attracting more and more attention recently due to their simple structures and efficient implementations. In this work, we develop our person search framework based on a classic one-stage anchor-free detector, thus making the whole framework simpler and faster.

Person Re-identification. Person re-id is also closely related to person search, aiming to learn identity embeddings from cropped person images. Traditional methods employed various handcrafted features [34], [5], [35] before the renaissance of deep learning. However, to pursue better performance, current re-id models are mostly based on deep learning. Some models employ structure/part information in the human body to learn more robust representations [36], [37], [38], [39], while others focus on learning better distance metrics [6], [40], [41], [42], [43], [44]. As person re-id usually lacks large-scale training data, data augmentation [45], [46], [47], [48] also becomes popular for tackling this task. Compared with detection which aims to learn common features of pedestrians, re-id needs to focus more on fine-grained details and unique features of each identity. Therefore, we propose to follow the “re-id first” principle to raise the priority of the re-id task, resulting in more discriminative identity embeddings for more accurate person search.

Person Search. Existing person search frameworks can be divided into two-step and one-step models. Two-step models first perform pedestrian detection and subsequently crop the detected people for re-id. Zheng et al [1] introduced the first two-step framework for person search and evaluated the combinations of different detectors and re-id models. Since then, several models [7], [8], [20], [49] have followed this pipeline. In [2], Xiao et al. proposed the first one-step person search framework based on Faster R-CNN. Specifically, a joint framework enabling end-to-end training of detection and re-id was proposed by stacking a re-id embedding layer after the detection features and proposing the Online Instance Matching (OIM) loss. So far, a number of improvements [9], [50], [10], [51], [11], [19], [12] have been made based on this framework. In general, two-step models may achieve better performance, while one-step models have the advantages of simplicity and efficiency. However, there is still room for improving one-step methods due to the aforementioned shortcomings of the two-stage anchor-based detectors they usually adopt. In this work,

we introduce the first anchor-free model to further improve the simplicity and efficiency of one-step models, without any sacrifice in accuracy.

Mutual Learning. Hinton et al [52] firstly introduced knowledge distillation in neural networks in 2015. Since then, it has been widely employed in various computer vision tasks to improve the capability of neural networks, e.g., image recognition [53], [54], [55], object detection [56], [57], [58], image segmentation [59], [60], [61], etc. In person search, several recent work have also adopted this approach. For example, QEEPS [11] and IGPN [62] employ the query person to find the most similar proposals in the RPN, which reduces the number of proposals and enhances the efficiency. In BiNet [19] and DKD [63], the features from cropped pedestrians are utilized to distill the re-id features in person search network, such that the robustness of the re-id features are improved. In this work, we design two feature learning branches in ROI-AlignPS. Instead of employing the query person or the cropped image, we investigate several mutual feature promotion strategies to improve the discriminative capability of the aggregated features.

III. FEATURE-ALIGNED PERSON SEARCH NETWORKS

In this section, we introduce the proposed anchor-free frameworks (i.e., AlignPS and ROI-AlignPS) for person search. First, we give an overview of the network architecture. Second, the proposed AFA module is elaborated with the aim of mitigating different levels of misalignment issues when transforming an anchor-free detector into a superior person search framework. Then, we present the designed loss function to obtain more discriminative features for person search. Finally, we present ROI-AlignPS, a model variant which further improves the performance of AlignPS.

A. Framework Overview

The basic framework of the proposed AlignPS is based on FCOS [21], one of the most popular one-stage anchor-free object detectors. Differently, we adhere to the “re-id first” principle to put emphasis on learning robust feature embeddings for the re-id subtask, which is crucial for enhancing the overall performance of person search.

As illustrated in Fig. 2, our model simultaneously localizes multiple people in the image and learns re-id embeddings for them. Specifically, an AFA module is developed to aggregate features from multi-level feature maps in the backbone network. To learn re-id embeddings, which is the key of our method, we directly take the flattened features from the output feature maps of AFA as the final embeddings, without any extra embedding layers. For detection, we employ the detection head from FCOS which is good enough for the detection subtask. The detection head consists of two branches, both of which contain four 3×3 conv layers. In the meantime, the first branch predicts regression offsets and centerness scores, while the second makes foreground/background classification. Finally, each location on the output feature map of AFA will be associated with a bounding box with classification and centerness scores, as well as a re-id feature embedding.

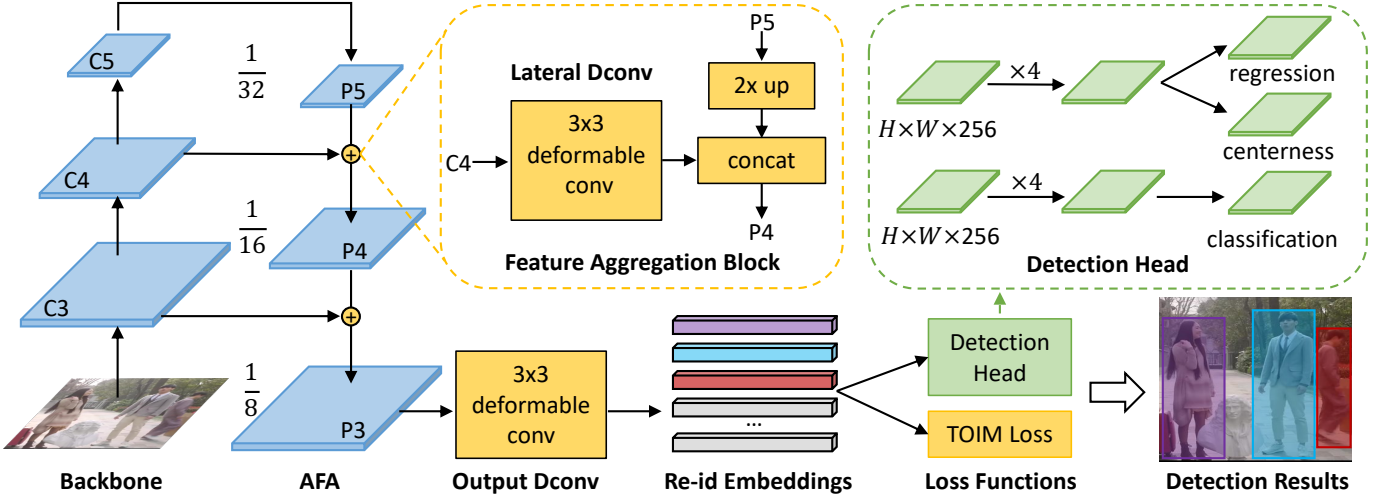


Fig. 2: Architecture of the proposed AlignPS framework, which shares the basic structure of FCOS [21]. The components in yellow are newly designed to accommodate FCOS for the task of person search. “Dconv” means deformable convolution.

B. Aligned Feature Aggregation

Following FPN [18], we make use of different levels of feature maps to learn detection and re-id features. As the key of our framework, the proposed AFA performs three levels of alignment, beyond the original FPN, to make the output re-id features more discriminative.

Scale Alignment. As shown in Fig. 3, the original FCOS model employs different levels of features to detect objects of different sizes. This significantly improves the detection performance since the overlapped ambiguous samples will be assigned to different layers. For the re-id task, however, the multi-level prediction could cause feature misalignment between different scales. In other words, when matching a person of different scales, re-id features are inconsistently taken from different levels of FPN, thus preventing the re-id features from being robust to scale variations. Furthermore, the people in the gallery set are of various scales, which could eventually make the multi-level model fail to find correct matches for the query person. Therefore, in our framework, we only make predictions based on a single layer of AFA, which explicitly addresses the feature misalignment caused by scale variations. Specifically, we employ the $\{C_3, C_4, C_5\}$ feature maps from the ResNet-50 backbone, and AFA sequentially outputs $\{P_5, P_4, P_3\}$, with strides of 32, 16, and 8, respectively. We only generate features from $\{P_3\}$, which is the largest output feature map, for both the detection and re-id subtasks, and $\{P_6, P_7\}$ are no longer generated as in the original FPN. Although this design may slightly influence the detection performance, we will show in Sec. IV-C that it achieves a good trade-off between the detection and re-id subtasks, which adapts well to the person search task.

Region Alignment. On the output feature map of AFA, each location perceives the information from the whole input image based on a large receptive field. Due to the lack of the ROI-Align operation as in Faster R-CNN, it is difficult for our anchor-free framework to learn more accurate features within the pedestrian bounding boxes, and thus leading to the issue of region misalignment. The re-id subtask is even more sensitive

to this issue as background features could greatly impact the discriminative capability of the learned features. In AlignPS, we address this issue from three perspectives. *First*, we replace the 1×1 conv layers in the lateral connections with 3×3 deformable conv layers. As the original lateral connections are designed to reduce the channels of feature maps, a 1×1 conv is enough. In our design, moreover, the 3×3 deformable conv enables the network to adaptively adjust the receptive field on the input feature maps, thus implicitly fulfilling region alignment. *Second*, we replace the “sum” operation in the top-down pathway with a “concatenation” operation, which can better aggregate multi-level features. *Third*, we again replace the 3×3 conv with a 3×3 deformable conv for the output layer of FPN, which further aligns the multi-level features to finally generate a more accurate feature map. The above three designs work seamlessly to address the region misalignment issue, and we notice that these simple designs are extremely effective when accommodating the basic anchor-free model for our person search task.

Task Alignment. Existing person search frameworks typically treat pedestrian detection as the primary task, i.e., re-id embeddings are just generated by stacking an additional layer after the detection features. A recent work [64] investigated a parallel structure by employing independent heads for the two tasks to achieve robust multiple object tracking results. In our task of person search, we find the inferior re-id features largely hinder the overall performance. Therefore, we opt for a different principle to align these two tasks by treating re-id as our primary task. Specifically, the output features of AFA are directly supervised with a re-id loss (which will be introduced in the following subsection), and then fed to the detection head. This “re-id first” design is based on two considerations. *First*, the detection subtask has been relatively well addressed by existing person search frameworks, which directly inherit the advantages from existing powerful detection frameworks. Therefore, learning discriminative re-id embeddings is our primary concern. As we discussed, re-id performance is more sensitive to region misalignment in an anchor-free framework.

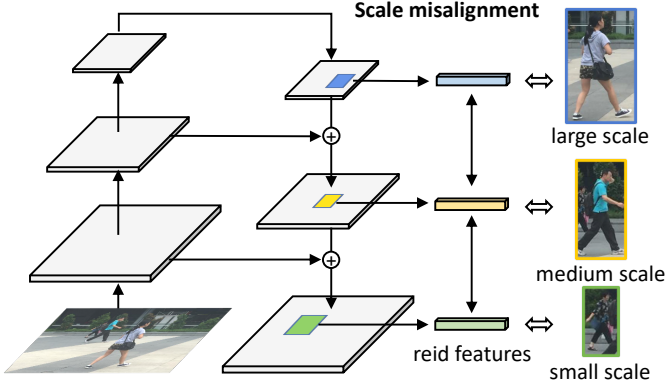


Fig. 3: Illustration of scale misalignment.

Therefore, it is desirable for the person search framework to be inclined towards the re-id subtask. We also show in our experiments that this design significantly improves the discriminative capability of the re-id embeddings, while having negligible impact on detection. *Second*, compared with “detection first” and parallel structures, the proposed “re-id first” structure does not require an extra layer to generate re-id embeddings, and is thus more efficient.

C. Triplet-Aided Online Instance Matching Loss

Existing works typically employ the OIM loss to supervise the training of the re-id subtask. Specifically, OIM stores the feature centers of all labeled identities in a lookup table (LUT), $V \in \mathbb{R}^{D \times L} = \{v_1, \dots, v_L\}$, which contains L feature vectors with D dimensions. Meanwhile, a circular queue $U \in \mathbb{R}^{D \times Q} = \{u_1, \dots, u_Q\}$ containing the features of Q unlabeled identities is maintained. At each iteration, given an input feature x with label i , OIM computes the similarity between x and all the features in the LUT and circular queue by $V^T x$ and $Q^T x$, respectively. The probability of x belonging to the identity i is calculated as:

$$p_i = \frac{\exp(v_i^T x / \tau)}{\sum_{j=1}^L \exp(v_j^T x / \tau) + \sum_{k=1}^Q \exp(u_k^T x / \tau)}, \quad (1)$$

where $\tau = 0.1$ is a hyperparameter that controls the softness of the probability distribution. The objective of OIM is to minimize the expected negative log-likelihood:

$$\mathcal{L}_{\text{OIM}} = -\mathbb{E}_x[\log p_t], \quad t = 1, 2, \dots, L. \quad (2)$$

Although OIM effectively employs both labeled and unlabeled samples, we still observe two limitations. First, the distances are only computed between the input features and the features stored in the lookup table and circular queue, while no comparisons are made between the input features. Second, the log-likelihood loss term does not give an explicit distance metric between feature pairs.

To improve OIM, we propose a specifically designed triplet loss. For each person in the input images, we employ the center sampling strategy as in [65]. As shown in Fig. 4, for each person, a set of features located around the person center are considered as positive samples. The objective is to pull

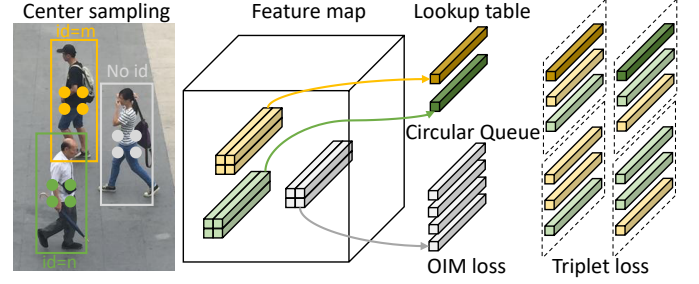


Fig. 4: Illustration of the Triplet-aided Online Instance Matching loss, where both the features from the input image and the lookup table are sampled to form the triplet.

the feature vectors from the same person close, and push the vectors from different people away. Meanwhile, the features from the labeled persons should be close to the corresponding features stored in the LUT, and away from the other features in the LUT.

More specifically, suppose we sample S vectors from one person; we get $X_m = \{x_{m,1}, \dots, x_{m,S}, v_m\}$ and $X_n = \{x_{n,1}, \dots, x_{n,S}, v_n\}$ as the candidate feature sets for the persons with identity labels m and n , respectively, where $x_{i,j}$ denotes the j -th feature of person i , and v_i is the i -th feature in the LUT. Given X_m and X_n , positive pairs can be sampled within each set, while negative pairs are sampled between the two sets. The triplet loss can be calculated as:

$$\mathcal{L}_{\text{tri}} = \sum_{\text{pos, neg}} [M + D_{\text{pos}} - D_{\text{neg}}], \quad (3)$$

where M denotes the distance margin, and D_{pos} and D_{neg} denote the Euclidean distances between the positive pair and the negative pair, respectively. Finally, the Triplet-aided OIM (TOIM) loss is the summation of these two terms:

$$\mathcal{L}_{\text{TOIM}} = \mathcal{L}_{\text{tri}} + \mathcal{L}_{\text{OIM}}. \quad (4)$$

D. AlignPS Augmented with Explicit Region Alignment

ROI-AlignPS. Although the proposed AFA module implicitly addresses the region misalignment issue with deformable convolution, it is still desirable to investigate the explicitly aligned learning schemes, e.g., cropping pedestrians for re-id learning [1], [7] and the ROI-Align operation in Faster R-CNN [2], [12]. However, learning based on the cropped pedestrians inevitably transfers the whole framework into a two-step model, which significantly downgrades the efficiency. In the meantime, the dense region proposals in Faster R-CNN also make it inefficient during inference. To address these issues, we propose a novel variant, which we name ROI-AlignPS, by augmenting our efficient anchor-free framework with an explicit region alignment module.

In the training phase, two parallel branches are simultaneously trained with implicit and explicit region alignment, respectively. As shown in Fig. 5a, the bottom branch is the original AlignPS model, while the top branch follows the architecture of NAE [12], which includes a RPN and ROI-Align operation to output the re-id features of a certain region.

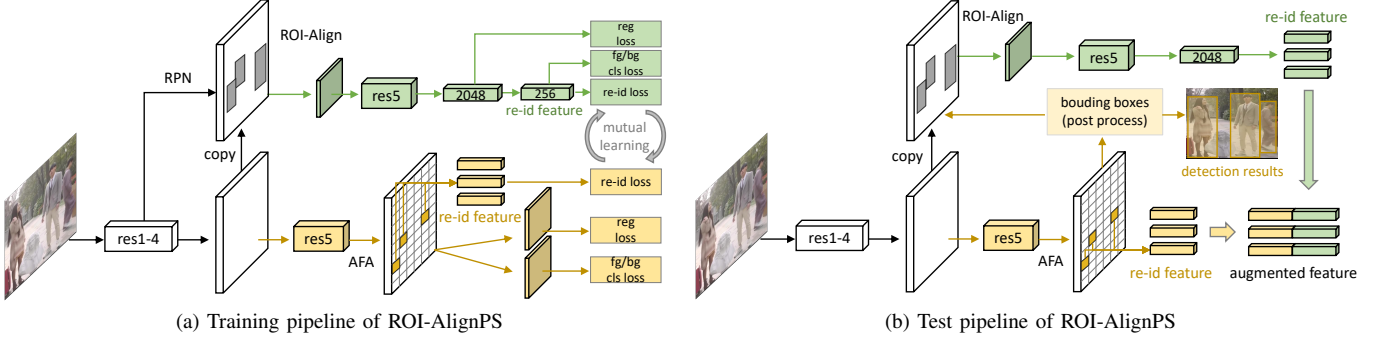


Fig. 5: The architecture of ROI-AlignPS, which can be regarded as AlignPS augmented with a Faster R-CNN branch.

Both branches share the ResNet-50 backbone from *res1* to *res4* layers, while having separate *res5* layers. In the test phase, as illustrated in Fig 5b, the bottom AlignPS branch first outputs bounding boxes and re-id features. Then, the processed bounding boxes are fed to the top branch to extract explicitly aligned re-id features corresponding to each person. Finally, the re-id features from these two branches are combined to generate a more robust representation.

Although ROI-AlignPS contains an anchor-free branch and an anchor-based branch, the dense anchors in Faster R-CNN are only involved in the training stage. During inference, RPN is discarded and only a small number of detected bounding boxes are processed with ROI-Align. In this way, ROI-AlignPS can still remain highly efficient during the inference stage. Moreover, the explicit and implicit region alignment strategies are complementary to each other. Specifically, the ROI-Align branch extracts re-id features based on bounding boxes, while AlignPS dynamically adapts to certain regions. Meanwhile, the performance of both methods may be influenced by occlusions and backgrounds. Therefore, combining these two kinds of features intuitively reduces such impacts, and the effectiveness of this design is validated in Sec IV-D.

Branch-Level Mutual Learning. Prior works [53], [55], [66], [67] found that feature interaction and knowledge distillation can benefit feature learning in neural networks. In our case, it would be desirable if the prediction of the two branches can reach a better consensus. To fulfill this, we investigate several branch interaction strategies as follows.

(1) *Mutual Information Maximization.* Mutual information is defined to measure the dependency between two random variables X and Y :

$$\mathcal{I}(X, Y) = \int_X \int_Y p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dx dy, \quad (5)$$

where $p(x, y)$ is the joint probability density function, $p(x)$ and $p(y)$ are the marginal probability density functions of X and Y , respectively. As demonstrated in prior works [68], [69], mutual information is able to depict the mutual dependence between X and Y , no matter how nonlinear the dependence is. Thus, mutual information is more flexible than the measurement of correlation. In this work, random variables X and Y represent the re-id features from the two branches in ROI-AlignPS, respectively. We aim to maximize

the mutual information between X and Y , such that the overall representation will focus on the identity information. As it is non-trivial to directly calculate the mutual information, we employ a neural estimator [70] to maximize the lower bound of the mutual information. Specifically,

$$\mathcal{I}_\Theta(X, Y) = \sup_{\theta \in \Theta} \mathbb{E}_{\mathbb{P}_{XY}}[T_\theta] - \log(\mathbb{E}_{\mathbb{P}_{X \otimes Y}}[e^{T_\theta}]), \quad (6)$$

where \mathbb{P}_{XY} is the joint distribution and $\mathbb{P}_{X \otimes Y}$ represents the marginal distribution, T_θ is a neural network parameterized by $\theta \in \Theta$. To maximize $\mathcal{I}_\Theta(X, Y)$, we minimize its negative value:

$$\mathcal{L}_{mi} = -\mathcal{I}_\Theta(X, Y). \quad (7)$$

(2) *KL-Divergence Minimization.* The Kullback-Leibler (KL) divergence is widely employed to measure the difference between two distributions, including successful applications in feature distillation in person search [19], [63]. In our case, we minimize the KL-divergence between identity predictions of the two branches. Suppose p^A and p^R denote the output probabilities of the AlignPS branch and the ROI-Align branch, respectively, the KL-divergence is calculated as:

$$\mathcal{L}_{kl} = \text{KL}(p^A \| p^R) = \sum_{i=1}^L p_i^A \log \frac{p_i^A}{p_i^R}, \quad (8)$$

where p_i^A and p_i^R are calculated from Eq. 1, and L is the number of identities in the training set. In this way, the output features could reach better prediction-level consensus.

(3) *Diversity Maximization.* Rather than pursuing a consensus between the two branches, we aim to diversify the features belonging to the same identity, to yield more robust ensemble results. To enhance diversity, we promote the output features of the two branches to be different. Specifically, suppose x_i^A and x_i^R denote the features of the i -th person from the AlignPS branch and ROI-Align branch, respectively, we aim to minimize the cosine similarity between the corresponding features:

$$\mathcal{L}_{dv} = \frac{1}{B} \sum_{i=1}^B \cos(x_i^A, x_i^R), \quad (9)$$

where B denotes the number of persons in a mini-batch, and \cos denotes the cosine similarity.

Discussions. A recent work, SeqNet [71], proposes a sequential architecture based on Faster R-CNN, which first




Fig. 6: Comparative results on CUHK-SYSU and PRW with different alignment strategies, i.e., scale alignment (SA), region alignment (RA), and task alignment (TA).

predicts the detection results, and then extracts re-id features with an additional ROI-Align branch. Although ROI-AlignPS shares certain spirits with SeqNet, there exist several significant differences. (1) The motivation of SeqNet is to improve the quality of proposals, such that the re-id features could be better learned. In ROI-AlignPS, rather than considering the detection issue, the ROI-Align branch serves as an explicitly aligned re-id feature extractor, to complement the features from the AlignPS branch. (2) Despite the strong performance, SeqNet relies on a post graph matching technique. In contrast, the re-id features in ROI-AlignPS could be mutually promoted and the ensemble results are highly discriminative without post-processing. (3) In the inference stage, ROI-AlignPS still runs in an anchor-free way. Consequently, our framework is more efficient than SeqNet. The comparative results between ROI-AlignPS and SeqNet can be found in Sec. IV-E.

IV. EXPERIMENTS

A. Datasets and Settings

CUHK-SYSU [2] is a large-scale person search dataset which contains 18,184 images, with 8,432 different identities and 96,143 annotated bounding boxes. The images come from two kinds of data sources (i.e., real street snaps and movies/TV), covering diverse scenes and including variations of viewpoints, lighting, resolutions, and occlusions. We utilize the standard training/test split, where the training set contains 5,532 identities and 11,206 images, and the test set contains 2,900 query persons and 6,978 images. This dataset also defines a set of protocols with gallery sizes ranging from 50 to 4,000. We report the results using the default gallery size of 100 unless otherwise specified.

PRW [1] was captured using six static cameras in a university campus. All the images are extracted from the surveillance videos, which consist of 11,816 video frames in total. Person identities and bounding boxes are manually annotated, resulting in 932 labeled persons with 43,110 bounding boxes. The dataset is split into a training set of 5,704 images with 482 different identities, and a test set of 2,057 query persons and 6,112 images. Results are reported based on this split.

Evaluation Metric. We employ the mean average precision (mAP) and top-1 accuracy to evaluate the performance for

person search. We also employ recall and average precision (AP) to measure the detection performance.

B. Implementation Details

We employ ResNet-50 [72] pretrained on ImageNet [73] as the backbone. We set the batch size to 4, and adopt the stochastic gradient descent (SGD) optimizer with weight decay of 0.0005. The initial learning rate is set to 0.001 and is reduced by a factor of 10 at epoch 16 and 22, with a total of 24 epochs. We use a warmup strategy for 300 steps. We employ a multi-scale training strategy, where the longer side of the image is randomly resized from 667 to 2000 during training, while zero padding is utilized to fit the images with different resolutions. For inference, we rescale the test images to a fixed size of 1500×900. Following [74], we add a focal loss [28] to the original OIM loss. All the experiments are implemented based on PyTorch [75] and MMDetection [76], with an NVIDIA Tesla V100 GPU. It takes around 38 and 24 hours for ROI-AlignPS to finish training on CUHK-SYSU and PRW, respectively.

C. Analysis of AlignPS

Baseline. We directly add a re-id head in parallel with the detection head to the FCOS model and take it as our baseline. As shown in Fig. 6, each of the alignment strategies brings notable improvements to the baseline, and combining all of them yields 23% and 26.1% improvements in mAP on CUHK-SYSU and PRW, respectively.

Scale Alignment. To evaluate the effects of scale alignment, we employ feature maps from different levels of AFA and report the results in Table I. Specifically, we evaluate the features from P_3 , P_4 , and P_5 with strides of 8, 16, and 32, respectively. As can be observed, features from the largest scale P_3 yield the best performance on both CUHK-SYSU and PRW, due to the fact that they absorb different levels of features from AFA, providing richer information for detection and re-id. Similar to FCOS, we also evaluate the performance by assigning people of different scales to different feature levels. We set the size ranges for $\{P_3, P_4\}$ as $[0, 128]$ and $[128, \infty]$, while the prediction ranges for $\{P_3, P_4, P_5\}$ are $[0, 128]$, $[128, 256]$, and $[256, \infty]$, respectively. We can see



(a) Deformable conv at lateral C_3 layer in AFA(b) Deformable conv at lateral C_4 layer in AFA

Fig. 7: Each image shows the sampling locations of two levels of 3×3 ($9^2 = 81$ points at each location) deformable filters: (a) Lateral deformable conv C_3 + Output deformable conv; (b) Lateral deformable conv C_4 + Output deformable conv. We illustrate different locations with different colors, while center locations of people are marked in green. Please zoom in for better visualization.

CUHK-SYSU	Detection		Re-id	
	Recall	AP	mAP	top-1
P_3	90.3	81.2	93.1	93.4
P_4	87.5	78.7	92.7	93.1
P_5	79.0	71.7	89.3	89.5
P_3, P_4	90.4	80.5	91.1	91.6
P_3, P_4, P_5	90.9	80.4	90.0	90.5
PRW	Detection		Re-id	
	Recall	AP	mAP	top-1
P_3	94.8	92.9	45.9	81.9
P_4	93.4	91.8	41.4	77.5
P_5	85.7	83.7	25.3	57.4
P_3, P_4	94.7	92.9	40.8	77.1
P_3, P_4, P_5	95.4	93.5	39.5	74.3

TABLE I: Comparative results on CUHK-SYSU and PRW by employing different levels of features. P_3 , P_4 , and P_5 are the feature maps with strides of 8, 16, and 32, respectively.

that these dividing strategies achieve slightly better detection results w.r.t. the recall rate. However, they bring back the scale misalignment issue to person re-id, resulting in worse performance compared with single-scale features. Also note that this issue is not well addressed with the multi-scale training strategy. All the above results demonstrate the necessity and effectiveness of the proposed scale alignment strategy.

Region Alignment. We conduct experiments with different combinations of lateral deformable conv, output deformable

Lateral dconv	Output dconv	Feature concat	CUHK-SYSU		PRW	
			mAP	top-1	mAP	top-1
			83.4	83.7	27.4	62.3
✓			90.6	90.8	39.4	74.7
	✓		91.4	91.9	40.6	77.0
		✓	84.0	84.1	29.1	62.4
✓	✓		91.8	92.2	42.8	80.1
✓		✓	90.7	91.0	39.6	75.3
	✓	✓	92.0	92.5	40.7	77.5
✓	✓	✓	93.1	93.4	45.9	81.9

TABLE II: Comparative results on CUHK-SYSU and PRW by employing different components in AFA for region alignment. “dconv” stands for deformable convolution.

conv and feature concatenation, and analyze how different region alignment components influence the overall performance. The results are reported in Table II. Without all these modules, the framework only achieves 83.4% and 27.4% in mAP on CUHK-SYSU, and PRW, respectively, which is 9.7% and 18.5% lower than the full model. The individual components of lateral deformable conv and output deformable conv improve the model by $\sim 7\%$ and $\sim 8\%$ on CUHK-SYSU, respectively. Feature concatenation also brings $\sim 1\%$ improvements. By combining two of the three components, we observe consistent improvements. Finally, employing all the three modules yields 93.1% in mAP and 93.4% in top-1 accuracy on CUHK-SYSU, significantly boosting the perfor-

CUHK-SYSU	Detection		Re-id	
	Recall	AP	mAP	top-1
T_1	87.5	79.0	80.3	79.2
T_2	89.1	78.6	77.1	75.9
T_3	90.1	81.4	80.7	80.2
AlignPS	90.3	81.2	93.1	93.4

PRW	Detection		Re-id	
	Recall	AP	mAP	top-1
T_1	94.3	92.5	35.5	67.4
T_2	96.9	92.8	31.8	64.3
T_3	94.5	92.6	35.8	69.1
AlignPS	94.8	92.9	45.9	81.9

TABLE III: Comparative results on CUHK-SYSU and PRW with different training structures.

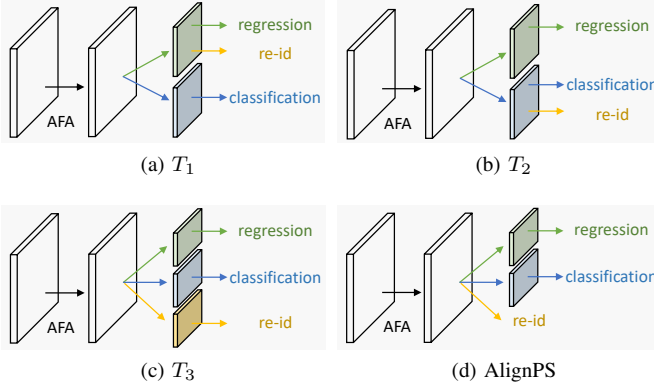


Fig. 8: Illustration of different structures for training the detection and re-id tasks.

mance. On PRW, we also observe consistent improvement by introducing these strategies. These ablation studies thoroughly demonstrate the effectiveness of region alignment.

To further illustrate how the deformable convolutions work in our framework, we visualize the learned offsets of the deformable filters in Fig. 7. We observe that the proposed framework is capable of learning adaptive receptive field according to the layout of the human body, and is robust to occlusion, crowding, and scale/illumination variations. We also observe that the lateral deformable conv in C_3 learns tighter offsets around the body center, while the offsets in the C_4 layer cover larger regions, which makes the two layers complementary to each other.

Task Alignment. Since person search aims to simultaneously address detection and re-id subtasks in a single framework, it is important to understand how different configurations of the two subtasks influence the overall task and which subtask should be paid more attention to. To this end, we design several structures to compare different training options (as shown in Fig. 8), the performance of which is summarized in Table III. As can be observed, the structures of T_1 and T_2 , where re-id features are shared with the regression and classification heads, respectively, yield significantly lower performance in re-id compared with our design. This indicates that the detection task takes advantage of the shared heads. As for T_3 where re-id and detection have independent feature heads, it achieves slightly better performance compared with

CUHK-SYSU	mAP	top-1	Δ mAP	Δ top-1
OIM	92.4	92.9	-	-
TOIM w/o LUT	92.8	93.2	+0.4	+0.3
TOIM w/ LUT	93.1	93.4	+0.7	+0.5

PRW	mAP	top-1	Δ mAP	Δ top-1
OIM	45.7	81.8	-	-
TOIM w/o LUT	45.8	81.8	+0.1	+0
TOIM w/ LUT	45.9	81.9	+0.2	+0.1

TABLE IV: Comparative results on CUHK-SYSU and PRW with different loss functions.

CUHK-SYSU	Deformable conv	mAP	top-1
ResNet-50	none	93.1	93.4
ResNet-50	res3	93.5	93.9
ResNet-50	res3 & res4	93.5	94.0
ResNet-50	res3 & res4 & res5	94.0	94.5

PRW	Deformable conv	mAP	top-1
ResNet-50	none	45.9	81.9
ResNet-50	res3	45.8	81.9
ResNet-50	res3 & res4	45.9	82.1
ResNet-50	res3 & res4 & res5	46.1	82.1

TABLE V: Comparative results on CUHK-SYSU and PRW with different deformable conv layers in the backbone model.

T_1 and T_2 , but still remarkably underperforms our design. These results indicate that our “re-id first” structure achieves the best task alignment among all these designs.

TOIM Loss. We evaluate the performance of our framework when adopting different loss functions and report the results in Table IV. We find that directly employing a triplet loss brings slight improvement. When employing the items in the LUT, the TOIM improves the mAP and top-1 accuracy on CUHK-SYSU by 0.7% and 0.5%, respectively. These two terms also slightly improve the performance on PRW. This indicates that it is beneficial to consider the relations between the input features and the features stored in the LUT.

Deformable Conv in the Backbone. As shown in Table V, inserting deformable convolutions into the backbone network has positive effects on our framework. However, the contribution of the deformable conv layers in the backbone network is less significant than the deformable conv layers in our AFA module, e.g., with all the res3 & res4 & res5 deformable conv layers, only $\sim 1\%$ and 0.2% improvements are observed on CUHK-SYSU and PRW, respectively. These results indicate that the proposed AFA works as the key module for successful feature alignment.

D. Analysis of ROI-AlignPS

We analyze the effectiveness of ROI-AlignPS by answering the following questions.

Is the ROI-Align branch necessary? In ROI-AlignPS, the ROI-Align branch extracts the explicitly aligned re-id features, complementary to the implicitly aligned features in the AlignPS branch. To evaluate the effectiveness of this design, we test each branch under different settings. (1) *Separate Training/Test*, i.e., two person search models are separately trained and evaluated. During inference, we also concatenate

Separate Training	CUHK-SYSU		PRW		Time (ms)
	mAP	top-1	mAP	top-1	
ROI-Align	92.0	92.3	43.1	80.7	83
AlignPS	93.1	93.4	45.9	81.9	61
Concatenation	93.0	93.6	44.5	81.3	144
Joint Training w/o RPN	CUHK-SYSU		PRW		Time (ms)
	mAP	top-1	mAP	top-1	
ROI-Align	90.4	91.3	43.9	81.5	83
AlignPS	93.2	92.8	44.9	80.0	61
Concatenation	94.3	94.8	48.6	83.2	75
ROI-Align*	93.4	93.8	47.5	84.5	75
Joint Training w/ RPN	CUHK-SYSU		PRW		Time (ms)
	mAP	top-1	mAP	top-1	
ROI-Align	93.3	93.9	47.8	83.3	83
AlignPS	94.0	94.1	46.4	81.2	61
Concatenation	95.0	95.3	50.3	84.3	75

TABLE VI: Comparative results with different structures. ROI-Align* denotes the model trained without re-id feature in the AlignPS branch, i.e., AlignPS concentrates on the detection task while ROI-Align focuses on re-id. This also yields an efficient anchor-free framework with fair performance.

the features from the ROI-Align model (with RPN) and the AlignPS model. (2) *Joint Training w/o RPN*, where the ROI-Align branch and the AlignPS branch are jointly trained, and both the training and test pipelines follow the structure in Fig. 5b, i.e., no RPN is employed during training and inference. (3) *Joint Training w/ RPN*, which denotes the proposed ROI-AlignPS framework, where the training pipeline follows Fig. 5a, and the test pipeline follows Fig. 5b. In this case, RPN is only employed in the ROI-Align branch during training. As shown in Table VI, although the separately trained models both obtain promising results, directly concatenating their re-id features does not yield better performance, which may be because the detection results of the two models are not perfectly aligned. In our joint training setting, in contrast, the branch-level features are well aligned. When the ROI-Align branch is trained without RPN, although the performance of each branch is similar to the separately trained models, the concatenated re-id features are notably improved by 1.3% and 4.1% w.r.t. mAP on CUHK-SYSU and PRW, respectively. Moreover, when the ROI-Align branch is trained with RPN, the performance of this branch is further improved, which in turn yields enhanced performance by concatenating the re-id features. These results not only indicate that the jointly training pipeline can better regularize features in the shared layers (i.e., res1-4), but also validate the effectiveness and necessity of the ROI-Align branch, which is complementary to AlignPS.

How does the ROI-Align branch influence the efficiency?

In ROI-AlignPS, the ROI-Align branch inevitably brings additional computational overhead. As shown in Table VI, in the inference stage, it takes 75 milliseconds (ms) to process an image with both AlignPS and ROI-Align, in comparison to 61 ms for AlignPS only. However, this is still less than the Faster R-CNN based ROI-Align model (83 ms), because our model does not require dense region proposals during inference. Therefore, ROI-AlignPS achieves significant performance improvement and remains highly efficient.

Mutual Learning	CUHK-SYSU		PRW	
	mAP	top-1	mAP	top-1
<i>None</i>	95.0	95.3	50.3	84.3
\mathcal{L}_{mi}	95.4	96.0	51.6	84.4
\mathcal{L}_{kl}	95.3	95.9	50.4	85.3
\mathcal{L}_{dv}	95.2	95.6	50.4	84.9
$\mathcal{L}_{mi} + \mathcal{L}_{kl}$	95.4	95.8	51.1	84.0
$\mathcal{L}_{mi} + \mathcal{L}_{dv}$	95.1	95.6	50.4	84.5
$\mathcal{L}_{kl} + \mathcal{L}_{dv}$	95.4	95.8	50.0	84.6
$\mathcal{L}_{mi} + \mathcal{L}_{kl} + \mathcal{L}_{dv}$	95.1	95.3	51.4	85.2

TABLE VII: Comparative results on CUHK-SYSU and PRW with different mutual learning loss functions.

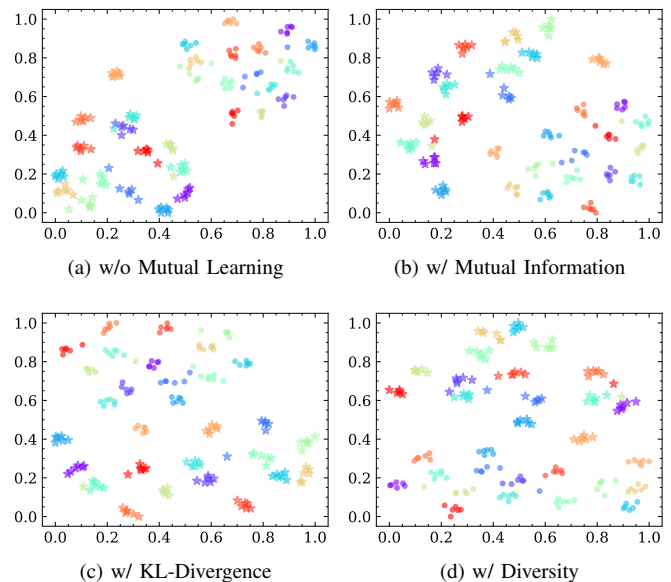


Fig. 9: The distribution of re-id features from 15 random identities. ‘*’ denotes the re-id features from the ROI-Align branch, and ‘.’ denotes the features from AlignPS branch. Different colors represent different identities.

Are the mutual learning strategies effective? As mentioned in Sec. III-D, we investigate several mutual learning strategies to enhance the representation capabilities of the re-id features in ROI-AlignPS. By applying these strategies, we observe 0.1%-1.3% performance gains in mAP on the two datasets, as shown in Table VII. However, combining several strategies together does not yield better performance, because the objectives of these strategies are partially overlapped. Furthermore, we visualize the distribution of re-id features based on t-SNE. As shown in Fig. 9, the features generated from the two branches are separable, which further validates the complementary property of these two branches. We also observe that the features with mutual learning strategies are more separable, where the overall performance with mutual information loss is slightly better than employing other strategies. Therefore, in the following, we denote ROI-AlignPS as the model trained with \mathcal{L}_{mi} .

E. Comparison to State-of-the-Art Methods

We compare our framework with state-of-the-art one-step models [2], [50], [9], [10], [74], [11], [19], [12], [77], [63], [71] and two-step models [7], [8], [20], [62], [49].



Fig. 10: Visualization of some difficult cases. The yellow bounding boxes denote the queries, while the green and red bounding boxes denote correct and incorrect top-1 matches, respectively.

Methods		CUHK-SYSU		PRW	
		mAP	top-1	mAP	top-1
one-step	OIM [2]	75.5	78.7	21.3	49.4
	IAN [50]	76.3	80.1	23.0	61.9
	NPSM [9]	77.9	81.2	24.2	53.1
	RCAA [10]	79.3	81.3	-	-
	CTXG [51]	84.1	86.5	33.4	73.6
	QEEPS [11]	88.9	89.1	37.1	76.7
	HOIM [74]	89.7	90.8	39.8	80.4
	BINet [19]	90.0	90.7	45.3	81.7
	NAE [12]	91.5	92.4	43.3	80.9
	NAE+ [12]	92.1	92.9	44.0	81.1
	PGA [77]	92.3	94.7	44.2	85.2
	DKD [63]	93.1	94.2	50.5	87.1
	SeqNet [71]	93.8	94.6	46.7	83.4
	SeqNet+CBGM[71]	94.8	95.7	47.6	87.6
	AlignPS	93.1	93.4	45.9	81.9
	ROI-AlignPS	95.4	96.0	51.6	84.4
two-step	DPM+IDE [1]	-	-	20.5	48.3
	CNN+MGTS [7]	83.3	83.9	32.8	72.1
	CNN+CLSA [8]	87.2	88.5	38.7	65.0
	FPN+RDRL [20]	93.0	94.2	42.9	70.2
	IGPN [62]	90.3	91.4	47.2	87.0
	OR [78]	92.3	93.8	52.3	71.5
	TCTS [49]	93.9	95.1	46.8	87.5

TABLE VIII: Comparison with state-of-the-art methods. The upper block lists the results of one-step models, while the lower block shows the results of two-step methods.

Results on CUHK-SYSU. As shown in Table VIII, ROI-AlignPS outperforms all the existing person search models. Notably, ROI-AlignPS outperforms the current best-performing SeqNet [71] by 0.6% and 0.3% in mAP and top-1 accuracy, respectively. Note that SeqNet requires a graph

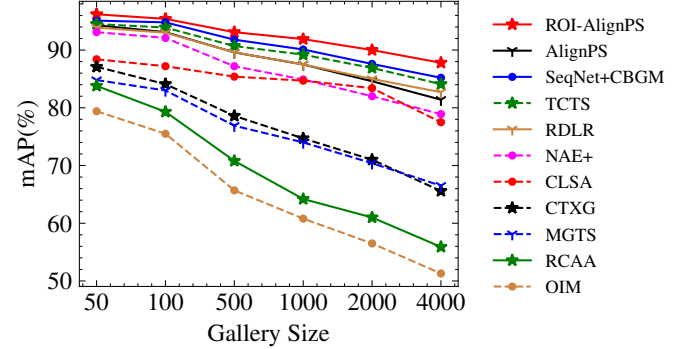


Fig. 11: Comparative results on CUHK-SYSU with different gallery sizes. Our models (AlignPS and ROI-AlignPS) are compared with both one-step models and two-step models.

matching strategy as post-processing to achieve its best performance, while ROI-AlignPS does not need such a process. We also observe from the table that our model outperforms all the two-step models, even though they employ two separate models for detection and re-id. In contrast, our model allows joint inference with a very simple structure, whilst running at a higher speed.

We visualize the results of AlignPS and ROI-AlignPS w.r.t. mAP with various gallery sizes and compare our model with both one-step and two-step models. Fig. 11 illustrates the detailed results, where ROI-AlignPS outperforms all the models by notable margins, in terms of all the gallery sizes.

Results on PRW. PRW contains less training data; therefore, all the models achieve worse performance on this dataset. Nevertheless, as can be observed from Table VIII, ROI-

Methods	Train Anchor	Test Anchor	GPU	Time (ms)
PGA [77]	✓	✓	Titan X	356
DKD [63]	✓	✓	1080 Ti	124
OIM [2]	✓	✓	V100	118
NAE+ [12]	✓	✓	V100	98
NAE [12]	✓	✓	V100	83
SeqNet [71]	✓	✓	V100	86
AlignPS	×	×	V100	61
ROI-AlignPS	✓	×	V100	75

TABLE IX: Runtime comparison of different models. Both AlignPS and ROI-AlignPS are anchor-free during inference.

AlignPS still outperforms all the one-step models in terms of mAP. We notice that DKD [63], PGA [77], IGP [62], and TCTS [49] achieve higher top-1 accuracy on PRW. These methods also seek to address the feature misalignment issue, but they either resort to feature distillation from the cropped instances [63], [62], [49], or apply attention mechanisms [77]. Differently, our model efficiently addresses this issue with the proposed AFA module.

Efficiency Comparison. Since different methods are evaluated with different GPUs, it is difficult to conduct a fair comparison of the efficiency among all the models. Here, we compare our method with OIM¹ [2], NAE/NAE+ [12] and SeqNet [71] on the same Tesla V100 GPU. All the test images are resized to 1500×900 before being fed to the networks. As shown in Table IX, our anchor-free AlignPS only takes 61 ms to process an image, which is 27% and 38% faster than NAE and NAE+, respectively. Meanwhile, ROI-AlignPS also runs faster than prior state-of-the-art models [12], [71]. This is because our models are anchor-free during inference, validating the advantage of the proposed framework.

Qualitative Results. Some qualitative results are illustrated in Fig. 10. We can observe that AlignPS and ROI-AlignPS are more robust in handling occlusions and scale/viewpoint variations, where OIM [2] and NAE [12] fail. A failure case is illustrated in the last row, where our models fail in distinguishing very tiny objects that share similar appearances. In our future work, we will work towards this direction.

V. CONCLUSION

In this paper, we propose an anchor-free approach to efficiently tackling the task of person search, by developing a person search framework based on an anchor-free detector. We design the aligned feature aggregation module to effectively address the scale, region, and task misalignment issues when accommodating the detector for the person search task. Furthermore, we propose to augment our anchor-free model with an ROI-Align branch, which additionally takes advantage of the anchor-based models. Extensive experiments demonstrate that the proposed framework not only outperforms existing person search methods but also runs at a higher speed.

REFERENCES

- [1] L. Zheng, H. Zhang, S. Sun, M. Chandraker, Y. Yang, and Q. Tian, “Person re-identification in the wild,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 3346–3355. **1, 3, 5, 7, 11**

¹We test the PyTorch implementation at https://github.com/serendip1ty/person_search.

- [2] T. Xiao, S. Li, B. Wang, L. Lin, and X. Wang, “Joint detection and identification feature learning for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 3376–3385. **1, 3, 5, 7, 10, 11, 12**
- [3] W. Ouyang and X. Wang, “Joint deep learning for pedestrian detection,” in *Int. Conf. Comput. Vis.*, 2013, pp. 2056–2063. **1**
- [4] S. Zhang, R. Benenson, and B. Schiele, “Citypersons: A diverse dataset for pedestrian detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 4457–4465. **1**
- [5] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani, “Person re-identification by symmetry-driven accumulation of local features,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2010, pp. 2360–2367. **1, 3**
- [6] E. Ahmed, M. J. Jones, and T. K. Marks, “An improved deep learning architecture for person re-identification,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 3908–3916. **1, 3**
- [7] D. Chen, S. Zhang, W. Ouyang, J. Yang, and Y. Tai, “Person search by separated modeling and A mask-guided two-stream CNN model,” *IEEE Trans. Image Process.*, vol. 29, pp. 4669–4682, 2020. **1, 3, 5, 10, 11**
- [8] X. Lan, X. Zhu, and S. Gong, “Person search by multi-scale matching,” in *Eur. Conf. Comput. Vis.*, vol. 11205, 2018, pp. 553–569. **1, 3, 10, 11**
- [9] H. Liu, J. Feng, Z. Jie, J. Karlekar, B. Zhao, M. Qi, J. Jiang, and S. Yan, “Neural person search machines,” in *Int. Conf. Comput. Vis.*, 2017, pp. 493–501. **1, 3, 10, 11**
- [10] X. Chang, P. Huang, Y. Shen, X. Liang, Y. Yang, and A. G. Hauptmann, “RCAA: relational context-aware agents for person search,” in *Eur. Conf. Comput. Vis.*, 2018, pp. 86–102. **1, 3, 10, 11**
- [11] B. Munjal, S. Amin, F. Tombari, and F. Galasso, “Query-guided end-to-end person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 811–820. **1, 3, 10, 11**
- [12] D. Chen, S. Zhang, J. Yang, and B. Schiele, “Norm-aware embedding for efficient person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 12 612–12 621. **1, 3, 5, 10, 11, 12**
- [13] S. Ren, K. He, R. B. Girshick, and J. Sun, “Faster R-CNN: towards real-time object detection with region proposal networks,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017. **1, 3**
- [14] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 779–788. **1, 3**
- [15] H. Law and J. Deng, “Corners: Detecting objects as paired keypoints,” in *Eur. Conf. Comput. Vis.*, 2018, pp. 765–781. **1, 3**
- [16] W. Liu, S. Liao, W. Ren, W. Hu, and Y. Yu, “High-level semantic feature detection: A new perspective for pedestrian detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 5187–5196. **1, 3**
- [17] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, “Centernet: Keypoint triplets for object detection,” in *Int. Conf. Comput. Vis.*, 2019, pp. 6568–6577. **1**
- [18] T. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie, “Feature pyramid networks for object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 936–944. **1, 4**
- [19] W. Dong, Z. Zhang, C. Song, and T. Tan, “Bi-directional interaction network for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 2836–2845. **1, 2, 3, 6, 10, 11**
- [20] C. Han, J. Ye, Y. Zhong, X. Tan, C. Zhang, C. Gao, and N. Sang, “Re-id driven localization refinement for person search,” in *Int. Conf. Comput. Vis.*, 2019, pp. 9813–9822. **1, 3, 10, 11**
- [21] Z. Tian, C. Shen, H. Chen, and T. He, “FCOS: fully convolutional one-stage object detection,” in *Int. Conf. Comput. Vis.*, 2019, pp. 9626–9635. **2, 3, 4**
- [22] Y. Yan, J. Li, J. Qin, S. Bai, S. Liao, L. Liu, F. Zhu, and L. Shao, “Anchor-free person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 7690–7699. **2**
- [23] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, “Deformable convolutional networks,” in *Int. Conf. Comput. Vis.*, 2017, pp. 764–773. **3**
- [24] Z. Cai and N. Vasconcelos, “Cascade R-CNN: delving into high quality object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 6154–6162. **3**
- [25] J. Pang, K. Chen, J. Shi, H. Feng, W. Ouyang, and D. Lin, “Libra R-CNN: towards balanced learning for object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 821–830. **3**
- [26] G. Song, Y. Liu, and X. Wang, “Revisiting the sibling head in object detector,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 11 560–11 569. **3**
- [27] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C. Fu, and A. C. Berg, “SSD: single shot multibox detector,” in *Eur. Conf. Comput. Vis.*, 2016, pp. 21–37. **3**

- [28] T. Lin, P. Goyal, R. B. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Int. Conf. Comput. Vis.*, 2017, pp. 2999–3007. 3, 7
- [29] J. Redmon and A. Farhadi, “YOLO9000: better, faster, stronger,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 6517–6525. 3
- [30] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li, “Single-shot refinement neural network for object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4203–4212. 3
- [31] X. Zhou, D. Wang, and P. Krähenbühl, “Objects as points,” *CoRR*, vol. abs/1904.07850, 2019. 3
- [32] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin, “Reppoints: Point set representation for object detection,” in *Int. Conf. Comput. Vis.*, 2019, pp. 9656–9665. 3
- [33] J. Li, S. Liao, H. Jiang, and L. Shao, “Box guided convolution for pedestrian detection,” in *ACM Int. Conf. Multimedia*, 2020, pp. 1615–1624. 3
- [34] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004. 3
- [35] D. Gray and H. Tao, “Viewpoint invariant pedestrian recognition with an ensemble of localized features,” in *Eur. Conf. Comput. Vis.*, 2008, pp. 262–275. 3
- [36] C. Su, J. Li, S. Zhang, J. Xing, W. Gao, and Q. Tian, “Pose-driven deep convolutional model for person re-identification,” in *Int. Conf. Comput. Vis.*, 2017, pp. 3980–3989. 3
- [37] Y. Sun, L. Zheng, Y. Yang, Q. Tian, and S. Wang, “Beyond part models: Person retrieval with refined part pooling (and A strong convolutional baseline),” in *Eur. Conf. Comput. Vis.*, 2018, pp. 501–518. 3
- [38] J. Miao, Y. Wu, P. Liu, Y. Ding, and Y. Yang, “Pose-guided feature alignment for occluded person re-identification,” in *Int. Conf. Comput. Vis.*, 2019, pp. 542–551. 3
- [39] Y. Yan, J. Qin, B. Ni, J. Chen, L. Liu, F. Zhu, W.-S. Zheng, X. Yang, and L. Shao, “Learning multi-attention context graph for group-based re-identification,” *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 3
- [40] A. Hermans, L. Beyer, and B. Leibe, “In defense of the triplet loss for person re-identification,” *CoRR*, vol. abs/1703.07737, 2017. 3
- [41] W. Chen, X. Chen, J. Zhang, and K. Huang, “Beyond triplet loss: A deep quadruplet network for person re-identification,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 1320–1329. 3
- [42] Y. Chen, X. Zhu, W. Zheng, and J. Lai, “Person re-identification by camera correlation aware feature augmentation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 2, pp. 392–408, 2018. 3
- [43] T. Wang, S. Gong, X. Zhu, and S. Wang, “Person re-identification by discriminative selection in video ranking,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 12, pp. 2501–2514, 2016. 3
- [44] X. Zhu, B. Wu, D. Huang, and W. Zheng, “Fast open-world person re-identification,” *IEEE Trans. Image Process.*, vol. 27, no. 5, pp. 2286–2300, 2018. 3
- [45] Y. Ge, Z. Li, H. Zhao, G. Yin, S. Yi, X. Wang, and H. Li, “FD-GAN: pose-guided feature distilling GAN for robust person re-identification,” in *Adv. Neural Inform. Process. Syst.*, 2018, pp. 1230–1241. 3
- [46] J. Liu, B. Ni, Y. Yan, P. Zhou, S. Cheng, and J. Hu, “Pose transferrable person re-identification,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 4099–4108. 3
- [47] L. Wei, S. Zhang, W. Gao, and Q. Tian, “Person transfer GAN to bridge domain gap for person re-identification,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 79–88. 3
- [48] Z. Zhong, L. Zheng, Z. Zheng, S. Li, and Y. Yang, “Camstyle: A novel data augmentation method for person re-identification,” *IEEE Trans. Image Process.*, vol. 28, no. 3, pp. 1176–1190, 2019. 3
- [49] C. Wang, B. Ma, H. Chang, S. Shan, and X. Chen, “TCTS: A task-consistent two-stage framework for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 11949–11958. 3, 10, 11, 12
- [50] J. Xiao, Y. Xie, T. Tillo, K. Huang, Y. Wei, and J. Feng, “IAN: the individual aggregation network for person search,” *Pattern Recognit.*, vol. 87, pp. 332–340, 2019. 3, 10, 11
- [51] Y. Yan, Q. Zhang, B. Ni, W. Zhang, M. Xu, and X. Yang, “Learning context graph for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 2158–2167. 3, 11
- [52] G. E. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *CoRR*, vol. abs/1503.02531, 2015. 3
- [53] Z. Li and D. Hoiem, “Learning without forgetting,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 12, pp. 2935–2947, 2018. 3, 6
- [54] X. Wang, T. Fu, S. Liao, S. Wang, Z. Lei, and T. Mei, “Exclusivity-consistency regularized knowledge distillation for face recognition,” in *Eur. Conf. Comput. Vis.*, 2020, pp. 325–342. 3
- [55] Z. Peng, Z. Li, J. Zhang, Y. Li, G. Qi, and J. Tang, “Few-shot image recognition with knowledge transfer,” in *Int. Conf. Comput. Vis.*, 2019, pp. 441–449. 3, 6
- [56] Q. Li, S. Jin, and J. Yan, “Mimicking very efficient network for object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 7341–7349. 3
- [57] G. Chen, W. Choi, X. Yu, T. X. Han, and M. Chandraker, “Learning efficient object detection models with knowledge distillation,” in *Adv. Neural Inform. Process. Syst.*, 2017, pp. 742–751. 3
- [58] X. Zhu, H. Hu, S. Lin, and J. Dai, “Deformable convnets V2: more deformable, better results,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 9308–9316. 3
- [59] R. T. Mullapudi, S. Chen, K. Zhang, D. Ramanan, and K. Fatahalian, “Online model distillation for efficient video inference,” in *Int. Conf. Comput. Vis.*, 2019, pp. 3572–3581. 3
- [60] Y. Liu, K. Chen, C. Liu, Z. Qin, Z. Luo, and J. Wang, “Structured knowledge distillation for semantic segmentation,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 2604–2613. 3
- [61] Y. Hou, Z. Ma, C. Liu, T. Hui, and C. C. Loy, “Inter-region affinity distillation for road marking segmentation,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 3
- [62] W. Dong, Z. Zhang, C. Song, and T. Tan, “Instance guided proposal network for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 2582–2591. 3, 10, 11, 12
- [63] X. Zhang, X. Wang, J. Bian, C. Shen, and M. You, “Diverse knowledge distillation for end-to-end person search,” in *AAAI*, 2021, pp. 3412–3420. 3, 6, 10, 11, 12
- [64] Y. Zhang, C. Wang, X. Wang, W. Zeng, and W. Liu, “Fairmot: On the fairness of detection and re-identification in multiple object tracking,” *CoRR*, vol. abs/2004.01888, 2020. 4
- [65] T. Kong, F. Sun, H. Liu, Y. Jiang, L. Li, and J. Shi, “Foveabox: Beyond anchor-based object detection,” *IEEE Trans. Image Process.*, vol. 29, pp. 7389–7398, 2020. 5
- [66] P. Hong, T. Wu, A. Wu, X. Han, and W.-S. Zheng, “Fine-grained shape-appearance mutual learning for cloth-changing person re-identification,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 10513–10522. 6
- [67] Y. Dai, X. Li, J. Liu, Z. Tong, and L.-Y. Duan, “Generalizable person re-identification with relevance-aware mixture of experts,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 16145–16154. 6
- [68] J. B. Kinney and G. S. Atwal, “Equitability, mutual information, and the maximal information coefficient,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 9, pp. 3354–3359, 2014. 6
- [69] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio, “Learning deep representations by mutual information estimation and maximization,” in *ICLR*, 2019. 6
- [70] M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, R. D. Hjelm, and A. C. Courville, “Mutual information neural estimation,” in *Int. Conf. Mach. Learn.*, 2018, pp. 530–539. 6
- [71] Z. Li and D. Miao, “Sequential end-to-end network for efficient person search,” in *AAAI*, 2021, pp. 2011–2019. 6, 10, 11, 12
- [72] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 770–778. 7
- [73] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li, “Imagenet: A large-scale hierarchical image database,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 248–255. 7
- [74] D. Chen, S. Zhang, W. Ouyang, J. Yang, and B. Schiele, “Hierarchical online instance matching for person search,” in *AAAI*, 2020, pp. 10518–10525. 7, 10, 11
- [75] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Adv. Neural Inform. Process. Syst.*, 2019, pp. 8024–8035. 7
- [76] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, “Mmdetection: Open mmlab detection toolbox and benchmark,” *CoRR*, vol. abs/1906.07155, 2019. 7
- [77] H. Kim, S. Joung, I.-J. Kim, and K. Sohn, “Prototype-guided saliency feature learning for person search,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 4865–4874. 10, 11, 12
- [78] H. Yao and C. Xu, “Joint person objectness and repulsion for person search,” *IEEE Trans. Image Process.*, vol. 30, pp. 685–696, 2021. 11