

PS-ARM: An End-to-End Attention-aware Relation Mixer Network for Person Search

Mustansar Fiaz¹[0000–0003–2289–2284], Hisham Cholakkal¹, Sanath Narayan²,
Rao Muhammad Anwer¹, and Fahad Shahbaz Khan¹

¹ Department of computer Vision, Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE.

([mustansar.fiaz](mailto:mustansar.fiaz@mbzuai.ac.ae), [hisham.cholakkal](mailto:hisham.cholakkal@mbzuai.ac.ae), [rao.anwer](mailto:rao.anwer@mbzuai.ac.ae), [fahad.khan](mailto:fahad.khan@mbzuai.ac.ae))@mbzuai.ac.ae

² Inception Institute of Artificial Intelligence, Abu Dhabi, UAE.

Abstract. Person search is a challenging problem with various real-world applications, that aims at joint person detection and re-identification of a query person from uncropped gallery images. Although, previous study focuses on rich feature information learning, it's still hard to retrieve the query person due to the occurrence of appearance deformations and background distractors. In this paper, we propose a novel attention-aware relation mixer (ARM) module for person search, which exploits the global relation between different local regions within RoI of a person and make it robust against various appearance deformations and occlusion. The proposed ARM is composed of a relation mixer block and a spatio-channel attention layer. The relation mixer block introduces a spatially attended spatial mixing and a channel-wise attended channel mixing for effectively capturing discriminative relation features within an RoI. These discriminative relation features are further enriched by introducing a spatio-channel attention where the foreground and background discriminability is empowered in a joint spatio-channel space. Our ARM module is generic and it does not rely on fine-grained supervisions or topological assumptions, hence being easily integrated into any Faster R-CNN based person search methods. Comprehensive experiments are performed on two challenging benchmark datasets: CUHK-SYSU and PRW. Our PS-ARM achieves state-of-the-art performance on both datasets. On the challenging PRW dataset, our PS-ARM achieves an absolute gain of 5% in the mAP score over SeqNet, while operating at a comparable speed. The source code and pre-trained models are available at ([this https URL](https://github.com/mustansarfiaz/PS-ARM)).

Keywords: Person Search · Transformer · Spatial attention · channel attention.

1 Introduction

Person search is a challenging computer vision problem where the task is to find a target query person in a gallery of whole scene images. The person search methods need to perform pedestrian detection [42,27,29] on the uncropped gallery

images and do re-identification (re-id) [44,25,8] of the detected pedestrians. In addition to addressing the challenges associated with these individual sub-tasks, both these tasks need to be simultaneously optimized within person search. Despite numerous real-world applications, person search is highly challenging due to the diverse nature of person detection and re-id sub-tasks within the person search problem.

Person search approaches can be broadly grouped into two-step [45,4,17] and one-step methods[37,39,6]. In two-step approaches, person detection and re-id are performed separately using two different steps. In the first step a detection network such as Faster R-CNN is employed to detect pedestrians. In the second step detected persons are first cropped and re-sized from the *input image*, then utilized in another independent network for the re-identification of the cropped pedestrians. Although two-step methods provide promising results, they are computationally expensive. Different to two-step methods, one step methods employ a unified framework where the backbone networks are shared for the detection and identifications of persons. For a given uncropped image, one-step methods predict the box coordinates and re-id features for all persons in that image. One-step person search approaches such as [6,24] generally extend Faster R-CNN object detection frameworks by introducing an additional branch to produce re-id feature embedding, and the whole network is jointly trained end-to-end. Such methods often struggle while the target person in the gallery images has large appearance deformations such as pose variation, occlusion, and overlapping background distractions within the region of interest (RoI) of a target person (see Figure. 1).

1.1 Motivation

To motivate our approach, we first distinguish two desirable characteristics to be considered when designing a Faster R-CNN based person search framework that is robust to appearance deformations (e.g. pose variations, occlusions) and background distractions occurring in the query person (see Figure. 1).

Discriminative Relation Features through Local Information Mixing: The position of different local person regions within an RoI can vary in case of appearance deformations such as pose variations and occlusions. This is likely to deteriorate the quality of re-id features, leading to inaccurate person matching. Therefore, a dedicated mechanism is desired that generates discriminative relation features by globally mixing relevant information from different local regions within an RoI. To ensure a straightforward integration into existing person search pipelines, such a mechanism is further expected to learn discriminative relation features without requiring fine-level region supervision or topological body approximations.

Foreground-Background Discriminability for Accurate Local Information Mixing: The quality of the aforementioned relation features rely on the assumption that the RoI region only contains foreground (person) information. However, in real-world scenarios the RoI regions are likely to contain unwanted background

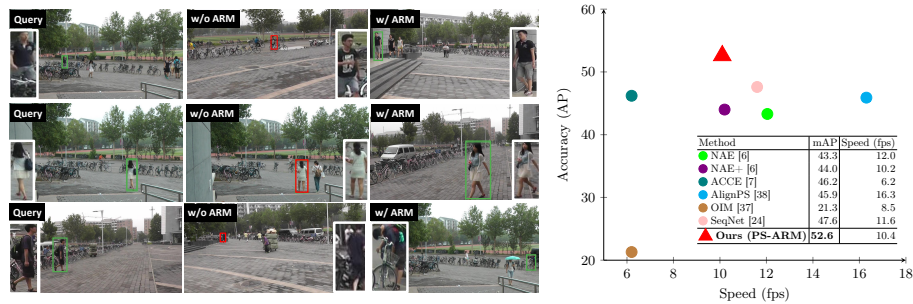


Fig. 1: **On the left:** Qualitative comparison showing different query examples and their corresponding top-1 matching results obtained with *and* without our ARM module in the same base framework. Here, true and false matching results are marked in green and red, respectively. These examples depict appearance deformations and distracting backgrounds in the gallery images for the query person. Our ARM module that explicitly captures discriminative relation features better handle the appearance deformations in these examples. **On the right:** Accuracy (AP) vs. speed (frames per second) comparison with state-of-the-art person search methods on PRW test set. All methods are reported with a Resnet50 backbone and speed is computed over V100 GPU. Our approach (PS-ARM) achieves an absolute mAP gain of 5% over SeqNet while operating at a comparable speed.

information due to less accurate bounding-box locations. Therefore, discriminability of the foreground from the background is essential for accurate local information mixing to obtain discriminative relation features. Further, such a FG/BG discrimination is expected to also improve the detection performance.

1.2 Contribution

We propose a novel end-to-end one-step person search method with the following novel contributions. We propose a novel attention-aware relation mixer (ARM) module that strives to capture global relation between different local person regions through global mixing of local information while simultaneously suppressing background distractions within an RoI. Our ARM module comprises a relation mixer block and a spatio-channel attention layer. The relation mixer block captures discriminative relation features through a spatially-attended spatial mixing and a channel-wise attended channel mixing. These discriminative relation features are further enriched by the spatio-channel attention layer performing foreground/background discrimination in a joint spatio-channel space. Comprehensive experiments are performed on two challenging benchmark datasets: CUHK-SYSU [37] and PRW [46]. On both datasets, our PS-ARM performs favourably against state-of-the-art approaches. On the challenging PRW benchmark, our PS-ARM achieves a mAP score of 52.6%. Our ARM module is

generic and can be easily integrated to any Faster R-CNN based person search methods. Our PS-ARM provides an absolute gain of 5% mAP score over SeqNet, while operating at a comparable speed (see Figure. 1), resulting in a mAP score of 52.6% on the challenging PRW dataset.

2 Related Work

Person search is a challenging computer vision problem with numerous real-world applications. As mentioned earlier, existing person methods can be broadly classified into two-step and one-step methods. Most existing two-step person search approaches address this problem by first detecting the pedestrians, followed by cropping and resizing into a fixed resolution before passing to the re-id network that identifies the cropped pedestrian [46,5,18,11,23]. These methods generally employ two different backbone networks for the detection and re-identification.

On the other hand, several one-step person search methods employ feature pooling strategies such as, RoIPooling or RoIAlign pooling to obtain a scale-invariant representation for the re-id sub-task. [5] proposed a two-step method to learn robust person features by exploiting person foreground maps using pre-trained segmentation network. Han et al. [18] introduced a bounding box refinement mechanism for person localization. Dong et al. [11] utilized the similarity between the query and query-like features to reduce the number of proposals for re-identification. Zhang et al. [46] introduced the challenging PRW dataset. A multi-scale feature pyramid was introduced in [23] for improving person search under scale variations. Wang et al. [35] proposed a method to address the inconsistency between the detection and re-id sub-tasks.

Most one-step person search methods [37,36,26,2,39,10,6,28,16,24] are developed based on Faster R-CNN object detector [31]. These methods generally introduce an additional branch to Faster R-CNN and jointly address the detection and Re-ID subtasks. One of the earliest Faster R-CNN based one-step person approach is [37], which proposed an online instance matching (OIM) loss. Xiao et al. [36] introduced a center loss to explore intra-class compactness. For generating person proposals, Liu et al. [26] introduced a mechanism to iteratively shrink the search area based on query guidance. Similarly, Chang et al. [2] used reinforcement learning to address the person search problem. Chang et al. [39] exploited complementary cues based on graph learning framework. Dont et al. [10] proposed Siamese based Bi-directional Interaction Network (BINet) to mitigate redundant context information outside the BBoxes. On the contrary, Chen et al. [6] proposed Norm Aware Embedding (NAE) to alleviate the conflict between person localization and re-identification by computing magnitude and angle of the embedded features respectively.

Chen et al. [3] developed a Hierarchical Online Instance Matching loss to guide the feature learning by exploiting the hierarchical relationship between detection and re-identification. A query-guided proposal network (QGPN) is proposed by Munjal et al. [28] to learn query guided re-identification score. H Li et al. [24] proposed a Sequential End-to-end Network (SeqNet) to refine the




Fig. 2: The overall architecture of the proposed PS-ARM framework. It comprises a person detection branch (shown in green) and person re-ID branch (shown in blue). The person detection branch predicts the initial box locations whereas the person re-id branch refines the box locations and perform a norm-aware embedding (NAE) to disentangle the detection and re-id. The focus of our design is the introduction of a novel *attention-aware relation Mixer* (ARM) module (shown in grey) to the detection and re-id branches. Our ARM module enriches standard RoI Align pooled features by capturing discriminative relation features between different local regions within an RoI. The resulting enriched features are used for box regression and classification in the detection branch, whereas these features are used to refine the box locations, along with generating a norm-aware embedding for box classification (person vs background) and re-id feature prediction in the re-id branch.

proposals by introducing Faster R-CNN as a proposal generator into the NAE pipeline to get refined features for detection and re-identification. The Faster R-CNN based one-step person search approaches often struggle while the target undergoes large appearance deformations or come across with distracting background objects within RoI. To address this, we propose a novel person search method, PS-ARM, where a novel ARM module is introduced to capture global relation between different local regions within an RoI. Our PS-ARM enables accurate detection and re-identification of person instances under under challenging scenarios such as pose variation and distracting backgrounds (See Figure. 1).

3 Method

3.1 Overall Architecture

Figure. 2 shows overall architecture of the proposed framework. It comprises a person detection branch (shown in green) followed by a person re-ID branch (shown in blue). The person detection branch follows the structure of standard Faster R-CNN, which comprises a ResNet backbone (*res1-res4*), a region proposal network (RPN), RoIAlign pooling, and a prediction head for box regression and classification. The person re-id branch takes the boxes predicted by the person detection branch as input and performs RoIAlign pooling on these predicted box locations. The resulting RoI Align pooled features are utilized to perform re-identification. We adopt norm-aware embedding (NAE) that is designed to

separate the detection and re-identification using shared feature representation. During inference, the person re-id branch takes only unique boxes (obtained by non-maximum suppression algorithm) from the person detection branch and performs a context bipartite graph matching for the re-id similar to [24]. The above-mentioned standard detection and re-id branches serve as a base network to which we introduce our novel attention-aware relation mixer (ARM) module that enriches the RoI features for accurate person search.

The focus of our design is the introduction of a novel ARM module (shown in grey). Specifically, we integrate our ARM module between the RoIAlign and convolution blocks (*res5*) in both the person detection and re-id branches of the base framework, without sharing the parameters between both branches. Our proposed ARM module strives to enrich standard RoI Align pooled features by capturing discriminative relation features between different local regions within an RoI through global mixing of local information. To ensure effective enrichment of RoI Align pooled features, we further introduce a foreground/background discrimination mechanism in our ARM module. Our ARM module strives to simultaneously improve both detection and re-id sub-tasks. Therefore, the output is passed to norm-aware embedding to decouple the features for the contradictory detection and re-id tasks. Furthermore, our ARM module is generic and can be easily integrated to other Faster R-CNN based person search methods. Next, we present the details of the proposed ARM module.

3.2 Attention-aware Relation Mixer (ARM) Module

Our ARM module is shown in Figure. 3. It comprises a relation mixer block and a spatio-channel attention layer. Our relation mixer block captures relation between different sub-regions (local regions) within an RoI. The resulting features are further enriched by our spatio-channel attention that attend to relevant input features in a joint spatio-channel space. Our ARM module takes RoIAlign pooled feature $\mathcal{F} \in \mathbb{R}^{C \times H \times W}$ as input. Here, H , W , C are the height, width and number of channels of the RoI feature. For computational efficiency, the number of channels are reduced to $c = C/4$ through a point (1×1) convolution layer before passing to relation mixer and spatio-channel attention blocks.



Fig. 3: The network structure of our ARM module. The module takes RoI Align pooled features as input and captures inter-dependency between different local regions, while simultaneously suppressing background distractions for the person search problem. To achieve this objective, ARM module comprises a relation mixer block and a joint spatio-channel attention layer.




Fig. 4: Structure of relation mixer block within our ARM module. It comprises a spatially attended spatial mixing operation where important local spatial regions will be emphasized using a spatial attention before globally mixing them across all spatial regions (tokens) within each channel using MLP-1 shared across all channels. Following this spatial mixing, we perform a channel attention to emphasize informative channels before globally mixing the channels for each local spatial region (token) using MLP-2 shared across all spatial regions.

Relation Mixer Block: As mentioned earlier, our relation mixer block is introduced to capture the relation between different sub-regions (local regions) within an RoI. This is motivated by the fact that the local regions of a person share certain ‘standard’ prior relationships among local regions, across RoIs of different person and it is desirable to explicitly learn these inter-dependencies without any supervision. One such module that can learn/encode such inter-dependencies, is the MLP-mixer [33] that performs spatial ‘token’ mixing followed by ‘point-wise’ feature refinement. Compared to other context aggregators [20,34,22], the mlp-mixer is more static, dense, and does not share parameters [13]. The core operation of the MLP-Mixer is transposed affinity matrix on a single feature group, which computes the affinity matrix with non-sharing W_{MLP-1} parameters as: $\mathcal{A} = (W_{MLP-1})^T$. To this end, MLP mixer contains a spatial mixer and a channel mixer. The spatial mixer comprise of a layer norm, skip connection and a token-mixing MLP with two fully-connected layers and a GELU nonlinearity. Similarly, the channel mixer employs a channel-mixing MLP, layer norm, skip connection and dropout. The MLP mixer conceptually acts as a persistent relationship memory that can learn and encode the prior relationships among the local regions of an object at a global level. To this end, we introduce our relation mixer comprising a spatially attended spatial mixer and a channel-wise attended channel mixer. our ARM module with residual connection not only enabled using MLP mixer for the first time in the problem of person search, but also provided impressive performance gain over the base framework.

Spatially attended Spatial Mixer: While learning the inter-dependencies of local RoI sub-regions using standard MLP mixer, the background regions are likely to get entangled with the foreground regions, thereby adversely affecting the resulting feature embedding used for the re-id and box predictions. In order to discriminate the irrelevant background information at token level, we introduce a spatial attention before performing token (spatial) mixing within our MLP mixer for emphasizing the foreground regions. In our spatial attention, we employ pooling operations along the channel axis, followed by convolution and

sigmoid layers to generate a 2D spatial attention weights $M_s \in \mathbb{R}^{1 \times H \times W}$. These attention weights are broadcasted along the channel dimension to generate the spatial attention $M'_s \in \mathbb{R}^{c \times H \times W}$. For a given feature $\mathcal{F}' \in \mathbb{R}^{c \times H \times W}$, we obtain the spatially attended feature map $\mathcal{F}'' = \mathcal{F}' \odot M'_s$. Here \odot denotes element-wise multiplication. These spatially attended features (\mathcal{F}'') are expected to discriminate irrelevant (background) spatial regions from the foreground. These features are (\mathcal{F}'') input to a shared multi-layer perceptron (MLP-1) for globally mixing local features (within \mathcal{F}'') across all spatial regions (tokens). Our spatially attended spatial mixing strives to achieve accurate spatial mixing and outputs the feature map Q (see Figure. 4).

Channel-wise attended Channel Mixer: To further prioritize the feature channels of Q that are relevant for detection and re-id of person instances, we introduce a channel attention before channel mixing. Our channel attention weights $M_c \in \mathbb{R}^{c \times 1}$ are generated through spatial pooling, fully connected (fc) and sigmoid layers, which are broadcasted along the spatial dimension to generate the channel attention weights $M'_c \in \mathbb{R}^{c \times H \times W}$. Similar to spatial attention, these channel weights are element-wise multiplied with the feature map to obtain channel-wise attended feature map. The resulting features are expected to emphasize only the channels that are relevant for effective channel-mixing within our relation mixing block. Our channel mixing employs another shared MLP (MLP-2) for global mixing of channel information. The final output of our relation mixer block results is feature maps $K \in \mathbb{R}^{c \times H \times W}$.

Spatio-channel Attention Layer: Our relation mixer block performs the mixing operations by treating the spatial and channel information in a disjoint manner. But, in many scenarios, all spatial regions within a channel and all channels at a given spatial location are not equally informative. Hence, it is desired to treat the entire spatio-channel information as a joint space. With this objective, we introduce a joint spatio-channel attention layer within our ARM module to further improve the foreground/background discriminability of RoIAlign pooled features. Our spatio-channel attention layer utilizes parameter-free 3D attention weights obtained based on [40] to modulate the 3D spatio-channel RoI pooled features. These spatio-channel attended features are aggregated with the relation mixer output to produce enriched features O for the person search task. These enriched features projected back to C channels ($\mathcal{H} \in \mathbb{R}^{C \times H \times W}$) and taken as input to the *res5* block.

In summary, within our ARM module, the relation mixer targets the global relation between different local regions within RoI and captures the discriminative relation features in disjoint spatial and channel spaces. The resulting features are further enriched by a spatio-channel attention that performs foreground/background discrimination in a joint spatio-channel space.

3.3 Training and Inference

For training and inference, we follow a strategy similar to [6,24]. Our PS-ARM is trained end-to-end with a loss formulation similar to [24]. That is, in the person detection branch, similar to Faster R-CNN, we employ Smooth-L1 and cross

entropy losses for box regression and classifications. For the person re-id branch, we employ three additional loss terms similar to [6] for regression, classification and re-ID. Both these branches are trained by utilizing an IoU threshold of 0.5 for selecting positive and negative samples.

During inference, we first obtain the re-id feature for a given query by using the provided bounding box. Then, for the gallery images, the predicted boxes and their re-id features are obtained from the re-id branch. Finally, employ cosine similarity between the re-id features to match a query person with an arbitrarily detected person in the galley.

4 Experiments

We perform the experiments on two person search datasets (*i.e.*, CUHK-SYSU [37]) and PRW [46] to demonstrate the effectiveness of our PS-ARM and compare it with the state-of-the-art methods.

4.1 Dataset and Evaluation Protocols

CUHK-SYSU [37]: is a large scale person search dataset with 96,143 person bounding boxes from a total of 18,184 images. The training and testing sets contains 11,206 images, 55,272 pedestrians, and 5,532 identities and test set includes 6,978 images, 40,871 pedestrians, and 2,900 identities. Instead of using full gallery during inference, different gallery sizes are used for each query from 50 to 4000. The default gallery size is set to 100.

PRW [46]: is composed of video frames recorded by six cameras that are being installed at different location in Tsinghua University. The dataset has a total 11,816 frames containing 43,110 person bounding boxes. In training set, 5,704 images are annotated with 482 identities. The test set has 2,057 frames are labelled as query persons while gallery set has 6,112 images. Hence, the gallery size of PRW dataset is notably larger compared to CUHK-SYSU gallery set.

Evaluation Protocol: We follow two standard protocol for person search performance evaluation of mean Average Precision (mAP) and top-1 accuracy. The mAP is computed by averaging over all queries with an intersection-over-union (IoU) threshold of 0.5. The top-1 accuracy is measured according to the IoU overlaps between the top-1 prediction and ground-truth with the threshold value set to 0.5.

Implementation Details: We used ResNet-50 as our backbone network. We followed [24] and utilized Stochastic Gradient Descent (SGD), set momentum and decay to 0.9 and 5×10^{-4} , respectively. We trained the model for 12 epochs over CUHK-SYSU dataset PRW dataset. During training, we used the batch-size of 3 with input size 900×1500 and set initial learning rate to 0.003 which is warmed up at first epoch and decayed by 0.1 at 8th epoch. During inference, the NMS threshold value is set to 0.4. The code is implemented in PyTorch [30]. The code and trained model will be publicly released.

Table 1: State-of-the-art comparison on CUHK and PRW test sets in terms of mAP and top-1 accuracy. On both datasets, our PS-ARM performs favourably against existing approaches. All the methods here utilize the same ResNet50 backbone. When compared with recently introduced SeqNet, our PS-ARM provides an absolute mAP gain of 5% on the challenging PRW dataset. Also, introducing our novel ARM module to a popular Faster R-CNN based approach (NAE [6]), provides an absolute mAP gain of 3.6%.

Method		CUHK-SYSU		PRW	
		mAP	top-1	mAP	top-1
Two-step	CLSA [23]	87.2	88.5	38.7	65.0
	IGPN [11]	90.3	91.4	42.9	70.2
	RDLR [18]	93.0	94.2	42.9	70.2
	MGTS [5]	83.0	83.7	32.6	72.1
	MGN+OR [41]	93.2	93.8	52.3	71.5
	TCTS [35]	93.9	95.1	46.8	87.5
End-to-end	OIM [37]	75.5	78.7	21.3	49.9
	RCAA [2]	79.3	81.3	-	-
	NPSM [26]	77.9	81.2	24.2	53.1
	IAN [36]	76.3	80.1	23.0	61.9
	QEEPS [28]	88.9	89.1	37.1	76.7
	CTXGraph [39]	84.1	86.5	33.4	73.6
	HOIM [3]	89.7	90.8	39.8	80.4
	BINet [10]	90.0	90.7	45.3	81.7
	AlignPS [38]	93.1	96.4	45.9	81.9
	PGSFL [21]	92.3	94.7	44.2	85.2
	DKD [43]	93.1	94.2	50.5	87.1
	NAE+ [6]	92.1	94.7	44.0	81.1
	PBNet [32]	90.5	88.4	48.5	87.9
	DIOIM [9]	88.7	89.6	36.0	76.1
	APNet [47]	88.9	89.3	41.2	81.4
	DMRN [19]	93.2	94.2	46.9	83.3
	CAUCPS [15]	81.1	83.2	41.7	86.0
	ACCE [7]	93.9	94.7	46.2	86.1
	NAE [6]	91.5	92.4	43.3	80.9
	SeqNet [24]	94.8	95.7	47.6	87.6
	Ours (NAE + ARM)	93.4	94.2	46.9	81.4
	Ours (PS-ARM)	95.2	96.1	52.6	88.1
	Ours (Cascaded PS-ARM)	-	-	53.1	88.3

4.2 Comparison with State-of-the-art Methods

Here, we compare our approach with state-of-the-art one-step and two-step person search methods in literature on two datasets: CUSK-SYSU and PRW.

CUHK-SYSU Comparison: Table. 1 shows the comparison of our PS-ARM with state-of-the-art two-step and single-step end-to-end methods with the gallery size of 100. Among existing two-step methods, MGN+OR [41] and TCTS [35] achieves mAP of 93.2 and 93.9, respectively. Among existing single-step end-to-end methods, SeqNet [24] and AlignPS [38] obtains mAP of 94.8%, 93.1% respectively.



Fig. 6: Qualitative comparison between the top-1 results obtained from SeqNet (row 2) and our PS-ARM (row3) for the same query input (row 1). Here, true and false matching results are marked in green and red, respectively. SeqNet provides inaccurate predictions due to the appearance deformations in these examples whereas our PS-ARM provides accurate predictions by explicitly capturing discriminative relation features within RoI.

To further analyse the benefits of our ARM module, we introduced the proposed ARM module in to a Faster R-CNN based method (NAE [6] method) after RoIAlign pooling. We observed that our ARM module can provide an absolute gains of 1.9% and 1.8% to the mAP and top-1 accuracies over NAE (see Table 1). Our PS-ARM outperforms all existing methods, and achieves a mAP score of 95.2. In terms of top-1 accuracy our method sets a state-of-the-art accuracy of 96.1%.

CUHK-SYSU dataset has different range of gallery sizes such as 50, 100, 500, 1000, 2000, and 4000. To further analyze our proposed method, we performed an experiment by varying the gallery size. Our mAP scores across different gallery size are compared with recent one-stage and two-stage methods as shown in Figure 5. The results shows that our PS-ARM provides consistent performance gain over other approaches across all gallery sizes.




Fig. 5: State-of-the-art comparison of existing methods over CUHK-SYSU dataset with varying gallery sizes. Dotted lines represent two-stage methods whereas solid lines represent one-stage methods. Our PS-ARM shows consistent improvement compared to other methods as the size of gallery increases.



Fig. 7: Qualitative results of our PS-ARM on challenging PRW dataset. The top-2 matching results for each query image is shown. Our PS-ARM accurately detect and re-identify the query person in both images.

PRW Comparison: Table 1 shows the state-of-the-art comparison on PRW dataset. Among the existing two stage methods, MGN+OR [41] achieves the best mAP score 52.3, but with a very low top-1 accuracy. While comparing the top-1 accuracy, TCTS [35] provides the best performance, but with a very low mAP score. To summarize, the performance of most two-step methods [41,5,18,14,11,23] are inferior either in mAP score or top-1 accuracy.

Among one-stage methods, NAE [6] and AlignPS [38], achieved mAP scores of 43.3% and 45.9%. These methods achieved top-1 accuracies of 80.9% and 81.9%. Among the other one-step methods SeqNet [24], PBNet [32], DMRN [19], and DKD [43] also performed well and obtain more than 46% mAP and have more than 86% top-1 accuracy.

To further analyze the effectiveness of our ARM module, we integrate our ARM module to NAE and achieved absolute mAP gain of 3.6% mAP, leading to an mAP score of 46.9%. We observe a similar performance gain over top-1 accuracy, resulting in top-1 score of 81.4%. We also introduced proposed ARM module in Han’s [16] method. Compared to the existing methods, [16] utilize a different approach, such as an RoI pooling of 24×8 size, instead of 14×14 . To this end, we modified our PS-ARM to adapt the setting of [16], resulting in an absolute gains of 2% and 1.3% improvement on PRW dataset and obtained 55.3% mAP and 89.0% top-1 scores, respectively.




Fig. 8: Person search and detection scores on PRW dataset with and without provided ground-truth detection boxes. The * indicates the results using ground-truth boxes.

Table 2: Ablation study over the PRW dataset by incrementally adding our novel contributions to the baseline. While introducing a MLP mixer to a baseline, both the detection and re-id performance increases over the baseline except top-1. The spatially attended spatial mixing and channel-wise attended channel mixing within our relation mixer captures discriminative relation features within RoI while suppressing distracting background features, hence provides superior re-id performance. Finally, our joint spatio-channel attention removes distracting backgrounds in a joint spatio-channel space, leading to improved detection and re-id performance.

Method	ReID		Detection	
	mAP	top-1	Recall	AP
Baseline	47.6	87.6	96.3	93.1
Baseline + MLP-Mixer	49.1	86.8	96.3	93.3
Baseline + Transformer	47.9	85.8	96.1	93.5
Baseline + Spatio-channel Attention Layer	48.1	86.2	95.2	93.0
Baseline + Spatial Mixing + Channel-wise Attended Channel Mixing	49.4	86.7	96.2	93.2
Baseline + Spatially Attended Spatial Mixing + Channel Mixing	49.5	86.9	96.5	93.4
Baseline + Relation Mixer	51.8	87.9	96.6	93.8
PS-ARM (Baseline + ARM)	52.6	88.1	97.1	93.9

Our PS-ARM achieve state-of-the-art performance compared the existing one-step and two-step methods. We achieve an mAP score of 52.6% and top-1 score of 88.1%.

Besides, similar to cascade RCNN [1], we extend our person search network by introducing an other person re-id branch, called Cascaded PS-ARM. This newly introduced branch takes refined bounding boxes from the Box2 as an input to perform RoIAlign pooling. This strategy further refines the detection and re-identification, producing improved mAP 53.1% and top-1 88.3 % scores. **Qualitative comparison:** Figure. 6 shows qualitative comparison between the SeqNet [24] (row 2) and our PS-ARM for the same query input (row 1). Here, true and false matching results are marked in green and red, respectively. The figure shows top-1 results obtained from both methods. It can be observed that SeqNet provides inaccurate top-1 predictions due to the appearance deformations. Our PS-ARM provides accurate predictions on these challenging examples by explicitly capturing discriminative relation features within RoI. Figure. 7 shows the qualitative results from our PS-ARM. Here we show the top-2 matching results for each query image. It can be seen that our PS-ARM can accurately detect and re-identify the query person in both gallery images.

4.3 Ablation study

Here, we perform the ablation study on the PRW dataset. Table. 2 shows the performance gain obtained by progressively integrating our novel contributions to the baseline. First we verify the effectiveness of the context aggregators including MLP-Mixer [33] and Transformer [12] within the proposed framework.

The experiment shows that choice of MLP-mixer is better. Moreover, we apply joint spatio-channel attention on the RoI feature maps which results in improved performance compared to baseline. Further, we investigate the introduction of spatially-attended spatial mixing and channel-wise attended channel mixing within our relation mixer which captures discriminative relation features within RoI while suppressing distracting background features. This resulted in superior re-id performance. Introducing our relation mixer comprising of a spatially attended spatial mixing and channel-wise attended channel mixing leads to an overall AP of 93.8 for detection and 51.8 mAP for re-id. To further complement the relation mixer that performs information mixing in the disjoint spatial and channel spaces, we introduce a joint spatio-channel attention. Our joint spatio-channel attention removes distracting backgrounds in a joint spatio-channel space, leading to improved detection and re-id performance by achieving 94.1 and 52.6, respectively.

Relation between Detection and ReID In Figure 8, we validate the effectiveness of the proposed PS-ARM to deal with the contradictory detection and ReID objectives. We compared our PS-ARM with the SOTA SeqNet [24] and NAE [6]. We notice that PS-ARM* and NAE+ARM* outperforms their counterparts provided the ground-truth boxes.

5 Conclusions

We propose a novel person search method named PS-ARM, that strives to capture global relation between different local regions within RoI of a person. The focus of our design is introduction of a novel ARM module, which effectively capturing the global relation within an RoI and make robust against occlusion. The relation mixer block introduces a spatially attended spatial mixing, a channel-wise attended channel mixing, and an input-output feature re-using for capturing discriminative relation features within an RoI. An additional spatio-channel attention layer is introduced within the ARM module to further enrich the discriminability between the foreground/background features in a joint spatio-channel space. Our ARM module is generic and it can be easily integrated to any Faster R-CNN based person search methods. Comprehensive experiments are performed on two benchmark datasets. We achieve state-of-the-art performance on both datasets, demonstrating the merits of our novel contributions.

References

1. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6154–6162 (2018)
2. Chang, X., Huang, P.Y., Shen, Y.D., Liang, X., Yang, Y., Hauptmann, A.G.: Rcaa: Relational context-aware agents for person search. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 84–100 (2018)

3. Chen, D., Zhang, S., Ouyang, W., Yang, J., Schiele, B.: Hierarchical online instance matching for person search. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 10518–10525 (2020)
4. Chen, D., Zhang, S., Ouyang, W., Yang, J., Tai, Y.: Person search via a mask-guided two-stream cnn model. *Proc. European Conference on Computer Vision* (2018)
5. Chen, D., Zhang, S., Ouyang, W., Yang, J., Tai, Y.: Person search via a mask-guided two-stream cnn model. In: *Proceedings of the european conference on computer vision (ECCV)*. pp. 734–750 (2018)
6. Chen, D., Zhang, S., Yang, J., Schiele, B.: Norm-aware embedding for efficient person search. In: *CVPR*. pp. 12615–12624 (2020)
7. Chen, S., Zhuang, Y., Li, B.: Learning context-aware embedding for person search. *arXiv preprint arXiv:2111.14316* (2021)
8. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: A deep quadruplet network for person re-identification. *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (2017)
9. Dai, J., Zhang, P., Lu, H., Wang, H.: Dynamic imposter based online instance matching for person search. *Pattern Recognition* **100**, 107120 (2020)
10. Dong, W., Zhang, Z., Song, C., Tan, T.: Bi-directional interaction network for person search. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2839–2848 (2020)
11. Dong, W., Zhang, Z., Song, C., Tan, T.: Instance guided proposal network for person search. In: *CVPR*. pp. 2585–2594 (2020)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *Proc. International Conference on Learning Representations* (2020)
13. Gao, P., Lu, J., Li, H., Mottaghi, R., Kembhavi, A.: Container: Context aggregation network. *arXiv preprint arXiv:2106.01401* (2021)
14. Girshick, R., Iandola, F., Darrell, T., Malik, J.: Deformable part models are convolutional neural networks. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 437–446 (2015)
15. Han, B.J., Ko, K., Sim, J.Y.: Context-aware unsupervised clustering for person search. *arXiv preprint arXiv:2110.01341* (2021)
16. Han, B.J., Ko, K., Sim, J.Y.: End-to-end trainable trident person search network using adaptive gradient propagation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 925–933 (2021)
17. Han, C., Ye, J., Zhong, Y., Tan, X., Zhang, C., Gao, C., Sang, N.: Re-id driven localization refinement for person search. *Proc. IEEE International Conference on Computer Vision* (2019)
18. Han, C., Ye, J., Zhong, Y., Tan, X., Zhang, C., Gao, C., Sang, N.: Re-id driven localization refinement for person search. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9814–9823 (2019)
19. Han, C., Zheng, Z., Gao, C., Sang, N., Yang, Y.: Decoupled and memory-reinforced networks: Towards effective feature learning for one-step person search. *arXiv preprint arXiv:2102.10795* (2021)
20. Kaiser, L., Gomez, A.N., Chollet, F.: Depthwise separable convolutions for neural machine translation. *arXiv preprint arXiv:1706.03059* (2017)
21. Kim, H., Joung, S., Kim, I.J., Sohn, K.: Prototype-guided saliency feature learning for person search. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4865–4874 (2021)

22. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
23. Lan, X., Zhu, X., Gong, S.: Person search by multi-scale matching. In: Proceedings of the European conference on computer vision (ECCV). pp. 536–552 (2018)
24. Li, Z., Miao, D.: Sequential end-to-end network for efficient person search. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2011–2019 (2021)
25. Liao, S., Hu, Y., Zhu, X., Li, S.Z.: Person re-identification by local maximal occurrence representation and metric learning. Proc. IEEE Conference on Computer Vision and Pattern Recognition (2015)
26. Liu, H., Feng, J., Jie, Z., Jayashree, K., Zhao, B., Qi, M., Jiang, J., Yan, S.: Neural person search machines. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 493–501 (2017)
27. Liu, W., Liao, S., Ren, W., Hu, W., Yu, Y.: High-level semantic feature detection: A new perspective for pedestrian detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5187–5196 (2019)
28. Munjal, B., Amin, S., Tombari, F., Galasso, F.: Query-guided end-to-end person search. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 811–820 (2019)
29. Pang, Y., Xie, J., Khan, M.H., Anwer, R.M., Khan, F.S., Shao, L.: Mask-guided attention network for occluded pedestrian detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4967–4975 (2019)
30. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
31. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
32. Tian, K., Huang, H., Ye, Y., Li, S., Lin, J., Huang, G.: End-to-end thorough body perception for person search. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12079–12086 (2020)
33. Tolstikhin, I.O., Houlsby, N., et al.: Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems* **34** (2021)
34. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Proc. *Advances in Neural Information Processing Systems* (2017)
35. Wang, C., Ma, B., Chang, H., Shan, S., Chen, X.: Tcts: A task-consistent two-stage framework for person search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11952–11961 (2020)
36. Xiao, J., Xie, Y., Tillo, T., Huang, K., Wei, Y., Feng, J.: Ian: the individual aggregation network for person search. *Pattern Recognition* **87**, 332–340 (2019)
37. Xiao, T., Li, S., Wang, B., Lin, L., Wang, X.: Joint detection and identification feature learning for person search. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3415–3424 (2017)
38. Yan, Y., Li, J., Qin, J., Bai, S., Liao, S., Liu, L., Zhu, F., Shao, L.: Anchor-free person search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7690–7699 (2021)
39. Yan, Y., Zhang, Q., Ni, B., Zhang, W., Xu, M., Yang, X.: Learning context graph for person search. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2158–2167 (2019)

40. Yang, L., Zhang, R.Y., Li, L., Xie, X.: Simam: A simple, parameter-free attention module for convolutional neural networks. In: International Conference on Machine Learning. pp. 11863–11874. PMLR (2021)
41. Yao, H., Xu, C.: Joint person objectness and repulsion for person search. IEEE Transactions on Image Processing **30**, 685–696 (2020)
42. Zhang, S., Yang, J., Schiele, B.: Occluded pedestrian detection through guided attention in cnns. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6995–7003 (2018)
43. Zhang, X., Wang, X., Bian, J.W., Shen, C., You, M.: Diverse knowledge distillation for end-to-end person search. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3412–3420 (2021)
44. Zheng, L., Zhang, H., Sun, S., Chandraker, M., Yang, Y., Tian, Q.: Person re-identification in the wild. Proc. IEEE Conference on Computer Vision and Pattern Recognition (2017)
45. Zheng, L., Zhang, H., Sun, S., Chandraker, M., Yang, Y., Tian, Q.: Person re-identification in the wild. Proc. IEEE Conference on Computer Vision and Pattern Recognition (2017)
46. Zheng, L., Zhang, H., Sun, S., Chandraker, M., Yang, Y., Tian, Q.: Person re-identification in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1367–1376 (2017)
47. Zhong, Y., Wang, X., Zhang, S.: Robust partial matching for person search in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6827–6835 (2020)