

# Domain Generalizable Person Search Using Unreal Dataset

Minyoung Oh\*, Duhyun Kim\*, and Jae-Young Sim†

Graduate School of Artificial Intelligence, UNIST, Republic of Korea  
 {mmyy2513, duhyunkim, jysim}@unist.ac.kr

## Abstract

Collecting and labeling real datasets to train the person search networks not only requires a lot of time and effort, but also accompanies privacy issues. The weakly-supervised and unsupervised domain adaptation methods have been proposed to alleviate the labeling burden for target datasets, however, their generalization capability is limited. We introduce a novel person search method based on the domain generalization framework, that uses an automatically labeled unreal dataset only for training but is applicable to arbitrary unseen real datasets. To alleviate the domain gaps when transferring the knowledge from the unreal source dataset to the real target datasets, we estimate the fidelity of person instances which is then used to train the end-to-end network adaptively. Moreover, we devise a domain-invariant feature learning scheme to encourage the network to suppress the domain-related features. Experimental results demonstrate that the proposed method provides the competitive performance to existing person search methods even though it is applicable to arbitrary unseen datasets without any prior knowledge and re-training burdens.

## Introduction

Person search is a technique to detect the person instances from the scene images first, and then find a query person among the detected instances. Recently, it has been drawing a lot of attention in various computer vision applications such as surveillance and life logging. In general, large datasets of labeled scene images, captured under diverse environments, are required to train the person search networks. However, collecting such datasets is a time-consuming task, and furthermore, it usually requires a great deal of effort to obtain the ground truth labels by human annotation such as the bounding boxes and identities of persons. In addition, real datasets including personal information often suffer from the privacy issues.

To reduce the burden of data labeling, attempts have been made such as weakly supervised learning (Han, Ko, and Sim 2021a; Han et al. 2021; Yan et al. 2022) and unsupervised domain adaptation (DA) (Li et al. 2022), whose concepts are compared in Figure 1. The weakly supervised methods assume that only the bounding box labels are given without the

\*These authors contributed equally.

†Corresponding author.

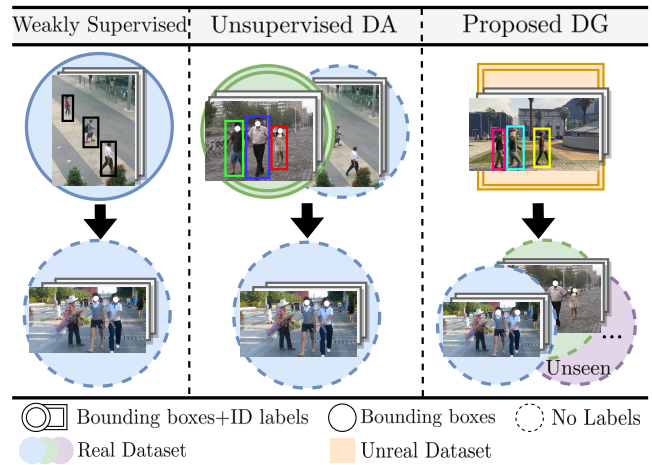


Figure 1: The proposed domain generalization concept compared to the weakly supervised and unsupervised domain adaptation methods. The upper and lower figures represent the training datasets and the test datasets, respectively.

ID labels in the training dataset. On the other hand, the unsupervised DA method considers source and target datasets, respectively, where the source dataset has the labels of both the bounding boxes and identities, but the target dataset has no labels at all. It uses the labeled source dataset and the unlabeled target dataset together for training. However, both approaches are not fully generalizable to be directly applied to arbitrary unseen datasets without additional training burdens, since they still require partial labels and/or need to re-train the networks for a given target dataset.

We propose a fully generalizable person search framework based on domain generalization (DG) from unreal dataset to arbitrary real datasets. In practice, we employ the unreal dataset of JTA (Joint Track Auto) (Fabbri et al. 2018), where the detailed labels were automatically annotated, as the only source dataset used for training. Then we test the trained network on arbitrary unseen target datasets captured in real environment. By using the unreal dataset, we are free from the time-consuming and labor-intensive labeling burdens as well as the privacy issues of real datasets. However, the knowledge transfer from an unreal dataset to the real




Figure 2: The characteristics of the real PRW (left) and unreal JTA (right) datasets. The identity labels of persons are shown at the top of the bounding boxes.

datasets suffers from huge domain gaps that usually degrade the performance of person search. Specifically, we observe that manually annotated datasets often include incorrectly labeled and/or unlabeled instances, as shown in Figure 2 (a). On the other hand, the automatically labeled unreal dataset always provides the correct labels even for some instances with degraded visibility due to severe occlusion, low contrast, or low resolution, as shown in Figure 2 (b). To alleviate the domain gaps of annotation between the unreal and real datasets, we estimate the fidelity of each person instance using the deep features, which is used for fidelity adaptive training. Moreover, we regard each sequence in the unreal training dataset as each domain, and force the network to learn the domain-invariant features while disentangling the domain-specific information from the ID-specific features.

The main contributions of this paper are as follows.

- To the best of our knowledge, we first propose a novel framework of generalizable person search where only an unreal dataset is used for training, and arbitrary unlabeled real datasets can be tested at the inference phase.
- We develop the fidelity adaptive training and domain-invariant feature learning to alleviate the domain gaps between the unreal and real datasets improving the generalization capability.
- We show that the proposed method provides the competitive performance to the existing weakly-supervised and unsupervised DA methods, even though it is free from the re-training burdens and privacy issues.

## Related Work

**Person Search** The supervised methods of person search have been proposed that require labor-intensive labeling burdens. Xiao et al. provided CUHK-SYSU (Xiao et al. 2017) dataset with the annotated ground truth labels of bounding boxes and identities. They proposed an end-to-end framework where the detection and re-identification networks are trained simultaneously. Zheng et al. introduced PRW (Zheng et al. 2017) dataset with the annotated labels. They reflected the detection confidence to improve the re-identification accuracy. Chen et al. decomposed the feature vector of each person instance into the norm and angle to overcome the conflict problem between the detection and re-identification

tasks (Chen et al. 2020b). Li and Miao performed the detection and re-identification progressively using the additional Faster-RCNN head (Li and Miao 2021). Han, Ko, and Sim adopted a part classifier network to prevent the overfitting and trained the network by weighting the detection confidence adaptively (Han, Ko, and Sim 2021b). Yu et al. tackled the occlusion problem by exchanging the tokens between the proposals based on the transformer (Yu et al. 2022).

To overcome the labeling burdens, the weakly supervised person search methods have been studied that assume only the bounding boxes are labeled without ID labels. Han, Ko, and Sim devised a context-aware clustering method using the uniqueness property that multiple persons in a certain scene image do not have the same ID, and the co-appearance property where the neighboring persons tend to appear simultaneously (Han, Ko, and Sim 2021a). Han et al. trained the network to yield more reliable results of re-identification by using both features from the scene-level proposals and the cropped bounding boxes (Han et al. 2021). Yan et al. used the context information to enhance the clustering accuracy (Yan et al. 2022). On the other hand, the unsupervised person search method was also introduced based on DA (Li et al. 2022), which uses the labeled source dataset and unlabeled target dataset together for training. However, the weakly supervised methods still need partial labels of target datasets, and the unsupervised method should re-train the network whenever a target dataset is newly given.

**Domain Generalization** The DG techniques aim to design robust networks when tested on any unseen dataset while using the limited training datasets. There have been three main ways to improve the generalization capability of image classification and segmentation: data augmentation (Qiao, Zhao, and Peng 2020; Wang et al. 2021; Volpi et al. 2018), meta-learning (Li et al. 2018; Balaji, Sankaranarayanan, and Chellappa 2018; Qiao and Peng 2021), and representation learning (Segu, Tonioni, and Tombari 2023; Motiian et al. 2017; Fan et al. 2021).

Recently, the DG techniques have been adopted for person re-identification tasks. Jin et al. normalized the style variations across the different domains and restored the lost ID-related information caused by the instance normalization (Jin et al. 2020). Choi et al. adopted the batch-instance normalization layers trained with the meta-learning strategy to avoid the overfitting to the source domain (Choi et al. 2021). Liu et al. defined a hybrid domain composed of the datasets from multiple domains, and trained the dataset in the hybrid domain with an ensemble of other batch normalization parameters to encourage the generalization capability (Liu et al. 2022). However, these methods were devised for re-identification and cannot be directly applied to the person search task combined with the addition person detection.

## Unreal Dataset

We used the unreal dataset of JTA (Fabbri et al. 2018) obtained from the photo-realistic video game Grand Theft Auto V, where the details for each person instance are automatically annotated such as the bounding boxes, identities, and keypoints. JTA provides about 450,000 images



Figure 3: Images from the unreal JTA dataset.

		JTA*	CUHK-SYSU	PRW
#Images	Train	10,049	11,206	5,134
	Test	4,426	6,978	6,112
#Persons	Train	175,035	55,272	16,243
	Test	74,382	40,871	25,062
#IDs	Train	10,912	5,532	482
	Test	1,480	2,900	450

Table 1: The specifications of datasets.

extracted from 512 video sequences with diverse characteristics, such as the background, viewpoint, and weather condition, as shown in Figure 3. We constructed the JTA\* dataset based on JTA for the purpose of person search by taking 256 sequences in the training category of JTA, which are then divided into 226 sequences for training and 30 sequences for test, respectively. We selected every tenth image from the training sequences and every sixth image from the test sequences, respectively. JTA\* is expected to serve as a more reliable training dataset, since many different identities can be used as the negative samples to improve the performance from the perspective of contrastive learning (Xiao et al. 2017; Chen et al. 2020c). Moreover, JTA\* has no incorrectly labeled or unlabeled instances at all with the help of automatic annotation. However, the unreal dataset does not completely capture the styles of real-world scenes in general, which makes it hard to transfer the knowledge learned from the unreal dataset to the real datasets. For example, the instances with severely degraded visibility tend to be undetected as persons in real datasets. Accordingly, using all the instances in unreal dataset for training may degrade the performance of person search when tested in real datasets. Therefore, we only used the person instances, where the numbers of occluded keypoints are less than 13, to exclude severely occluded instances from training. Table 1 shows the specifications of JTA\* that exhibits relatively larger numbers of person identities and instances compared to the other existing real datasets of CUHK-SYSU and PRW.

## Method

We train the unreal dataset JTA\* as the only source dataset, and test the trained network on arbitrary real target datasets based on the DG framework for person search. To alleviate the domain gaps between the unreal and real datasets,

we first estimate the fidelity of person instance in JTA\* by extracting the deep features. Then we use the estimated fidelity to adaptively train the network suppressing the influence of degraded person instances which are difficult to be identified. Furthermore, we also improve the generalization capability of network by disentangling the domain and ID-specific features to reduce the dependency on the domain information. Figure 4 shows the overall framework of the proposed method.

## Fidelity Adaptive Training

When all the instances with degraded visibility in the unreal dataset are used for training, the network tends to overfit to the source dataset and thus yields low performance on real target datasets. We may remove such degraded instances from the training dataset by using the automatically annotated information, e.g., the keypoints with occlusion information and the size of bounding box. However, it is not trivial to set a criterion for the fidelity of instance in terms of the performance of person search. Moreover, some of the degraded instances in the training dataset may help to improve the robustness of the network to identify the challenging person instances in real datasets. In order to strike a balance between suppressing the effect of degraded instances and improving the robustness of network, we estimate the fidelity of person instances which is then used to train the network adaptively.

**Fidelity Estimation.** We basically estimate the fidelity as the visibility or quality of the image. To this end, we use the BRISQUE (Mittal, Moorthy, and Bovik 2012) which measures the naturalness of image. The high values of BRISQUE score represent severely distorted or noisy images. Figure 5 compares the distributions of BRISQUE scores computed over the bounding box images of person instances among the three datasets of CUHK-SYSU, PRW, and JTA\*. Whereas most of the person instances in the real datasets of CUHK-SYSU and PRW have lower BRISQUE scores than 60, many instances in the unreal dataset of JTA\* exhibit relatively higher scores. For example, we see blurred and/or low contrast images at the scores around 70.

We train the fidelity estimation network, composed of the four convolutional layers and a fully-connected layer, by minimizing the fidelity estimation loss  $\mathcal{L}_{\text{fid}}$ .

$$\mathcal{L}_{\text{fid}} = \frac{1}{|\Omega_p|} \sum_{i \in \Omega_p} \{\eta_i - \exp(-b_i/\tau_{\text{fid}})\}^2, \quad (1)$$

where  $\eta_i$  denotes the estimated fidelity of the  $i$ -th person instance,  $b_i$  is the BRISQUE score measured on the  $i$ -th person instance,  $\tau_{\text{fid}}$  is a hyperparameter, and  $\Omega_p$  denotes the index set of the predicted person instances among all the proposals in a batch. Note that we do not determine the fidelity directly from the BRISQUE score, but we estimate the fidelity values by extracting the deep features. It means that even the person instances with similar BRISQUE scores may be assigned largely different fidelity values according to their actual appearance or visibility. Therefore, the proposed network flexibly learns the relationship between the level of naturalness of image and the actual fidelity.





to learn the domain-invariant and ID-specific features, respectively. We regard each sequence in the unreal dataset as a unique domain assuming that it represents different characteristics such as the background, viewpoint, and weather condition. We employ the attention encoder network composed of the global average pooling and the convolutional layers to learn a channel attention vector. As shown in Figure 4, two attention vectors of  $\mathbf{m}_i^{\text{id}} \in \mathbb{R}^c$  and  $\mathbf{m}_i^{\text{dom}} \in \mathbb{R}^c$ , where  $c$  is the number of channels, are independently extracted from the  $i$ -th person instance.

**Domain-Guided Feature Normalization.** We multiply each element in  $\mathbf{m}_i^{\text{id}}$  and  $\mathbf{m}_i^{\text{dom}}$  with each feature map at the corresponding channel to extract the ID-specific feature map,  $\mathbf{X}_i^{\text{id}} \in \mathbb{R}^{c \times h \times w}$ , and the domain-specific feature map,  $\mathbf{X}_i^{\text{dom}} \in \mathbb{R}^{c \times h \times w}$ , respectively, where  $h$  and  $w$  indicate the height and width of the feature map. To improve the domain-agnostic ID discriminative capability, we additionally normalize  $\mathbf{X}_i^{\text{id}}$  by using the statistics of  $\mathbf{X}_i^{\text{dom}}$  such that

$$\tilde{\mathbf{X}}_i^{\text{id}} = \frac{\mathbf{X}_i^{\text{id}} - \mu(\mathbf{X}_i^{\text{dom}})}{\sigma(\mathbf{X}_i^{\text{dom}})}, \quad (5)$$

where  $\tilde{\mathbf{X}}_i^{\text{id}}$  is the result of the domain-guided normalization (DGN), and  $\mu$  and  $\sigma$  denote the operations to calculate the mean and standard deviation at each channel of feature map, respectively. Whereas  $\tilde{\mathbf{X}}_i^{\text{id}}$  is mapped into  $\mathbf{x}_i^{\text{id}}$  by the head network both at the training and test phases together,  $\mathbf{X}_i^{\text{dom}}$  is fed into the domain encoder network at the training phase only, yielding a domain-specific feature vector  $\mathbf{x}_i^{\text{dom}}$ .

**Domain Separation Loss.** Note that, when both  $\mathbf{m}_i^{\text{id}}$  and  $\mathbf{m}_i^{\text{dom}}$  become identical to each other, the DGN operation becomes equivalent to the instance normalization (Ulyanov, Vedaldi, and Lempitsky 2016). To exploit the benefit of DGN by learning the distinct features from each other, we suggest a domain separation loss given by

$$\mathcal{L}_{\text{sep}} = \exp(-\text{MMD}^2(\{\mathbf{x}_i^{\text{dom}}\}, \{\mathbf{x}_i^{\text{id}}\})), \text{ for } \forall i \in \Omega_p, \quad (6)$$

where  $\text{MMD}(A, B)$  means the mean maximum discrepancy between two sets of  $A$  and  $B$  (Gretton et al. 2012). By maximizing the difference of distribution between the sets of  $\mathbf{x}_i^{\text{dom}}$  and  $\mathbf{x}_i^{\text{id}}$ , we force them construct unique distributions with respect to ID and domain, respectively.

**Domain Feature Update.** To extract a representative  $\mathbf{x}_i^{\text{dom}}$  for its own domain, we first build a new domain look-up table (DLT) and update the DLT such that

$$\mathbf{f}_{s_i}^{\text{dom}} \leftarrow w_{\text{dom}} \mathbf{f}_{s_i}^{\text{dom}} + (1 - w_{\text{dom}}) \mathbf{x}_i^{\text{dom}}, \quad (7)$$

where  $\mathbf{f}_s^{\text{dom}}$  is the feature vector corresponding to the  $s$ -th element in DLT that represents the domain characteristics of the  $s$ -th sequence in the unreal training dataset,  $s_i$  is the ground truth sequence label where  $\mathbf{x}_i^{\text{dom}}$  belongs to, and  $w_{\text{dom}}$  is a momentum parameter. The updated feature vector  $\mathbf{f}_s^{\text{dom}}$  is also normalized to have the length of 1. Based on the DLT, we introduce the domain loss  $\mathcal{L}_{\text{dom}}$  as follows.

$$\mathcal{L}_{\text{dom}} = -\frac{1}{|\Omega_p|} \sum_{i \in \Omega_p} \log \frac{\exp(\langle \mathbf{f}_{s_i}^{\text{dom}}, \mathbf{x}_i^{\text{dom}} \rangle / \tau_{\text{dom}})}{\sum_{j=1}^D \exp(\langle \mathbf{f}_j^{\text{dom}}, \mathbf{x}_i^{\text{dom}} \rangle / \tau_{\text{dom}})}, \quad (8)$$

where  $D$  is the size of DLT, and  $\tau_{\text{dom}}$  is a temperature parameter. By maximizing the cosine similarity of  $\mathbf{x}_i^{\text{dom}}$  to  $\mathbf{f}_{s_i}^{\text{dom}}$  while minimizing that to the others,  $\mathcal{L}_{\text{dom}}$  enhances the domain-specific representation of  $\mathbf{x}_i^{\text{dom}}$ .

## Experimental Results

### Experimental Setup

**Datasets.** We used the unreal dataset of JTA\* only for training, and used the real datasets of CUHK-SYSU (Xiao et al. 2017) and PRW (Zheng et al. 2017) for testing. CUHK-SYSU consists of various scene images captured with a moving camera and the frames selected from the movies and TV shows. PRW includes the images captured by 6 fixed cameras with different locations and viewing directions. The dataset specifications are summarized in Table 1.

**Evaluation Measures.** We used the Precision and Recall to evaluate the detection performance, and used the mean Average Precision (mAP) and Top-1 scores for re-identification performance. Only the proposals with larger than 0.5 IoU to the ground truth bounding boxes are used to evaluate the Top-1 scores.

**Implementation Details.** We adopted the end-to-end person search network (Chen et al. 2020b) with the adaptive gradient weighting function (Han, Ko, and Sim 2021b) as a baseline, where the detection and re-identification networks are trained simultaneously. We used PyTorch for all experiments with a single NVIDIA RTX-3090 GPU. We used the ImageNet pre-trained ResNet50 as our backbone network for a fair comparison. We set the batch size to 4 and used the SGD optimizer with a momentum of 0.9. The warm-up learning rate scheduler linearly increases the learning rate from 0 to 0.003 during the first epoch, and the learning rate decays by multiplying 0.1 every third epoch. We empirically set the weights of losses to 10 and 0.1 for  $\mathcal{L}_{\text{fid}}$  and  $\mathcal{L}_{\text{dom}}$ , respectively, and 1 otherwise.  $\tau_{\text{fid}} = 200$  in (1),  $\tau_{\text{id}} = 1/30$  in (4),  $\tau_{\text{dom}} = 1$  in (8),  $w_{\text{id}} = 2/3$  in (3), and  $w_{\text{dom}} = 2/3$  in (7). During the training phase, we applied the Resize and HorizontalFlip transformations with the probability of 0.5.

### Performance Comparison

Note that the proposed framework of domain generalizable person search is first introduced in this paper, and there is no existing method fairly comparable to the proposed one. Instead, we compared the quantitative performance of the proposed method and the existing person search methods with different experimental settings including the supervised, weakly-supervised, and unsupervised DA methods, as shown in Table 2. Whereas all the compared existing methods use the target test dataset for training in any way, the proposed method does not access test datasets at all during training. Nevertheless, the proposed method provides the comparable performance to the existing methods and even surpasses several supervised and weakly-supervised methods on both target datasets.

In addition, we also compared the DG performance of several supervised methods and the proposed one in Table 3. We trained all the compared networks by using the JTA\*

Method	CUHK		PRW	
	mAP	Top-1	mAP	Top-1
OIM (Xiao et al. 2017)	75.5	78.7	21.3	49.9
HOIM (Chen et al. 2020a)	89.7	90.8	39.8	80.4
NAE (Chen et al. 2020b)	92.1	92.9	44.0	81.1
OIMNet++ (Lee et al. 2022)	93.1	93.9	46.8	83.9
SeqNet (Li and Miao 2021)	94.8	95.7	47.6	87.6
PSTR (Cao et al. 2022)	93.5	95.0	49.5	87.8
DMRNet++ (Han et al. 2022)	94.4	95.5	51.0	86.8
COAT (Yu et al. 2022)	94.2	94.7	53.3	87.4
AGWF (Han, Ko, and Sim 2021b)	93.3	94.2	53.3	87.7
CGPS (Yan et al. 2022)	80.0	82.3	16.2	87.6
R-SiamNet (Han et al. 2021)	86.0	87.1	21.4	75.2
CUCPS (Han, Ko, and Sim 2021a)	81.1	83.2	41.7	86.0
Unsupervised DA (Li et al. 2022)	77.6	79.6	34.7	80.6
Proposed DG	76.1	78.4	25.5	79.4

Table 2: Comparison of the quantitative performance. The supervised and weakly-supervised methods are grouped in the first and second categories, respectively.

Method	CUHK-SYSU				PRW			
	mAP	Top-1	AP	Recall	mAP	Top-1	AP	Recall
HOIM	38.5	42.5	57.1	81.4	12.2	37.8	71.4	92.4
NAE	40.8	44.9	57.1	69.2	14.1	42.1	65.6	80.0
SeqNet	62.3	65.1	56.3	64.5	19.2	74.7	77.2	88.3
OIMNet++	66.3	69.0	60.4	69.7	19.8	74.0	74.8	84.7
COAT	61.4	64.7	57.0	60.2	22.6	76.9	81.2	87.9
Proposed	<b>76.1</b>	<b>78.4</b>	<b>72.3</b>	<b>87.3</b>	<b>25.5</b>	<b>79.4</b>	<b>84.8</b>	<b>96.0</b>

Table 3: Comparison of the DG performance. All the methods were trained by using the JTA\* dataset only.

dataset only. We see that the proposed method achieves a much higher performance of DG compared with the existing methods. Consequently, the experimental results demonstrate that the proposed method is a promising technique for person search which is completely free from the burden of time-consuming and labor-intensive labeling as well as the privacy issues.

### Ablation Study

We validated the effectiveness of the proposed fidelity adaptive training (FAT) and domain-invariant feature learning (DIL), respectively. Table 4 demonstrates that each of FAT and DIL improves not only the re-identification performance but also the detection performance.

**Effect of Resize Transformation.** There are huge differences in the size and aspect ratio of image between CUHK-SYSU and JTA\* datasets. While the image size and aspect ratio of JTA\* are fixed to  $1920 \times 1080$  and 1.78, respectively, CUHK-SYSU dataset has very diverse image sizes and aspect ratios. This discrepancy between the source and target datasets can be another source of domain gap. Therefore, when training the JTA\* dataset, we applied the resize transformation with 0.5 probability to prevent the model

Method	CUHK-SYSU				PRW			
	mAP	Top-1	AP	Recall	mAP	Top-1	AP	Recall
Baseline	66.7	71.0	64.6	78.5	20.9	76.0	77.8	92.3
w/ FAT	75.8	<b>78.5</b>	68.3	87.3	21.5	77.9	82.3	95.9
w/ DIL	69.3	72.7	66.8	80.2	24.8	<b>79.8</b>	79.8	92.7
Proposed	<b>76.1</b>	78.4	<b>72.3</b>	<b>87.3</b>	<b>25.5</b>	79.4	<b>84.8</b>	<b>96.0</b>

Table 4: Ablation study of the proposed method.

Method	Resize	CUHK-SYSU		PRW	
		mAP	Top-1	mAP	Top-1
Baseline	✓	62.6	66.7	<b>21.0</b>	<b>76.9</b>
		<b>66.7</b>	<b>71.0</b>	20.9	76.0
Proposed	✓	72.3	74.9	25.0	<b>80.8</b>
		<b>76.1</b>	<b>78.4</b>	<b>25.5</b>	79.4

Table 5: Effect of the resize transformation.

FWDL	FGCL	FWFU	CUHK-SYSU		PRW	
			mAP	Top-1	mAP	Top-1
			66.7	71.0	20.9	76.0
✓			75.5	78.0	21.4	77.9
	✓		71.2	74.4	21.1	77.5
		✓	71.7	75.1	21.3	77.4
✓	✓	✓	<b>75.8</b>	<b>78.5</b>	<b>21.5</b>	<b>77.9</b>
†✓	†✓	†✓	74.4	77.5	21.3	77.8

Table 6: Ablation study of fidelity-weighted detection loss (FWDL), fidelity-guided confidence loss (FGCL), and fidelity-weighted feature update (FWFU). †✓ indicates that we use the pre-defined ground-truth fidelity instead of the learned fidelity.

from overfitting to the source dataset with the fixed image size. Table 5 shows that the resize transformation keeps the performance from degradation caused by the size difference in CUHK-SYSU dataset. On the other hand, most of the images in PRW dataset have the same size to that of JTA\*, and thus the resize transformation does not yield a significant performance gain when applied on PRW dataset.

**Effect of Fidelity Adaptive Training.** Table 6 shows the detailed results of ablation study for the three schemes of FAT: fidelity-weighted detection loss (FWDL), fidelity-guided confidence loss (FGCL), and fidelity-weighted feature update (FWFU), where we see that each scheme contributes to the performance gain from the baseline. In addition, as shown in Table 6, we also evaluated the performance of using the three schemes together without fidelity estimation (FE) by replacing the learned fidelity with the pre-defined ground-truth fidelity. We see that using the proposed fidelity values further improves the performance compared with using the static pre-defined fidelity values.

Figure 6 also demonstrates the effectiveness of the proposed FE by showing the estimated fidelity values and the detection confidence scores. The three examples of person




Figure 6: The fidelity and detection confidence estimated by the proposed fidelity adaptive training. The three instances have similar ground-truth fidelity values around 0.8. However, they are assigned different fidelity values according to their actual appearance or visibility. The initial scores of the detection confidence are also changed accordingly.

$\mathcal{L}_{\text{dom}}$	DGN	$\mathcal{L}_{\text{sep}}$	CUHK-SYSU		PRW	
			mAP	Top-1	mAP	Top-1
			66.7	71.0	20.9	76.0
✓			66.9	70.4	21.1	77.5
✓	✓		69.0	72.1	23.2	77.0
✓	✓	✓	<b>69.3</b>	<b>72.7</b>	<b>24.8</b>	<b>79.8</b>

Table 7: Ablation study of  $\mathcal{L}_{\text{dom}}$ , domain-guided normalization (DGN), and  $\mathcal{L}_{\text{sep}}$ .

Method	CUHK-SYSU		PRW	
	mAP	Top-1	mAP	Top-1
Instance Norm.	68.2	70.6	22.0	75.4
Domain-guided Norm.	<b>69.0</b>	<b>72.1</b>	<b>23.2</b>	<b>77.0</b>

Table 8: The performance of the domain guided normalization compared to the instance normalization.

instances have similar BRISQUE scores and hence similar pre-defined fidelity values around 0.8. However, we observe their appearance and visibility are different from one another, which actually affect the performance of person search. For example, the first instance exhibits a relatively clear appearance of person, and yields a much higher learned fidelity value than the ground-truth one by using the proposed FAT. On the contrary, the third instance has relatively degraded visibility with blur and low contrast. In such a case, the network assigns a low fidelity value than the ground-truth according to its visibility. The detection confidence score is also forced to be a lower value compared to that of the baseline, by the confidence loss  $\mathcal{L}_{\text{con}}$  in (2) in FAT.

**Effect of Domain Invariant Feature Learning.** Table 7 compares the performance by incorporating the domain loss ( $\mathcal{L}_{\text{dom}}$ ), domain-guided normalization (DGN), and domain separation loss ( $\mathcal{L}_{\text{sep}}$ ), respectively. We see that every method improves the performance. Note that the instance normalization is widely employed for DG that serves to alleviate the style variations between different domains. We also conducted an experiment to see the effect of DGN compared to the instance normalization. Table 8 shows the results where we see that the proposed DGN outperforms the instance normalization on both target datasets. It means that



Figure 7: Comparison of the qualitative performance. Query person images (left) and the Top-5 matching results of the baseline (middle) and the proposed method (right). The true and false matching results are depicted in blue and red, respectively. The camera IDs are indicated in yellow.

the proposed DGN preserves more useful information for person re-identification while alleviating the information associated with domain variations more effectively.

Note that DIL achieves a relatively high performance gain on the PRW dataset compared to the CUHK-SYSU dataset. Whereas the images in CUHK-SYSU are captured by a single camera within relatively short time durations, the images in PRW are captured by 6 different cameras possibly comprising 6 different domains. Therefore, it becomes more challenging in the PRW dataset to find the person instances having the same ID across different domains. Accordingly, a relatively high impact of DIL is observed in PRW where the domain-related features are suppressed while the ID-specific features are exploited. Figure 7 verifies the cross-domain discriminative capability of the proposed method by showing the Top-5 matching results to the query images in PRW dataset. The true and false matching results are depicted in blue and red, respectively, and the camera IDs are indicated in yellow. The baseline method tends to match the person instances from the same camera to the query with high similarity values, and usually fails to find the correct persons across different domains. On the contrary, the proposed method effectively alleviates the camera-dependent information, and therefore, successfully finds the persons across different cameras.

## Conclusion

In this paper, we introduced a novel framework of domain generalizable person search that uses an automatically labeled unreal dataset only for training to avoid the time-consuming and labor-intensive data labeling and the privacy issues in real datasets. To alleviate the domain gaps between the unreal and real datasets, we trained an end-to-end network by estimating the fidelity of person instances simultaneously. We also devised the domain-invariant feature learning scheme to encourage the network to suppress the domain-specific information while learning the ID-related features more faithfully. Experimental results showed that the proposed method achieves the competitive performance compared to the existing person search methods, even though it is applicable to arbitrary unseen datasets without any prior knowledge of the target domain and additional re-training burdens.

## Acknowledgments

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-01336, Artificial Intelligence Graduate School Program(UNIST)) and (No.2021-0-02068, Artificial Intelligence Innovation Hub).

## References

- Balaji, Y.; Sankaranarayanan, S.; and Chellappa, R. 2018. Metareg: Towards domain generalization using meta-regularization. In *NIPS*, 1006–1016.
- Cao, J.; Pang, Y.; Anwer, R. M.; Cholakkal, H.; Xie, J.; Shah, M.; and Khan, F. S. 2022. Pstr: End-to-end one-step person search with transformers. In *CVPR*, 9458–9467.
- Chen, D.; Zhang, S.; Ouyang, W.; Yang, J.; and Schiele, B. 2020a. Hierarchical online instance matching for person search. In *AAAI*, 10518–10525.
- Chen, D.; Zhang, S.; Yang, J.; and Schiele, B. 2020b. Norm-aware embedding for efficient person search. In *CVPR*, 12615–12624.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020c. A simple framework for contrastive learning of visual representations. In *PMLR*, 1597–1607.
- Choi, S.; Kim, T.; Jeong, M.; Park, H.; and Kim, C. 2021. Meta batch-instance normalization for generalizable person re-identification. In *CVPR*, 3425–3435.
- Fabbi, M.; Lanzi, F.; Calderara, S.; Palazzi, A.; Vezzani, R.; and Cucchiara, R. 2018. Learning to detect and track visible and occluded body joints in a virtual world. In *ECCV*, 430–446.
- Fan, X.; Wang, Q.; Ke, J.; Yang, F.; Gong, B.; and Zhou, M. 2021. Adversarially adaptive normalization for single domain generalization. In *CVPR*, 8208–8217.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A Kernel Two-Sample Test. *JMLR*, 13(25): 723–773.
- Han, B.-J.; Ko, K.; and Sim, J.-Y. 2021a. Context-aware unsupervised clustering for person search. In *BMVC*.
- Han, B.-J.; Ko, K.; and Sim, J.-Y. 2021b. End-to-end trainable trident person search network using adaptive gradient propagation. In *ICCV*, 925–933.
- Han, C.; Su, K.; Yu, D.; Yuan, Z.; Gao, C.; Sang, N.; Yang, Y.; and Wang, C. 2021. Weakly supervised person search with region siamese networks. In *CVPR*, 12006–12015.
- Han, C.; Zheng, Z.; Su, K.; Yu, D.; Yuan, Z.; Gao, C.; Sang, N.; and Yang, Y. 2022. Dmrnet++: Learning discriminative features with decoupled networks and enriched pairs for one-step person search. *TPAMI*, 45(6): 7319–7337.
- Jin, X.; Lan, C.; Zeng, W.; Chen, Z.; and Zhang, L. 2020. Style normalization and restitution for generalizable person re-identification. In *CVPR*, 3143–3152.
- Lee, S.; Oh, Y.; Baek, D.; Lee, J.; and Ham, B. 2022. Oimnet++: Prototypical normalization and localization-aware learning for person search. In *ECCV*, 621–637.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 3490–3497.
- Li, J.; Yan, Y.; Wang, G.; Yu, F.; Jia, Q.; and Ding, S. 2022. Domain adaptive person search. In *ECCV*, 302–318.
- Li, Z.; and Miao, D. 2021. Sequential end-to-end network for efficient person search. In *AAAI*, 2011–2019.
- Liu, J.; Huang, Z.; Li, L.; Zheng, K.; and Zha, Z.-J. 2022. Debiased batch normalization via gaussian process for generalizable person re-identification. In *AAAI*, 1729–1737.
- Mittal, A.; Moorthy, A. K.; and Bovik, A. C. 2012. No-reference image quality assessment in the spatial domain. *TIP*, 21(12): 4695–4708.
- Motiiian, S.; Piccirilli, M.; Adjero, D. A.; and Doretto, G. 2017. Unified deep supervised domain adaptation and generalization. In *ICCV*, 5715–5725.
- Qiao, F.; and Peng, X. 2021. Uncertainty-guided model generalization to unseen domains. In *CVPR*, 6790–6800.
- Qiao, F.; Zhao, L.; and Peng, X. 2020. Learning to learn single domain generalization. In *CVPR*, 12556–12565.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 91–99.
- Segu, M.; Tonioni, A.; and Tombari, F. 2023. Batch normalization embeddings for deep domain generalization. *PR*, 135: 109115.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*.
- Volpi, R.; Namkoong, H.; Sener, O.; Duchi, J. C.; Murino, V.; and Savarese, S. 2018. Generalizing to unseen domains via adversarial data augmentation. In *NIPS*, 5339–5349.
- Wang, Z.; Luo, Y.; Qiu, R.; Huang, Z.; and Baktashmotlagh, M. 2021. Learning to diversify for single domain generalization. In *ICCV*, 834–843.
- Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2017. Joint detection and identification feature learning for person search. In *CVPR*, 3415–3424.
- Yan, Y.; Li, J.; Liao, S.; Qin, J.; Ni, B.; Lu, K.; and Yang, X. 2022. Exploring visual context for weakly supervised person search. In *AAAI*, 3027–3035.
- Yu, R.; Du, D.; LaLonde, R.; Davila, D.; Funk, C.; Hoogs, A.; and Clipp, B. 2022. Cascade transformers for end-to-end person search. In *CVPR*, 7267–7276.
- Zheng, L.; Zhang, H.; Sun, S.; Chandraker, M.; Yang, Y.; and Tian, Q. 2017. Person re-identification in the wild. In *CVPR*, 1367–1376.