

Debiased Batch Normalization via Gaussian Process for Generalizable Person Re-identification

Jiawei Liu^{1*}, Zhipeng Huang^{1*}, Liang Li², Kecheng Zheng¹, Zheng-Jun Zha^{1†}

¹ University of Science and Technology of China

² Institute of Computing Technology, Chinese Academy of Sciences

{hzp1104,zkcys001}@mail.ustc.edu.cn, {jwliu6,zhazj}@ustc.edu.cn, liang.li@ict.ac.cn

Abstract

Generalizable person re-identification aims to learn a model with only several labeled source domains that can perform well on unseen domains. Without access to the unseen domain, the feature statistics of the batch normalization (BN) layer learned from a limited number of source domains is doubtlessly biased for unseen domain. This would mislead the feature representation learning for unseen domain and deteriorate the generalization ability of the model. In this paper, we propose a novel Debiased Batch Normalization via Gaussian Process approach (GDNorm) for generalizable person re-identification, which models the feature statistic estimation from BN layers as a dynamically self-refining Gaussian process to alleviate the bias to unseen domain for improving the generalization. Specifically, we establish a lightweight model with multiple set of domain-specific BN layers to capture the discriminability of individual source domain, and learn the corresponding parameters of the domain-specific BN layers. These parameters of different source domains are employed to deduce a Gaussian process. We randomly sample several paths from this Gaussian process served as the BN estimations of potential new domains outside of existing source domains, which can further optimize these learned parameters from source domains, and estimate more accurate Gaussian process by them in return, tending to real data distribution. Even without a large number of source domains, GDNorm can still provide debiased BN estimation by using the mean path of the Gaussian process, while maintaining low computational cost during testing. Extensive experiments demonstrate that our GDNorm effectively improves the generalization ability of the model on unseen domain.

Introduction

Person re-identification (Re-ID) (Ye et al. 2021; Luo et al. 2019; Liu et al. 2016, 2019) aims to identify a person-of-interest across non-overlapping camera networks under different times and locations. It has drawn extensive research attention in recent years, due to its significant role in various practice scenarios, such as object tracking, activity analysis and smart retail *etc* (Wang 2013; Caba Heilbron et al. 2015; Martini, Paolanti, and Frontoni 2020). Along with the success of deep learning, large amount of sophisticated person

Re-ID methods have been proposed and achieved promising performances in a fully supervised manner.

However, when applying these trained models to previously unseen domains, supervised approaches suffer from significant performance degradation due to style discrepancies across domains. To this end, unsupervised domain adaptation (UDA) has been introduced in person Re-ID. UDA approaches employ unlabeled samples from target domain to adapt the pre-trained model on the labeled source domain to target domain (Bak, Carr, and Lalonde 2018). Although UDA approaches are more practical than the full supervised approaches, they still require data collection and model update for the target domain, resulting in additional computational cost. Moreover, samples from a target domain are usually not unavailable in advance.

Compared to UDA, domain generalization (DG) is a promising solution for real-world applications. Domain generalization aims at learning models that are generalizable to unseen target domain, without requiring access to the data and annotations from target domain as well as model updating, *i.e.*, fine-tune or adaptation. Most existing DG methods (Chattopadhyay, Balaji, and Hoffman 2020; Huang et al. 2020) assume that the source and target domains share the same label space with a fixed of classes, which are designed for a classification task. In contrast, domain generalization for person Re-ID is an open-set retrieval task, having different and variable number of identities between source and target domains. Therefore, it is difficult to achieve satisfying generalization capability when the existing DG approaches are directly applied to person Re-ID.

Existing generalizable person Re-ID approaches can be divided into two categories: single model methods (Zhao et al. 2021; Choi et al. 2021; Song et al. 2019) and ensemble learning based methods (Dai et al. 2021; Seo et al. 2020; Zhuang et al. 2020; Mancini et al. 2018). The former collects all source domain data and trains a single model on them to extract the shared domain-invariant representations. They mainly design customized loss functions (*e.g.*, Maximum Mean Discrepancy regularization to align the distributions among different domain), specific architectures (*e.g.*, integrating Batch Normalization and Instance Normalization in the models to alleviate the domain discrepancy due to appearance style variations), or ad-hoc training policies (*e.g.*, model-agnostic meta learning to mimic real train-test

*Equal Contribution

†Corresponding Author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

domain shift). However, these methods discard the specific characteristics of individual source domains and neglect the inherent relevance of unseen target domain with respect to source domains, result in unsatisfying generalization capability. Such domain-specific characteristics can provide adequately discriminative and meaningful information and infer the distribution of unseen target domain, thus greatly improving the model’s generalizability.

Ensemble learning based methods attempt to train domain-specific models for each source domain by exploiting domain specific characteristics, and combine multiple domain-specific models (*e.g.*, branches, classifiers or experts) to enhance the generalization ability. These methods assume that any sample can be considered as a comprehensive sample from multi-source domains, such that the overall prediction result can be seen as the superposition of multi-domain networks. Nevertheless, these works have two main drawbacks: (1) The number of source domains in DG person Re-ID task is extremely limited, directly ensembling them results in a high bias to the underlying hyper-distribution (real data distribution) on domains. The biased issue injury the generalization ability of the model; (2) Ensemble learning significantly increases the computational cost due to training and testing multiple neural networks. When there exists a large amount of source domains, the computational cost of the ensemble learning based method is unacceptable. Thus, how to efficiently utilize the limited number of source domains to improve the generalization is the major issue for generalizable person re-identification.

In this work, we propose a novel Debiased Batch Normalization via Gaussian Process approach (GDNorm) for DG person Re-ID, which models the feature statistic estimation from BN layers as a dynamically self-refining Gaussian process to alleviate the bias issue on unseen domain for improving the generalization. GDNorm establishes a lightweight model with multiple set of domain-specific BN layers for source domains to learn the corresponding population statistics and rescaling parameters of the domain-specific BN layers. The parameters of the domain-specific BN layers from different source domains are assumed to fit a Gaussian process, and act as an approximate distribution of the domain-specific model for each source domain. We randomly sample paths (the parameters of the BN layers) from the deduced Gaussian process as the BN estimations of potential new domains, which can further optimize these parameters, and estimate more accurate Gaussian process by them in return, tending to the underlying hyper-distribution (real data distribution) on domains. Such self-refining mechanism guarantees that even with a limited number of source domains, GDNorm can still provide debiased BN estimation on unseen domain. In the testing phase, GDNorm directly exploits the mean path of the Gaussian process as the optimal BN estimating for target domain, greatly reducing the computational cost as compared to ensemble learning. Extensive experiments have demonstrated the effectiveness of the proposed GDNorm when testing on the unseen domain.

The main contributions of this paper are three-fold: (1) We propose a novel Debiased Batch Normalization via Gaussian Process approach (GDNorm) for generalizable person

Re-ID. (2) We design a novel debiased Batch Normalization, which models the feature statistic estimation of BN as a dynamically self-refining Gaussian process, alleviating the bias issue of the BN estimation on unseen domain. (3) We propose a lightweight model for generalizable person Re-ID, which exploits the mean path of the Gaussian process as the optimal debiased BN parameters, without requiring extra computational cost as compared to ensemble learning.

Related Work

DG Person Re-identification

Existing DG person Re-ID methods (Zhao et al. 2021; Choi et al. 2021; Song et al. 2019; Lin, Li, and Kot 2020; Chen et al. 2021; Jin et al. 2020; Tamura and Murakami 2019) can be divided into two categories: single model methods and ensemble learning based methods. For the single model methods, Chen *et al.* (Chen et al. 2021) proposed a Dual Distribution Alignment Network with dual-level constraints, *i.e.*, a domainwise adversarial feature learning and an identity-wise similarity enhancement, which maps pedestrian images into a domain-invariant feature space. Choi *et al.* (Choi et al. 2021) designed learnable batch-instance normalization layers, which prevents the model from overfitting to the source domains by the simulation of unsuccessful generalization scenarios in meta-learning pipeline. Lin *et al.* (Lin, Cheng, and Wang 2020) present an episodic learning scheme, which advances meta-learning algorithm to exploit the labeled data from source domain and learns domain-invariant features without observing target domain data.

Ensemble Learning Based DG

Ensemble learning based methods ensemble multiple domain specific models like experts or classifiers to enhance the generalization ability of the overall network, which are applied to general DG and GD person Re-ID tasks. Mancini *et al.* (Mancini et al. 2018) proposed to fuse the predictions from different domain-specific classifiers with learnable weight and a domain predictor. D’Innocente *et al.* (D’Innocente and Caputo 2018) proposed domain-specific layers of different source domains and learned the linear aggregation of these layers to represent a test sample. Zhou *et al.* (Zhou et al. 2020) proposed a domain adaptive ensemble learning method, which improves the generalization of a multi-expert network by explicitly training the ensemble to solve the target task. Dai *et al.* (Dai et al. 2021) proposed a method called the relevance-aware mixture of experts (RaMoE), by using an effective voting-based mixture mechanism to dynamically leverage source domains’ diverse characteristics to improve the generalization.

Normalization in DA and DG

Normalization techniques in deep neural networks are designed for regularizing trained models and improving their generalization performance. Recently, several methods on domain adaptation (DA) and domain generalization (DG) discovered the relationship between domain gap and normalization operation. For example, Jin *et al.* (Jin et al. 2021) proposed a Style Normalization and Restitution module,

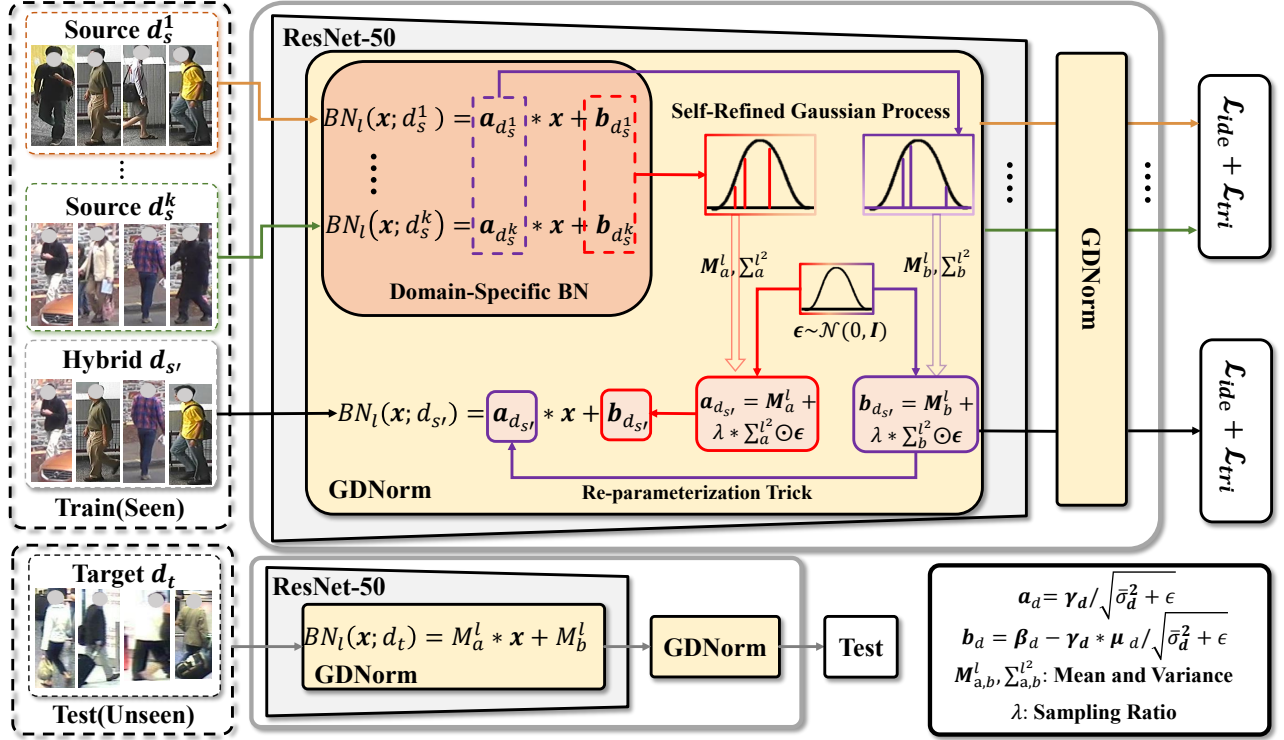


Figure 1: The overall architecture of the proposed GDNorm. It contains a lightweight model with multiple set of domain-specific BN layers for capturing the specific feature statistics information, and a dynamically self-refining Gaussian process estimation for learning debiased BN parameters on unseen domain.

which utilizes the Instance Normalization (IN) layers to filter out style variations and compensates for the identity-relevant features discarded by IN layers. Seo *et al.* (Seo et al. 2020) proposed to leverage instance normalization to optimize the trade-off between cross-category variance and domain invariance, which is desirable for domain generalization in unseen domains. Zhuang *et al.* (Zhuang et al. 2020) proposed camera-based batch normalization (CBN) to force the images of all cameras to fall onto the same subspace and to shrink the distribution gap between any camera pair.

Method

Bias in Generalizable Person Re-ID

For DG person Re-ID, we assume that we have K source domains (datasets) $\mathcal{D}_S = \{d_s^1, d_s^2, \dots, d_s^K\}$ in the training stage. Each domain $d_s^k = \{(\mathbf{x}_i^k, y_i^k)\}_{i=1}^{N_k}$ has its own label space (person IDs among different datasets are non-overlapping), where N_k is the total number of pedestrian images in the source domain d_s^k , each image \mathbf{x}_i^k is associated with an identity label y_i^k . During the testing stage, the trained model is frozen and directly deployed to a previously unseen dataset (target domain d_t) without additional model updating. More formally, we denote \mathcal{X} as the input pedestrian images space, and denote \mathcal{Y} as the label space. A domain (dataset) in DG person Re-ID can be denoted as the aggregation of data sampled from a specific joint distribution $d^k = (\mathbf{x}_i, y_i)_{i=1}^{N_k} \sim P_{xy}^{d^k}$. Their joint distributions of source and target domains are completely different. DG person Re-ID aims to learn a generalizable model $\bar{P}(Y|X)$

which achieves the smallest prediction error on unknown target domain without fine-tuning:

$$\min_{\bar{P}} \mathbb{E}_{(\mathbf{x}, y) \sim P_{xy}^{d_t}} [\ell(\bar{P}(Y|X = \mathbf{x}), y)] \quad (1)$$

where ℓ denotes a loss function on \mathcal{Y} . Due to the fact that the target domain is totally unknown (not even unsupervised data), it is required to measure the average risk over all possible target domains. Assuming both source and target domain distributions follow an underlying hyper-distribution $\mathcal{H}: P_{xy} \sim \mathcal{H}$. Then the target prediction error (risk) is reformulated as the following:

$$\mathcal{E} := \mathbb{E}_{P_{xy} \sim \mathcal{H}} \mathbb{E}_{(\mathbf{x}, y) \sim P_{xy}} [\ell(\bar{P}(Y|X = \mathbf{x}), y)] \quad (2)$$

Existing ensemble learning based DG methods (Mancini et al. 2018; Segu, Tonioni, and Tombari 2020; D’Innocente and Caputo 2018; Zhou et al. 2020) simply regard the hyper-distribution \mathcal{H} as a static and discrete distribution with respect to the limited source domains: $\{P_{\mathcal{H}}(P_{xy} = P_{xy}^{d_s^k}) = \pi^{d_s^k}\}_{k=1}^K$. Under this approximate hypothesis, they exploit the relationship between multiple source domains by designing specific network architecture and training strategies to improve generalization. They assume that any sample can be regarded as an integrated sample of the multiple source domains, such that the overall prediction result can be seen as the superposition of the multiple domain networks. Thus, the approximate optimal solution for Eq. 2 is then computed as the following:

$$\bar{P}(Y|X) = \sum_{k=1}^K \pi^{d_s^k} \bar{P}(Y|X, P_{xy}^{d_s^k}) \quad (3)$$

where $\overline{P}(Y|X, P_{xy}^{d_s^k})$ represents the k -th source domain-specific model, and π^k denotes the mixture ratio. Nevertheless, these methods employ such simple discrete, static distribution to represent the complex hyper-distribution \mathcal{H} , leading to high bias. First, the discrete distribution considers all possible domains under \mathcal{H} as a mixture of a limited number of source domains. Only using a small number of domains could definitely not express the complex unknown \mathcal{H} . Second, the static distribution is not able to reduce and adjust the deviation to the real hyper-distribution \mathcal{H} . In addition, ensemble model leads to a heavy increasing of computational cost, due to maintaining multiple domain specific networks in the testing stage.

Unlike these methods, we rethink the hyper-distribution \mathcal{H} as a dynamically self-refining Gaussian process. In theory, a linear combination of multiple Gaussian distributions or Gaussian process can fit any type of distribution (Bond, Hoeffler, and Temple 2001; Yu, Yang, and Li 2019), which is a better choice to serve as the hyper-distribution \mathcal{H} . Based on this, we design a self-refining mechanism to refine and debias the estimated gaussian process, towards obtaining the ideal unbiased optimal solution of Eq. 2:

$$\begin{aligned} \overline{P}_{unbiased}(Y|X) &= \int_{P_{xy} \sim \mathcal{H}} P(P_{xy}|X) \overline{P}(Y|X, P_{xy}) \\ &= \mathbb{E}_{\mathcal{H}|x} \overline{P}(Y|X, P_{xy}) \end{aligned} \quad (4)$$

Specifically, we propose a novel Debaised Batch Normalization via Gaussian Process approach (GDNorm) for generalizable person Re-ID, as illustrated in Fig. 1. GDNorm establishes a lightweight model $\overline{P}(Y|X, P_{xy})$ with different set of domain-specific BN layers, which is based on the popular ResNet-50 model (He et al. 2016). The expected domain specific models for each source domain share all the training parameters of the lightweight model apart from the learned parameters θ_{d_s} from the domain-specific BN layers. GDNorm utilizes these parameters to serve as the distribution of the domain-specific model, $\overline{P}(Y|X = x, P_{xy}^{d_s})$ s.t. $P_{xy}^{d_s} \sim \mathcal{H}$, which follows a Gaussian process. It then randomly samples paths (the parameters of domains-specific BN layers) from this Gaussian process to obtain the BN estimations of potential new domains outside of the training source domains. The sampled paths can in return, facilitate the Gaussian process close to (debias to) the distribution of the hyper-distribution by end-to-end optimizing. In the testing phase, unlike the previous ensemble learning based methods, GDNorm directly employs the mean path of the learned debaised Gaussian process as optimal batch normalization layers' parameters on unseen target domain for model inference, without requiring extra computational cost.

Domain-Specific Batch Normalization

BN is a widely-used training technique in deep networks, which normalizes the activations from each domain to the same reference distribution by domain-specific normalization statistics (Chang et al. 2019). To capture the specific characteristics of individual domain, a lightweight model

with different set of domain-specific BN layers is designed, which can capture the source domain distribution $P_{xy}^{d_s}$.

Concretely, given a batch size of pedestrian images from a certain source domain d_s at the training stage, a domain-specific BN layer normalizes the activations by matching their first and second order moments ($\mu_{d_s}, \sigma_{d_s}^2$) to a reference Gaussian distribution with zero mean and unitary variance $\mathcal{N}(0, 1)$. And then transforms channel-wise whitened activations by using rescaling parameters. The formulation of a BN layer is defined as the following:

$$BN(\mathbf{x}; d_s) = \gamma_{d_s} \frac{\mathbf{x} - \mu_{d_s}}{\sqrt{\sigma_{d_s}^2 + \epsilon}} + \beta_{d_s} \quad (5)$$

where \mathbf{x} refers to an input activation from the domain d_s , $\epsilon > 0$ is a small constant to avoid numerical problem. γ_{d_s} and β_{d_s} are the learnable rescaling parameters, which are optimized by samples from the specific source domain. μ_{d_s} and $\sigma_{d_s}^2$ are the domain-specific batch statistics for the domain d_s , which are calculated as the following:

$$\begin{aligned} \mu_{d_s} &= \frac{\sum_{b,i,j} \mathbf{x}[b, i, j]}{B \cdot H \cdot W} \\ \sigma_{d_s}^2 &= \frac{\sum_{b,i,j} (\mathbf{x}[b, i, j] - \mu_{d_s})^2}{B \cdot H \cdot W} \end{aligned} \quad (6)$$

where B denotes the batch size, H and W is the height and width size of the input activations. The domain-specific batch statistics are updated by using moving average strategy as the following:

$$\begin{aligned} \overline{\mu}_{d_s} &= m \overline{\mu}_{d_s} + (1 - m) \mu_{d_s} \\ \overline{\sigma}_{d_s}^2 &= m \overline{\sigma}_{d_s}^2 + (1 - m) \sigma_{d_s}^2 \end{aligned} \quad (7)$$

where $\overline{\mu}_{d_s}, \overline{\sigma}_{d_s}^2$ reflect domain-specific population statistics of this source domain across all batch data. The expected domain-specific model with the parameters θ_{d_s} of the domain-specific BN layers for each source domain is then learned as the following:

$$\begin{aligned} \theta_{d_s} &= \{(\gamma_{d_s}^l, \beta_{d_s}^l, \overline{\mu}_{d_s}^l, \overline{\sigma}_{d_s}^l)\}_{l=1}^L \\ \overline{P}(Y|X, P_{xy}^{d_s}) &= \overline{P}(Y|X, \theta_{d_s}) \end{aligned} \quad (8)$$

where L is the number of batch normalization layers in our lightweight model for each source domain.

Debias Gaussian Process

The learned parameters θ_{d_s} of the domain-specific BN layers contain specific feature statistics information of the corresponding domain d_s . Therefore, the distribution of the parameters θ_{d_s} can be represented as the hyper-distribution \mathcal{H} on domains. We assume that the distributions of the parameters θ_{d_s} in the model follow a Gaussian process (A collection of Gaussian distributions along different depth position BN layers in the model). GDNorm deduces a dynamical Gaussian process from the learned parameters θ_{d_s} of the domain-specific BN layers to fit the hyper-distribution \mathcal{H} . To reduce the bias between the Gaussian process and the hyper-distribution, it designs a self-refining mechanism to

randomly sample several paths from the estimated Gaussian process to serve as the BN estimations of potential new domains outside of source domains, which can further optimize these parameters, and infer more accurate Gaussian process by them in return, tending to real hyper-distribution.

For simplicity, we reformulate the BN operation into a linear transformation as the following:

$$\begin{aligned} BN(\mathbf{x}; d_s) &= \mathbf{a}_{d_s} \mathbf{x} + \mathbf{b}_{d_s} \\ \mathbf{a}_{d_s} &= \frac{\gamma_{d_s}}{\sqrt{\sigma_{d_s}^2 + \epsilon}} \\ \mathbf{b}_{d_s} &= \beta_{d_s} - \gamma_{d_s} \frac{\bar{\boldsymbol{\mu}}_{d_s}}{\sqrt{\sigma_{d_s}^2 + \epsilon}} \end{aligned} \quad (9)$$

GDNorm calculates the mean \mathbf{M} and variance Σ^2 of the learned parameters in each domain-specific BN layer for all source domains to obtain a collection of Gaussian distributions along different depth position BN layers in the model.

$$\begin{aligned} \mathbf{M}(\mathbf{a}^l) &= \frac{\sum_{k=1}^K \mathbf{a}_{d_s^k}^l}{K} & \mathbf{M}(\mathbf{b}^l) &= \frac{\sum_{k=1}^K \mathbf{b}_{d_s^k}^l}{K} \\ \Sigma^2(\mathbf{a}^l) &= \frac{\sum_{k=1}^K (\mathbf{a}_{d_s^k}^l - \mathbf{M}(\mathbf{a}^l))^2}{K} \\ \Sigma^2(\mathbf{b}^l) &= \frac{\sum_{k=1}^K (\mathbf{b}_{d_s^k}^l - \mathbf{M}(\mathbf{b}^l))^2}{K} \end{aligned} \quad (10)$$

These Gaussian distributions constitute a Gaussian process whose random variables are indexed by the different depth position: $GP = \{(\mathbf{M}^l, \Sigma^{l^2})\}_{l=1}^L$.

To refine the estimated Gaussian process unbiased to the real hyper-distribution, GDNorm randomly samples several paths from the Gaussian process as the BN estimations of potential domains for further training and updating these learned parameters. It then employs a reparameterization trick (Blundell et al. 2015) to bypass the discrete sampling process, which allows end-to-end optimizing:

$$\mathbf{a}_{s'}^l, \mathbf{b}_{s'}^l = \mathbf{M}^l + \lambda \Sigma^{l^2} \odot \boldsymbol{\epsilon} \quad (11)$$

where $\{\mathbf{a}_{s'}^l, \mathbf{b}_{s'}^l\}_{l=1}^L$ is a randomly sample path from the Gaussian process, corresponding to a potential new domain $d_{s'}$. $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and λ is the sampling ratio. The randomly sampled path is used as the BN parameters for the potential new domain. A batch of images sampled equally from all existing source domains as the samples of the potential new domain to train the model, which is supervised by a identity loss and a triplet loss (Schroff, Kalenichenko, and Philbin 2015). The parameters of the domain-specific BN layer for the existing source domains will be updated by this optimizing process, which can be used to deduce more accurate Gaussian process by them in return, tending to real hyper-distribution. In the testing phase, GDNorm directly utilizes the mean path of the final debiased Gaussian process $\{(\mathcal{M}(\mathbf{a}^l), \mathcal{M}(\mathbf{b}^l))\}_{l=1}^L$ to serve as every BN layers in the model, runs the trained model only once for inference, without requiring extra computational cost as compared to ensemble learning.

Experiments

Datasets and Evaluation Settings

Protocol-1: Following the previous methods (Song et al. 2019; Jia, Ruan, and Hospedales 2019; Tamura and Murakami 2019), we employ the existing Re-ID benchmarks to evaluate the Re-ID model’s generalization ability, where the existing large-scale Re-ID datasets are viewed as multiple source domains, and the small-scale Re-ID datasets are used as unseen target domains. As shown in Tab. 2, source domains include CUHK02 (Li and Wang 2013), CUHK03 (Li et al. 2014), Market-1501 (Zheng et al. 2015), DukeMTMC (Zheng, Zheng, and Yang 2017) and CUHK-SYSU (Xiao et al. 2017). Target domains contain VIPeR (Gray and Tao 2008), PRID (Hirzer et al. 2011), GRID (Loy, Xiang, and Gong 2009) and iLIDS (Zheng, Gong, and Xiang 2009). All training sets and testing sets in the source domains are used for model training. The four small-scale Re-ID datasets are tested respectively, where the final performances are obtained by the average of 10 repeated random splits of testing sets.

Protocol-2: Considering that the image quality of the small-scale Re-ID datasets is quite poor, the performances on these datasets can not precisely reflect the generalization ability of a model in real scenarios. The previous methods (Zhao et al. 2021; Dai et al. 2021) thus set a new protocol (*i.e.*, leave-one-out setting) for four large-scale Re-ID datasets. Specifically, four large-scale Re-ID datasets (Market-1501 (Zheng et al. 2015), DukeMTMC (Zheng, Zheng, and Yang 2017), CUHK03 (Li et al. 2014) and MSMT17 (Wei et al. 2018)) are divided into two parts: three datasets as the source domains for training and the remaining one as the target domain for testing. For simplicity, in the following sections, we denote Market1501 as M, DukeMTMC as D, CUHK02 as C2, CUHK03 as C3, MSMT17 as MT and CUHK-SYSU as CS. The two different evaluation settings are shown in Tab. 2.

Evaluation metrics: We follow the common evaluation metrics person Re-ID, *i.e.*, mean Average Precision (mAP) and Cumulative Matching Characteristic (CMC) at Rank- k .

Implementation Details

We adopt ResNet50 (He et al. 2016) pretrained on Imagenet as our backbone. Following previous method (Luo et al. 2019), the last residual layer’s stride size is set to 1. A generalized mean Pooling (GeM) (Radenović, Tolia, and Chum 2018) with a batch normalization layer is used after the backbone to obtain the Re-ID features. Images are resized to 384×128 , and the training batch size is set to 128, including 8 identities and 16 images per identity. For data augmentation, we use random flipping, random cropping and color jittering. We train the model for 60 epochs. The learning rate is initialized as 3.5×10^{-4} and divided by 10 at 40th epochs, weight decay is 5×10^{-4} . λ in Eq. 11 is set to 0.6. For our Baseline, we combine all the source domains into a hybrid dataset, and employ the same settings for the loss function, data augmentation strategy and the backbone, and then train the model without GDNorm. To speed up the training process and increase memory efficiency, we

Method	Type/Source	VIPeR (V)		PRID (P)		GRID (G)		iLIDS (I)		Mean	
		R1	mAP	R1	mAP	R1	mAP	R1	mAP	R1	mAP
ImpTrpLoss (Cheng et al. 2016)	S/Target	47.8	-	22.0	-	-	-	60.4	-	-	-
JLML (Li, Zhu, and Gong 2017)	S/Target	50.2	-	-	-	37.5	-	-	-	-	-
SSM (Bai, Bai, and Tian 2017)	S/Target	53.7	-	-	-	27.2	-	-	-	-	-
TJ-AIDL (Wang et al. 2018)	UDA/M	38.5	-	26.8	-	-	-	-	-	-	-
MMFA (Lin et al. 2018)	UDA/M	39.1	-	35.1	-	-	-	-	-	-	-
UDML (Peng et al. 2016)	UDA/Comb1	31.5	-	24.2	-	-	-	49.3	-	-	-
SyRI (Bak, Carr, and Lalonde 2018)	UDA/Comb2	43.0	-	43.0	-	-	-	56.5	-	-	-
CDEL (Lin, Cheng, and Wang 2020)	DG/MS	38.5	-	57.6	-	33.0	-	62.3	-	47.9	-
DIMN (Song et al. 2019)	DG/MS	51.2	60.1	39.2	52.0	29.3	41.1	70.2	78.4	47.5	57.9
AugMining (Tamura et al. 2019)	DG/MS	49.8	-	34.3	-	46.6	-	76.3	-	51.8	-
DualNorm (Jia et al. 2019)	DG/MS	53.9	58.0	60.4	64.9	41.4	45.7	74.8	78.5	57.6	61.8
DDAN (Chen et al. 2021)	DG/MS	56.5	60.8	62.9	67.5	46.2	50.9	78.0	81.2	60.9	65.1
RaMoE (Dai et al. 2021)	DG/MS	56.6	64.6	57.7	67.3	46.8	54.2	85.0	90.2	61.5	69.1
DIR-ReID (Zhang et al. 2021)	DG/MS	58.3	62.9	71.1	75.6	47.8	52.1	74.4	78.6	62.9	67.3
GDNorm (Ours)	DG/MS	66.1	74.1	72.6	79.9	55.4	63.8	81.3	87.2	68.9	76.3

Table 1: Performance (%) comparison with the state-of-the-art methods under Protocol-1.

Setting	Training Data	Testing Data
Protocol-1	M+D+C2+C3+CS	PRID, GRID, VIPeR, iLIDs
Protocol-2	D+C3+MT	M
	M++C3+MT	D
	M+D+MT	C3
	M+D+C3	MT

Table 2: Different evaluation settings of Protocol-1 and Protocol-2.

E	S	D+MT+C3→M			
		R1	R5	R10	mAP
×	×	78.5	91.4	94.5	55.2
✓	×	83.2	92.9	95.3	62.5
✓	✓	86.5	95.2	96.9	68.2

Table 4: Ablation studies on the effectiveness of each component of GDNorm under Protocol-2.

D+MT+C3→M	R1	R5	R10	mAP
<i>QACConv</i> ₅₀ * (Liao et al. 2020)	65.7	-	-	35.6
<i>M</i> ³ <i>L</i> (Zhao et al. 2021)	75.9	-	-	50.2
RaMoE (Dai et al. 2021)	82.0	91.4	94.4	56.5
GDNorm (Ours)	86.5	95.2	96.9	68.2
M+MT+C3→D	R1	R5	R10	mAP
<i>QACConv</i> ₅₀ * (Liao et al. 2020)	66.1	-	-	47.1
<i>M</i> ³ <i>L</i> (Zhao et al. 2021)	69.2	-	-	51.1
RaMoE (Dai et al. 2021)	73.6	85.3	88.4	56.9
GDNorm (Ours)	78.2	87.8	90.4	63.8
D+MT+M→C3	R1	R5	R10	mAP
<i>QACConv</i> ₅₀ * (Liao et al. 2020)	23.5	-	-	21.0
<i>M</i> ³ <i>L</i> (Zhao et al. 2021)	33.1	-	-	32.1
RaMoE (Dai et al. 2021)	34.6	54.3	64.6	33.5
GDNorm (Ours)	48.6	70.3	79.3	47.9
D+M+C3→MT	R1	R5	R10	mAP
<i>QACConv</i> ₅₀ * (Liao et al. 2020)	24.3	-	-	7.5
<i>M</i> ³ <i>L</i> (Zhao et al. 2021)	33.0	-	-	12.9
RaMoE (Dai et al. 2021)	34.1	46.0	51.8	13.5
GDNorm (Ours)	48.1	63.1	68.5	20.4

Table 3: Performance (%) comparison with the state-of-the-art methods under Protocol-2.

use automatic mixed-precision training (Micikevicius et al. 2017) in the entire process. All experiments are conducted on a single NVIDIA Titan XP GPU.

Comparison to state-of-the-art Methods

Comparison under the Protocol-1. As shown in Tab. 1, ‘MS’ is the multiple source datasets under Protocol-1 (M, D, C2, C3 and CS). ‘Comb1’ is the leave-one-out setting for VIPeR, PRID, CUHK01, iLIDS and CAVIAR datasets. ‘Comb2’ is the combination of C3, D and synthetic datasets.

The comparison methods are mainly divided into three groups: supervised learning (S), unsupervised domain adaptation (UDA), and domain generalization (DG). Supervised methods suffer from the over-fitting issue due to the extremely small scale of training set. In contrast, our GDNorm method achieves better performance than all the supervised methods without access to labeled target data. Furthermore, we can observe that even if unlabeled target samples can be utilized for UDA methods, our GDNorm is still superior to these UDA methods without using any target data. Moreover, the mean performance on four target domains of our GDNorm method improves RaMoE (Dai et al. 2021) by 7.4% R-1 accuracy and 7.2% mAP, and DIR-ReID (Zhang et al. 2021) 6.0% R-1 accuracy and 6.0% mAP, respectively. Previous DG person Re-ID methods directly adopt the model trained on the source domains to unseen target domain without considering the biased issue. Compared to these methods, our proposed method finds the debiased feature statistics from an accurate estimated Gaussian process and significantly enhances the generalization ability.

Comparison under the Protocol-2. As shown in Tab. 3, we compare the proposed GDNorm with *QACConv*₅₀ (Liao and Shao 2020), *M*³*L* (Zhao et al. 2021) and RaMoE (Dai et al. 2021) under Protocol-2. GDNorm outperforms the performances of these methods by a large margin. Specifically,

Testing Paths	D+MT+C3→M				
	R1	R5	R10	mAP	Time
Single Path (D)	75.5	89.3	92.1	48.0	0.09s / batch
Single Path (MT)	80.8	90.8	92.2	56.4	0.09s / batch
Single Path (C3)	78.9	91.4	93.7	54.2	0.09s / batch
Mean Path	86.5	95.2	96.9	68.2	0.09s / batch

Table 5: Evaluation of the influence of different paths as well as the inference time under Protocol-2.

λ	D+MT+C3→M			
	R1	R5	R10	mAP
0.1	83.9	93.6	95.5	63.2
0.5	86.3	94.4	96.2	67.7
0.6	86.5	95.2	96.9	68.2
0.7	85.8	94.2	95.9	67.3
1	84.3	92.8	93.2	64.4

Table 6: Ablation study on sample ratio λ under Protocol-2.

our method improves the second best RaMoE by 4.5% R1 accuracy and 11.7% mAP on Market-1501. When testing on DukeMTMC, GDNorm improves the second best RaMoE by 4.6% and 6.9% in terms of R1 accuracy and mAP. When testing on CUHK03, GDNorm outperforms RaMoE by 10.0% R1 accuracy and 10.4% mAP. GDNorm also improves RaMoE by 14.0% R-1 accuracy and 6.9% mAP on MSMT17. The performances on these four large-scale Re-ID datasets have demonstrated the strong domain generalization of our GDNorm through debiased BN estimation from the accurate Gaussian process.

Ablation Study

Effectiveness of components in GDNorm. To investigate the effectiveness of each component in DGNorm, we conduct ablation studies in Tab. 4. ‘E’ denotes the model directly trained with domain-specific BN layers and tested with the mean parameters of the domain specific BN layers across all source domains. ‘S’ means employing the mean path from the learned Gaussian process as the debiased BN parameters. As shown in Tab. 4, the performance of ‘E’ outperforms the baseline by 5.2% R1 accuracy and 8.3% mAP. It indicates that the domain-specific BN layers can effectively capture the specific feature statistics information from the corresponding domains, which can be integrated to make full use of all the effective information of source domains for enhancing the generalization. On this basis, the model with the dynamically self-refining Gaussian process improves the performance by 2.8% R1 accuracy and 4.7% mAP, which demonstrates the effectiveness of the learned debiased feature statistics of BN layers, tending to real hyper-distribution.

Analysis of the influence of different paths. We analyze the performance of using different paths from the estimated Gaussian process as the optimal debias parameters of BN layers for the target domain. As shown in Fig. 2, we randomly sample paths from the learned Gaussian process to serve as BN layers in the model and test their performance. We change λ in Eq. 11 from 0.1 to 1, and randomly sample 100 paths at each λ as 100 different models. Their results are marked with black dots. The grey region is the area




Figure 2: Performance comparison between the mean path and other sampled paths under Protocol-2.

between the best sampled path’s model and the worst sampled path’s model at each λ . The horizontal line in red represents the performance of the “Mean Path”. From the results in the Fig. 2, few number of models generated by randomly sampled paths perform slightly better than the “Mean Path”, which indicates that these models more closely match the target domain. But from the perspective of average risk, “Mean Path” achieves the best performance. Moreover, as shown in Tab. 5, “Single Path (D)/(MT)/(C3)” means inference with the parameters of domain-specific BN layers learned from D/MT/C3 domain-specific model, respectively. “Mean Path” means to directly use the mean path of the Gaussian process to serve as every BN layer in the model. “Single Path (D)/(MT)/(C3)” are inferior to “Mean Path”. It indicates that our GDNorm makes use of all the domain-specific models’ features to improve the generalization ability. GDNorm only takes 0.09 seconds per query batch images to inference, which is just one-third time cost of the ensemble model with three domain-specific networks, showing the high efficiency of our GDNorm.

Effectiveness of the sampling ratio λ . The results in Tab. 6 show the influence of different sampling ratio λ in Eq. 11 during the training stage. As shown in Tab. 6, we can observe that when λ increases from 0.1 to 0.6, the proposed method obtains 2.4% and 5.0% improvements in terms of R-1 accuracy and mAP, respectively. When n_d increases from 0.6 to 1, the performance decreases by 2.2% and 3.8% R-1 in terms of accuracy and mAP, respectively. GDNorm obtains the best performance when λ is set to 0.6.

Conclusion

In this paper, we propose a novel Debiased Batch Normalization via Gaussian Process approach (GDNorm) for generalizable person re-identification. GDNorm models the feature statistic estimation (the parameters) from the domain-specific BN layers as a dynamically self-refining Gaussian process to fit the real hyper-distribution of domains for improving the generalization ability on unseen target domain. It employs the self-refining mechanism to further optimize these learned parameters of domain-specific BN layers for source domains, and estimates more accurate Gaussian process by them in return, tending to real data distribution. In addition, GDNorm directly exploits the mean path of the Gaussian process as the optimal BN estimating for unseen target domain, without requiring no extra computational cost as compared to ensemble learning. Extensive experiments show the proposed GDNorm achieves the state-of-the-art performance on multiple standard benchmarks.

Acknowledgments

This work was supported by the National Key R&D Program of China under Grant 2020AAA0105702, National Natural Science Foundation of China (NSFC) under Grant U19B2038 and 62106245, the University Synergy Innovation Program of Anhui Province under Grants GXXT-2019-025, and the Fundamental Research Funds for the Central Universities under Grant WK2100000021.

References

- Bai, S.; Bai, X.; and Tian, Q. 2017. Scalable person re-identification on supervised smoothed manifold. In *CVPR*.
- Bak, S.; Carr, P.; and Lalonde, J.-F. 2018. Domain adaptation through synthesis for unsupervised person re-identification. In *ECCV*.
- Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; and Wierstra, D. 2015. Weight uncertainty in neural network. In *ICML*.
- Bond, S. R.; Hoeffler, A.; and Temple, J. R. 2001. GMM estimation of empirical growth models. *SSRN*.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*.
- Chang, W.-G.; You, T.; Seo, S.; Kwak, S.; and Han, B. 2019. Domain-specific batch normalization for unsupervised domain adaptation. In *CVPR*.
- Chattopadhyay, P.; Balaji, Y.; and Hoffman, J. 2020. Learning to balance specificity and invariance for in and out of domain generalization. In *ECCV*.
- Chen, P.; Dai, P.; Liu, J.; Zheng, F.; Tian, Q.; and Ji, R. 2021. Dual Distribution Alignment Network for Generalizable Person Re-Identification. In *AAAI*.
- Cheng, D.; Gong, Y.; Zhou, S.; Wang, J.; and Zheng, N. 2016. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *CVPR*.
- Choi, S.; Kim, T.; Jeong, M.; Park, H.; and Kim, C. 2021. Meta Batch-Instance Normalization for Generalizable Person Re-Identification. In *CVPR*.
- Dai, Y.; Li, X.; Liu, J.; Tong, Z.; and Duan, L.-Y. 2021. Generalizable person re-identification with relevance-aware mixture of experts. In *CVPR*.
- D’Innocente, A.; and Caputo, B. 2018. Domain generalization with domain-specific aggregation modules. In *GCPA*.
- Gray, D.; and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hirzer, M.; Belezni, C.; Roth, P. M.; and Bischof, H. 2011. Person re-identification by descriptive and discriminative classification. In *SCIA*.
- Huang, Z.; Wang, H.; Xing, E. P.; and Huang, D. 2020. Self-challenging improves cross-domain generalization. *arXiv preprint arXiv:2007.02454*.
- Jia, J.; Ruan, Q.; and Hospedales, T. M. 2019. Frustratingly Easy Person Re-Identification: Generalizing Person Re-ID in Practice. In *BMVC*.
- Jin, X.; Lan, C.; Zeng, W.; and Chen, Z. 2021. Style Normalization and Restitution for Domain Generalization and Adaptation. *arXiv preprint arXiv:2101.00588*.
- Jin, X.; Lan, C.; Zeng, W.; Chen, Z.; and Zhang, L. 2020. Style normalization and restitution for generalizable person re-identification. In *CVPR*.
- Li, W.; and Wang, X. 2013. Locally aligned feature transforms across views. In *CVPR*.
- Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deep-reid: Deep filter pairing neural network for person re-identification. In *CVPR*.
- Li, W.; Zhu, X.; and Gong, S. 2017. Person Re-Identification by Deep Joint Learning of Multi-Loss Classification. In *IJ-CAI*.
- Liao, S.; and Shao, L. 2020. Interpretable and Generalizable Person Re-identification with Query-Adaptive Convolution and Temporal Lifting. In *ECCV*.
- Lin, C.-S.; Cheng, Y.-C.; and Wang, Y.-C. F. 2020. Domain Generalized Person Re-Identification via Cross-Domain Episodic Learning. *arXiv preprint arXiv:2010.09561*.
- Lin, S.; Li, C.-T.; and Kot, A. C. 2020. Multi-Domain Adversarial Feature Generalization for Person Re-Identification. *TIP*.
- Lin, S.; Li, H.; Li, C.; and Kot, A. C. 2018. Multi-task Mid-level Feature Alignment Network for Unsupervised Cross-Dataset Person Re-Identification. In *BMVC*.
- Liu, J.; Zha, Z.-J.; Chen, D.; Hong, R.; and Wang, M. 2019. Adaptive transfer network for cross-domain person re-identification. In *CVPR*.
- Liu, J.; Zha, Z.-J.; Tian, Q.; Liu, D.; Yao, T.; Ling, Q.; and Mei, T. 2016. Multi-scale triplet cnn for person re-identification. In *ACMMM*, 192–196.
- Loy, C. C.; Xiang, T.; and Gong, S. 2009. Multi-camera activity correlation analysis. In *CVPR*.
- Luo, H.; Jiang, W.; Gu, Y.; Liu, F.; Liao, X.; Lai, S.; and Gu, J. 2019. A strong baseline and batch normalization neck for deep person re-identification. *IEEE Trans Multimedia*.
- Mancini, M.; Bulò, S. R.; Caputo, B.; and Ricci, E. 2018. Best sources forward: domain generalization through source-specific nets. In *ICIP*.
- Martini, M.; Paolanti, M.; and Frontoni, E. 2020. Open-world person re-identification with RGBD camera in top-view configuration for retail applications. *IEEE Access*.
- Micikevicius, P.; Narang, S.; Alben, J.; Diamos, G.; Elsen, E.; Garcia, D.; Ginsburg, B.; Houston, M.; Kuchaiev, O.; Venkatesh, G.; et al. 2017. Mixed precision training. *arXiv preprint arXiv:1710.03740*.
- Peng, P.; Xiang, T.; Wang, Y.; Pontil, M.; Gong, S.; Huang, T.; and Tian, Y. 2016. Unsupervised cross-dataset transfer learning for person re-identification. In *CVPR*.
- Radenović, F.; Toliás, G.; and Chum, O. 2018. Fine-tuning CNN image retrieval with no human annotation. *TPAMI*.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *CVPR*.

Segu, M.; Tonioni, A.; and Tombari, F. 2020. Batch normalization embeddings for deep domain generalization. *arXiv preprint arXiv:2011.12672*.

Seo, S.; Suh, Y.; Kim, D.; Kim, G.; Han, J.; and Han, B. 2020. Learning to optimize domain specific normalization for domain generalization. In *ECCV*.

Song, J.; Yang, Y.; Song, Y.-Z.; Xiang, T.; and Hospedales, T. M. 2019. Generalizable person re-identification by domain-invariant mapping network. In *CVPR*.

Tamura, M.; and Murakami, T. 2019. Augmented hard example mining for generalizable person re-identification. *arXiv preprint arXiv:1910.05280*.

Wang, J.; Zhu, X.; Gong, S.; and Li, W. 2018. Transferable joint attribute-identity deep learning for unsupervised person re-identification. In *CVPR*.

Wang, X. 2013. Intelligent multi-camera video surveillance: A review. *Pattern Recognit Lett*.

Wei, L.; Zhang, S.; Gao, W.; and Tian, Q. 2018. Person transfer gan to bridge domain gap for person re-identification. In *CVPR*.

Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2017. Joint detection and identification feature learning for person search. In *CVPR*.

Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; and Hoi, S. C. 2021. Deep learning for person re-identification: A survey and outlook. *TPAMI*.

Yu, Y.; Yang, L.; and Li, H. 2019. An adaptive model recognition and construction method for RSSI fingerprint-based localization. *Meas Sci Technol*.

Zhang, Y.-F.; Zhang, H.; Zhang, Z.; Li, D.; Jia, Z.; Wang, L.; and Tan, T. 2021. Learning Domain Invariant Representations for Generalizable Person Re-Identification. *arXiv preprint arXiv:2103.15890*.

Zhao, Y.; Zhong, Z.; Yang, F.; Luo, Z.; Lin, Y.; Li, S.; and Sebe, N. 2021. Learning to Generalize Unseen Domains via Memory-based Multi-Source Meta-Learning for Person Re-Identification. In *CVPR*.

Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. In *ICCV*.

Zheng, W.-S.; Gong, S.; and Xiang, T. 2009. Associating Groups of People. In *BMVC*.

Zheng, Z.; Zheng, L.; and Yang, Y. 2017. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*.

Zhou, K.; Yang, Y.; Qiao, Y.; and Xiang, T. 2020. Domain adaptive ensemble learning. *arXiv preprint arXiv:2003.07325*.

Zhuang, Z.; Wei, L.; Xie, L.; Zhang, T.; Zhang, H.; Wu, H.; Ai, H.; and Tian, Q. 2020. Rethinking the distribution gap of person re-identification with camera-based batch normalization. In *ECCV*.