

PartMix: Regularization Strategy to Learn Part Discovery for Visible-Infrared Person Re-identification

Minsu Kim¹, Seungryong Kim², Jungin Park¹, Seongheon Park¹, Kwanghoon Sohn^{1,3*}
¹ Yonsei University ² Korea University ³ Korea Institute of Science and Technology (KIST)
 {minsukim320, newrun, sam121796, khsohn}@yonsei.ac.kr seungryong_kim@korea.ac.kr

Abstract

Modern data augmentation using a mixture-based technique can regularize the models from overfitting to the training data in various computer vision applications, but a proper data augmentation technique tailored for the part-based Visible-Infrared person Re-Identification (VI-ReID) models remains unexplored. In this paper, we present a novel data augmentation technique, dubbed **PartMix**, that synthesizes the augmented samples by mixing the part descriptors across the modalities to improve the performance of part-based VI-ReID models. Especially, we synthesize the positive and negative samples within the same and across different identities and regularize the backbone model through contrastive learning. In addition, we also present an entropy-based mining strategy to weaken the adverse impact of unreliable positive and negative samples. When incorporated into existing part-based VI-ReID model, PartMix consistently boosts the performance. We conduct experiments to demonstrate the effectiveness of our PartMix over the existing VI-ReID methods and provide ablation studies.

1. Introduction

Person Re-Identification (ReID), aiming to match person images in a query set to ones in a gallery set captured by non-overlapping cameras, has recently received substantial attention in numerous computer vision applications, including video surveillance, security, and persons analysis [64, 76]. Many ReID approaches [1, 2, 25, 26, 32, 41, 59, 73, 78] formulate the task as a visible-modality retrieval problem, which may fail to achieve satisfactory results under poor illumination conditions. To address this, most surveillance systems use an infrared camera that can capture the scene even in low-light conditions. However,

*Corresponding author

This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (NRF2021R1A2C2006703).

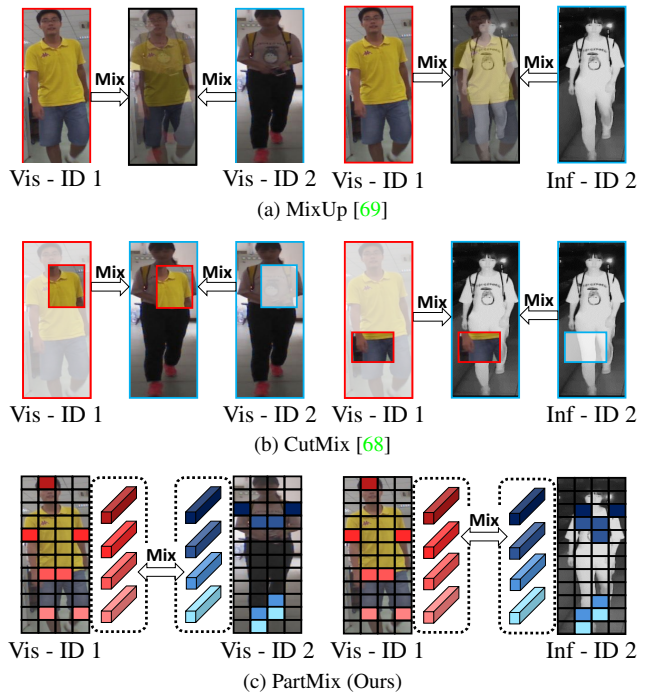


Figure 1. **Comparison of data augmentation methods for VI-ReID.** (a) MixUp [69] using a global image mixture and (b) CutMix [68] using a local image mixture can be used to regularize a model for VI-ReID, but these methods provide limited performances because they yield unnatural patterns or local patches with only background or single human part. Unlike them, we present (c) PartMix using a part descriptor mixing strategy, which boosts the VI-ReID performance (Best viewed in color).

directly matching these infrared images to visible ones for ReID poses additional challenges due to an inter-modality variation [60, 62, 65].

To alleviate these inherent challenges, Visible-Infrared person Re-Identification (VI-ReID) [5–7, 12, 24, 35, 48, 54, 60–62, 66] has been popularly proposed to handle the large intra- and inter-modality variations between visible images and their infrared counterparts. Formally, these approaches first extract a person representation from whole

visible and infrared images, respectively, and then learn a modality-invariant feature representation using feature alignment techniques, *e.g.*, triplet [5, 7, 24, 60–62] or ranking criterion [12, 66], so as to remove the inter-modality variation. However, these global feature representations solely focus on the most discriminative part while ignoring the diverse parts which are helpful to distinguish the person identity [53, 63].

Recent approaches [53, 57, 63] attempted to further enhance the discriminative power of person representation for VI-ReID by capturing diverse human body parts across different modalities. Typically, they first capture several human parts through, *e.g.*, horizontal stripes [63], clustering [53], or attention mechanisms [57] from both visible and infrared images, extract the features from these human parts, and then reduce inter-modality variation in a part-level feature representation. Although these methods reduce inter-modality variation through the final prediction (*e.g.*, identity probability), learning such part detector still leads to overfitting to the specific part because the model mainly focuses on the most discriminative part to classify the identity, as demonstrated in [4, 15, 31, 52]. In addition, these parts are different depending on the modality, it accumulates errors in the subsequent inter-modality alignment process, which hinders the generalization ability on unseen identity in test set.

On the other hand, many data augmentation [19, 22, 38, 45, 68, 69] enlarge the training set through the image mixture technique [69]. They typically exploit the samples that linearly interpolate the global [38, 44, 69] or local [19, 68] images and label pairs for training, allowing the model to have smoother decision boundaries that reduce overfitting to the training samples. This framework also can be a promising solution to reduce inter-modality variation by mixing the different modality samples to mitigate overfitting to the specific modality, but directly applying these techniques to part-based VI-ReID models is challenging in that they inherit the limitation of global and local image mixture methods (*e.g.*, ambiguous and unnatural patterns, and local patches with only background or single human part). Therefore, the performance of part-based VI-ReID with these existing augmentations would be degraded.

In this paper, we propose a novel data augmentation technique for VI-ReID task, called **PartMix**, that synthesizes the part-aware augmented samples by mixing the part descriptors. Based on the observation that learning with the unseen combination of human parts may help better regularize the VI-ReID model, we randomly mix the inter- and intra-modality part descriptors to generate positive and negative samples within the same and across different identities, and regularize the model through the contrastive learning. In addition, we also present an entropy-based mining strategy to weaken the adverse impact of unreliable posi-

tive and negative samples. We demonstrate the effectiveness of our method on several benchmarks [33, 55]. We also provide an extensive ablation study to validate and analyze components in our model.

2. Related work

Person ReID. Person Re-Identification (ReID) aims to search a target person from a large gallery set, where the images are captured from non-overlapping visible camera views. With the advent of deep convolutional neural networks (CNNs), to solve this task, existing works [1, 3, 46] encourage the person representation within the same identity to be similar through feature-level constraint, including triplet constraint [46], quadruplet constraint [3], or group consistency constraint [1]. However, since these methods learn features from the whole person image, they often suffer from intra-modality variation caused by the human pose variation and part occlusions [41]. Thus, there have been many efforts to focus on extracting human body parts that can provide fine-grained person image descriptions through uniform partitions [8, 41, 42, 47, 75] or attention mechanism [25, 28, 29, 59, 74, 77].

Visible-Infrared Person ReID. Most surveillance systems deploy infrared images to achieve satisfactory results under poor illumination conditions [20, 55]. However, directly applying person re-id methods suffer from the different distribution between modalities, typically known as modality discrepancy [60, 62, 65]. To alleviate this, Visible-Infrared person re-id (VI-ReID) has been popularly studied to match a person between visible and infrared images with challenges posed by large intra-modality variation and inter-modality variations. Wu *et al.* [55] first introduced a large-scale VI-ReID dataset, named SYSU-MM01, and proposed a deep zero-padding strategy to explore modality-specific structure in a one-stream network. Recently, modality invariant feature learning based VI-ReID has been proposed to project the features from different modalities into the same feature space. Formally, these methods extract features from visible and infrared person images and then reduce the inter-modality discrepancy by feature-level constraints, such as triplet constraint [7, 60–62], or ranking constraint [12, 66]. However, these methods usually focus only on the most discriminative part rather than the diverse parts which are helpful to distinguish different persons [53, 63]. Therefore, several methods have been proposed to align inter-modality discrepancy in a fine-grained manner, exploiting horizontal stripes [63], or inter-modality nuances [53, 57]. However, the part detector can be easily overfitted to the specific part. Moreover, these part maps are not semantically consistent across modalities, showing limited performance.

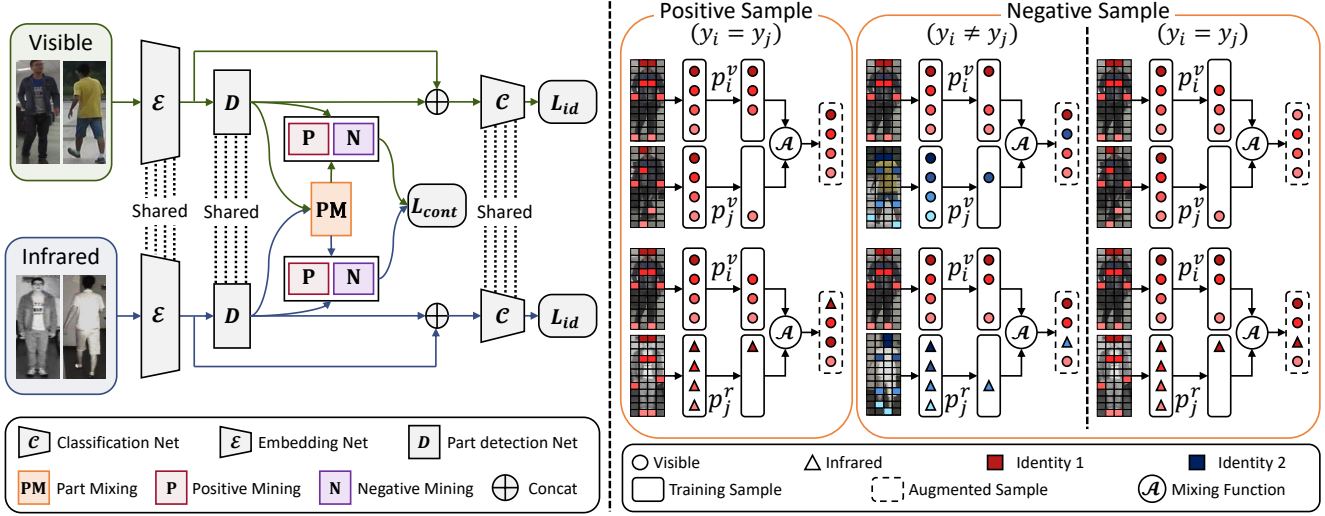


Figure 2. **Overview of our PartMix-based framework for VI-ReID.** The visible and infrared images are fed into the embedding network to extract features and obtain part descriptors through the global average pooling layer. The part descriptors from visible and infrared modalities are fed into the part mixing module to synthesize the positive and negative samples. The reliable positive and negative samples are selected through the positive and negative mining module for contrastive learning.

Data Augmentation. Data augmentation has been widely explored in various computer vision tasks [13, 21, 39, 44, 58, 69] to prevent deep neural network from overfitting to the training data. Specifically, mixture-based methods, including global image mixture [38, 44, 69], e.g., MixUp, and local image mixture [19, 68], e.g., CutMix, are dominant regularization strategies. These methods formally synthesize the virtual training samples by linearly interpolating the images and the corresponding labels. It results in smoother decision boundaries that reduce overfitting to the training data. On the other hand, some methods [22, 38] leveraged the different image mixtures for contrastive-based unsupervised learning. Recently, there exist several efforts to adapt the global image mixture to VI-ReID tasks [16, 71], but there still remains the problem of unnatural patterns, which hinder the localization ability [43, 68] for part-based VI-ReID. Unlike the methods above, we present for the first time a part-level augmentation for part-based VI-ReID, overcoming the limitations of existing data augmentation methods [38, 44, 68, 69].

3. Proposed Method

3.1. Preliminaries and Problem Formulation

Let us denote visible and infrared image sets as $\mathcal{X}^v = \{x_i^v\}_{i=1}^{N^v}$ and $\mathcal{X}^r = \{x_i^r\}_{i=1}^{N^r}$, where x_i^v and x_i^r are images, and N^v and N^r are the number of visible and infrared images, and they are unpaired. For visible and infrared image sets, the corresponding identity label sets are defined such that $\mathcal{Y}^v = \{y_i^v\}_{i=1}^{N^v}$ and $\mathcal{Y}^r = \{y_i^r\}_{i=1}^{N^r}$, whose label candidates are shared. The objective of VI-ReID is to

learn modality-invariant person descriptors, denoted by d^v and d^r , for matching persons observed from visible and infrared cameras. For simplicity, we denote visible and infrared modalities as $t \in \{v, r\}$ unless stated.

Most recent state-of-the-art methods are based on part-based person representation [53, 57, 63] that aims to extract discriminative human parts information and use them to enhance the robustness of person representation against human pose variation across the modality. These methods typically involve discovering diverse and discriminative human parts in an attention mechanism, and generating the person descriptor by assembling a global descriptor and part descriptors for retrieving visible (or infrared) images according to the given infrared (or visible) images.

Specifically, given visible and infrared images, the feature map for each modality is computed through an embedding network $\mathcal{E}(\cdot)$ such that $f^t = \mathcal{E}(x^t)$. The part detector $\mathcal{D}(\cdot)$ then produces human parts, followed by sigmoid function $\sigma(\cdot)$, to output part map probability, denoted by $\{m^t(k)\}_{k=1}^M = \sigma(\mathcal{D}(f^t))$, where M is the number of part maps. The part descriptors are then formulated as follows:

$$p^t = [p^t(k)]_{k=1}^M = [\text{GAP}(m^t(k) \odot f^t)]_{k=1}^M, \quad (1)$$

where $\text{GAP}(\cdot)$ denotes a global average pooling, \odot is an element-wise multiplication, and $[\cdot]$ is a concatenate operation. Note that they apply element-wise multiplication between $m^t(k)$ and each channel dimension in f^t . They finally concatenate the global descriptor g^t such that $l^t = \text{GAP}(f^t)$ and part descriptors p^t to obtain person descriptor d^t for matching the persons observed from visible and

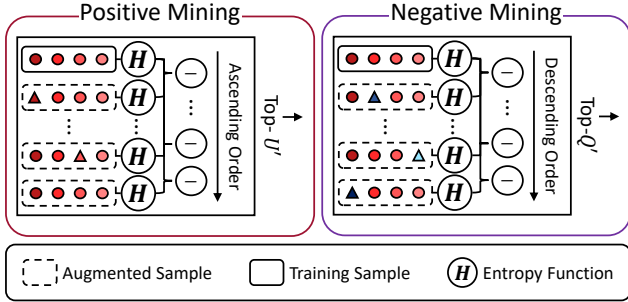


Figure 3. **Illustration of entropy-based mining strategy.** It computes the difference between the pairwise entropy for positive and negative samples. These uncertainty value sets are sorted in ascending and descending order to select reliable positive and negative samples.

infrared cameras such that

$$d^t = [g^t, p^t]. \quad (2)$$

To train such a model, since only identity labels are available, they adopted a cross-entropy loss between the identity probabilities and ground-truth identities. In addition, they also adopted several loss functions, including knowledge distillation [57] or metric learning [53, 57, 63] loss, to learn modality invariance in a part-level feature representation.

While these losses let the network focus on human parts across modalities and reduce inter-modality variations by aligning person descriptors within the same identity, the part detectors learned by these methods have been often overfitted to the specific part, as exemplified in Fig. 4. In addition, these learned parts may not be semantically consistent across modalities [14, 17, 23]. Therefore, they fail to localize discriminative and semantically-consistent human parts across the visual and infrared modalities, thus showing limited performance.

3.2. Overview

To overcome the limitations of previous part-based person representation methods [53, 57, 63], our PartMix, which we introduce in the remainder of this section, accounts for the observation that learning with the part-aware augmented samples across both inter- and intra-modality can help better regularize the model. Unlike conventional methods [19, 38, 44, 68, 69] that exploit image mixture techniques that yield unnatural patterns, and local patches with only background or single human part, we present a novel part mixing strategy to synthesize augmented samples by mixing partial descriptors across the modalities, and use them to synthesize positive and negative samples to maximize the similarities of positive pairs and minimize the similarities of negative pairs through the contrastive learning objective. It helps to regularize the model and mitigate the overfitting to the specific part and modality, improving the generalization

capability of the model. Furthermore, to eliminate the unreliable positive and negative samples, we present an entropy-based mining strategy that can help guide a representation to be more discriminative.

3.3. Part Mixing for Data Augmentation

One of the most straightforward ways of leveraging regularization to better learn part discovery may be to utilize existing augmentation techniques, e.g., using global [38, 44, 69] or local [19, 68] image mixtures. These existing strategies, however, are difficult to be directly applied to part-based VI-ReID methods [53, 57, 63] due to the following two respects. Firstly, directly applying the global image mixture methods [38, 44, 69] suffers from the locally ambiguous and unnatural pattern, and thus mixed sample confuses the model, especially for localization as demonstrated in [43, 68]. Secondly, mixing the local image region [19, 68] without part annotation may contain only a background or single human part, and thus it may cause performance degradation for part-based models which require diverse and discriminative discovery of human parts for distinguishing the person identity.

To overcome these, we present a novel data augmentation technique tailored to part-based methods, called PartMix, that mixes the part descriptors extracted from different person images. By mixing the part descriptors rather than the images, we can synthesize the augmented samples with diverse combinations of human parts. Concretely, we first collect the part descriptors in the visible and infrared modalities of the mini-batch, denoted as the descriptor bank $P^t = \{p_1^t, p_2^t, \dots, p_{N^t}^t\}$. We then mix the part descriptors through part mix operation across the inter-modality $\mathcal{A}(p_i^v(u), p_j^r(h))$ and intra-modality $\mathcal{A}(p_i^t(u), p_j^t(h))$ sequentially as follows:

$$\begin{aligned} &\mathcal{A}(p_i^v(u), p_j^r(h)) \\ &= [p_1^v(1), \dots, p_i^v(u-1), p_j^r(h), p_i^v(u+1), \dots, p_i^v(M)], \\ &\mathcal{A}(p_i^t(u), p_j^t(h)) \\ &= [p_1^t(1), \dots, p_i^t(u-1), p_j^t(h), p_i^t(u+1), \dots, p_i^t(M)], \end{aligned} \quad (3)$$

where h, u denote the randomly sampled indexes of part descriptors p^t . Note that we exclude the global descriptor g^t in the part mixing above because it contains all human body parts information.

3.4. Sample Generation for Contrastive Learning

Existing image mixture-based methods [19, 38, 44, 68, 69] generate the training samples by linearly interpolating the images and the corresponding labels. These approaches, however, only synthesize the samples with the combination of identities in the training set, and thus they have limited generalization ability on the VI-ReID task where identities

in the testing set are different from the training set. To alleviate this, we present a sample generation strategy that can synthesize positive and negative samples with the unseen combination of human parts (*i.e.* the unseen identity). In the following section, we explain how to achieve positive bank $B_i^{+,t}$ and negative bank $B_i^{-,t}$ in detail. For simplicity, only visible samples are described as an example.

Positive Samples. Our first insight is that the combination of the human parts of the persons with the same identity has to be consistent. To this end, we design positive samples that mix the same part information between the person images within the same identity. Specifically, we mix the part descriptors with the same identity using (3). Each positive sample for visible modality is denoted as

$$b_i^{+,v} = [\mathcal{A}(p_i^v(k), p_j^r(k)), \mathcal{A}(p_i^v(k), p_j^v(k))], \text{ if } y_i = y_j. \quad (4)$$

Note that we only mix the part-level descriptor within the same identity (*i.e.* $y_i = y_j$).

Negative Sample. The positive samples in (4) encourage the part-level feature within the same identity across the inter- and intra-modality to be invariant. This, however, does not guarantee that the model can distinguish the person identity with different combinations of human parts, and localize the diverse human parts within the same person.

To overcome this, we design the negative samples that encourage the model to distinguish the person identity when the combination of human parts is different and localize diverse human parts in each identity. We mix the part descriptor within the same and different identity using (3) as

$$b_i^{-,v} = \begin{cases} [\mathcal{A}(p_i^v(k), p_j^r(h)), \mathcal{A}(p_i^v(k), p_j^v(h))], & \text{if } y_i = y_j \\ [\mathcal{A}(p_i^v(k), p_j^r(k)), \mathcal{A}(p_i^v(k), p_j^v(k))], & \text{if } y_i \neq y_j \end{cases}, \quad (5)$$

where k and h denote the different indexes of part descriptors. Note that our negative samples cover the unseen combination of human parts in the training set. Therefore, these samples can be seen as out-of-distribution negative samples that can provide supportive information for improving the generalization capability of the model as investigated in [9, 40].

3.5. Entropy-based Mining Strategy

Even though the proposed sample generation strategy through part mixing yields comparable performance to some extent (which will be discussed in experiments), it relies on the assumption that there are different human part information (e.g. clothes, hairstyle) for each person. However, several persons with different identities share a similar appearance (e.g., wearing similar clothes). Therefore, simply mixing these samples makes false negative samples that still have the same combination of human parts. Moreover, the false positive samples can be synthesized by randomly

mixing human parts that have different semantic meanings, especially in the early training stage.

To overcome this, inspired by the uncertainty-based approaches [10, 37], we present an entropy-based mining strategy for eliminating the false positive and false negative samples. We exploit the difference between the pairwise entropy of the identity prediction that can be an uncertainty measure for positive and negative samples, where the reliable samples are determined as a positive pair with a smaller entropy difference and a negative pair with a larger entropy difference. We first obtain identity probability for each sample through the part-level identity classifier $\mathcal{C}_p(\cdot)$. We then compute the difference between the pairwise entropy for positive and negative samples as follows:

$$\begin{aligned} h_i^{+,v} &= [|H(\mathcal{C}_p(p_i^v)) - H(\mathcal{C}_p(b_i^{+,v}(j)))|]_{j=1}^U, \\ h_i^{-,v} &= [|H(\mathcal{C}_p(p_i^v)) - H(\mathcal{C}_p(b_i^{-,v}(j)))|]_{j=1}^Q, \end{aligned} \quad (6)$$

where $H(\cdot)$ is entropy function [37], and U, Q are the number of positive and negative pairs. These uncertainty value sets for positive $h_i^{+,v}$ and negative $h_i^{-,v}$ are then sorted in ascending and descending order, respectively. We select the top U' and Q' samples for positive bank $B_i^{+,v} = [b_i^{+,v}(j)]_{j=1}^{U'}$ and negative bank $B_i^{-,v} = [b_i^{-,v}(j)]_{j=1}^{Q'}$, respectively.

3.6. Loss Functions

In this section, we describe several loss functions to train our network. We newly present contrastive regularization loss $\mathcal{L}_{\text{cont}}$ and part ID loss \mathcal{L}_{aid} to regularize the model through positive and negative samples.

Contrastive Regularization Loss. We aim to maximize the similarities between positive pairs and minimize similarities between negative pairs. To this end, inspired by [34, 56], we adopt the contrastive regularization loss that jointly leverages positive and negative samples to regularize the model, and thus the model can overcome the limitation of strong reliance on supervision signal. We first define the similarity between part descriptors and positive (negative) samples, defined as $s_{i,j}^+ = \text{sim}(p_i, b_i^+(j))$ ($s_{i,k}^- = \text{sim}(p_i, b_i^-(k))$), where $\text{sim}(\cdot, \cdot)$ denotes a similarity function. The contrastive loss can be written as follows:

$$\mathcal{L}_{\text{cont}} = \sum_{i=1}^{N^v+N^r} -\log \frac{\sum_{j=1}^{U'} \exp(s_{i,j}^+/\tau)}{\sum_{j=1}^{U'} \exp(s_{i,j}^+/\tau) + \sum_{k=1}^{Q'} \exp(s_{i,k}^-/\tau)}, \quad (7)$$

where τ is a scaling temperature parameter.

Part ID Loss. To estimate accurate entropy values for the augmented samples, we adopt part ID loss for an additional

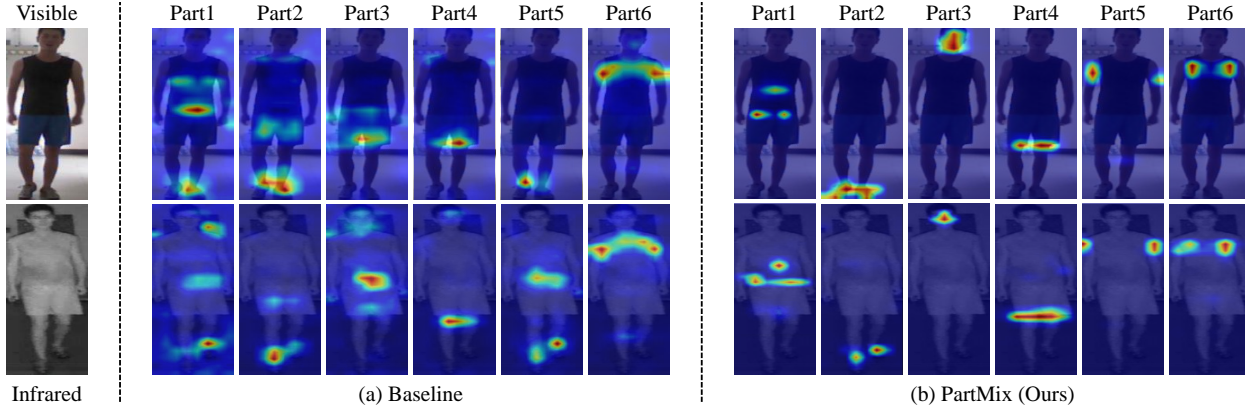


Figure 4. **Visualization of part detection results by (a) baseline [57] and (b) baseline with PartMix (Ours).** The baseline method is easily overfitted to the specific part, while our method can capture diverse and discriminative human parts.

part classifier that provides identity scores from part descriptors. It allows the model to eliminate the unreliable positive and negative samples, and learn discriminative and semantically consistent part discovery, simultaneously.

$$\mathcal{L}_{\text{aid}} = -\frac{1}{N} \sum_{i=1}^N y_i \log(\mathcal{C}_p(p_i)). \quad (8)$$

Total Loss. Following the baseline [57], we also adopt modality learning loss \mathcal{L}_{ML} , modality specific ID loss \mathcal{L}_{sid} , center cluster loss \mathcal{L}_{cc} , and identity classification loss \mathcal{L}_{id} . The detailed losses are described in the supplementary material. The total loss function of our approach can be written as $\mathcal{L} = \mathcal{L}_{\text{id}} + \mathcal{L}_{\text{cc}} + \lambda_{\text{sid}}\mathcal{L}_{\text{sid}} + \lambda_{\text{ML}}\mathcal{L}_{\text{ML}} + \lambda_{\text{aid}}\mathcal{L}_{\text{aid}} + \lambda_{\text{cont}}\mathcal{L}_{\text{cont}}$, where λ_{sid} , λ_{ML} , λ_{aid} , and λ_{cont} are weights that control the importance of each loss.

3.7. Discussion

Most recent trends in inter-domain scenarios (*e.g.*, domain adaptation [67, 70]) exploited counterfactual intervention to learn domain invariant knowledge, improving generalization capabilities. These frameworks consist of two main components: generating “counterfactual samples” by changing the domain (*e.g.*, style) and using them in the model training for “intervention”. Our PartMix satisfies these conditions, as we synthesize the modality-mixed samples by changing the part descriptor across the modality and training the whole model with these samples. Therefore, we can interpret our PartMix as a counterfactual intervention for inter-modality part-discovery, where the part mixing module can be viewed as a “counterfactual” sample generator and the contrastive regularization as an “intervention”. This allows our model to learn modality invariant part representation by encouraging the part discovery to be invariant under different interventions.

4. Experiments

4.1. Experimental Setup

In this section, we comprehensively analyze and evaluate our PartMix on several benchmarks [33, 55]. First, we analyze the effectiveness of our PartMix and comparison with other regularization methods. We then evaluate our method compared to the state-of-the-art methods for VI-ReID. In the experiment, we utilize MPANet [57] as our baseline model. Additional implementation details will be explained in the supplementary material.

Dataset. We evaluate our method on two benchmarks, SYSU-MM01 [55] and RegDB [33]. Firstly, SYSU-MM01 dataset [55] is a large-scale VI-ReID dataset. This dataset contains 395 identities with 22,258 visible images acquired by four cameras and 11,909 near-infrared images acquired by two cameras for the training set. The testing set contains 96 identities with 3,803 near-infrared images in the query, and 301 and 3,010 visible images in the gallery for single-shot and multi-shot, respectively. Secondly, RegDB dataset [33] contains 4,120 visible and infrared paired images with 412 person identities, where each person has 10 visible and 10 far-infrared images. Following [33], we randomly split the dataset for training and testing sets, where each set contains non-overlapping 206 person identities between the sets.

Evaluation Protocols. For SYSU-MM01 [55] benchmark, we follow the evaluation protocol in [55]. We test our model in all-search and indoor-search settings, where the gallery sets for the former include images captured by all four visible cameras, and the latter includes two indoor ones. For RegDB [33] benchmark, we evaluate our model on infrared to visible and visible to infrared setting, where the former retrieve infrared images from visible ones, and the latter retrieves visible ones from infrared ones. For both benchmarks, we adopt the cumulative matching character-

Methods	SYSU-MM01			
	Single-shot		Multi-shot	
	Rank-1	mAP	Rank-1	mAP
Base	70.58	68.24	75.58	62.91
Base+IntraPM	72.72	69.84	76.82	64.64
Base+InterPM	75.61	71.79	78.64	67.72
Base+IntraPM + InterPM	75.86	72.71	79.05	68.80
Ours	77.78	74.62	80.54	69.84

Table 1. Ablation study for the different components of our method on the SYSU-MM01 dataset [55].




Figure 5. Ablation study for part mixing with different numbers of the mixed part (left) and number of part maps (right).

istics (CMC) and mean average precision (mAP) as evaluation metrics.

4.2. Ablation Study

In this section, we conduct ablation analyses to investigate the effectiveness of our framework. In this ablation study, all experiments are conducted on SYSU-MM01 dataset [55].

Effectiveness of Each Component. Here we analyze the key components of PartMix, including intra-modality part mixing (IntraPM), inter-modality part mixing (InterPM), and entropy-based mining (EM). For the fair comparison, we utilize MPANet [57] as our baseline (Base) for all experiments. As summarized in Table 1, each component helps to boost performance. Starting from the baseline, adding IntraPM improves the performance which indicates that IntraPM effectively mitigates intra-modality variation (e.g., human pose variation) thanks to samples synthesized by mixing part information extracted from diverse human poses. When the InterPM is added, we observe a significant improvement, which confirms that the InterPM effectively synthesizes the unseen combination of human parts across the modalities, and thus mitigates the overfitting to the specific parts and modality, simultaneously. We also note that EM also brings the performance gain for VI-ReID by mining the reliable samples among initial positive and negative samples.

Number of Mixing Parts. In our experiment, we mix the part descriptors B times using (3). Note that we set the number of part maps as 6 in this experiment. We analyze the quantitative comparison with a different number of mixed parts B as shown in Fig. 5. The result shows that ours with

Methods	SYSU-MM01			
	Single-shot		Multi-shot	
	Rank-1	mAP	Rank-1	mAP
Base	70.58	68.24	75.58	62.91
Base+MixUp [69]	51.48	46.25	58.30	38.48
Base+Manifold [44]	71.25	67.74	76.72	62.39
Base+CutMix [68]	73.35	70.69	77.03	64.76
Ours	77.78	74.62	80.54	69.84

Table 2. Ablation study for comparison with other regularization methods on the SYSU-MM01 dataset [55].

various numbers of mixed parts, which shows the effectiveness of our PartMix. We consider ours with $B = 0$ as the baseline [57]. The result with $B = 6$ is that positive and negative samples are set to the counterpart modality samples within the same and different identities, respectively. The performance has improved as B is increased from 1, and after achieving the best performance at $B = 2$ the rest showed slightly improved performance. The lower performance with the larger number of mixed parts indicates that easy samples in which the combination of human parts is not significantly different from samples with a different identity in mini-batch can actually be less effective in regularization. Since the result with $B = 2$ has shown the best performance on the SYSU-MM01 dataset [55] and RegDB dataset [33], we set B as 2 for all experiments.

Number of Part Maps. In Fig. 5, we evaluate our model with the different number of part maps. In all experiments, we set the number of mixed parts to 1/3 of the number of part maps. The results show that our PartMix consistently boosts the performance of mAP in single-shot all-search on the SYSU-MM01 dataset. For $M = 1$, we exploit training samples with the same and different identities in the mini-batch as positive and negative samples, respectively. We consider ours with $M = 1$ as contrastive learning with the global descriptor. It shows that contrastive learning effectively regularizes the model, and thus the model can mitigate modality discrepancy. Specifically, as M is increased from 4, the performance of our model converges to high mAP. These results indicate that our PartMix can consistently capture diverse and discriminative human parts with only a small number of part maps. Since the result with $M = 6$ has shown the best performance of mAP in single-shot all-search on the SYSU-MM01 dataset, we set $M = 6$ for the remaining experiments.

4.3. Comparison to Other Regularization Methods

In this section, we validate the effectiveness of our PartMix through the comprehensive comparison with other regularization methods, including MixUp [69], Manifold MixUp [44], and CutMix [68]. Table 2 shows PartMix significantly outperforms all other regularization methods [44, 68, 69]. Interestingly, the MixUp method [69]

Method	SYSU-MM01 [55]								RegDB [33]			
	All-Search				Indoor-Search				infrared to visible		visible to infrared	
	Single-Shot		Multi-Shot		Single-Shot		Multi-Shot		Rank-1	mAP	Rank-1	mAP
	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
Two-stream [48]	11.65	12.85	16.33	8.03	15.60	21.49	22.49	13.92	-	-	-	-
One-stream [48]	12.04	13.67	16.26	8.59	16.94	22.95	22.62	15.04	-	-	-	-
Zero-Padding [48]	14.80	15.95	19.13	10.89	20.58	26.92	24.43	18.86	16.7	17.9	17.8	18.9
cmGAN [5]	26.97	27.80	31.49	22.27	31.63	42.19	37.00	32.76	-	-	-	-
D ² RL [50]	28.90	29.20	-	-	-	-	-	-	-	-	43.4	44.1
JSIA-ReID [49]	38.10	36.90	45.10	29.50	43.80	52.90	52.70	42.70	48.1	48.9	48.5	49.3
AlignGAN [48]	42.40	40.70	51.50	33.90	45.90	54.30	57.10	45.30	56.3	53.4	57.9	53.6
DFE [11]	48.71	48.59	54.63	42.14	52.25	59.68	59.62	50.60	68.0	66.7	70.2	69.2
XIV-ReID [24]	49.92	50.73	-	-	-	-	-	-	62.3	60.2	-	-
CMM+CML [27]	51.80	51.21	56.27	43.39	54.98	63.70	60.42	53.52	59.8	60.9	-	-
SIM [18]	56.93	60.88	-	-	-	-	-	-	75.2	78.3	74.7	75.2
CoAL [51]	57.22	57.20	-	-	63.86	70.84	-	-	74.1	69.9	-	-
DG-VAE [36]	59.49	58.46	-	-	-	-	-	-	-	-	73.0	71.8
cm-SSFT [30]	61.60	63.20	63.40	62.00	70.50	72.60	73.00	72.40	71.0	71.7	72.3	72.9
SMCL [54]	67.39	61.78	72.15	54.93	68.84	75.56	79.57	66.57	83.05	78.57	83.93	79.83
MPANet [57]	70.58	68.24	75.58	62.91	76.74	80.95	84.22	75.11	82.8	80.7	83.7	80.9
MSCLNet [72]	76.99	71.64	-	-	78.49	81.17	-	-	83.86	78.31	84.17	80.99
Ours	77.78	74.62	80.54	69.84	81.52	84.38	87.99	79.95	84.93	82.52	85.66	82.27

Table 3. **Quantitative evaluation on SYSU-MM01 dataset [55] and RegDB dataset [33].** For evaluation, we measure Rank-1 accuracy(%) and mAP(%). Our results show the best results in terms of Rank-1 accuracy and mAP.

highly degrades the performance for part-based VI-ReID. It demonstrates that simply applying MixUp to the part-based VI-ReID degrades the localization ability of the model due to ambiguous and unnatural patterns in mixed images, and thus the model fails to distinguish the different person identities as done in the literature [43, 68]. On the other hand, Manifold MixUp [44] shows slightly improved rank-1 accuracy, but achieves lower performance in mAP than the baseline. The result shows that it still inherits the limitation of the global mixture model [38, 44, 69]. Although CutMix [68] achieves improved performance than the baseline [57], it achieves relatively lower performance than ours by 3.93% and 5.08% mAP in single-shot all-search and multi-shot all-search on the SYSU-MM01 dataset. It demonstrated that our method effectively alleviates the overfitting to the specific part and modality in part-based VI-ReID. Based on all these evaluation and comparison results, we can confirm the effectiveness of our methods.

4.4. Comparison to Other Methods

In this section, we evaluate our framework through comparison to state-of-the-art methods for VI-ReID, including Two-stream [48], One-stream [48], Zero-Padding [48], cmGAN [5], D²RL [50], JSIA-ReID [49], AlignGAN [48], DFE [11], XIV-ReID [24], CMM+CML [27], SIM [18], CoAL [51], DG-VAE [36], cm-SSFT [30], SMCL [54], MPANet [57], and MSCLNet [72].

Results on SYSU-MM01 dataset. We evaluate our PartMix on SYSU-MM01 benchmark [55] as provided in Table 3. PartMix achieves the Rank-1 accuracy of 77.78% and

mAP of 74.62% in all-search with single-shot mode, improving the Rank-1 accuracy by 0.79% and mAP by 2.98% over the MSCLNet [72]. In indoor-search with single-shot mode, our PartMix outperforms the MSCLNet [72] by Rank-1 accuracy of 3.03% and mAP of 3.21%.

Results on RegDB dataset. We also evaluate our method on RegDB benchmark [33]. As shown in Table 3, PartMix records state-of-the-art results with the Rank-1 accuracy of 84.93% and mAP of 82.52% in infrared to visible and the Rank-1 accuracy of 85.66% and mAP of 82.27% in visible to infrared mode. Our PartMix outperforms the Rank-1 accuracy by 1.07% and mAP by 4.21% in infrared to visible mode and the Rank-1 accuracy by 1.49% and mAP by 1.28% in visible to infrared mode over the MSCLNet [72].

5. Conclusion

In this paper, we have presented a novel data augmentation technique, called PartMix, that generates part-aware augmented samples by mixing the part descriptors. We introduce a novel sample generation method to synthesize the positive and negative samples and an entropy-based mining strategy to select reliable positive and negative samples to regularize the model through the contrastive objective. We have shown that PartMix achieves state-of-the-art performance over the existing methods on several benchmarks.

Acknowledgements. This research was supported by the Yonsei Signature Research Cluster Program of 2022 (2022-22-0002) and the KIST Institutional Program (Project No.2E32283-23-064).

References

- [1] Dapeng Chen, Dan Xu, Hongsheng Li, Nicu Sebe, and Xiaogang Wang. Group consistent similarity learning via deep crf for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8649–8658, 2018. 1, 2
- [2] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abdnnet: Attentive but diverse person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8351–8361, 2019. 1
- [3] Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 403–412, 2017. 2
- [4] Junsuk Choe and Hyunjung Shim. Attention-based dropout layer for the weakly-supervised object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2219–2228, 2019. 2
- [5] Pingyang Dai, Rongrong Ji, Haibin Wang, Qiong Wu, and Yuyu Huang. Cross-modality person re-identification with generative adversarial training. In *IJCAI*, volume 1, page 6, 2018. 1, 2, 8
- [6] Xing Fan, Wei Jiang, Hao Luo, and Weijie Mao. Modality-transfer generative adversarial network and dual-level unified latent representation for visible thermal person re-identification. *The Visual Computer*, pages 1–16, 2020. 1
- [7] Zhanxiang Feng, Jianhuang Lai, and Xiaohua Xie. Learning modality-specific representations for visible-infrared person re-identification. *IEEE Transactions on Image Processing*, 29:579–590, 2019. 1, 2
- [8] Yang Fu, Yunchao Wei, Yuqian Zhou, Honghui Shi, Gao Huang, Xinchao Wang, Zhiqiang Yao, and Thomas Huang. Horizontal pyramid matching for person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8295–8302, 2019. 2
- [9] Jonas Geiping, Gowthami Somepalli, Ravid Shwartz-Ziv, Andrew Gordon Wilson, Tom Goldstein, and Micah Goldblum. How much data is augmentation worth? In *ICML 2022: Workshop on Spurious Correlations, Invariance and Stability*. 5
- [10] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. *Advances in neural information processing systems*, 17, 2004. 5
- [11] Yi Hao, Nannan Wang, Xinbo Gao, Jie Li, and Xiaoyu Wang. Dual-alignment feature embedding for cross-modality person re-identification. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 57–65, 2019. 8
- [12] Yi Hao, Nannan Wang, Jie Li, and Xinbo Gao. Hsme: Hypersphere manifold embedding for visible thermal person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8385–8392, 2019. 1, 2
- [13] Dan Hendrycks, Norman Mu, Ekin Dogus Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*, 2019. 3
- [14] Luwei Hou, Yu Zhang, Kui Fu, and Jia Li. Informative and consistent correspondence mining for cross-domain weakly supervised object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9929–9938, 2021. 4
- [15] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. *Advances in Neural Information Processing Systems*, 31, 2018. 2
- [16] Zhipeng Huang, Jiawei Liu, Liang Li, Kecheng Zheng, and Zheng-Jun Zha. Modality-adaptive mixup and invariant decomposition for rgb-infrared person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1034–1042, 2022. 3
- [17] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5001–5009, 2018. 4
- [18] Mengxi Jia, Yunpeng Zhai, Shijian Lu, Siwei Ma, and Jian Zhang. A similarity inference metric for rgb-infrared cross-modality person re-identification. In *IJCAI*, pages 1026–1032, 2021. 8
- [19] Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song. Puzzle mix: Exploiting saliency and local statistics for optimal mixup. In *International Conference on Machine Learning*, pages 5275–5285. PMLR, 2020. 2, 3, 4
- [20] Minsu Kim, Sunghun Joung, Kihong Park, Seungryong Kim, and Kwanghoon Sohn. Unpaired cross-spectral pedestrian detection via adversarial feature learning. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1650–1654. IEEE, 2019. 2
- [21] Dogyoon Lee, Jaeha Lee, Junhyeop Lee, Hyeongmin Lee, Minhyeok Lee, Sungmin Woo, and Sangyoun Lee. Regularization strategy for point cloud via rigidly mixed sample. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15900–15909, 2021. 3
- [22] Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee. *i*-mix: A domain-agnostic strategy for contrastive representation learning. In *International Conference on Learning Representations*, 2020. 2, 3
- [23] Dong Li, Jia-Bin Huang, Yali Li, Shengjin Wang, and Ming-Hsuan Yang. Weakly supervised object localization with progressive domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3512–3520, 2016. 4
- [24] Diangang Li, Xing Wei, Xiaopeng Hong, and Yihong Gong. Infrared-visible cross-modal person re-identification with an x modality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4610–4617, 2020. 1, 2, 8
- [25] Wei Li, Xiatian Zhu, and Shaogang Gong. Harmonious attention network for person re-identification. In *Proceedings*

- of the *IEEE/CVF conference on computer vision and pattern recognition*, pages 2285–2294, 2018. 1, 2
- [26] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z Li. Person re-identification by local maximal occurrence representation and metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2197–2206, 2015. 1
- [27] Yongguo Ling, Zhun Zhong, Zhiming Luo, Paolo Rota, Shaozi Li, and Nicu Sebe. Class-aware modality mix and center-guided metric learning for visible-thermal person re-identification. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 889–897, 2020. 8
- [28] Hao Liu, Jiashi Feng, Meibin Qi, Jianguo Jiang, and Shuicheng Yan. End-to-end comparative attention networks for person re-identification. *IEEE Transactions on Image Processing*, 26(7):3492–3506, 2017. 2
- [29] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 350–359, 2017. 2
- [30] Yan Lu, Yue Wu, Bin Liu, Tianzhu Zhang, Baopu Li, Qi Chu, and Nenghai Yu. Cross-modality person re-identification with shared-specific feature transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13379–13389, 2020. 8
- [31] Jinjie Mai, Meng Yang, and Wenfeng Luo. Erasing integrated learning: A simple yet effective approach for weakly supervised object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8766–8775, 2020. 2
- [32] Tetsu Matsukawa, Takahiro Okabe, Einoshin Suzuki, and Yoichi Sato. Hierarchical gaussian descriptor for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1363–1372, 2016. 1
- [33] Dat Tien Nguyen, Hyung Gil Hong, Ki Wan Kim, and Kang Ryoung Park. Person recognition system based on a combination of body images from visible light and thermal cameras. *Sensors*, 17(3):605, 2017. 2, 6, 7, 8
- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 5
- [35] Hyunjong Park, Sanghoon Lee, Junhyup Lee, and Bumsub Ham. Learning by aligning: Visible-infrared person re-identification using cross-modal correspondences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12046–12055, 2021. 1
- [36] Nan Pu, Wei Chen, Yu Liu, Erwin M Bakker, and Michael S Lew. Dual gaussian-based variational subspace disentanglement for visible-infrared person re-identification. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2149–2158, 2020. 8
- [37] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 5
- [38] Zhiqiang Shen, Zechun Liu, Zhuang Liu, Marios Savvides, Trevor Darrell, and Eric Xing. Un-mix: Rethinking image mixtures for unsupervised visual representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2216–2224, 2022. 2, 3, 4, 8
- [39] Connor Shorten and Taghi M Khoshgofaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019. 3
- [40] Abhishek Sinha, Kumar Ayush, Jiaming Song, Burak Uzkent, Hongxia Jin, and Stefano Ermon. Negative data augmentation. *arXiv preprint arXiv:2102.05113*, 2021. 5
- [41] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, pages 480–496, 2018. 1, 2
- [42] Rahul Rama Varior, Bing Shuai, Jiwen Lu, Dong Xu, and Gang Wang. A siamese long short-term memory architecture for human re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 135–153. Springer, 2016. 2
- [43] Shashanka Venkataramanan, Ewa Kijak, Laurent Amsaleg, and Yannis Avrithis. Alignmixup: Improving representations by interpolating aligned features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19174–19183, 2022. 3, 4, 8
- [44] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *International Conference on Machine Learning*, pages 6438–6447. PMLR, 2019. 2, 3, 4, 7, 8
- [45] Devesh Walawalkar, Zhiqiang Shen, Zechun Liu, and Marios Savvides. Attentive cutmix: An enhanced data augmentation approach for deep learning based image classification. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3642–3646. IEEE, 2020. 2
- [46] Guangrun Wang, Liang Lin, Shengyong Ding, Ya Li, and Qing Wang. Dari: Distance metric and representation integration for person verification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016. 2
- [47] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 274–282, 2018. 2
- [48] Guan’an Wang, Tianzhu Zhang, Jian Cheng, Si Liu, Yang Yang, and Zengguang Hou. Rgb-infrared cross-modality person re-identification via joint pixel and feature alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3623–3632, 2019. 1, 8
- [49] Guan-An Wang, Tianzhu Zhang, Yang Yang, Jian Cheng, Jianlong Chang, Xu Liang, and Zeng-Guang Hou. Cross-modality paired-images generation for rgb-infrared person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12144–12151, 2020. 8

- [50] Zhixiang Wang, Zheng Wang, Yinqiang Zheng, Yung-Yu Chuang, and Shin'ichi Satoh. Learning to reduce dual-level discrepancy for infrared-visible person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 618–626, 2019. 8
- [51] Xing Wei, Diangang Li, Xiaopeng Hong, Wei Ke, and Yihong Gong. Co-attentive lifting for infrared-visible person re-identification. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1028–1037, 2020. 8
- [52] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1568–1576, 2017. 2
- [53] Ziyu Wei, Xi Yang, Nannan Wang, and Xinbo Gao. Flexible body partition-based adversarial learning for visible infrared person re-identification. *IEEE Transactions on Neural Networks and Learning Systems*, 2021. 2, 3, 4
- [54] Ziyu Wei, Xi Yang, Nannan Wang, and Xinbo Gao. Syncretic modality collaborative learning for visible infrared person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 225–234, 2021. 1, 8
- [55] Ancong Wu, Wei-Shi Zheng, Hong-Xing Yu, Shaogang Gong, and Jianhuang Lai. Rgb-infrared cross-modality person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5380–5389, 2017. 2, 6, 7, 8
- [56] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Contrastive learning for compact single image dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10551–10560, 2021. 5
- [57] Qiong Wu, Pingyang Dai, Jie Chen, Chia-Wen Lin, Yongjian Wu, Feiyue Huang, Bineng Zhong, and Rongrong Ji. Discover cross-modality nuances for visible-infrared person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4330–4339, 2021. 2, 3, 4, 6, 7, 8
- [58] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. *Advances in Neural Information Processing Systems*, 33:6256–6268, 2020. 3
- [59] Wenjie Yang, Houjing Huang, Zhang Zhang, Xiaotang Chen, Kaiqi Huang, and Shu Zhang. Towards rich feature discovery with class activation maps augmentation for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1389–1398, 2019. 1, 2
- [60] Mang Ye, Xiangyuan Lan, Qingming Leng, and Jianbing Shen. Cross-modality person re-identification via modality-aware collaborative ensemble learning. *IEEE Transactions on Image Processing*, 29:9387–9399, 2020. 1, 2
- [61] Mang Ye, Xiangyuan Lan, Jiawei Li, and Pong Yuen. Hierarchical discriminative learning for visible thermal person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 1, 2
- [62] Mang Ye, Xiangyuan Lan, Zheng Wang, and Pong C Yuen. Bi-directional center-constrained top-ranking for visible thermal person re-identification. *IEEE Transactions on Information Forensics and Security*, 15:407–419, 2019. 1, 2
- [63] Mang Ye, Jianbing Shen, David J Crandall, Ling Shao, and Jiebo Luo. Dynamic dual-attentive aggregation learning for visible-infrared person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 229–247. Springer, 2020. 2, 3, 4
- [64] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):2872–2893, 2021. 1
- [65] Mang Ye, Jianbing Shen, and Ling Shao. Visible-infrared person re-identification via homogeneous augmented trimodal learning. *IEEE Transactions on Information Forensics and Security*, 16:728–739, 2020. 1, 2
- [66] Mang Ye, Zheng Wang, Xiangyuan Lan, and Pong C Yuen. Visible thermal person re-identification via dual-constrained top-ranking. In *IJCAI*, volume 1, page 2, 2018. 1, 2
- [67] Zhongqi Yue, Qianru Sun, Xian-Sheng Hua, and Hanwang Zhang. Transporting causal mechanisms for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8599–8608, 2021. 6
- [68] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019. 1, 2, 3, 4, 7, 8
- [69] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. 1, 2, 3, 4, 7, 8
- [70] Xiheng Zhang, Yongkang Wong, Xiaofei Wu, Juwei Lu, Mohan Kankanhalli, Xiangdong Li, and Weidong Geng. Learning causal representation for training cross-domain pose estimator via generative interventions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11270–11280, 2021. 6
- [71] Yukang Zhang, Yan Yan, Yang Lu, and Hanzi Wang. Towards a unified middle modality learning for visible-infrared person re-identification. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 788–796, 2021. 3
- [72] Yiyuan Zhang, Sanyuan Zhao, Yuhao Kang, and Jianbing Shen. Modality synergy complement learning with cascaded aggregation for visible-infrared person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 462–479. Springer, 2022. 8
- [73] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Densely semantically aligned person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 667–676, 2019. 1

- [74] Liming Zhao, Xi Li, Yueting Zhuang, and Jingdong Wang. Deeply-learned part-aligned representations for person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3219–3228, 2017. [2](#)
- [75] Feng Zheng, Cheng Deng, Xing Sun, Xinyang Jiang, Xiaowei Guo, Zongqiao Yu, Feiyue Huang, and Rongrong Ji. Pyramidal person re-identification via multi-loss dynamic training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8514–8522, 2019. [2](#)
- [76] Liang Zheng, Yi Yang, and Alexander G Hauptmann. Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984*, 2016. [1](#)
- [77] Meng Zheng, Srikrishna Karanam, Ziyang Wu, and Richard J Radke. Re-identification with consistent attentive siamese networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5735–5744, 2019. [2](#)
- [78] Kuan Zhu, Haiyun Guo, Zhiwei Liu, Ming Tang, and Jinqiao Wang. Identity-guided human semantic parsing for person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 346–363. Springer, 2020. [1](#)