

Hi-CMD: Hierarchical Cross-Modality Disentanglement for Visible-Infrared Person Re-Identification

Seokeon Choi Sumin Lee Youngeun Kim Taekyung Kim Changick Kim
Korea Advanced Institute of Science and Technology, Daejeon, Korea
{seokeon, suminlee94, youngeunkim, tkkim93, changick}@kaist.ac.kr

Abstract

Visible-infrared person re-identification (VI-ReID) is an important task in night-time surveillance applications, since visible cameras are difficult to capture valid appearance information under poor illumination conditions. Compared to traditional person re-identification that handles only the intra-modality discrepancy, VI-ReID suffers from additional cross-modality discrepancy caused by different types of imaging systems. To reduce both intra- and cross-modality discrepancies, we propose a Hierarchical Cross-Modality Disentanglement (Hi-CMD) method, which automatically disentangles ID-discriminative factors and ID-excluded factors from visible-thermal images. We only use ID-discriminative factors for robust cross-modality matching without ID-excluded factors such as pose or illumination. To implement our approach, we introduce an ID-preserving person image generation network and a hierarchical feature learning module. Our generation network learns the disentangled representation by generating a new cross-modality image with different poses and illuminations while preserving a person's identity. At the same time, the feature learning module enables our model to explicitly extract the common ID-discriminative characteristic between visible-infrared images. Extensive experimental results demonstrate that our method outperforms the state-of-the-art methods on two VI-ReID datasets. The source code is available at: <https://github.com/bismex/HiCMD>.

1. Introduction

Person re-identification (ReID) aims to match a specific person across multiple non-overlapping camera views. Due to its usefulness in security and surveillance systems, person ReID has been of great research interest in recent years. Existing ReID methods mainly treat visible images captured by single-modality cameras, and depend on human appearance for RGB-RGB matching [35, 42, 45, 25, 36, 26]. However, visible light cameras can not capture all the ap-

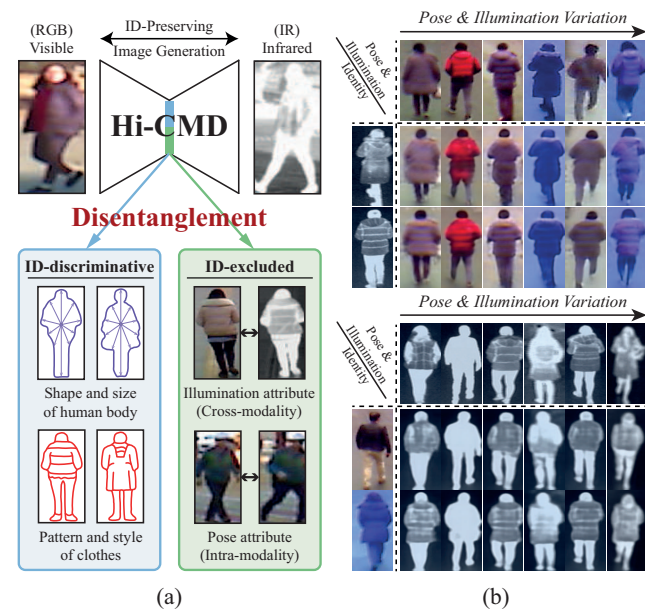


Figure 1. (a) Illustration of our Hierarchical Cross-Modality Disentanglement (Hi-CMD) concept. Our Hi-CMD method aims to hierarchically disentangle ID-discriminative factors (e.g. body shape and clothes pattern) and ID-excluded factors (e.g. pose and illumination) from RGB-IR images. (b) Examples of ID-preserving Person Image Generation (ID-PIG). The images in each row show that pose and illumination attributes can be changed while maintaining the identity information. Best viewed in color.

pearance characteristics of a person under poor illumination conditions. For these conditions, most surveillance cameras automatically switch from visible to the infrared mode in dark environments [33, 34]. After all, it becomes essential to consider visible-infrared person re-identification (VI-ReID). The goal of VI-ReID is to match pedestrians observed from visible and infrared cameras with different spectra.

Compared to the traditional ReID task that only has the intra-modality discrepancy problem, VI-ReID encounters the additional cross-modality discrepancy problem resulting from the natural difference between the reflectivity of

the visible spectrum and the emissivity of the thermal spectrum [23]. Eventually, the coexistence of intra- and cross-modality discrepancies leads to a critical situation where the intra-class distance is larger than the inter-class distance in VI-ReID [38, 31]. In this situation, most studies [33, 37, 39, 2, 38, 9] have attempted to reduce both discrepancies with feature-level constraints like the traditional ReID methods [28, 27]. It is difficult to eliminate the intractable discrepancies successfully using only feature-level constraints, since the illumination and pose attributes are entangled in a single image. More recent research [31] has attempted to bridge the cross-modality gap using an image-level constraint. However, they only translate an infrared (or visible) image into its visible (or infrared) counterpart without considering intra-modality discrepancy, despite the insufficient amount of cross-view paired training data.

To mitigate the coexisting intra- and cross-modality discrepancies at the same time, we propose a novel Hierarchical Cross-Modality Disentanglement (Hi-CMD) method, as shown in Fig. 1 (a). The goal of our approach is to hierarchically disentangle ID-excluded factors (*i.e.* pose and illumination) and ID-discriminative factors (*i.e.* body shape and clothes pattern) from cross-modality images using image-level constraints. To this end, we introduce the ID-preserving Person Image Generation (ID-PIG) network. The ID-PIG network focuses on learning the ID-excluded feature representation by replacing some latent vectors in a pair of cross-modality images. As a result, ID-PIG can transform pose and illumination attributes while preserving the identity information of a person as visualized in Fig. 1 (b). The visualization results of the ID-PIG network show that unnecessary information (*i.e.* pose or illumination attribute) can be separated from entangled representations.

Besides, we introduce the Hierarchical Feature Learning (HFL) module coupled with the ID-PIG network. This module enables the encoders in the generator to extract the common ID-discriminative factors explicitly, which is robust to pose and illumination variations. It also implicitly helps our model separate pose and illumination from RGB-IR images, which improves image generation quality. Finally, the ID-discriminative feature is used for solving the cross-modality image retrieval problem in VI-ReID. Note that we train a whole network in an end-to-end manner without pose-related supervision (*e.g.* 3D skeletal pose, keypoint heatmaps, and pose-guided parsing) compared to the existing pose generation methods [18, 22, 6, 45, 26].

Our main contributions can be summarized as follows:

- We propose a Hierarchical Cross-Modality Disentanglement (Hi-CMD) method. It is an efficient learning structure that extracts pose- and illumination-invariant features for cross-modality matching. To the best of our knowledge, this is the first work to disentangle ID-discriminative factors and ID-excluded factors simul-

taneously from cross-modality images in VI-ReID.

- The proposed ID-preserving Person Image Generation (ID-PIG) network makes it possible to change the pose and illumination attributes while maintaining the identity characteristic of a specific person. Exploring person attributes through ID-PIG demonstrates the effectiveness of our disentanglement approach.
- Extensive experimental results show that our novel framework outperforms the state-of-the-art methods on two VI-ReID datasets. The visualization results of the ID-PIG network demonstrate the overwhelming performance of our proposed method.

2. Related Work

Visible-infrared person re-identification. The visible-infrared person re-identification (VI-ReID) is about matching cross-modality images under different illumination conditions. The VI-ReID task is challenging due to cross-modality variation in addition to intra-modality variation. At the beginning of the study, most of the work has focused on how to design a feature embedding network such as a deep zero-padding network [33] and a two-stream CNN network [37]. Recently, adversarial learning [2] or metric learning methods [39, 9, 38] are applied to learn the feature representation of heterogeneous person images involving intra-modality and cross-modality variations. However, it is difficult to overcome pixel-level differences resulting from illumination or pose variations by feature-level constraints alone because of the insufficient data. In contrast to most existing feature-level approaches, our Hi-CMD method focuses on an image-level approach by combining an image generation technique with the VI-ReID task to effectively bridge both cross-modality and intra-modality gaps.

Person re-identification based on image generation. Recently, image generation methods by Generative Adversarial Networks (GANs) [8] have drawn a lot of attention in person ReID. Most of the existing work is categorized in two ways as follows: pose transfer [18, 22, 6] and style transfer [32, 19, 5, 31]. The work of the former approach points out that the existing datasets do not provide sufficient pose coverage to learn a pose-invariant representation. Accordingly, it addresses this issue by using pose-rich data augmentation. However, since this work is designed for a single-modality environment, it is difficult to apply the pose-guided methods directly to the VI-ReID task.

Another image generation approach in ReID is to reduce the domain gap between different camera domains based on unsupervised domain adaptation [32, 19, 5]. Most methods focus on transforming styles while maintaining the structural information of the person. In a similar approach, Wang *et al.* [31] translate an infrared (or visible) image into its visible (or infrared) counterpart to reduce the cross-modality

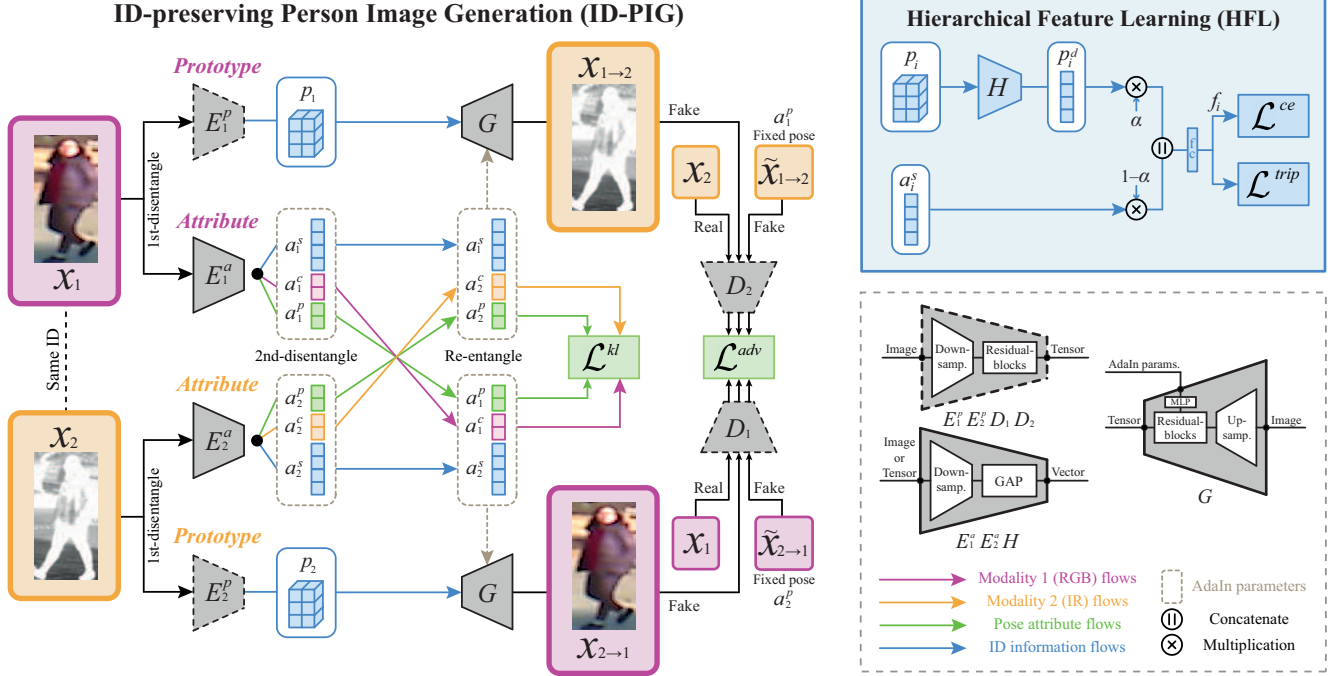


Figure 2. The framework of our Hi-CMD method. The entire framework includes two important components: the ID-preserving Person Image Generation (ID-PIG) network and the Hierarchical Feature Learning (HFL) module. Our goal is to disentangle ID-discriminative factors and ID-excluded factors from cross-modality images. Reconstruction losses are shown in Fig. 3. Best viewed in color.

discrepancy in the VI-ReID task. However, since most style transfer methods in ReID do not care about the lack of cross-view paired training data, view-invariant representations are hardly exploited.

In short, the above image generation methods in ReID handle only structural information or only convert the image style at the image-level. Unlike single-modality person re-identification, it is more important to consider intra-modality and cross-modality characteristics simultaneously in VI-ReID. To this end, we attempt to alleviate the cross-modality and intra-modality discrepancies at the same time by applying a novel hierarchical disentanglement approach even without pose supervision.

Disentangled representation learning for recognition.

The goal of disentangled representation learning is to extract explanatory factors from diverse data variation for generating a meaningful representation. Recently, considerable attention has been focused on learning disentangled representations in various fields [12, 7, 13, 24]. Also in the recognition task, several studies have tried to disentangle the identity-related information and the identity-irrelevant information from an image (e.g. pose, viewpoint, age, and other attributes) [41, 29, 1, 30]. Among them, some previous works in the single-modality person re-identification task have been conducted with a purpose of disentangling foreground, background, and pose factors [20], or extracting illumination-invariant features [40]. Note that the VI-

ReID task is particularly challenging to disentangle the common identity information and the remaining attributes from RGB-IR images due to the coexistence of cross-modality and intra-modality discrepancies. To deal with both pose and illumination attributes simultaneously, we introduce a novel hierarchical disentanglement approach. To the best of our knowledge, this is the first work to disentangle ID-discriminative factors and ID-excluding factors (*i.e.* pose and illumination attributes) from RGB-IR images in the VI-ReID task.

3. Proposed Method

3.1. Problem Definition and Overview

Problem definition. We denote the visible image and the infrared image as $x_1 \in \mathbb{R}^{H \times W \times 3}$ and $x_2 \in \mathbb{R}^{H \times W \times 3}$ respectively, where H and W are the height and the width of images. Each of images x_1 and x_2 corresponds to an identity label $y \in \{1, 2, \dots, N\}$, where N is the number of person identities. In the training stage, a feature extraction network $\phi(\cdot)$ is trained with the cross-modality image sets \mathcal{X}_1 and \mathcal{X}_2 . In the testing stage, given a query image with one modality, a ranking list within the gallery set of the other modality is calculated. The distance between two feature vectors $\phi(x_1)$ and $\phi(x_2)$ is computed by the Euclidean distance.

Framework overview. In the VI-ReID task, the most challenging issue is that both cross- and intra-modality discrepancies coexist between visible and infrared images. To address this issue effectively, we propose a novel Hierarchical Cross-Modality Disentanglement (Hi-CMD) method. Our Hi-CMD method aims to disentangle ID-discriminative factors and ID-excluded factors from cross-modality images to reduce the cross-modality and intra-modality discrepancies at the same time. To achieve this goal, we introduce two key components, the ID-preserving Person Image Generation (ID-PIG) network and the Hierarchical Feature Learning (HFL) module, as shown in Fig. 2.

3.2. Identity Preserving Person Image Generation

Hierarchical representation. We present the hierarchical representation of person images for VI-ReID. As illustrated in Fig. 2, our ID-PIG network consists of two disentanglement stages. In the first stage, we design a prototype encoder E_i^p and an attribute encoder E_i^a for each modality ($i = 1$ for visible images and $i = 2$ for infrared images). These encoders E_i^p and E_i^a map x_i to the corresponding prototype code p_i and attribute code a_i , respectively. The prototype code p_i is a tensor containing the underlying form of a person appearance such as clothes pattern and body shape. On the other hand, the attribute code a_i is a vector including clothes style and changeable attributes depending on the situation such as pose and illumination. In the second stage, the attribute code is divided into three types of codes once more as $a_i = [a_i^s; a_i^c; a_i^p]$, which includes a style attribute code a_i^s , an illumination attribute code a_i^c , and a pose attribute code a_i^p . The illumination and pose attribute codes a_i^c and a_i^p correspond to the cross-modality variation and the intra-modality variation, respectively. Note that we refer to the visual difference caused by different RGB and IR cameras as the illumination attribute. For clarity, both codes a_i^c and a_i^p can be integrated into an ID-excluded attribute code as $a_i^{ex} = [a_i^c; a_i^p]$. In summary, ID-excluded factors involves the illumination attribute code a_i^c and the pose attribute code a_i^p , while ID-discriminative factors corresponds to the style attribute code a_i^s and the prototype code p_i . This assumption is different from the case of single-modality re-identification where color information is the key clue as discussed in [43].

Disentangling ID-excluded factors. In the image generation process, our primary strategy is to synthesize a pair of cross-modality images by swapping the ID-excluded factors of two images with the same ID. Since two cross-modality images share the same ID characteristic, we can apply the image reconstruction loss between the translated image and the cross-modality image. Formally, this cross-modality reconstruction loss is formulated as follows:

$$\mathcal{L}_{recon1}^{cross} = \mathbb{E}_{\substack{x_1 \sim p_{data}(x_1) \\ x_2 \sim p_{data}(x_2)}}} [\|x_1 - G(p_2, a_2^s, a_1^{ex})\|_1], \quad (1)$$

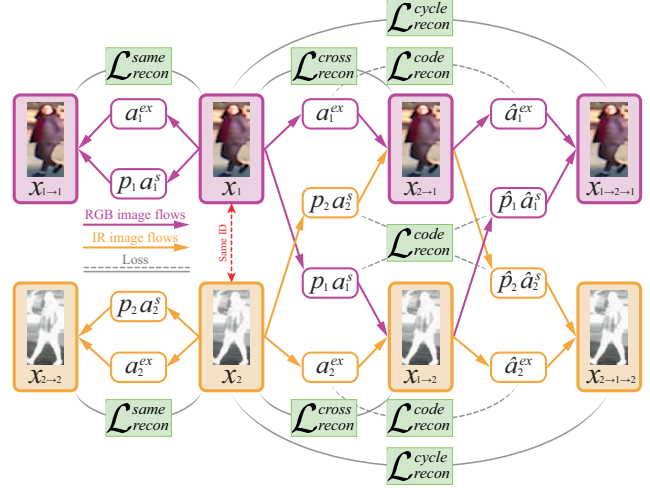


Figure 3. Illustration of the proposed reconstruction losses.

where $p_i = E_i^p(x_i)$, $[a_i^s; a_i^{ex}] = E_i^a(x_i)$, and G denotes a decoder. The l_1 loss encourages the generator to create sharp images. From this cross-modality reconstruction loss, the generator learns how to encode and decode the ID-excluded factors. To be clear, we only represent the loss for one modality as $\mathcal{L}_{recon1}^{cross}$. Another loss $\mathcal{L}_{recon2}^{cross}$ is defined by changing the index of modalities.

Reconstruction loss. We propose three additional reconstruction losses to improve the generation quality further, as expressed in Fig. 3. In addition to the loss of reconstructing images of different modalities, we apply a loss to reconstruct images of the same modality. This same-modality reconstruction loss plays a key role in regularization in the generation network, which is formulated as

$$\mathcal{L}_{recon1}^{same} = \mathbb{E}_{x_1 \sim p_{data}(x_1)} [\|x_1 - G(p_1, a_1^s, a_1^{ex})\|_1], \quad (2)$$

where $p_i = E_i^p(x_i)$ and $[a_i^s; a_i^{ex}] = E_i^a(x_i)$.

We also apply the cycle consistency to regularize the ill-posed unsupervised image-to-image translation problem [44]. The cycle reconstruction loss is formulated as follows:

$$\mathcal{L}_{recon1}^{cycle} = \mathbb{E}_{\substack{x_1 \sim p_{data}(x_1) \\ x_2 \sim p_{data}(x_2)}}} [\|x_1 - G(\hat{p}_1, \hat{a}_1^s, \hat{a}_1^{ex})\|_1], \quad (3)$$

where \hat{p}_1 , \hat{a}_1^s , and \hat{a}_1^{ex} denote the reconstructed prototype code, the reconstructed style attribute code, and the reconstructed ID-excluded attribute code, respectively. \hat{p}_1 , \hat{a}_1^s , and \hat{a}_1^{ex} are obtained from $E_2^p(G(p_1, a_1^s, a_2^{ex}))$, $E_2^a(G(p_1, a_1^s, a_2^{ex}))$, and $E_1^a(G(p_2, a_2^s, a_1^{ex}))$ respectively, where $p_i = E_i^p(x_i)$ and $[a_i^s; a_i^{ex}] = E_i^a(x_i)$.

Besides, we apply a code reconstruction loss as follows:

$$\begin{aligned} \mathcal{L}_{recon1}^{code} = & \mathbb{E}_{\substack{x_1 \sim p_{data}(x_1) \\ x_2 \sim p_{data}(x_2)}}} [\|a_1^s - \hat{a}_1^s\|_1] \\ & + \mathbb{E}_{\substack{x_1 \sim p_{data}(x_1) \\ x_2 \sim p_{data}(x_2)}}} [\|a_1^{ex} - \hat{a}_1^{ex}\|_1]. \end{aligned} \quad (4)$$

\hat{a}_1^s and \hat{a}_1^{ex} are obtained from $E_2^a(G(p_1, a_1^s, a_2^{ex}))$ and $E_1^a(G(p_2, a_2^s, a_1^{ex}))$, where $p_i = E_i^p(x_i)$ and $[a_i^s; a_i^{ex}] =$

$E_i^a(\mathbf{x}_i)$. This loss includes both assumptions that ID-discriminative factors should be preserved during the cross-modality reconstruction process as well as ID-excluded factors should be maintained during the same modality reconstruction process. The overall losses for reconstruction are expressed as follows:

$$\mathcal{L}^{recon} = \lambda_1 \mathcal{L}_{recon}^{cross} + \lambda_2 \mathcal{L}_{recon}^{same} + \lambda_3 \mathcal{L}_{recon}^{cycle} + \lambda_4 \mathcal{L}_{recon}^{code}, \quad (5)$$

where $\lambda_t, t \in \{1, 2, 3, 4\}$ controls the relative importance of four losses. $\mathcal{L}_{recon}^{cross}$ indicates the sum of $\mathcal{L}_{recon1}^{cross}$ and $\mathcal{L}_{recon2}^{cross}$. Other losses $\mathcal{L}_{recon}^{same}$, $\mathcal{L}_{recon}^{cycle}$, $\mathcal{L}_{recon}^{code}$ are calculated in the same manner.

KL divergence loss. To help the attribute encoder E_a learn more informative representations, we utilize the Kullback-Leibler (KL) divergence loss. This loss encourages the ID-excluded attribute representation to be as close to a prior Gaussian distribution as follows:

$$\mathcal{L}_1^{kl} = \mathbb{E}_{\mathbf{x}_1 \sim p(\mathbf{x}_1)} [\mathcal{D}_{KL}(\mathbf{a}_1^{ex} \| N(0, 1))], \quad (6)$$

where $\mathcal{D}_{KL}(p||q) = -\int p(z) \log \frac{p(z)}{q(z)} dz$ and $[\mathbf{a}_1^s; \mathbf{a}_1^{ex}] = E_1^a(\mathbf{x}_1)$. By limiting the distribution range of the cross-modality and intra-modality characteristics, this KL divergence loss enables ID-excluded attribute codes to change continuously in the latent space. \mathcal{L}_2^{kl} is defined in a similar manner and $\mathcal{L}^{kl} = \mathcal{L}_1^{kl} + \mathcal{L}_2^{kl}$.

Adversarial loss. Since generating realistic images is crucial for image-to-image translation, we apply an adversarial loss [8] by using the cross-reconstructed images with different modalities. Two discriminators D_1, D_2 corresponding visible and infrared domains are employed for adversarial training. In the case of modality 1, the RGB discriminator D_1 distinguishes the real image \mathbf{x}_1 and the fake image $G(\mathbf{p}_2, \mathbf{a}_2^s, \mathbf{a}_1^{ex})$ used for cross-modality reconstruction. The generator tries to synthesize a more realistic RGB image to fool the discriminator. Accordingly, the ID-excluded attribute code \mathbf{a}_1^{ex} is encouraged to include the modality characteristic of RGB.

Furthermore, we introduce a new strategy to distinguish the cross-modality characteristic (*i.e.* illumination) and the intra-modality characteristic (*i.e.* pose). As mentioned above, the ID-excluded attribute code can be divided into two attribute codes as $\mathbf{a}_1^{ex} = [\mathbf{a}_1^c; \mathbf{a}_1^p]$. Our idea is to swap only the illumination attribute leaving the pose attribute unchanged as $G(\mathbf{p}_2, \mathbf{a}_2^s, \mathbf{a}_1^c, \mathbf{a}_2^p)$. By feeding this image to the RGB discriminator, the modality characteristic is concentrated only on the illumination attribute code \mathbf{a}_1^c . The remaining intra-modality characteristic across RGB-IR images is collected in the pose attribute code \mathbf{a}_2^p . Adversarial losses are employed to play the minimax game, which is formulated as follows:

$$\begin{aligned} \mathcal{L}_1^{adv} = & \mathbb{E}_{\substack{\mathbf{x}_1 \sim p_{data}(\mathbf{x}_1) \\ \mathbf{x}_2 \sim p_{data}(\mathbf{x}_2)}} [\log (1 - D_1(G(\mathbf{p}_2, \mathbf{a}_2^s, \mathbf{a}_1^c, \mathbf{a}_1^p)))] \\ & + \mathbb{E}_{\substack{\mathbf{x}_1 \sim p_{data}(\mathbf{x}_1) \\ \mathbf{x}_2 \sim p_{data}(\mathbf{x}_2)}} [\log (1 - D_1(G(\mathbf{p}_2, \mathbf{a}_2^s, \mathbf{a}_1^c, \mathbf{a}_2^p)))] \quad (7) \\ & + \mathbb{E}_{\mathbf{x}_1 \sim p(\mathbf{x}_1)} [\log D_1(\mathbf{x}_1)], \end{aligned}$$

where $\mathbf{p}_i = E_i^p(\mathbf{x}_i)$ and $[\mathbf{a}_i^s; \mathbf{a}_i^c; \mathbf{a}_i^p] = E_i^a(\mathbf{x}_i)$. The generator is trained to minimize (7) while the discriminator attempts to maximize it. Especially, the parameters of the discriminator are updated when the parameters of the generator are fixed. \mathcal{L}_2^{adv} is defined in a similar way and $\mathcal{L}^{adv} = \mathcal{L}_1^{adv} + \mathcal{L}_2^{adv}$.

3.3. Hierarchical Feature Learning

As illustrated in Fig. 2, our Hierarchical Feature Learning (HFL) module is coupled with ID-PIG by sharing the prototype and attribute encoders. This module enables both encoders to extract the common ID-discriminative factors between RGB-IR images. At the same time, this feature learning process implicitly assists in separating intra- and cross-modality characteristics from cross-modality images and enhances the quality of image generation.

Re-entangling ID-discriminative factors. We introduce the ID-discriminative feature by concatenating the style attribute code and the prototype code to distinguish person identities. Compared to using one of the two codes, the combination of both codes with different characteristics encourages the network to learn rich representations of a person’s identity. Given the prototype tensor \mathbf{p}_i from the prototype encoder E_i^p , the feature embedding network H projects it to the ID-discriminative prototype code \mathbf{p}_i^d , where $\mathbf{p}_i^d = H(\mathbf{p}_i)$. We then concatenate the ID-discriminative prototype code \mathbf{p}_i^d and the style attribute code \mathbf{a}_i^s with a learnable parameter $\alpha \in [0, 1]$, which is expressed as $\mathbf{d}_i^{comb} = [\alpha \cdot \mathbf{p}_i^d; (1 - \alpha) \cdot \mathbf{a}_i^s]$. Then, the combined code \mathbf{d}_i^{comb} is feed into a fully connected layer. In the testing phase, we use the output \mathbf{f} of the fully connected layer for cross-modality retrieval with the Euclidean distance.

Alternate sampling strategy. We form a set of training feature vectors $\mathbf{f} \in \mathcal{F}$ by selecting various types of style attribute codes and prototype codes alternately. This alternate sampling strategy improves the discrimination ability by overcoming the lack of diversity in the training dataset. We alternately combine style attribute codes \mathbf{a}^s and prototype codes \mathbf{p}^d extracted from the original images $\mathbf{x}_1, \mathbf{x}_2$ and the cross-reconstructed images $\mathbf{x}_{1 \rightarrow 2} = G(\mathbf{p}_1, \mathbf{a}_1^s, \mathbf{a}_2^c, \mathbf{a}_2^p)$, $\mathbf{x}_{2 \rightarrow 1} = G(\mathbf{p}_2, \mathbf{a}_2^s, \mathbf{a}_1^c, \mathbf{a}_1^p)$. Note that the attribute and prototype codes for combination must be of the same person.

Cross-entropy loss. Given a set of training feature vectors with the identity annotation $\{\mathbf{f}_i, y_i\}$, we use the cross-entropy loss for ID-discriminative learning, which is expressed as follows:

$$\mathcal{L}^{ce} = \mathbb{E}_{\mathbf{f} \in \mathcal{F}, y \sim \mathcal{Y}} [-\log(p(y|\mathbf{f}))], \quad (8)$$

where $p(y|\mathbf{f})$ indicates the predicted probability of a sampled feature vector \mathbf{f} belonging to the identity y .

Triplet loss. For similarity learning, we also employ the triplet loss. The triplet loss is expressed as follows:

$$\mathcal{L}^{trip} = \sum_{\mathbf{f}^a, \mathbf{f}^p, \mathbf{f}^n \in \mathcal{F}} [d(\mathbf{f}^a, \mathbf{f}^p) - d(\mathbf{f}^a, \mathbf{f}^n) + m]_+, \quad (9)$$

where \mathbf{f}^a , \mathbf{f}^p , and \mathbf{f}^n indicate anchor, positive, and negative samples. $d(\cdot, \cdot)$ is the Euclidean distance, m is a margin parameter, and $[z]_+ = \max(z, 0)$. For each sample \mathbf{f}^a in the set \mathcal{F} , we select the hardest positive sample \mathbf{f}^p and the hardest negative samples \mathbf{f}^n within the batch in the same way as [11]. The triplet loss forces intra-class samples closer and inter-class samples farther. As a result, the cross-entropy and triplet losses help the encoder to clearly disentangle ID-discriminative factors and ID-excluded factors from RGB-IR images.

End-to-end training. As a summary, the overall loss for our Hi-CMD method is expressed as follows:

$$\mathcal{L} = \mathcal{L}^{recon} + \lambda_{kl} \mathcal{L}^{kl} + \lambda_{adv} \mathcal{L}^{adv} + \lambda_{ce} \mathcal{L}^{ce} + \lambda_{trip} \mathcal{L}^{trip}, \quad (10)$$

where λ_{kl} , λ_{adv} , λ_{ce} , and λ_{trip} are hyperparameters to control the relative importance of loss terms. We train the whole network to optimize the total loss in an end-to-end manner. For adversarial learning, we alternatively train the parameters of discriminators and the remaining parameters.

3.4. Discussion

We compare Hi-CMD with the most related disentanglement approach DG-Net [43]. DG-Net is similar to our proposed Hi-CMD in that both methods combine an image generation network with a discriminative learning module in an end-to-end manner. However, the disentangled elements and the feature vectors used for person ReID are entirely different. While DG-Net decomposes each RGB image into appearance and structure codes, our Hi-CMD hierarchically disentangles ID-discriminative factors and ID-excluded factors including pose and illumination attributes from RGB-IR images. In addition, DG-Net only uses an appearance code where color information is vital to distinguish people. However, since this factor is not feasible in the VI-ReID task, we manage the ID-discriminative information between RGB-IR images by hierarchical disentanglement. This hierarchical approach is more useful for extracting the common ID-discriminative feature.

4. Experiments

4.1. Datasets and Settings

Datasets. Extensive experiments were conducted on two widely used VI-ReID datasets, RegDB [21] and SYSU-MM01 [33]. We followed the RegDB evaluation protocol in [37, 39] and the SYSU-MM01 evaluation protocol

Datasets	RegDB [21]			SYSU-MM01 [33]		
	Methods	R=1	R=10	mAP	R=1	R=10
HOG [3]	13.49	33.22	10.31	2.76	18.25	4.24
LOMO [15]	0.85	2.47	2.28	1.75	14.14	3.48
MLBP [16]	2.02	7.33	6.77	2.12	16.23	3.86
GSM [17]	17.28	34.47	15.06	5.29	33.71	8.00
SVDNet [27]	17.24	34.12	19.04	14.64	53.28	15.17
PCB [28]	18.32	36.42	20.13	16.43	54.06	16.26
One stream [33]	13.11	32.98	14.02	12.04	49.68	13.67
Two stream [33]	12.43	30.36	13.42	11.65	47.99	12.85
Zero padding [33]	17.75	34.21	18.90	14.80	54.12	15.95
TONE [37]	16.87	34.03	14.92	12.52	50.72	14.42
TONE+HCML[37]	24.44	47.53	20.80	14.32	53.16	16.16
BCTR [39]	32.67	57.64	30.99	16.12	54.90	19.15
BDTR [39]	33.47	58.42	31.83	17.01	55.43	19.66
eBDTR(alex) [38]	34.62	58.96	33.46	22.42	64.61	24.11
eBDTR(resnet) [38]	31.83	56.12	33.18	27.82	67.34	28.42
cmGAN [2]	-	-	-	26.97	67.51	27.80
D2RL [31]	43.40	66.10	44.10	28.90	70.60	29.20
HSME [9]	41.34	65.21	38.82	18.03	58.31	19.98
D-HSME [9]	50.85	73.36	47.00	20.68	62.74	23.12
Ours (Hi-CMD)	70.93	86.39	66.04	34.94	77.58	35.94

Table 1. Comparison with the state-of-the-arts on RegDB and SYSU-MM01 datasets. Re-identification rates (%) at rank R and mAP (%). 1st and 2nd best results are indicated by **red** and **blue** color, respectively.

in [33]. The RegDB dataset consists of 2,060 visible images and 2,060 far-infrared images with 206 identities for training. The testing set contains 206 identities with 2,060 visible images for the query and 2,060 far-infrared images for the gallery. We repeated 10 trials with a random split to achieve statistically stable results. The SYSU dataset contains 22,258 visible images and 11,909 near-infrared images of 395 identities for training. The testing set includes 96 identities with 3,803 near-infrared images for the query and 301 visible images as the gallery set. The SYSU dataset is collected by six cameras (four visible and two near-infrared), including indoor and outdoor environments. We adopted the most challenging *single-shot all-search* mode and repeated the above evaluation 10 trials with a random split of the gallery and probe set.

Evaluation metrics. Two popular evaluation metrics are adopted: Cumulative Matching Characteristic (CMC) and mean Average Precision (mAP). The rank- k identification rate in the CMC curve indicates the cumulative rate of true matches in the top- k position. The other evaluation metric is the mean average precision (mAP), considering person re-identification as a retrieval task.

Implementation details. Our method is implemented with the Pytorch framework on an NVIDIA Titan Xp GPU. Visible and infrared images are resized to $256 \times 128 \times 3$. Each mini-batch contains 4 pairs of visible and infrared images with different identities. The reconstruction parameters $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ in (5) were set to 50, 50, 50, 10, respectively. The parameters $\lambda_{kl}, \lambda_{adv}, \lambda_{ce}, \lambda_{trip}$ in (10) were set to 1, 20, 1, 1, respectively. We used Stochastic Gradient

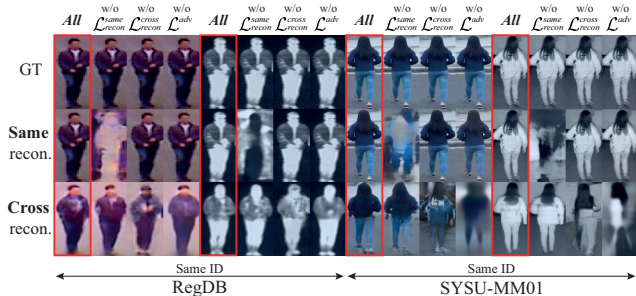


Figure 4. Qualitative comparison between image generation networks for various loss combinations on RegDB and SYSU-MM01. Zoom in for best view.

Descent with the learning rate 0.001 and momentum 0.9 to optimize the HFL module. We adopted the Adam optimizer [14] with the learning rate 0.0001 for the ID-PIG network. The ID-PIG framework is modified based on MUNIT [12] and the feature embedding network H is based on ResNet-50 [10] pretrained on ImageNet [4]. For more details, please refer to the supplementary material.

4.2. Comparison with State-of-the-art Methods

Comparison with conventional methods. We compare our method with conventional methods, which are not designed for VI-ReID. Feature extraction (HOG [3], LOMO [15], MLBP [16]), cross-domain matching (GSM [17]), and RGB-based person ReID (SVDNET [27], PCB [28]) methods are included for comparison. Table 1 shows that all methods have relatively poor performance. Although the PCB method achieves high performance in single-modality person ReID, a significant performance drop is inevitable in the VI-ReID task. Note that a pixel-level difference between visible and infrared images is challenging to deal with at the feature-level representation.

Comparison with state-of-the-arts. We compare our method with the state-of-the-art methods in VI-ReID. The competing methods include feature learning frameworks (one-stream, two-stream, zero-padding [33], and TONE [37]), ranking losses (BCTR [39], BDTR [39], eBDTR [38]), metric learning (HCML [37], HSME [9], D-HSME [9]), reducing distribution divergence (cmGAN [2]), and image generation (D2RL [31]) methods. Our model achieves 70.93% rank-1 identification rate and 66.04% mAP score on the RegDB dataset [21], and 34.94% rank-1 identification rate and 35.94% mAP score on the SYSU-MM01 dataset [33]. Our method significantly outperforms the state-of-the-art VI-ReID methods on both the RegDB and SYSU-MM01 datasets. This comparison indicates the effectiveness of our disentanglement approach for bridging the cross-modality and intra-modality gaps. Moreover, this improvement in performance can be analyzed by the visualization of ID-discriminative factors, which is discussed in Section 4.4.

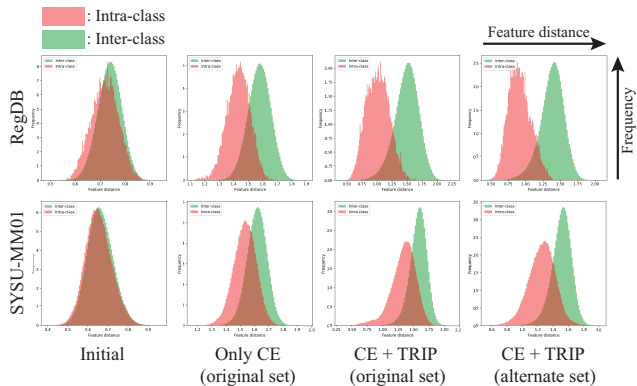


Figure 5. The distribution of the Euclidean distance between cross-modality (RGB-IR) features. The intra-class and inter-class distances are indicated by red and green color, respectively.

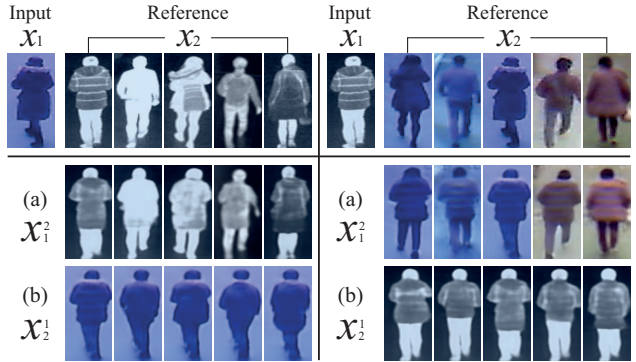
Methods			RegDB		SYSU-MM01	
Input set	Loss	Feature	R=1	mAP	R=1	mAP
Original	CE	A+P	36.36	33.47	18.65	19.49
Original	CE+TRIP	A	15.33	15.37	6.05	7.74
Original	CE+TRIP	P	49.02	45.75	22.51	23.73
Original	CE+TRIP	A+P	53.25	49.53	29.19	30.53
Alternate	CE+TRIP	A+P	70.93	66.04	34.94	35.94

Table 2. Component analysis of our HFL module on RegDB and SYSU-MM01 datasets. A and P represent the style attribute code and the prototype code respectively, which are used to train HFL.

4.3. Further Evaluations and Analysis

Impact of image generation losses. We performed an ablation study of our ID-preserving Person Image Generation (ID-PIG) network. To evaluate image generation losses qualitatively, we compare four variations of ID-PIG: 1) our best model with all components; 2) removing the reconstruction loss $\mathcal{L}_{recon}^{same}$; 3) removing the disentanglement loss $\mathcal{L}_{recon}^{cross}$; 4) removing the adversarial loss \mathcal{L}^{adv} . The network structure and training strategy remain the same for all settings. Figure 4 shows the results of this image translation experiment. The samples are randomly selected from the testing set. We observe that the generated images can contain unpleasant artifacts such as blurriness or color shifts if one kind of loss is excluded from the training process. On the other hand, the results generated by our ID-PIG network with all components show more a realistic and sharp appearance regardless of the modalities are changed or not.

Effectiveness of hierarchical feature learning. We study the several variants of the Hierarchical Feature Learning (HFL) module on both datasets to demonstrate the effectiveness of our hierarchical disentanglement approach. Figure 5 shows the distribution of the Euclidean distances between RGB-IR images from the testing set. Compared to the initial state, the HFL module based on an alternate sampling strategy minimizes the intra-class distance and maximizes the inter-class distance simultaneously. Moreover, as shown in Table 2, the use of the alternate sampling strategy



$$\begin{aligned} \mathcal{X}_1^2 &= G(p_1, a_1^s, a_2^e, a_2^p) \rightarrow \text{ID-discriminative} [\mathcal{X}_1] + \text{ID-excluded} [\mathcal{X}_2] \\ \mathcal{X}_2^1 &= G(p_2, a_2^s, a_1^e, a_1^p) \rightarrow \text{ID-discriminative} [\mathcal{X}_2] + \text{ID-excluded} [\mathcal{X}_1] \end{aligned}$$

Figure 6. Examples of person image generation across two different swapping manners: (a) swapping ID-excluded factors; (b) swapping ID-discriminative factors.

significantly improves performance compared to learning with the original image set. Our HFL method also achieves much higher performance than learning the style attribute code or the prototype code alone. These results prove that diverse code combinations by an alternate sampling strategy improve the discrimination ability and reduce the cross-modality gap significantly.

4.4. Person Attribute Exploration

Disentangling ID-discriminative and ID-excluded factors. In this section, we present image generation results by our ID-PIG network. To demonstrate that ID-discriminative factors and ID-excluded factors are clearly disentangled from RGB-IR images, we conducted two experiments: 1) swapping the ID-excluded factors of x_1 and x_2 with preserving ID-discriminative factors; 2) swapping the ID-discriminative factors of x_1 and x_2 with maintaining ID-excluded factors. The images in Fig. 6 (a) are changed to the pose and illumination attributes of the reference images, while the clothes and patterns of the input image are retained. The images in Fig. 6 (b) are synthesized with the clothes of the reference images, while maintaining the pose and illumination attributes of the input image. The visualization of ID-PIG proves that the shapes, patterns, and styles of clothes are important factors for cross-modality matching, which means that our novel hierarchical disentanglement approach can extract the common ID-discriminative feature effectively.

Disentangling illumination and pose attributes. To demonstrate that our proposed system can manipulate pose and illumination attributes in ID-excluded factors independently, we conducted an experiment that changes illumination attribute codes while maintaining pose attribute codes, as shown in Fig. 7. Note that our adversarial loss enables the network to distinguish pose and illumination attributes




Figure 7. Examples of ID-excluded factor manipulation.




Figure 8. Examples of person image generation by linear interpolation of ID-excluded factors.

in (7). This experiment shows that the pose can be transformed into other poses without any supervision of human pose estimation, unlike conventional pose generation methods [18, 22].

Interpolation of ID-excluded attributes. In this part, we validate that ID-excluded attributes of the generated images continuously change in the latent space. We extracted the ID-excluded attribute codes from a pair of images and obtained a series of attribute codes by linear interpolation. As shown in Fig. 8, the pose and illumination are changed smoothly from left to right. This experiment ensures that our ID-PIG can synthesize a lot of unseen poses or illuminations.

5. Conclusion

In this work, we have proposed a novel hierarchical cross-modality disentanglement method for VI-ReID. In contrast to previous studies, the proposed model reduces both cross- and intra-modality discrepancies at the same time by disentangling ID-discriminative factors and ID-excluded factors from visible-infrared images. Furthermore, our novel ID-preserving person image generation network provides resources to analyze cross-modality matching results and alleviates the problem of the insufficient data. Quantitative and qualitative evaluation on VI-ReID datasets demonstrates the superiority of our proposed method over the state-of-the-art methods.

References

- [1] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. Towards open-set identity preserving face synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6713–6722, 2018. 3
- [2] Pingyang Dai, Rongrong Ji, Haibin Wang, Qiong Wu, and Yuyu Huang. Cross-modality person re-identification with generative adversarial training. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 677–683, 2018. 2, 6, 7
- [3] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. 2005. 6, 7
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009. 7
- [5] Weijian Deng, Liang Zheng, Qixiang Ye, Guoliang Kang, Yi Yang, and Jianbin Jiao. Image-image domain adaptation with preserved self-similarity and domain-dissimilarity for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 994–1003, 2018. 2
- [6] Yixiao Ge, Zhuowan Li, Haiyu Zhao, Guojun Yin, Shuai Yi, Xiaogang Wang, et al. Fd-gan: Pose-guided feature distilling gan for robust person re-identification. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 1222–1233, 2018. 2
- [7] Abel Gonzalez-Garcia, Joost van de Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 1287–1298, 2018. 3
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680, 2014. 2, 5
- [9] 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 (AAAI)*, volume 33, pages 8385–8392, 2019. 2, 6, 7
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 7
- [11] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 6
- [12] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018. 3, 7
- [13] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, 2019. 3
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [15] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2197–2206, 2015. 6, 7
- [16] Shengcai Liao and Stan Z Li. Efficient psd constrained asymmetric metric learning for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3685–3693, 2015. 6, 7
- [17] Liang Lin, Guangrun Wang, Wangmeng Zuo, Xiangchu Feng, and Lei Zhang. Cross-domain visual matching via generalized similarity measure and feature learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(6):1089–1102, 2016. 6, 7
- [18] Jinxian Liu, Bingbing Ni, Yichao Yan, Peng Zhou, Shuo Cheng, and Jianguo Hu. Pose transferrable person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4099–4108, 2018. 2, 8
- [19] Jiawei Liu, Zheng-Jun Zha, Di Chen, Richang Hong, and Meng Wang. Adaptive transfer network for cross-domain person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7202–7211, 2019. 2
- [20] Liqian Ma, Qianru Sun, Stamatis Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 99–108, 2018. 3
- [21] Dat Nguyen, Hyung Hong, Ki Kim, and Kang Park. Person recognition system based on a combination of body images from visible light and thermal cameras. *Sensors*, 17(3):605, 2017. 6, 7
- [22] Xuelin Qian, Yanwei Fu, Tao Xiang, Wenxuan Wang, Jie Qiu, Yang Wu, Yu-Gang Jiang, and Xiangyang Xue. Pose-normalized image generation for person re-identification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 650–667, 2018. 2, 8
- [23] M Saquib Sarfraz and Rainer Stiefelwagen. Deep perceptual mapping for cross-modal face recognition. *International Journal of Computer Vision (IJCV)*, 122(3):426–438, 2017. 2
- [24] Krishna Kumar Singh, Utkarsh Ojha, and Yong Jae Lee. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6490–6499, 2019. 3
- [25] Jifei Song, Yongxin Yang, Yi-Zhe Song, Tao Xiang, and Timothy M Hospedales. Generalizable person re-identification by domain-invariant mapping network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 719–728, 2019. 1
- [26] Sijie Song, Wei Zhang, Jiaying Liu, and Tao Mei. Unsupervised person image generation with semantic parsing trans-

- formation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2357–2366, 2019. 1, 2
- [27] Yifan Sun, Liang Zheng, Weijian Deng, and Shengjin Wang. Svdnet for pedestrian retrieval. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3800–3808, 2017. 2, 6, 7
- [28] 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. 2, 6, 7
- [29] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1415–1424, 2017. 3
- [30] Hao Wang, Dihong Gong, Zhifeng Li, and Wei Liu. Decorrelated adversarial learning for age-invariant face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3527–3536, 2019. 3
- [31] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 618–626, 2019. 2, 6, 7
- [32] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 79–88, 2018. 2
- [33] 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 International Conference on Computer Vision (ICCV)*, pages 5380–5389, 2017. 1, 2, 6, 7
- [34] Xuezhi Xiang, Ning Lv, Zeting Yu, Mingliang Zhai, and Abdulmotaleb El Saddik. Cross-modality person re-identification based on dual-path multi-branch network. *IEEE Sensors Journal*, 2019. 1
- [35] Qize Yang, Hong-Xing Yu, Ancong Wu, and Wei-Shi Zheng. Patch-based discriminative feature learning for unsupervised person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3633–3642, 2019. 1
- [36] 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1389–1398, 2019. 1
- [37] Mang Ye, Xiangyuan Lan, Jiawei Li, and Pong C Yuen. Hierarchical discriminative learning for visible thermal person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018. 2, 6, 7
- [38] 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*, 2019. 2, 6, 7
- [39] Mang Ye, Zheng Wang, Xiangyuan Lan, and Pong C Yuen. Visible thermal person re-identification via dual-constrained top-ranking. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1092–1099, 2018. 2, 6, 7
- [40] Zelong Zeng, Zhixiang Wang, Zheng Wang, Yung-Yu Chuang, and Shin’ichi Satoh. Illumination-adaptive person re-identification. *arXiv preprint arXiv:1905.04525*, 2019. 3
- [41] Ziyuan Zhang, Luan Tran, Xi Yin, Yousef Atoum, Xiaoming Liu, Jian Wan, and Nanxin Wang. Gait recognition via disentangled representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4710–4719, 2019. 3
- [42] Meng Zheng, Srikrishna Karanam, Ziyuan Wu, and Richard J Radke. Re-identification with consistent attentive siamese networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5735–5744, 2019. 1
- [43] Zhedong Zheng, Xiaodong Yang, Zhiding Yu, Liang Zheng, Yi Yang, and Jan Kautz. Joint discriminative and generative learning for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2138–2147, 2019. 4, 6
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2223–2232, 2017. 4
- [45] Zhen Zhu, Tengting Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2347–2356, 2019. 1, 2