

On Exploring Pose Estimation as an Auxiliary Learning Task for Visible-Infrared Person Re-identification

Yunqi Miao^a, Nianchang Huang^b, Xiao Ma^c, Qiang Zhang^b and Jungong Han^{d,*}

^aThe Warwick Manufacturing Group (WMG), University of Warwick, Coventry, CV4 7AL, United Kingdom

^bThe School of Mechano-Electronic Engineering, Xidian University, Xi'an, 710071, China

^cNottingham Business School, Nottingham Trent University, Nottingham, NG1 4FQ, United Kingdom

^dThe Department of Computer Science, University of Sheffield, Sheffield, S1 4DP, United Kingdom

ARTICLE INFO

Keywords:

person re-identification, pose estimation, visible-infrared, auxiliary learning task

ABSTRACT

Visible-infrared person re-identification (VI-ReID) has been challenging due to the existence of large discrepancies between visible and infrared modalities. Most pioneering approaches reduce intra-modality variations and inter-modality discrepancies by learning modality-shared features. However, an explicit modality-shared cue, i.e., body keypoints, has not been fully exploited in VI-ReID. Additionally, existing feature learning paradigms imposed constraints on either global features or partitioned feature stripes, which neglect the prediction consistency of global and part features. To address the above problems, we exploit Pose Estimation as an auxiliary learning task to assist VI-ReID in an end-to-end framework. By jointly training these two tasks in a mutually beneficial manner, our model learns higher quality ID-related features. On top of it, the learnings of global features and local features are seamlessly synchronized by Hierarchical Feature Constraint (HFC), where the former supervises the latter using the knowledge distillation strategy. Experimental results on two benchmark VI-ReID datasets show that the proposed method consistently improves state-of-the-art methods by significant margins. Specifically, our method achieves nearly 20% mAP improvements against the state-of-the-art method on the RegDB dataset. Our intriguing findings highlight the usage of auxiliary task learning in VI-ReID. Our source code is available at https://github.com/yoqim/Pose_VIReID.

1. Introduction

Person re-identification (ReID) aims at retrieving the same identity across multiple disjoint cameras, which has gained much attention from the recent computer vision community [16, 45]. Most existing ReID methods focus on the matching between visible images, which are generally collected under good illumination conditions [6, 1, 9, 53, 25]. However, those systems seem impractical because visible images cannot provide sufficient discriminatory information in poor lighting environments, *e.g.*, at night. To this end, Visible-Infrared person Re-identification (VI-ReID) emerges as an alternative to performing the retrieval between visible (RGB) images and infrared (IR) counterparts, thus enabling the day-to-night person re-identification.

However, VI-ReID is a challenging problem due to large intra-class variations and modality discrepancies across different cameras. The former refers to identity's appearance differences within a modality caused by poses, clothes, viewpoints, *etc.* While, the latter denotes intrinsic differences between visible and infrared images caused by the spectrum of cameras. To reduce both discrepancies, the central research question in this field has always been seeking better ways to extract discriminative features for identity retrieval, which are ID-related and modality-invariant [40, 45, 55, 44, 46, 32].

Despite the vigorous development of this field, we realize that most algorithms extract identity features in a heuristics manner with NO common view of what sort of features are specifically helpful for VI-ReID. To tackle this problem, we visualize the features extracted by several representative methods in Fig. 1, aiming to investigate how visual feature extractions evolved over the years to improve the performance of VI-ReID systems. Concretely, ZeroPad [40], as the first work in VI-ReID, extracts features from random regions in an image. Additionally, for a specific identity, the features derived from two modalities share NO commonality. As a result, the mean Average Precision (mAP) score is far from satisfactory. Later, TSLFN [55] horizontally partitions the backbone feature maps (global) into several stripes (local) and employs local-level constraints on each of them. Clearly, features derived by TSLFN cover more parts of the human body, compared to the global-level constraint-based method, *e.g.*, ZeroPad. That might be the reason why its performance is significantly superior to that of ZeroPad. In the meantime, a baseline for VI-ReID (AGW [45]) inserts non-local attention blocks during the feature extraction, which enforces the features to be extracted from identity's body instead of backgrounds. It reveals that *attention-aware* features help to increase performance. Recently, a state-of-the-art approach (DDAG [44]) integrates both local-level and global-level constraints into an end-to-end framework. Compared to previous works, the features derived by DDAG are more fine-grained, which are not only shared by two modalities but distinguishable for different identities. Based on the above observations, a conclusion can be drawn: *as more features from the human region and attentive body*

*Corresponding author



Yunqi.Miao.1@warwick.ac.uk (Yunqi Miao);

nchuang@stu.xidian.edu.cn (Nianchang Huang); xiao.ma@ntu.ac.uk (Xiao Ma); qzhang@xidian.edu.cn (Qiang Zhang); jungonghan77@gmail.com (Jungong Han)

ORCID(s):




Figure 1: Visualization of features derived by ZeroPad [40], TSLFN [55], AGW [45], DDAG [44], and our method SYSU-MM01 [40]. For each method, the RGB image of an identity is shown on the left and the IR image is on the right. The publication year and mAP score (%) at *All-search* setting of each method are also reported. For a fair comparison, the mAP scores of TSLFN, AGW and our method are given by models trained merely with vanilla identity loss and triplet loss.

parts are extracted, the identity retrieval performance improves consistently. Benefiting from this conclusion, we aim to extract more features from body skeleton joints, which are not only ID-related but immune to modality changes. Therefore, we achieve over 10% mAP improvements against DDAG on the challenging dataset - SYSU-MM01 [40] at *All-search* setting.

Having depicted the visual features that are conducive to VI-ReID, the next question is: how can we extract them effectively? It is noted that body skeleton points are explicit modality-shared cues and the features describing certain skeleton points are ID-related. In light of this, in the paper, we aim to facilitate the extraction of discriminative features for identity retrieval with pose estimation. However, making effective use of the pose information for cross-modality ReID does not seem easy, though it has recently appeared to be exploited in some single-modality ReID works [48, 34, 52, 49, 35, 26], where visible images are exclusively involved. Earlier methods [48, 34, 52] utilize detected body keypoint to segment [48] or align [34, 52] body regions in order to cope with human pose changes. After the calibration of body parts, different body parts need to be stitched, which usually yields visually unrealistic transformed images. If we move to the cross-modality setting, the transformed errors would be further magnified due to the huge discrepancy between the two modalities. Alternatively, another group of methods employ the body keypoints information to refine ID-related feature maps either by means of highlighting discriminative body regions [49, 26] or complementing human appearance features [35]. Although pose-assisted features are proved to improve the feature discriminability under the single-modality setting, they are rigidly based on the outputs of off-the-shelf pose estimators. However, such a blind trust in the pre-trained pose estimators will lead to a poor re-identification performance if the gap between the source domain and the target domain of the pose estimator is huge. Therefore, employing pose information in VI-ReID is extremely challenging due to the massive gap between visible images (source domain) and infrared images (target domain). To highlight the problem, we extend a state-of-the-art pose-guided single-modality person ReID method [35] to VI-ReID and the results turn out that the best performance in

the mAP score only achieves 42.19%, which is far from satisfaction due to the inadequate usage of the pose information (more results and comparisons are provided in **Experiments** section).

To solve the above problems, we propose a two-stream VI-ReID framework, where modality-shared features for identity retrieval are extracted by means of learning an auxiliary task (pose estimation) and the main task (person ReID) jointly. Unlike previous works, which rely dramatically on the off-the-shelf pose estimators, body keypoint features are adaptively adjusted to facilitate person ReID in our work. Specifically, apart from ID-related constraints, an extra constraint is imposed on the pose estimation branch, ensuring that not only body skeleton points are precisely estimated but also the ID-related information are fully embedded in feature maps, *i.e.*, at both local and global levels. Additionally, despite significant improvements obtained by the horizontal-divided feature constraints [36] in VI-ReID [55, 44], the learning of individual striped features is generally independent, which means its discriminability consistency with global backbone features is neglected. To this end, we propose a Hierarchical Feature Constraint (HFC), which bonds the learnings of global features and local ones via the knowledge distillation strategy. Concretely, predictions of backbone features serve as “soft-target” to supervise the learning of partitioned feature stripes, hence preserving the discriminability consistency of global features and local ones.

In summary, the contributions made in our work are mainly three-fold:

- A novel two-stream framework for VI-ReID is proposed, where the pose estimation acts as an auxiliary learning task to facilitate VI-ReID. To learn body keypoint information enhanced ID-related features, both pose estimation and ReID constraints are imposed on the pose estimation branch.
- Instead of imposing feature constraints on local feature stripes only, Hierarchical Feature Constraint (HFC) is proposed to ensure the discriminability consistency of global features and local ones via the knowledge distillation strategy.

- The proposed method performs far better than the state-of-the-art methods on two benchmark datasets: SYSU-MM01 [40] and RegDB [27].

2. Related work

2.1. Auxiliary tasks in person ReID

Research has shown that semantic information such as body parts, human pose *etc.*, can significantly facilitate the person ReID task [16, 45]. Therefore, recent works utilize auxiliary tasks to improve the performance of ReID models, which can be mainly categorized as attribute-guided methods [20, 37], segmentation-guided methods [15, 33, 54], and pose-guided methods [5, 22, 30, 48, 34, 52, 49, 35].

Attribute information provides complementary details of identities for person ReID. Ling *et al.* [20] propose a multi-task learning framework, which attempts to improve the discriminability of identity features by embedding attribute information. Tay *et al.* [37] enhance the identity features with attribute attention maps, where class-sensitive activation regions of various attributes, *i.e.*, clothing color, hair, gender *etc.* are emphasized.

Segmentation-guided methods are based on pixel-level body parts segmentation, which improve the discriminability of identity features by masking out backgrounds [33] or leverage fine-grained local features from discriminative body regions [15, 54]. However, both attribute- and segmentation-assisted methods require extensive additional annotations, which are too expensive to obtain in real-world applications. Additionally, since color information is not available in infrared images, the usage of color-related attributes is restricted in VI-ReID, thereby limiting performance improvement.

The pose estimation is employed to facilitate person ReID from mainly two aspects: 1) generating person images with various poses to augment training data [5, 22, 30], and 2) aligning body parts [48, 34, 52, 49, 35, 26]. With the help of Generative Adversarial Networks (GAN), Qian *et al.* [30] synthesize eight new images for an identity with a set of target canonical poses. In addition to generating visually preferable samples, Liu *et al.* [22] propose a guider module to ensure that the generated samples have discriminative power for ReID. Apart from image generation, pose information is employed to handle the problem of feature misalignment in ReID [48, 34, 52]. Zhao *et al.* [48] segment the human body into seven regions by human landmark information and combine the representations over them for identity retrieval. Zheng *et al.* [52] align pedestrians to a standard pose by stitching the segmented body regions with affine transformations. However, such an unnatural stitch destroys the authenticity of human and requires an elaborately designed following fusion method to fuse the local features. Instead of using detected body regions rigidly, Zhao *et al.* [49] and Miao *et al.* [26] employ body keypoint maps to refine image feature maps, where discriminative body parts for person ReID are emphasised. However, both approaches are rigidly based on the outputs of off-the-shelf pose estimators,

which may generate unreliable information because of the gap between the source domain and the target domain. To avoid the blind trust in the pre-trained pose estimator, Suh *et al.* [35] take an on-the-fly pose estimator as an individual branch to derive pose features, which are then aggregated with the appearance features from the other branch via a bilinear pooling layer. Such combined training encourages the pose branch to learn features that are beneficial to the person ReID task. While pose features are learned to adapt to the ReID task in the work, only ID-related constraints are imposed on the fused features, which ignores the quality of pose features from the individual branch. Given the above concerns, the auxiliary pose estimation task designed for single-modality ReID cannot perform well in VI-ReID. In other words, how to enable the pose information to facilitate VI-ReID has not been elaborately investigated.

Therefore, in the paper, we propose a novel pose estimation assisted framework for VI-ReID, where body skeleton points cues are learned under both pose and identity guidance, which are then employed to enhance the discriminative identity feature extraction for the cross-modality person re-identification.

2.2. Feature constraints in VI-ReID

Visible-Infrared Person ReID (VI-ReID) aims to match person images cross the visible modality and the infrared modality, which faces the challenges of both intra-modality discrepancy and inter-modality discrepancy. As a pioneering VI-ReID work, Zero-Pad [40] proposes to use the zero-padding input strategy to evolve domain-specific structures automatically in a one-stream network. To improve the discriminability of identity features, most existing works impose image-level constraints [38, 39, 17, 13], feature-level constraints [42, 43, 10, 24, 45, 55, 21, 14] or both [44, 47] during training.

The image-level constraints aim to alleviate the inter-modality discrepancy by generating cross-modality fake images via generative models, such as GAN [8]. For example, AlignGAN [38] employs a pixel alignment module to generate fake infrared images from real RGB images. D²RL [39] adopts a sub-network to translate the infrared images to their visible counterparts and vice versa. Instead of generating between the visible and the infrared modality, X-Modal [17] introduces an intermediate modality to bridge the gap between the visible modality and the infrared modality.

In terms of feature-level constraints, existing works apply either the global-level feature constraints on backbone convolutional features or the local-level feature constraints on partitioned feature stripes. As a representative of global-level constraints, HCML [42] proposes a two-stream network, which jointly optimizes modality-specific and modality-shared metrics. Based on the idea, a bi-directional top-ranking loss is then introduced in eDBTR [43] to incorporate the above two constraints. Similarly, HSME [10] uses the sphere softmax function to build a hypersphere embedding space where the intra-modality variations and

the cross-modality variations are constrained. With a shared-specific feature transfer, cm-SSFT [24] enables the flow between modality-shared and modality-specific features, which serve as compensation to each other. Alternatively, AGW [45] presents a weighted regularized triplet loss to embed the neighboring relationship of images from two modalities in a common feature space. Inspired by the competitive performance of the Part-based Convolutional Baseline (PCB) model [36] in single-modality ReID, recent VI-ReID studies impose local-level feature constraints on feature stripes obtained by partitioning backbone convolutional features. a kickoff work TSLFN [55] presents a center-based loss that pulls the centers of RGB features and IR features of a given identity closer. On top of it, Liu *et al.* [21] present a hetero-center triplet loss, where the feature centers of different identities are pushed away.

Recently, DDAG [44] and GLMC [47] propose to impose the image-level and feature-level constraints simultaneously and achieve state-of-the-art performance. Specifically, DDAG mines cross-modality global-level and intra-modality part-level contextual cues. GLMC learns multi-granularity representations based on global and local features to leverage both coarse and fine-grained information. Apart from feature-level and image-level constraints, LbA [29] exploits pixel-wise features by establishing dense correspondences between RGB images and IR images and enforcing semantically similar local regions to be closely embedded. Although both global-level and local-level constraints are considered, those constraints are imposed independently without taking the discriminability consistency of them into account.

To this end, we propose a Hierarchical Feature Constraint (HFC) to bond the global feature learning with the local ones, where predictions of global features supervise the learning of part features via the knowledge distillation strategy.

3. Proposed method

We propose a novel pose estimation assisted framework for VI-ReID, which aims to learn modality-shared and ID-related features for identity retrieval. The framework of the proposed method is shown in Fig. 2. As can be seen, our model mainly consists of four components: Modality-specific module, Modality-shared module, Pose Estimation branch, and ReID branch. Details about these components will be discussed in the following subsections.

3.1. Modality-specific module and Modality-shared module

Following previous works [43, 21, 55], ResNet50 [11] is exploited as a backbone feature extractor to provide discriminative features for both pose estimation and ReID tasks. Specifically, Modality-specific module consists of two blocks (“*Conv Block1~2*”), which adopt the structures of shallow convolution block (*layer0*) and the first res-convolution block (*layer1*) of ResNet50, respectively. Note

that the parameters of Modality-specific module for RGB and IR modalities are separately updated. Then, modality-specific features of two modalities are projected into a shared feature space by the Modality-shared module (“*Conv Block3*”), which adopts the structure of the second res-convolution block (*layer2*) of ResNet50.

Mathematically, given a RGB image I_{RGB} and an IR image I_{IR} , modality-specific features $\mathbf{F}_m$ and modality-shared features $\mathbf{F}_S$ can be obtained by,

$$\begin{aligned}\mathbf{F}_m &= \text{ConvB2}(\text{ConvB1}(I_m, \theta_m^1), \theta_m^2), \quad m \in \{RGB, IR\}, \\ \mathbf{F}_S &= \text{ConvB3}([\mathbf{F}_{RGB}, \mathbf{F}_{IR}]_b, \theta^3),\end{aligned}\quad (1)$$

where $[\cdot, \cdot]_b$ represents the feature concatenation along the data dimension. $\text{ConvB1}(*, \theta_m^1)$, $\text{ConvB2}(*, \theta_m^2)$, and $\text{ConvB3}(*, \theta^3)$ denote the convolution blocks of Modality-specific module and Modality-shared module with corresponding parameters θ_m^1 , θ_m^2 and θ^3 , respectively.

3.2. Pose Estimation branch

3.2.1. Body keypoint features extraction

Given the observation that modality-shared features are beneficial to VI-ReID, a Pose Estimation branch is integrated as an auxiliary to extract modality-shared features. The structure of our Pose Estimation branch is shown in Fig. 2. Specifically, a convolutional layer (“*Conv1*”) and a deconvolutional layer (“*DConv2*”) are employed to extract high-level features, *i.e.*, $\mathbf{F}_{p_1}$ and $\mathbf{F}_{p_2}$, and restore the resolution of feature maps to that of ground-truth body keypoint heatmaps, which are denoted as follows,

$$\mathbf{F}_{p_1} = \text{Conv1}(\mathbf{F}_S, \theta_p^1), \quad \mathbf{F}_{p_2} = \text{DConv2}(\mathbf{F}_{p_1}, \theta_p^2), \quad (2)$$

where $\text{Conv1}(*, \theta_p^1)$ and $\text{DConv2}(*, \theta_p^2)$ represents a 3×3 convolutional layer with parameters θ_p^1 and θ_p^2 , respectively. Note that both layers are followed by a ReLU activation function, which are omitted in equations for simplicity.

Subsequently, a Refinement Module [28] is used to extract refined body keypoint features and predict body keypoint heatmaps. Specifically, the Refinement Module consists of a U-Shaped Block, three Refine Blocks and two convolutional layers. The U-Shaped Block and Refine Blocks are employed to extract refined features. On top of it, the convolutional layers are applied for the heatmap estimation. Mathematically, given high-level features $\mathbf{F}_{p_2}$, refined keypoint features $\mathbf{F}_R$ and predicted heatmaps $\hat{\mathbf{H}}$ can be respectively obtained by,

$$\mathbf{F}_R = \text{RM}_F(\mathbf{F}_{p_2}, \theta_{RM}^F), \quad (3)$$

$$\hat{\mathbf{H}} = \text{RM}_H(\mathbf{F}_R, \theta_{RM}^H), \quad (4)$$

where $\text{RM}_F(*, \theta_{RM}^F)$ and $\text{RM}_H(*, \theta_{RM}^H)$ denote the refined feature extraction stage with parameters θ_{RM}^F , and the heatmap estimation stage with parameters θ_{RM}^H in the Refinement Module, respectively.




Figure 2: The framework of the proposed method.

3.2.2. Body keypoint features transferring

To exploit body keypoint features derived by Pose Estimation branch in the ReID branch, a convolutional layer (“Conv3”) and a convolutional block (“Conv Block6”) are employed to deal with the mismatch in terms of the resolution and the channel number of feature maps.

Specifically, the refined keypoint features $\mathbf{F}_R$ are firstly downsampled by “Conv3” so that the derived $\mathbf{F}'_R$ has the identical resolution as $\mathbf{F}_{p1}$, which is denoted as,

$$\mathbf{F}'_R = \text{Conv3}(\mathbf{F}_R, \theta^3_p), \quad (5)$$

where $\text{Conv3}(*, \theta^3_p)$ represents the 3×3 convolutional layer with a stride of 2 and parameters θ^3_p . On top of that, “Conv Block6” aligns the channel number of feature maps from the Pose Estimation branch with that from the ReID branch. Therefore, the final body keypoint features $\mathbf{F}_P$ can be derived as follows,

$$\mathbf{F}_P = \text{ConvB6}(\mathbf{F}'_R + \mathbf{F}_{p1}, \theta^4_p), \quad (6)$$

where $\text{ConvB6}(*, \theta^4_p)$ denotes the convolutional block with the parameters θ^4_p , consisting of a 3×3 convolutional layer with a stride of 2 and a 1×1 convolutional layer. Note that all convolutional layers are followed by a ReLU activation function.

3.2.3. Body keypoint features integration

To highlight the body keypoint regions in the features output by the ReID branch, the final keypoint features $\mathbf{F}_P$ are employed to generate the body keypoint masks $\mathbf{M}$, i.e.,

$$\mathbf{M} = \text{sigmoid}(\mathbf{F}_P). \quad (7)$$

The values of the masks are regularized by the sigmoid function to $[0, 1]$, which serve as soft attention maps to refine the identity features from the ReID branch.

3.3. ReID branch

3.3.1. Global-level feature extraction

Apart from the Pose Estimation branch that aids modality-shared features extraction, a ReID branch is employed to extract ID-related features. As can be seen from Fig. 2, the ReID branch mainly consists of 2 convolutional blocks (“Conv Block4~5”). Following previous works [43, 45, 55], “Conv Block4”, “Conv Block5” follow the structures of the third and fourth res-convolution block (*layer3*, *layer4*) of ResNet50 [11], respectively. Mathematically, identity features $\mathbf{F}_{id_2}$ are extracted by,

$$\mathbf{F}_{id_2} = \text{ConvB5}(\text{ConvB4}(I_m, \theta^1_{ID}), \theta^2_{ID}), m \in \{RGB, IR\}, \quad (8)$$

where $\text{ConvB4}(*, \theta^1_{ID})$ and $\text{ConvB5}(*, \theta^2_{ID})$ denote a convolutional block with the parameters θ^1_{ID} and θ^2_{ID} , respectively. Then, identity features $\mathbf{F}_{id_2}$ are refined by the body keypoint masks $\mathbf{M}$ derived by Pose Estimation branch by performing the element-wise product operation $\odot$ to obtain the final identity features $\mathbf{F}_{ID}$, i.e.,

$$\mathbf{F}_{ID} = \mathbf{F}_{id_2} \odot \mathbf{M}. \quad (9)$$

3.3.2. Local-level feature partition

Since part features can offer fine-grained information for identity identification, PCB models [36] are exploited in the proposed framework for local feature learning. Following [44, 55], convolutional features $\mathbf{F}_{ID}$ from ReID branch are firstly partitioned into P horizontal stripes and then transferred to feature vectors via Global Average Pooling (GAP) before being sent to the corresponding PCB model, which can be formulated as follow,

$$\mathbf{F}_{ID_1}, \dots, \mathbf{F}_{ID_P} = \text{GAP}(\text{Part}(\mathbf{F}_{ID})), \quad (10)$$

where $\text{Part}(\cdot)$, $\text{GAP}(\cdot)$ denote the horizontal partition and GAP, respectively.




Figure 3: Illustration of Hierarchical Feature Constraint (HFC). GAP represents Global Average Pooling.

As can be seen from Fig. 2, a PCB model consists of a Fully-Connected (FC) layer and a classifier. The former reduces the dimensions of feature vectors from 2048-dim to 512-dim, and the latter is employed for identity prediction. For the i -th PCB model, the fine-grained part features f_{ID_i} are obtained by,

$$f_{ID_i} = FC_{ID_i}(\mathbf{F}_{ID_i}, \theta_{f_{cid}}^i), \quad i = \{1, \dots, P\}, \quad (11)$$

where $FC_{ID_i}(*, \theta_{f_{cid}}^i)$ denotes the FC layer with the parameters $\theta_{f_{cid}}^i$. P is the number of PCB models, which is empirically set as 6 in the paper.

The local feature learning is also performed on convolutional features $\mathbf{F}_P$ from Pose Estimation branch to obtain the corresponding fine-grained part features f_{P_i} . During inference, fine-grained part features are concatenated for identity retrieval, i.e.,

$$f_{ID} = [f_{ID_1}, \dots, f_{ID_P}]_c, \quad f_P = [f_{P_1}, \dots, f_{P_P}]_c, \quad (12)$$

$$f_{ALL} = [f_{ID}, f_P]_c, \quad (13)$$

where $[\cdot, \dots, \cdot]_c$ represents the feature concatenation along the channel dimension.

3.4. Hierarchical Feature Constraint

To ensure the discriminability consistency of global and local features, Hierarchical Feature Constraint (HFC) is proposed to bond the learnings of global features and local ones. The structure of HFC is illustrated in Fig. 3, which is inspired by the Teacher-Student learning spirit in Knowledge Distillation (KD) [12]. As can be seen, instead of introducing an additional pre-trained teacher model, the predictions of convolutional features are employed as “soft-targets” to provide an *extra* supervision for “Student” models, i.e., PCB models of Pose Estimation branch and ReID branch.

Specifically, convolutional features of Pose Estimation branch and ReID branch are firstly concatenated along channel dimension and then transferred to feature vectors by Global Average Pooling. Then, a PCB model (“PCB_T”) is

employed to obtain “soft-targets”, i.e., $P_T = \{p_T^i\}_{i=1}^N$. N refers to the number of training images. Formally, given an image I_i with the identity label y_i , p_T^i can be obtained as follows,

$$p_T^i = p(y_T^i = y_i | I_i) = \frac{\exp(y_T^i)}{\sum_{k=1}^{N_{id}} \exp(y_T^k)}, \quad (14)$$

where $y_T^k = w_{t_k}^T f_T^i$, f_T^i is the fine-grained feature of the i -th image output by the FC layer in “PCB_T”. w_{t_k} indicates the parameter of the classifier in “PCB_T” for the k -th identity. N_{id} is the number of identities in the whole training set. For the j -th PCB model of Pose Estimation branch and ReID branch, the corresponding probability predictions, i.e., $P_{P_j} = \{p_{P_j}^i\}_{i=1}^N$ and $P_{ID_j} = \{p_{ID_j}^i\}_{i=1}^N$, can be calculated in the same way, respectively.

In order to supervise the local feature learning with the global one, KD loss L_{KD} is employed to reduce the distance between two prediction distributions, i.e., P_T and $P_{ID(P)_j}$. Given a mini-batch with M images, L_{KD} is formulated as follows,

$$L_{KD} = \frac{1}{M} \sum_{i=1}^M \sum_{j=1}^P (KL(p_T^i, p_{ID_j}^i) + KL(p_T^i, p_{P_j}^i)), \quad (15)$$

where $KL(p, q)$ measures the Kullback-Leibler divergence between distribution p and distribution q . P denotes the number of PCB models of each branch.

3.5. Loss functions

Batch sampling Following [55, 44], an online batch sampling strategy is adopted during training. Specifically, D identities are randomly selected at each iteration. For each identity, K RGB images and K IR images are then randomly selected to build the mini-batch. Therefore, the batch size $M = 2 * DK$. In the paper, we set $D = 8$, $K = 4$, and $M = 64$ during training.

Pose estimation loss To encourage Pose Estimation branch to learn modality-shared features, pose estimation loss L_{pose}

is introduced to minimize pixel-wise Euclidean distances between ground-truth body keypoint heatmaps and the predicted ones. In the paper, ground-truth heatmaps are derived by a pose estimation model [28], which is pre-trained on the LIP dataset [18]. Formally, L_{pose} across a mini-batch is defined by,

$$L_{pose} = \frac{1}{M} \sum_{i=1}^M \sum_{x,y} (H_i(x, y) - \hat{H}_i(x, y))^2, \quad (16)$$

where $H_i(x, y)$, $\hat{H}_i(x, y)$ represent the pixel value at (x, y) position of the i -th ground truth and the predicted body keypoint heatmap, respectively.

Identity loss To extract ID-related features, identity loss is performed on each PCB model of both Pose Estimation branch and ReID branch. Given the probability predictions P_{P_j} and P_{ID_j} given by the j -th PCB model of Pose Estimation branch and ReID branch, respectively, identity loss L_{id} across a mini-batch is formulated as,

$$L_{id} = -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^P (\log p_{P_j}^i + \log p_{ID_j}^i), \quad (17)$$

where P denotes the number of PCB models of each branch.

Hetero-center triplet loss To reduce both intra- and inter-modality discrepancies, hetero-center triplet (HC-tri) loss [21] is also employed for local feature learning. Similar to identity loss, HC-tri loss is also performed on each PCB model. Specifically, for each PCB model, the centers of fine-grained part features of the i -th identity in a given mini-batch from RGB modality $c_{i,j}^{RGB}$ and IR modality $c_{i,j}^{IR}$ can be computed as follows,

$$c_{i,j}^{RGB} = \frac{1}{K} \sum_{j=1}^K f_{i,j}^{RGB}, \quad c_{i,j}^{IR} = \frac{1}{K} \sum_{j=1}^K f_{i,j}^{IR}, \quad i = \{1, \dots, D\}, \quad (18)$$

where $f_{i,j}^{RGB(IR)}$ represents the fine-grained part feature of the j -th RGB (IR) image in the given mini-batch. Therefore, HC-tri loss for the z -th PCB model of Pose Estimation branch is computed as follows,

$$L_{hctri}^{P_z} = \sum_{i=1}^D [\rho + \|c_i^{RGB} - c_i^{IR}\|_2 - \min_{\substack{m \in \{RGB, IR\} \\ j \neq i}} \|c_i^{RGB} - c_j^m\|_2]_+ + \sum_{i=1}^D [\rho + \|c_i^{IR} - c_i^{RGB}\|_2 - \min_{\substack{m \in \{RGB, IR\} \\ j \neq i}} \|c_i^{IR} - c_j^m\|_2]_+, \quad (19)$$

where $[x]_+ = \max(x, 0)$. ρ refers to the margin value, which is empirically set as 0.3. HC-tri loss of the z -th PCB model of ReID branch, i.e., $L_{hctri}^{ID_z}$, can be calculated in the same way. Therefore, the overall HC-tri loss across the given mini-batch can be derived as follows,

$$L_{hctri} = \sum_{j=1}^P (L_{hctri}^{P_j} + L_{hctri}^{ID_j}). \quad (20)$$

The overall objective for training is defined as,

$$L = L_{id} + \beta L_{hctri} + \lambda L_{pose} + \gamma L_{KD}, \quad (21)$$

where β , λ , and γ are the weighting factors to balance each loss term, which are empirically set as 0.1, 5, 1 in the paper, respectively.

4. Experiments

4.1. Datasets, evaluation metrics and implementation details

Datasets Two benchmark datasets (SYSU-MM01 [40] and RegDB [27]) are employed to evaluate the performance of the proposed method.

SYSU-MM01 [40] consists of images captured by 6 cameras, including 2 IR cameras and 4 RGB ones (2 outdoors and 2 indoors). The training set contains 395 persons, with 22,258 RGB images and 11,909 IR images. The test set contains 96 persons, with 3,803 IR images for query and 301 randomly selected RGB images as the gallery. Following [40], two evaluation modes are conducted: *All-search* and *Indoor-search*. For *Indoor-search* mode, images collected by indoor RGB cameras are exclusively selected to build the gallery set. For *All-search* mode, images are randomly selected from all RGB cameras to form the gallery set.

RegDB [27] contains 412 identities, with 206 for training and 206 for testing. Each identity has 10 RGB and 10 IR images. Two evaluation modes are employed: *Visible-Thermal* and *Thermal-Visible*. The former refers to searching for corresponding IR images with a RGB image and vice versa. The dataset is randomly split into 10 training/testing trials. The evaluation results are given by averaging the performances over the 10 trials.

Evaluation metrics Following the standard evaluation protocol given by [42, 55, 45], Cumulative Matching Characteristics (CMC) curve and mean Average Precision (mAP) are adopted as evaluation metrics. Here, CMC reports the probabilities of the targeted identity occurring at top- r in the ranking list, i.e., “Rank- r ” accuracy. mAP measures the overall retrieval performance when multiple matching cases occur in the gallery set.

Implementation details The experiments are deployed on an NVIDIA GeForce 2080Ti GPU with Pytorch. Following most existing works [44, 55, 45], all input images are resized to 288×144. Random cropping, random erasing, and horizontal flipping are adopted for data augmentation. The parameters of Modality-specific module, Modality-shared module, and ReID branch are initialized by ResNet50 [11] pre-trained on ImageNet. Other parameters are initialized by Xavier initialization [7]. We adopt the SGD optimizer with a weight decay of 0.0005 for optimization. The learning rate is initialized as 0.01 and decays by 0.5 at every 20 epochs. The training process iterates for 100 epochs in total.

4.2. Comparison with state-of-the-arts

We extensively compare our algorithm with the current State-Of-The-Art (SOTA) methods on SYSU-MM01 [40]

Table 1

Comparison with the state-of-the-art methods on **SYSU-MM01** [40]. Rank-r ($r = 1, 10, 20$) accuracy(%) and mAP(%) are reported. $Ours_{ID}$ and $Ours_{ALL}$ denote the features used for evaluation are obtained from ReID branch and both branches, respectively.

Method	Venue	All-search				Indoor-search			
		Rank-1	Rank-10	Rank-20	mAP	Rank-1	Rank-10	Rank-20	mAP
Zero-Pad [40]	ICCV2017	14.80	54.12	71.33	15.95	20.58	68.38	85.79	26.92
HCML [42]	AAAI2018	14.32	53.16	69.17	16.16	24.52	73.25	86.73	30.08
cmGAN [3]	IJCAI2018	26.97	67.51	80.56	31.49	31.63	77.23	89.18	42.19
eDBTR [43]	TIFS2019	27.82	67.34	81.34	28.42	32.46	77.42	89.62	42.46
HSME [10]	AAAI2019	20.68	32.74	77.95	23.12	-	-	-	-
D ² RL [39]	CVPR2019	28.90	70.60	82.40	29.20	-	-	-	-
MSR [4]	TIP2019	37.35	83.40	93.34	38.11	39.64	89.29	97.66	50.88
AlignGAN [38]	ICCV2019	42.40	85.00	93.70	40.70	45.90	87.60	94.40	54.30
TSLFN [55]	Neuro2020	56.96	91.50	96.82	54.95	59.74	92.07	96.22	64.91
AGW [45]	Arxiv2020	47.50	-	-	47.65	54.17	-	-	62.97
X-Modal [17]	AAAI2020	49.92	89.79	95.96	50.73	-	-	-	-
MACE [41]	TIP2020	51.64	87.25	94.44	50.11	57.35	93.02	97.47	64.79
DDAG [44]	ECCV2020	54.75	90.36	95.81	53.02	61.02	94.06	98.41	67.98
cm-SSFT [24]	CVPR2020	61.60	89.20	93.90	63.20	70.50	94.90	97.70	72.60
NFS [2]	CVPR2021	56.91	91.34	96.52	55.45	62.69	96.53	99.07	69.79
CICL [51]	AAAI2021	57.2	94.3	98.4	59.3	66.6	98.8	99.7	74.7
GLMC [47]	TNNLS2021	64.37	93.90	97.53	63.43	67.35	98.10	99.77	74.02
LbA [29]	ICCV2021	55.41	-	-	54.14	58.46	-	-	66.33
DFA [46]	Neuro2022	56.98	90.81	96.13	55.10	66.34	94.91	98.57	71.43
TCOM [32]	Neuro2023	63.92	94.39	98.42	60.71	68.35	97.37	98.90	73.08
$Ours_{ID}$	-	65.82	94.53	98.23	64.52	71.74	94.57	97.60	74.54
$Ours_{ALL}$	-	71.21	95.35	98.81	67.15	72.55	97.15	98.60	77.05

and RegDB [27]. The SOTA methods include pioneering ones (Zero-Pad [40] and HCML [42]), GAN-based ones (cmGAN [3], AlignGAN [38], and D²RL [39]), middle modality based ones (X-Modal [17], cm-SSFT [24], TCOM [32]), feature constraints based ones (eBDTR [43], HSME [10], MSR [4], TSLFN [55], AGW [45], and GLMC [47]), dual-level feature alignment based ones (DDAG [44], MACE [41], CICL [51], LbA [29], and DFA [46]).

Evaluations on SYSU-MM01 Table 1 reports the performance of our model and State-Of-The-Art (SOTA) methods on SYSU-MM01 [40]. “ $Ours_{ID}$ ” and “ $Ours_{ALL}$ ” refer to the retrieval performance with features f_{ID} (Equ. 12) and f_{ALL} (Equ. 13), respectively. It can be seen that both “ $Ours_{ID}$ ” and “ $Ours_{ALL}$ ” outperform SOTA approaches on ALL evaluation metrics in both *All-search* and *Indoor-search* evaluation modes. Compared to the second-best method (GLMC [47]), our model achieves performance improvements by 6.84%/3.72% in terms of Rank-1 accuracy/mAP at *All-search* setting, and 5.20%/3.03% at *Indoor-search* setting.

Evaluations on RegDB The evaluation results on RegDB [27] are reported in Table 2. It can be observed that the proposed model obtains surprisingly good results in both “*Visible-Thermal*” and “*Thermal-Visible*” modes. Specifically, the performance of $Ours_{ID}$ exceeds GLMC [47] by

6.46% and 5.64% in terms of the mAP score in two evaluation modes, respectively. The improvements achieve 7.56% and 6.77% when $Ours_{ALL}$ is employed for identity retrieval. The comparison results also prove that, with the proposed Hierarchical Feature Constraint (HFC), each feature stripe is embedded with modality-shared and ID-related information. Such advantage leads to a satisfactory identity accuracy even when features with less dimensions are used for identity retrieval.

4.3. Comparison with pose-guided single-modality person ReID

As aforementioned, employing pose information in VI-ReID is extremely challenging due to the massive gap between visible images (source domain) and infrared images (target domain). Although there are some pose-guided single-modality person ReID works, how to enable pose information to facilitate VI-ReID has been unrevealed. In the section, several experiments are conducted on SYSU-MM01 [40] to explore the potential of pose estimation in VI-ReID.

Firstly, VI-ReID is considered as the single-modality person ReID task with two different types of images. For each modality, a modality-specific pose-assisted feature extractor is trained. During the inference, features of the query

Table 2

Comparison with the state-of-the-art methods on **RegDB** [27]. Rank- r ($r = 1, 10, 20$) accuracy(%) and mAP(%) are reported. $Ours_{ID}$ and $Ours_{ALL}$ denote the features used for evaluation are obtained from ReID branch and both branches, respectively.

Method	Visible-Thermal				Thermal-Visible			
	Rank-1	Rank-10	Rank-20	mAP	Rank-1	Rank-10	Rank-20	mAP
Zero-Pad [40]	17.75	34.21	44.35	18.90	16.63	34.68	44.25	17.82
HCML [42]	24.44	47.53	56.78	20.80	21.70	45.02	55.58	22.24
eDBTR [43]	34.62	58.96	68.72	33.46	34.21	58.74	68.64	32.49
HSME [10]	50.85	73.36	81.66	47.00	50.15	72.40	81.07	46.16
D ² RL [39]	43.40	66.10	76.30	44.10	-	-	-	-
MSR [4]	48.43	70.32	79.95	48.67	-	-	-	-
AlignGAN [38]	57.9	-	-	53.6	56.3	-	-	53.4
X-Modal [17]	62.21	83.13	91.72	60.18	-	-	-	-
DDAG [44]	69.34	86.19	91.49	63.46	68.06	85.15	90.31	61.80
AGW [45]	70.05	-	-	66.37	-	-	-	-
MACE [41]	72.37	88.40	93.59	69.09	72.12	88.07	93.07	68.57
cm-SSFT [24]	72.3	-	-	72.9	71.0	-	-	71.7
NFS [2]	80.54	91.96	95.07	72.10	77.95	90.45	93.62	69.79
CICL [51]	78.8	-	-	69.4	77.9	-	-	69.4
GLMC [47]	91.84	97.86	98.98	81.42	91.12	97.86	98.69	81.06
LbA [29]	74.17	-	-	67.64	72.43	-	-	65.46
DFA [46]	73.05	87.92	91.52	67.76	70.53	85.82	90.51	63.93
TCOM [32]	87.04	96.12	98.20	80.40	83.20	93.15	97.18	76.73
$Ours_{ID}$	92.14	98.16	99.22	87.88	91.36	97.57	98.88	86.70
$Ours_{ALL}$	93.35	98.61	99.42	88.98	92.72	98.79	99.36	87.83

Table 3

Comparison with SOTA pose-guided single-modality person ReID methods on SYSU-MM01 [40] in terms of Rank-1 accuracy(%) and mAP(%) at *All-search* setting. The dimension of features used for the evaluation is set as 2560 for all experiments. $Ours_{ID}^*$ means P in Equ. (11) is set as 5. Impl. means experiments are implemented with the official source code provided.

Methods	Rank-1	mAP
PGFA(Impl.) [26]	10.04	11.45
PABR(Impl.) [35]	40.73	42.19
$Ours_{ID}^*$	62.82	61.90

set (IR images) and gallery set (RGB images) are derived by the corresponding feature extractor to conduct the cross-modality person retrieval. Taking a state-of-the-art (SOTA) pose-guided single-modality person ReID work, *i.e.*, PGFA [26], as an example, where human landmarks obtained by an off-the-shelf pose estimator are rigidly used to generate attention maps to highlight body regions, the performance is reported in the first row of Table 3. As can be seen, although the same pose estimator is used for two modalities, retrieving identities across modalities with modality-specific pose-assisted features can only achieve 11.45% in terms of the mAP score. The poor performance demonstrates that

the off-the-shelf pose estimators cannot handle the huge gap between the data distribution of the two modalities. In this case, the pose information cannot serve as effective modality-shared cues in VI-ReID.

To make the best of the pose information in VI-ReID, we applied the two-stream network proposed by a single-modality person ReID approach, *i.e.*, PABR [35], where the ReID branch and the pose branch extract appearance and pose features, respectively. Then two kinds of features are fused for identity retrieval. To adapt PABR to VI-ReID, the cross-modality hard triplet loss [50] is used to replace the single-modality one, where the hardest cross-modality triplets are also considered. The performance of PABR in VI-ReID is shown in the second row of Table 3. Compared to PGFA, PABR achieves a higher Rank-1 accuracy (40.73%) and mAP score (42.19%). In addition to numerical results, we also visualize the attentive feature maps output by the pose branch of PABR [35] by means of Grad-CAM [31] in Fig. 4(a), to have a better understanding of where pose features are extracted. As can be seen, due to the lack of adequate guidance on the pose branch during training, the extracted pose features only locate the whole body coarsely.

To solve the problem, as shown in Fig. 2, the body key-points generated by a pre-trained pose estimator serve only as the guidance of Pose Estimation branch in our method. By applying the *pose estimation loss* during training, more fine-grained pose features can be obtained. The performance




Figure 4: Visualization of attentive feature maps output by the pose branch of PABR [35] and the proposed method on SYSU-MM01 [40]. For each person (group), the RGB image is shown on the left and the IR image is on the right.

of the proposed method is shown in the third row of Table 3. It can be seen that, with above improvements, our method outperforms PGFA and PABR by a large margin in terms of the Rank-1 accuracy (62.92%) and the mAP score (61.90%). Additionally, the visualization of the attentive feature maps output by our Pose Estimation branch is shown Fig. 4(b). In comparison to PABR, our method focuses on more body details, such as shoulders and feet, which can serve as distinctive cues for VI-ReID.

4.4. Ablation studies

An ablation analysis is conducted on SYSU-MM01 [40] and RegDB [27] to verify the effectiveness of the proposed components in our model. Firstly, a “Baseline” model is trained under the supervision of identity loss L_{id} and HCTri loss L_{hctri} , which only consists of Modality-specific module, Modality-shared module, and ReID branch. The evaluation results are shown in the first row in Table 4. Based on “Baseline”, Pose Estimation branch (PEB), pose estimation loss (L_{pose}), and Hierarchical Feature Constraint (HFC) are gradually applied, which results are illustrated in the following rows in Table 4. Note that the reported results are in the *All-search* mode for SYSU-MM01 while in the “Visible-Thermal” mode for RegDB.

Effectiveness of Pose Estimation branch (PEB) As can be seen from the 2nd row in Table 4, by integrating the PEB branch that is pre-trained on a pose estimation dataset, the framework yields an increase of approximately 2% and 4% in terms of the mAP score on SYSU-MM01 and RegDB, respectively, when f_{ID} is applied for identity retrieval. Likewise, the improvements exceed 6.5% when employing f_{ALL} , which clearly demonstrates the effectiveness of our PEB branch.

Table 4

Ablation studies on SYSU-MM01 [40] and RegDB [27]. “PEB”, “HFC” refer to Pose Estimation branch and Hierarchical Feature Constraint, respectively.

	Components			SYSU-MM01		RegDB	
	PEB	L_{pose}	HFC	Rank-1	mAP	Rank-1	mAP
Baseline				57.03	56.21	85.76	77.21
$Ours_{ID}$	✓			60.19	58.96	87.39	81.75
	✓	✓		63.90	62.79	89.25	84.16
	✓	✓	✓	65.82	64.52	92.14	87.88
$Ours_{ALL}$	✓			65.05	62.78	89.53	84.12
	✓	✓		68.84	65.94	91.96	87.08
	✓	✓	✓	71.21	67.15	93.35	88.98

Apart from L_{id} and L_{hctri} , L_{pose} is further employed on such a structure to ensure the body skeleton points are precisely estimated. The results for f_{ID} and f_{ALL} are shown in the 3rd and 6th rows, respectively. Specifically, an increase of 3.83% and 2.41% can be found in terms of the mAP score on two benchmark datasets for f_{ID} . A similar increase, *i.e.*, 3.16% and 2.96%, can be obtained if using f_{ALL} .

Effectiveness of Hierarchical Feature Constraint (HFC)

It can be observed from the 4th row that HFC yields a satisfactory improvement for f_{ID} in terms of the mAP score, which are 1.73% and 3.72% on SYSU-MM01 and RegDB, respectively. A similar improvement can also be found for f_{ALL} (the 7th row), which are around 2% for both datasets. The boost of performance proves that the proposed HFC encourages the information flow between global and local features by introducing extra supervision on part-level features.

4.5. Impact of Loss coefficient

In this section, we aim to investigate an optimal combination of weight factors to balance multiple loss terms. According to the paper, the overall objective L for training is defined as,

$$L = L_{id} + \beta L_{hctri} + \lambda L_{pose} + \gamma L_{KD}. \quad (22)$$

β , λ , and γ adjust the contribution of Hetero-center triplet loss, Pose estimation loss, and KD loss, respectively. For a fair comparison, we only change one factor at a time while keeping the rest fixed. Subsequently, models are trained by objectives with different combinations of weight factors. The corresponding mAP scores on the test set are depicted in Fig. 5.

It can be observed that no matter which combination is employed, the proposed framework could achieve a stable performance within 50 epochs. Additionally, according to the experimental results, the optimal combination $\{\beta, \lambda, \gamma\} = \{0.1, 5, 1\}$ is chosen during the training.




Figure 5: Parameter analysis on RegDB [27]. Best viewed in color.




Figure 6: Visualization of gradient feature maps output by ReID branch on SYSU-MM01 [40]. Top: RGB images. Bottom: IR images. The corresponding mAP scores are reported in brackets.

4.6. Visualization

Apart from quantitative results, we also visualize the gradient feature maps output by ReID branch with Grad-CAM [31], in order to examine where the features are extracted.

Gradient feature maps output by “Baseline” and the proposed framework (“Ours”) are illustrated in Fig. 6(a) and Fig. 6(b), respectively. Feature maps are grouped by modality. Specifically, RGB images are illustrated at the top row while IR images are at the bottom. As can be seen, the “Baseline” model extracts identity features from random regions rigidly. Although covering the partial body, the features are less identity-related since the extraction is imposed coarsely. In terms of our method, identity-related features are extracted from regions that are close to body skeleton joints. Such fine-grained features not only can be used to differentiate different identities but are highly immune to viewpoint, accessories, and modality variations. Taking the girl in the first image as an example, although the upper body is covered by the bag, our method is able to extract

effective features from joint-related regions. The superiority of our method in fine-grained feature extraction leads to better performance.

Therefore, the visualization results not only intuitively reveal the reason why our method performs better, but also show the potential of the pose estimation task in the field of VI-ReID.

5. Conclusion

We have proposed a novel two-stream VI-ReID framework, where modality-shared and ID-related features for identity retrieval are extracted by means of learning an auxiliary task (pose estimation) and the main task (person ReID) simultaneously. By imposing pose estimation and ReID constraints on the Pose Estimation branch at the same time, both modality-shared and ID-related are fully embedded on each feature stripe. Apart from learning discriminative features at the local level, we also propose a Hierarchical Feature Constraint to bond the learning of global features with local

ones by employing the knowledge distillation strategy to ensure discriminability consistency. The proposed framework achieves new state-of-the-art on VI-ReID benchmarks in terms of rank-1 accuracy and the mAP score. Inspired by the recent advance in network architecture, such as Transformer, CapsuleNet [23], and the feature fusion strategies, such as Bilinear Pooling [19] and cross attention, we will explore their potential in the proposed framework in the future.

References

- [1] Chen, T., Ding, S., Xie, J., Yuan, Y., Chen, W., Yang, Y., Ren, Z., Wang, Z., 2019. Abd-net: Attentive but diverse person re-identification, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 8351–8361.
- [2] Chen, Y., Wan, L., Li, Z., Jing, Q., Sun, Z., 2021. Neural feature search for rgb-infrared person re-identification, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 587–597.
- [3] Dai, P., Ji, R., Wang, H., Wu, Q., Huang, Y., 2018. Cross-modality person re-identification with generative adversarial training, in: IJ-CAI, p. 2.
- [4] Feng, Z., Lai, J., Xie, X., 2019. Learning modality-specific representations for visible-infrared person re-identification. IEEE Trans. Image Process. 29, 579–590.
- [5] Ge, Y., Li, Z., Zhao, H., Yin, G., Yi, S., Wang, X., Li, H., 2018a. Fd-gan: Pose-guided feature distilling gan for robust person re-identification. arXiv preprint arXiv:1810.02936 .
- [6] Ge, Y., Li, Z., Zhao, H., Yin, G., Yi, S., Wang, X., et al., 2018b. Fd-gan: Pose-guided feature distilling gan for robust person re-identification. Proc. Advances in Neural Inf. Process. Syst. 31.
- [7] Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks, in: Proc. Conf. Artif. Intell. and Stat., pp. 249–256.
- [8] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2020. Generative adversarial networks. Commun. ACM 63, 139–144.
- [9] Guo, J., Yuan, Y., Huang, L., Zhang, C., Yao, J.G., Han, K., 2019. Beyond human parts: Dual part-aligned representations for person re-identification, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 3642–3651.
- [10] Hao, Y., Wang, N., Li, J., Gao, X., 2019. Hsme: hypersphere manifold embedding for visible thermal person re-identification, in: Proc. AAAI Conf. Artificial Intell., pp. 8385–8392.
- [11] He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 770–778.
- [12] Hinton, G., Vinyals, O., Dean, J., 2015. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 .
- [13] Huang, N., Liu, J., Miao, Y., Zhang, Q., Han, J., 2023. Deep learning for visible-infrared cross-modality person re-identification: A comprehensive review. Information Fusion 91, 396–411.
- [14] Huang, N., Liu, K., Liu, Y., Zhang, Q., Han, J., 2022. Cross-modality person re-identification via multi-task learning. Pattern Recognition 128.
- [15] Kalayeh, M.M., Basaran, E., Gökmen, M., Kamasak, M.E., Shah, M., 2018. Human semantic parsing for person re-identification, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 1062–1071.
- [16] Leng, Q., Ye, M., Tian, Q., 2019. A survey of open-world person re-identification. IEEE Trans. Circuit Syst. Video Technol. 30, 1092–1108.
- [17] Li, D., Wei, X., Hong, X., Gong, Y., 2020. Infrared-visible cross-modal person re-identification with an x modality, in: Proc. AAAI Conf. Artificial Intell., pp. 4610–4617.
- [18] Liang, X., Gong, K., Shen, X., Lin, L., 2018. Look into person: Joint body parsing & pose estimation network and a new benchmark. IEEE Trans. Pattern Anal. Mach. Intell. , 871–885.
- [19] Lin, T.Y., RoyChowdhury, A., Maji, S., 2017. Bilinear convolutional neural networks for fine-grained visual recognition. IEEE Trans. Pattern Anal. Mach. Intell. 40, 1309–1322.
- [20] Ling, H., Wang, Z., Li, P., Shi, Y., Chen, J., Zou, F., 2019. Improving person re-identification by multi-task learning. Neurocomputing 347, 109–118.
- [21] Liu, H., Tan, X., Zhou, X., 2020. Parameter sharing exploration and hetero-center triplet loss for visible-thermal person re-identification. IEEE Trans. Multimedia .
- [22] Liu, J., Ni, B., Yan, Y., Zhou, P., Cheng, S., Hu, J., 2018. Pose transferrable person re-identification, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 4099–4108.
- [23] Liu, Y., Zhang, D., Zhang, Q., Han, J., 2022. Part-object relational visual saliency. IEEE Trans. Pattern Anal. Mach. Intell. 44, 3688–3704.
- [24] Lu, Y., Wu, Y., Liu, B., Zhang, T., Li, B., Chu, Q., Yu, N., 2020. Cross-modality person re-identification with shared-specific feature transfer, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 13379–13389.
- [25] Lv, Y., Gu, Y., Xinggao, L., 2020. The dilemma of trihard loss and an element-weighted trihard loss for person re-identification. Proc. Advances in Neural Inf. Process. Syst. 33, 17391–17402.
- [26] Miao, J., Wu, Y., Liu, P., Ding, Y., Yang, Y., 2019. Pose-guided feature alignment for occluded person re-identification, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 542–551.
- [27] Nguyen, D.T., Hong, H.G., Kim, K.W., Park, K.R., 2017. Person recognition system based on a combination of body images from visible light and thermal cameras. Sensors 17, 605.
- [28] Osokin, D., 2019. Global context for convolutional pose machines. arXiv preprint arXiv:1906.04104 .
- [29] Park, H., Lee, S., Lee, J., Ham, B., 2021. Learning by aligning: Visible-infrared person re-identification using cross-modal correspondences, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 12046–12055.
- [30] Qian, X., Fu, Y., Xiang, T., Wang, W., Qiu, J., Wu, Y., Jiang, Y.G., Xue, X., 2018. Pose-normalized image generation for person re-identification, in: Proc. Eur. Conf. Comp. Vis., pp. 650–667.
- [31] Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 618–626.
- [32] Si, T., He, F., Li, P., Gao, X., 2023. Tri-modality consistency optimization with heterogeneous augmented images for visible-infrared person re-identification. Neurocomputing 523, 170–181.
- [33] Song, C., Huang, Y., Ouyang, W., Wang, L., 2018. Mask-guided contrastive attention model for person re-identification, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 1179–1188.
- [34] Su, C., Li, J., Zhang, S., Xing, J., Gao, W., Tian, Q., 2017. Pose-driven deep convolutional model for person re-identification, in: Proc. IEEE Int. Conf. Comp. Vis., pp. 3960–3969.
- [35] Suh, Y., Wang, J., Tang, S., Mei, T., Lee, K.M., 2018. Part-aligned bilinear representations for person re-identification, in: Proc. Eur. Conf. Comp. Vis., pp. 402–419.
- [36] Sun, Y., Zheng, L., Yang, Y., Tian, Q., Wang, S., 2018. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline), in: Proc. Eur. Conf. Comp. Vis., pp. 480–496.
- [37] Tay, C.P., Roy, S., Yap, K.H., 2019. Aanet: Attribute attention network for person re-identifications, in: Proc. IEEE Conf. Comp. Vis. Patt. Recogn., pp. 7134–7143.
- [38] Wang, G., Zhang, T., Cheng, J., Liu, S., Yang, Y., Hou, Z., 2019a. Rgb-infrared cross-modality person re-identification via joint pixel

- and feature alignment, in: *Proc. IEEE Int. Conf. Comp. Vis.*, pp. 3623–3632.
- [39] Wang, Z., Wang, Z., Zheng, Y., Chuang, Y.Y., Satoh, S., 2019b. Learning to reduce dual-level discrepancy for infrared-visible person re-identification, in: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pp. 618–626.
 - [40] Wu, A., Zheng, W.S., Yu, H.X., Gong, S., Lai, J., 2017. Rgb-infrared cross-modality person re-identification, in: *Proc. IEEE Int. Conf. Comp. Vis.*, pp. 5380–5389.
 - [41] Ye, M., Lan, X., Leng, Q., Shen, J., 2020a. Cross-modality person re-identification via modality-aware collaborative ensemble learning. *IEEE Trans. Image Process.* 29, 9387–9399.
 - [42] Ye, M., Lan, X., Li, J., Yuen, P., 2018. Hierarchical discriminative learning for visible thermal person re-identification, in: *Proc. AAAI Conf. Artificial Intell.*
 - [43] Ye, M., Lan, X., Wang, Z., Yuen, P.C., 2019. Bi-directional center-constrained top-ranking for visible thermal person re-identification. *IEEE Trans. Inf. Foren. and Sec.* 15, 407–419.
 - [44] Ye, M., Shen, J., J. Crandall, D., Shao, L., Luo, J., 2020b. Dynamic dual-attentive aggregation learning for visible-infrared person re-identification, in: *Proc. Eur. Conf. Comp. Vis.*, Springer. pp. 229–247.
 - [45] Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., Hoi, S.C., 2021. Deep learning for person re-identification: A survey and outlook. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 2872–2893.
 - [46] Yin, J., Ma, Z., Xie, J., Nie, S., Liang, K., Guo, J., 2022. Dual-granularity feature alignment for cross-modality person re-identification. *Neurocomputing* 511, 78–90.
 - [47] Zhang, L., Du, G., Liu, F., Tu, H., Shu, X., 2021. Global-local multiple granularity learning for cross-modality visible-infrared person reidentification. *IEEE Trans. Neural Networks Learn. Syst.* .
 - [48] Zhao, H., Tian, M., Sun, S., Shao, J., Yan, J., Yi, S., Wang, X., Tang, X., 2017a. Spindle net: Person re-identification with human body region guided feature decomposition and fusion, in: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pp. 1077–1085.
 - [49] Zhao, L., Li, X., Zhuang, Y., Wang, J., 2017b. Deeply-learned part-aligned representations for person re-identification, in: *Proc. IEEE Int. Conf. Comp. Vis.*, pp. 3219–3228.
 - [50] Zhao, Y.B., Lin, J.W., Xuan, Q., Xi, X., 2019. Hpiln: a feature learning framework for cross-modality person re-identification. *IET Image Process.* 13, 2897–2904.
 - [51] Zhao, Z., Liu, B., Chu, Q., Lu, Y., Yu, N., 2021. Joint color-irrelevant consistency learning and identity-aware modality adaptation for visible-infrared cross modality person re-identification, in: *Proc. AAAI Conf. Artificial Intell.*, pp. 3520–3528.
 - [52] Zheng, L., Huang, Y., Lu, H., Yang, Y., 2019a. Pose-invariant embedding for deep person re-identification. *IEEE Trans. Image Process.* 28, 4500–4509.
 - [53] Zheng, Z., Yang, X., Yu, Z., Zheng, L., Yang, Y., Kautz, J., 2019b. Joint discriminative and generative learning for person re-identification, in: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pp. 2138–2147.
 - [54] Zhu, K., Guo, H., Liu, Z., Tang, M., Wang, J., 2020a. Identity-guided human semantic parsing for person re-identification, in: *Proc. Eur. Conf. Comp. Vis.*, pp. 346–363.
 - [55] Zhu, Y., Yang, Z., Wang, L., Zhao, S., Hu, X., Tao, D., 2020b. Hetero-center loss for cross-modality person re-identification. *Neurocomputing* 386, 97–109.