

# Modality-Aware Bias Mitigation and Invariance Learning for Unsupervised Visible-Infrared Person Re-Identification

Menglin Wang<sup>1</sup>, Xiaojin Gong<sup>2\*</sup>, Jiachen Li<sup>2</sup>, Genlin Ji<sup>1</sup>

<sup>1</sup>School of Computer and Electronic Information, Nanjing Normal University, China

<sup>2</sup>College of Information Science and Electronic Engineering, Zhejiang University, China

lynnwang6875@gmail.com, gongxj@zju.edu.cn

## Abstract

Unsupervised visible-infrared person re-identification (USVI-ReID) aims to match individuals across visible and infrared cameras without relying on any annotation. Given the significant gap across visible and infrared modality, estimating reliable cross-modality association becomes a major challenge in USVI-ReID. Existing methods usually adopt optimal transport to associate the intra-modality clusters, which is prone to propagating the local cluster errors, and also overlooks global instance-level relations. By mining and attending to the visible-infrared modality bias, this paper focuses on addressing cross-modality learning from two aspects: bias-mitigated global association and modality-invariant representation learning. Motivated by the camera-aware distance rectification in single-modality re-ID, we propose modality-aware Jaccard distance to mitigate the distance bias caused by modality discrepancy, so that more reliable cross-modality associations can be estimated through global clustering. To further improve cross-modality representation learning, a ‘split-and-contrast’ strategy is designed to obtain modality-specific global prototypes. By explicitly aligning these prototypes under global association guidance, modality-invariant yet ID-discriminative representation learning can be achieved. While conceptually simple, our method obtains state-of-the-art performance on benchmark VI-ReID datasets and outperforms existing methods by a significant margin, validating its effectiveness.


**Code** — <https://github.com/Terminator8758/BMIL>

## Introduction

Visible-infrared person re-identification (Wu et al. 2017) is the task of matching the same person from visible to infrared camera or vice versa. Due to the potential application in night-time surveillance and person retrieval, this task has received increasing research interest. However, it is costly and time-consuming to annotate the cross-modality image relations, therefore unsupervised visible-infrared re-identification (USVI-ReID) has gained its popularity in recent years (Liang et al. 2021; Yang et al. 2022a; Wu and Ye 2023; Chen et al. 2023; Yin et al. 2024; Yang, Chen, and Ye 2023a, 2024; Shi et al. 2024b,a; Teng et al. 2025).

\*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) Intra and inter-modality ratio in Top-30 NNs



(b) T-SNE visualization of visible (circle shape) and infrared (triangle shape) image features

Figure 1: Illustration of KNN distribution and feature visualization, using intra-modality (stage-1) trained model.

Given the substantial modality gap between visible and infrared images, most existing methods formulate USVI-ReID as a two-stage learning process. The first stage focuses on intra-modality learning to recognize identities from the same modality. This is typically done by iterative clustering and prototypical contrast learning, inspired by previous unsupervised single-modality re-ID methods (Fan et al. 2018; Ge et al. 2020; Dai et al. 2022; Wang et al. 2021a; Cho et al. 2022; Wu et al. 2022; Zou et al. 2023). The second stage deals with cross-modality matching and learning, where the key challenge is how to estimate reliable cross-modality association, under the presence of large modality discrepancy. A common strategy is to use optimal transport (e.g. Hungarian matching) to predict the correspondence among visi-

ble and infrared clusters, then enforce a cross-modality contrastive loss based on the matching result. However, this strategy propagates the intra-modality cluster errors while matching, and also overlooks the global instance relationship, reducing the efficiency of cross-modality learning.

On the other hand, a naive modality-agnostic global clustering is also not ideal. Due to the modality gap, intra-modality images generally exhibit higher similarity compared to cross-modality images, leading to a modality bias that negatively influences the pairwise distance distribution. Specifically, unsupervised clustering typically adopts pairwise Jaccard distance (Zhong, Zheng, and Li 2017), which entails the  $K$ -nearest neighbor computation. With the modality bias, the retrieved KNNs are dominated by intra-modality instances, as observed in Figure 1(a). The biased neighbors further exacerbate the intra- and inter-modality distance gap, causing the global clustering to be unreliable and fail to associate cross-modality instances. Some recent efforts (Yang, Hu, and Hu 2025) try to select a larger number of KNNs and separately fuse visible and infrared features of the original KNNs. However, this does not address the inherently biased neighbor distribution, therefore has limited effect in rectifying the cross-modal distance.

To this end, we propose modality-aware global clustering to mitigate the modality bias and release the potential of global clustering. Inspired by the camera-aware distance rectification (Chen et al. 2024) in single modality re-ID, we mitigate the bias in Jaccard distance by independently retrieving  $K$ -nearest neighbors within each modality. These modality-balanced neighbors are incorporated into two key steps of Jaccard distance, *i.e.* reciprocal neighbor computation and local query expansion, ensuring a fair contribution of both intra- and inter-modality neighbor instances in distance computation. As a result, the pairwise distance is modulated and less impacted by modality bias, thereby increasing the likelihood of associating cross-modality instances using global clustering.

While the cross-modality association bias is mitigated, there also exists issues in cross-modality representation learning. As identified by previous methods (Zou et al. 2023; Shi et al. 2024b), utilizing the centroid feature as the cluster prototype only represents the commonality, but fails to capture the variance and diversity of a cluster. The cluster variance becomes a more notable issue in cross-modality scenario, as the features of visible and infrared samples can have distinct distributions within the same global cluster. This is also supported by our visualizations in Figure 1(b), where images of the same modality are more likely to form a tight clique. Therefore, it is suboptimal to represent the global cluster with one unified centroid prototype. More informative prototypes are needed to precisely describe the modality-mixed cluster distribution.

Leveraging the modality label as a natural prior, we propose a *split-and-contrast* strategy to split each cluster into modality-specific sub-clusters, and design modality-specific global prototypes to capture fine-grained dynamics within global clusters. By explicitly aligning the modality-specific prototypes using a multi-positive contrastive loss, the model is enforced to learn representations invariant to modality dis-

crepancy. Together with the intra-modality prototypes, the representation learning is simultaneously guided by locally and globally associated prototypes, leading to a model that is robust to modality variation, yet discriminative at identity recognition.

To summarize, the contributions of this paper are:

- We present a simple yet effective method for USVI-ReID, targeting at cross-modality learning from the perspective of bias mitigation and invariance learning.
- To mitigate the distance bias caused by modality discrepancy, modality-aware Jaccard distance is proposed. Meanwhile, modality-specific global prototypes are designed to learn modality-invariant representations.
- On benchmark datasets, we perform extensive experiments and achieve state-of-the-art performance, demonstrating the effectiveness and superiority of our method.

## Related Work

**Supervised Visible-Infrared Re-ID.** The goal of supervised visible-infrared re-ID (VI-ReID) is to learn discriminative feature to recognize a person across modalities, by exploiting the given annotation. The VI-ReID task and the SYSU-MM01 dataset are initially proposed by Wu *et al.* (Wu et al. 2017). Since cross-modality alignment is important for tackling the modality discrepancy, existing VI-ReID methods can be broadly categorized into image-level or feature-level alignment. For image-level alignment, D2RL (Wang et al. 2019) generate cross-modality images using generative models. CAJ (Ye et al. 2021) and CAJ+ (Ye et al. 2024) design random channel exchange to generate color-irrelevant images, improving the robustness against color variations. For feature-level alignment, MACE (Ye et al. 2020a) presents a middle-level sharable two-stream network to handle the modality discrepancy in both feature and classifier level. MAUM (Liu et al. 2022) utilizes historically accumulated proxies to align cross-modality features. PartMix (Kim et al. 2023) randomly mixes cross-modality part-level features to regularize the model. FMCNet (Zhang et al. 2022) and MUN (Yu et al. 2023a) generate intermediate modality feature to implicitly reduce the modality gap. Some recent methods (Hu, Yang, and Ye 2024; Yu et al. 2025) also exploit text modality to enhance visual feature learning. Although supervised VI-ReID performance has been significantly improved, the reliance on heavy annotation presents a limitation to its practical applicability.

**Unsupervised Visible-Infrared Re-ID.** Due to the task characteristic, current USVI-ReID approaches usually combine intra-modality and cross-modality learning for method design. Driven by the success of single-modality unsupervised re-ID methods (Dai et al. 2022; Ge et al. 2020; Wang et al. 2021a; Cho et al. 2022; Zou et al. 2023), intra-modality learning usually follows an iterative clustering and prototypical contrastive learning paradigm, to gradually boost the model performance. On this basis, some methods focus on intra-modality learning by designing multiple prototypes to represent cluster diversity or variance (Shi et al. 2024b,a), utilizing metric loss to assist representation learning (Yin et al. 2024), or considering neighbor information (Teng

et al. 2025; Yin et al. 2024) to refine the pseudo label. For cross-modality learning, the key issue is to predict reliable cross-modality association, and design efficient representation learning thereon. To this end, some methods estimate cross-modality correspondence by counting the number of reliable instance pairs in cross-modality clusters (Yang et al. 2022a), designing label propagation based on normalized similarity (Yang, Chen, and Ye 2023a), exploiting optimal transport (Wang et al. 2022) and graph matching (Wu and Ye 2023), or assigning instance weight according to cluster and neighbor correlation (Teng et al. 2025).

Despite the progress in USVI-ReID, one notable issue is the *accumulation of cluster noise* when matching the intra-modality clusters. After associating the local clusters, intra-modality cluster noise is still kept and renders the association unreliable. What’s more, associating the clusters ignores the global similarity relation among visible and infrared instances. In this paper, we aim to address this by proposing modality-aware global clustering to obtain more reliable association among instances.

**Debiased and Invariance Learning.** Learning unbiased model from biased data has been a long-studied topic in visual recognition. Data bias such as image style, color or texture, can bring spurious correlations that negatively influence model learning. To address the bias, some methods (Wang et al. 2021b; Wad et al. 2022) turn to causal intervention such as predicting data partition as effective backdoor adjustment. Regularization based debiasing methods are also popular for their lower complexity. For example, Invariant Risk Minimization (IRM) (Arjovsky et al. 2019) try to reduce the loss variance among different biases subsets, so as to achieve invariance to bias. EnD (Tartaglione, Barbano, and Grangetto 2021) and BiasCon (Hong and Yang 2021) bring closer the representation of positive samples in case of different biases, while push apart negative samples sharing the same bias attribute. In general re-ID task, the camera bias also brings a spurious correlation to model. To reduce its influence in similarity measure, (Chen et al. 2024) proposes camera-aware distance rectification that enhances the reliability of Jaccard distance. Compared to general re-ID, the visible-infrared modality discrepancy becomes a new source of bias in VI re-ID. In a similar spirit to (Chen et al. 2024), we propose modality-aware Jaccard distance to effectively mitigate the bias and facilitate cross-modal clustering. Moreover, from the invariance learning perspective, we design a modality-invariant loss to explicitly optimize the modality-biased representations, thus better combat the modality confounder in model learning.

## Methodology

### Overview

In USVI-ReID task, the overall aim is to learn a robust and discriminative model in an unsupervised setting, to match an identity across visible and infrared images. To achieve the goal, this paper addresses cross-modality learning from the perspective of mitigating the modality bias. Our method focuses on two complementary aspects: bias-mitigated global cross-modality association and modality-invariant represen-

tation learning. An overview of the proposed framework is illustrated in Figure 2. The rest of this section is organized as follows: First, we introduce the preliminary information of the USVI-reID task, then we introduce the intra-modality learning. Afterwards, we present the cross-modality learning, where the proposed global association and invariant representation learning will be detailed.

### Intra-Modality Learning

Suppose a cross-modality dataset is given as  $\mathcal{D} = \{\mathcal{D}_v, \mathcal{D}_r\}$ , where  $\mathcal{D}_v = \{x_1^v, x_2^v, \dots, x_{N_1}^v\}$  denotes the visible training set composed of  $N_1$  images, and  $\mathcal{D}_r = \{x_1^r, x_2^r, \dots, x_{N_2}^r\}$  is the infrared training set of  $N_2$  images. As previous methods, we use a two-stream backbone  $f_\theta$  to extract features for visible and infrared images. The backbone has separate initial convolution blocks for visible and infrared images, while the rest blocks are shared for both modalities.

Due to the image discrepancy in visible and infrared modalities, the backbone model is initially weak at matching cross-modality images. Therefore, a two-stage learning procedure is usually conducted, with intra-modality learning in the first stage to warm-up the model. For intra-modality learning, an iterative clustering and prototypical contrastive learning paradigm is adopted in each modality.

Consider visible modality as an example. After obtaining a  $d$ -dimensional image feature  $f_\theta(x^v)$  from the backbone  $f_\theta$ , DBSCAN clusters all visible images into  $C^v$  clusters. Then for each cluster, a representative prototype is constructed and initialized as the cluster mean feature. A memory bank for visible modality is then formed as the collection of all clusters’ prototypes, i.e.  $M^v \in R^{C^v \times d}$ . During mini-batch training, when the backbone parameter is updated by back-propagation, the prototype memory is updated by online image features:

$$M^v[j] \leftarrow \mu M^v[j] + (1 - \mu) f_\theta(x_i^v), \quad (1)$$

where  $M^v[j]$  is the  $j$ -th entry of the memory bank,  $\mu$  is the updating rate.  $x_i^v$  is an image belonging to the  $j$ -th cluster.

To recognize different identities within the same modality, each visible training image  $x_i^v$  is contrasted with all the visible prototypes using the InfoNCE loss (Gutmann and Hyvärinen 2010) as follow:

$$\mathcal{L}_{intra}^v = - \sum_{i=1}^{N_b} \log \frac{\exp(M^v[\tilde{y}_i]^T f_\theta(x_i^v)/\tau)}{\sum_{j=1}^{C^v} \exp(M^v[j]^T f_\theta(x_i^v)/\tau)}, \quad (2)$$

where  $N_b$  is the batch size,  $\tilde{y}_i$  is the pseudo label of image  $x_i^v$  from clustering,  $\tau$  is the temperature. When optimizing the prototypical contrastive loss in Eq. (2), the image is drawn closer to its belonging prototype, while pushed away from all the other prototypes, thus achieving discriminative representation learning in the same modality.

For the infrared modality, the clustering, infrared prototype memory and the loss  $\mathcal{L}_{intra}^r$  can be formulated similarly. The total intra-modality loss is then computed as:

$$\mathcal{L}_{intra} = \mathcal{L}_{intra}^v + \mathcal{L}_{intra}^r. \quad (3)$$

**Dealing with over-clustering issue.** In real-world applications, it is common that a greater number of visible images

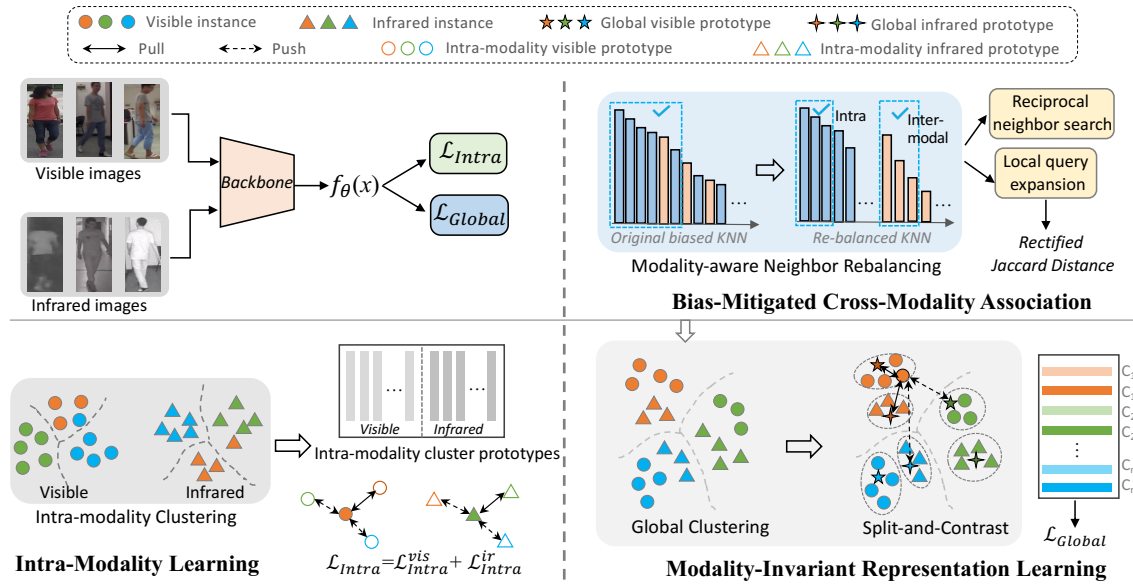


Figure 2: An overview of the proposed modality-aware learning framework. For cross-modality learning, a bias-mitigating global association is proposed, featuring the modality-aware neighbor rebalancing strategy to rectify the pairwise distance and overcome the modality-induced bias. The rectified distance is utilized by global clustering to obtain modality-mixed clusters. A split-and-contrast design captures the variance of global clusters, then a multi-positive contrastive loss is designed to facilitate the learning of modality-invariant yet ID-discriminative representation.

are collected compared to infrared images. When applying merging-based clustering algorithms like DBSCAN on the visible modality, over-clustering tends to appear, *i.e.* much more clusters are predicted than the actual identity number. To alleviate over-clustering, we utilize a **subset clustering strategy** (Jin et al. 2022) to randomly sample a subset of images for clustering. By sampling at a fixed ratio (*e.g.* 0.5), the average number of images per identity is reduced, making it easier to cluster the sampled images. In addition, each epoch samples a different subset of images due to randomness, making sure the full image set gets covered during training. Our experiments show that the subset clustering reduces over-clustering while also saves computation cost.

### Bias-Mitigated Cross-Modality Association

Based on prototypes obtained from intra-modality clustering, a common way to establish cross-modality correspondence is to match the cross-modality prototypes through optimal transport algorithms such as Hungarian (Kuhn 1955) or Sinkhorn algorithm (Cuturi 2013). However, this two-step matching process is prone to accumulating the label noise of intra-modality clustering. Through cluster matching, intra-modality noise is falsely propagated across modality.

**The modality bias.** With the perception of instance-level similarity relation, one might consider global clustering as a suitable choice for cross-modality association. However, modality-induced bias becomes a main obstacle limiting the effectiveness of global clustering. The recognition of cross-modality images is confounded by the modality-specific context, such as image style, color and texture. As a result, a spurious correlation is formed between modality informa-

tion and image similarity, negatively impacting the distance measure which is vital to clustering quality.

To this end, we propose bias-mitigated cross-modality association by rectifying the pairwise distance of cross-modality images. For image clustering in typical unsupervised re-ID, Jaccard distance (2017) is often adopted as the distance measure given its consideration of neighbor relations. In cross-modality scenario however, due to modality bias, the nearest neighbors are dominated by intra-modality images, as observed in Figure 1(a). The imbalanced neighbor composition only strengthens the interaction of intra-modality neighbors, but ignores the cross-modality ones, making Jaccard distance weak in cross-modality association. Next, we revisit the key steps of Jaccard distance, then introduce our proposal for bias mitigation.

**Review of Jaccard distance.** Nearest neighbors play a crucial role in Jaccard distance computation. Given the features of all training images,  $k_1$ -nearest neighbors  $N(x_i, k_1)$  are first retrieved for each query using Euclidean or Cosine distance. These KNNs are then involved in two key steps of Jaccard distance, *i.e.*  $K$ -reciprocal nearest neighbor retrieval (KRNN) and local query expansion (LQE).

(1)  **$K$ -Reciprocal Nearest Neighbor Retrieval:** For sample  $x_i$ , its  $k$ -reciprocal nearest neighbor set is computed as:  $R(x_i, k_1) = \{x_j | x_j \in N(x_i, k_1) \wedge x_i \in N(x_j, k_1)\}$ . To increase the recall of positive samples, the reciprocal neighbors are expanded by adding the  $\frac{1}{2}k_1$ -reciprocal nearest neighbors of each candidate in  $R(x_i, k_1)$ .

(2) **Local Query Expansion:** To obtain a more robust feature, this step uses the  $k_2$ -nearest neighbors to compute an averaged feature as the local expansion result.

**Bias-mitigated distance calibration.** To rectify the biased distance and ensure fair contribution of intra and inter-modality neighbors, we propose a modality-aware distance rectification. Essentially, we consider a balanced number of modality-specific neighbors, and incorporate them into Jaccard distance. We do so by altering the following two basic steps of Jaccard distance computation:

(1) **Initial  $K$ -nearest neighbor rectification.** We independently search intra- and inter-modality  $\frac{1}{2}k_1$ -nearest neighbors. Denote  $N^{intra}(x_i, k)$  and  $N^{inter}(x_i, k)$  as the intra- and inter-modality  $k$ -nearest neighbors, respectively. The rectified  $k_1$ -nearest neighbor set is represented as:

$$N^*(x_i, k_1) = N^{intra}(x_i, k_1/2) \cup N^{inter}(x_i, k_1/2), \quad (4)$$

and then sorted according to the query-to-neighbor distance:

$$\bar{N}^*(x_i, k_1) = \text{Sort}(N^*(x_i, k_1)). \quad (5)$$

The rectified neighbors  $\bar{N}^*(x_i, k_1)$  are then adopted to retrieve  $k$ -reciprocal nearest neighbors.

(2) **Balanced local query expansion.** Instead of averaging over the distance distribution of the  $k_2$ -nearest neighbors, we compute the expanded distance using the union of intra- and inter-modality  $\frac{1}{2}k_2$ -nearest neighbors:

$$\overline{Dist}(x_i) = \frac{1}{k_2} \sum_{j \in N^*(x_i, k_2)} \text{Dist}(x_j). \quad (6)$$

**Discussion.** A few recent methods (Yang, Hu, and Hu 2025; Shi et al. 2024a) also consider an improved Jaccard distance for global clustering, but with limited gains. The key difference between their method and ours is that they only consider the modality imbalance in local query expansion. In contrast, our modality-balanced calibration strategy directly mitigates the biased KNN, achieving more thorough rectification of cross-modality distance. The distance distribution in Figure 3 further validates the effect of our method. Besides, when camera label is available, our method can be complementary to CA-Jaccard (Chen et al. 2024), allowing jointly handling modality and camera biases.

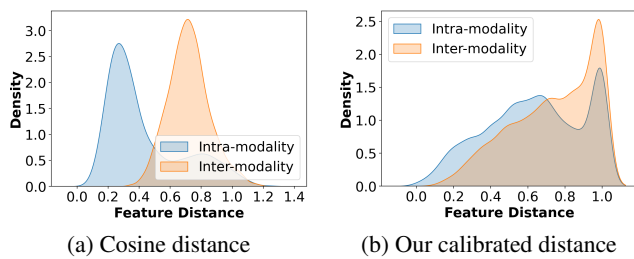


Figure 3: Intra-cluster feature distance distribution on SYSU-MM01. Features are extracted with intra-modality trained model.

## Modality-Invariant Representation Learning

Unlike intra-modality clusters, the global clusters contain a mixture of visible and infrared images, therefore demonstrate a higher level of intra-cluster variance. As observed in Figure 1(b), images of different modalities form distinct

cliques in feature space. Even when associated into the same cluster, the image features may still exhibit a “two-peak” cosine distance distribution (Figure 3(a)).

**Modality-aware global prototypes.** To preserve and model the intra-cluster characteristic, we leverage the available modality label as a natural prior to design modality-aware prototypes. Specifically, we propose a *split-and-contrast* strategy to capture the fine-grained cluster distribution: Each global cluster is split into modality-specific sub-clusters, and a fine-grained prototype is associated to each subcluster. The prototypes are initialized as the subcluster mean feature, stored in a memory bank  $\mathcal{K} \in R^{C \times d}$ , and updated similarly as Eq. (1). In this way, clusters containing mixed-modality images will be represented with two modality-specific global prototypes, preserving the variance of global clusters.

**Modality-invariant contrastive learning.** In a similar spirit to Invariant Risk Minimization (Arjovsky et al. 2019), we would like to reduce the response variance of different modalities, so as to achieve invariance to modality bias. To this end, we propose a unified modality-invariant loss to jointly optimize the similarity relationship between query image and the modality-specific global prototypes.

Our proposed loss generalizes InfoNCE to multi-positive setting. For a query visible image from a mixed global cluster, there exist two positive prototypes, one visible and one infrared. Using a multi-positive contrastive loss as below, we ensure the query feature to be equally close to its infrared and visible positive prototypes, thus explicitly aligning the modality-specific features for invariance learning:

$$\mathcal{L}_{glob}^v = - \sum_{i=1}^{N_b} \frac{1}{|P(z_i)|} \sum_{p \in P(z_i)} \log \frac{\exp(\mathcal{K}[p]^T f_\theta(x_i^v)/\tau)}{\sum_{j \in S(x_i^v)} \exp(\mathcal{K}[j]^T f_\theta(x_i^v)/\tau)}, \quad (7)$$

where  $z_i$  is the global cluster label of image  $x_i^v$ , and  $P(z_i)$  is the index set of prototypes associated to cluster  $z_i$ .  $S(x_i^v) = P(z_i) \cup \text{NegK}(x_i^v)$  is the index union of positive and Top- $k$  hard negative prototypes.

The global loss  $\mathcal{L}_{global}^r$  for infrared modality can be formulated likewise. Then the average of the two losses is taken as the full global loss:

$$\mathcal{L}_{global} = 0.5(\mathcal{L}_{glob}^v + \mathcal{L}_{glob}^r). \quad (8)$$

## Experiments

### Datasets and Settings

**Datasets and evaluation protocols.** Following existing methods, we perform experiments on two benchmark VI-ReID datasets: SYSU-MM01 (Wu et al. 2017) and RegDB (Nguyen et al. 2017). SYSU-MM01 is a large-scale dataset containing 491 identities captured by 6 cameras. The training set consists of 22,258 visible images and 11,909 infrared images from 395 identities. The rest 96 identities constitute the test set. RegDB contains 4,120 images from 412 identities, captured under one visible and one infrared camera. For RegDB, 10 random train/test splits are conducted, and the average accuracy of 10 splits is reported.

Settings			SYSU-MM01				RegDB			
			All Search		Indoor Search		Visible2Thermal		Thermal2Visible	
Type	Method	Venue	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP
SVI-ReID	DDAG (Ye et al. 2020b)	ECCV'20	54.8	53.0	61.0	68.0	69.4	63.5	68.1	61.8
	AGW (Ye et al. 2022)	TPAMI'21	47.5	47.7	54.2	63.0	70.1	66.4	70.5	65.9
	CAJ (Ye et al. 2021)	ICCV'21	69.9	66.9	76.3	80.4	85.0	79.1	84.8	77.8
	MPANet (Wu et al. 2021)	CVPR'21	70.6	68.2	76.7	81.0	83.7	80.9	82.8	80.7
	DART (Yang et al. 2022b)	CVPR'22	68.7	66.3	72.5	78.2	83.6	75.7	82.0	73.8
	FMCNet (Zhang et al. 2022)	CVPR'22	66.3	62.5	68.2	74.1	89.1	84.4	88.4	83.9
	LUPI (Alehdaghi et al. 2022)	ECCV'22	71.1	67.6	82.4	82.7	88.0	82.7	86.8	81.3
	DEEN (Zhang and Wang 2023)	CVPR'23	74.7	71.8	80.3	83.3	91.1	85.1	89.5	83.4
	PartMix (Kim et al. 2023)	CVPR'23	77.8	74.6	81.5	84.4	85.7	82.3	84.9	82.5
	MUN (Yu et al. 2023b)	ICCV'23	76.2	73.8	79.4	82.1	95.2	87.2	91.9	85.0
RLE (Lei et al. 2024)	NeurIPS'24	75.4	72.4	84.7	87.0	92.8	88.6	91.0	86.6	
CSDN (Yu et al. 2025)	TMM'25	75.2	71.8	82.0	85.0	89.0	84.7	88.2	82.8	
SSVI-ReID	OTLA (Wang et al. 2022)	ECCV'22	48.2	43.9	47.4	56.8	49.9	41.8	49.6	42.8
	TAA (Yang et al. 2023)	TIP'23	48.8	42.3	50.1	56.0	62.2	56.0	63.8	56.5
	DPIS (Shi et al. 2023)	ICCV'23	58.4	55.6	63.0	70.0	62.3	53.2	61.5	52.7
USVI-ReID	H2H (Liang et al. 2021)	TIP'21	30.2	29.4	-	-	23.8	18.9	-	-
	OTLA (Wang et al. 2022)	ECCV'22	29.9	27.1	29.8	38.8	32.9	29.7	32.1	28.6
	ADCA (Yang et al. 2022a)	MM'22	45.5	42.7	50.6	59.1	67.2	64.1	68.5	63.8
	NGLR (Cheng et al. 2023b)	MM'23	50.4	47.4	53.5	61.7	85.6	76.7	82.9	75.0
	MBCCM (Cheng et al. 2023a)	MM'23	53.1	48.2	55.2	62.0	83.8	77.9	82.8	76.7
	CCLNet (Chen et al. 2023)	MM'23	54.0	50.2	56.7	65.1	69.9	65.5	70.2	66.7
	PGM (Wu and Ye 2023)	CVPR'23	57.3	51.8	56.2	62.7	69.5	65.4	69.9	65.2
	GUR (2023b)	ICCV'23	61.0	57.0	64.2	69.5	73.9	70.2	75.0	69.9
	MMM (Shi et al. 2024a)	ECCV'24	61.6	57.9	64.4	70.4	89.7	80.5	85.8	77.0
	PCLHD (Shi et al. 2024b)	NeurIPS'24	64.4	58.7	69.5	74.4	84.3	80.7	82.7	78.4
	PCLHD† (Shi et al. 2024b)	NeurIPS'24	<u>65.9</u>	<u>61.8</u>	<u>70.3</u>	<u>74.9</u>	89.6	83.7	87.0	80.9
	RPNR (Yin et al. 2024)	MM'24	65.2	60.0	68.9	74.4	90.9	84.7	90.1	<u>83.2</u>
Ours	-	<b>67.1</b>	<b>63.1</b>	<b>75.0</b>	<b>78.6</b>	<b>94.3</b>	<b>89.1</b>	<b>93.6</b>	<b>88.5</b>	

Table 1: Comparison with state-of-the-art methods on SYSU-MM01 and RegDB dataset. PCLHD† denotes the combination of MMM and PCLHD. The best unsupervised performance is marked in **Bold** and the second best is marked in underline.

In line with the standard evaluation settings, Cumulative Matching Characteristics (CMC) and mean Average Precision (mAP) are utilized for accuracy measure. Specifically, Rank-1 accuracy is reported, along with mAP.

**Implementation details.** Based on ResNet-50, AGW (Ye et al. 2022) is adopted as the backbone network and initialized with ImageNet-pretrained weights. Following existing methods, the training is conducted in two stages. The first stage focuses on intra-modality learning and the second stage further adds cross-modality learning. During training, Adam optimizer is utilized, with an initial learning rate of  $3.5e-3$  and divided by 10 every 20 epochs. Each stage is trained for 50 epochs. Batch size is 128 for both visible and infrared modality. Based on the clustering results per epoch, each mini-batch randomly samples 8 pseudo identities and 16 instances per identity. DBSCAN is utilized for clustering both within and cross-modality. Following PGM (Wu and Ye 2023), the *eps* value for DBSCAN is set as 0.6 for SYSU-MM01 and 0.3 for RegDB. The memory update rate  $\mu$  is set as 0.1, and temperature  $\tau$  is 0.05. For more details and the pseudo Algorithm, please refer to the *Appendix*.

### Comparison with the State-of-the-arts

To validate the effectiveness of our proposed method, we compare with state-of-the-art unsupervised, semi-supervised and full-supervised VI-ReID methods on SYSU-MM01 and RegDB. The results are presented in Table 1.

**Comparison with USVI-ReID methods.** As shown in Ta-

ble 1, our method achieves superior accuracy compared to SoTA unsupervised methods while retaining its simplicity. On SYSU-MM01 dataset, our method improves the previous best Rank-1 Acc by 1.2% and 4.7% for *All* and *Indoor* search, respectively. On RegDB, our methods is also competitive, exceeding RPNR by 3.4% and 4.8% on average Rank-1 and mAP. The results demonstrate that the potential of global clustering has been overlooked by previous USVI-ReID methods. Bias-mitigated distance measure and invariant representation learning are both important aspects for enhancing the performance of USVI-ReID.

**Comparison with SVI-ReID and SSVI-ReID methods.** Supervised VI-ReID setting has been extensively studied over the last several years. From Table 1, we observe that USVI-ReID methods are catching up on the performance even without any annotation. Specifically, our proposed unsupervised method achieves comparable accuracy with some of the supervised VI-ReID methods, such as CAJ, DART and FMCNet. In addition, compared to the three semi-supervised VI-ReID methods listed in Table 1, our unsupervised method shows higher performance, proving the potential of learning from unlabeled data.

### Ablation Study

In Table 2, we provide the ablation study to analyze the impact of each proposed component.

**Effectiveness of bias-mitigated global clustering.** Compared to the global baseline  $\mathcal{M}2$ ,  $\mathcal{M}3$  simply replaces the

	Components			All Search		Indoor Search	
	SC	BMGC	MIRL	R1	mAP	R1	mAP
$\mathcal{M}1$	<i>Intra Baseline</i>			39.5	38.9	47.1	55.5
$\mathcal{M}2$	<i>Global Baseline</i>			54.9	51.5	62.9	68.9
$\mathcal{M}3$		✓		64.9	60.0	68.0	73.5
$\mathcal{M}4$	✓	✓		64.8	61.0	73.9	77.4
$\mathcal{M}5$	✓		✓	61.1	58.3	72.1	76.1
$\mathcal{M}6$	✓	✓	✓	<b>67.1</b>	<b>63.1</b>	<b>75.0</b>	<b>78.6</b>

Table 2: Ablation study on SYSU-MM01. ‘SC’: Random Subset Clustering on visible modality. ‘BMGC’: Bias-Mitigated Global Clustering. ‘MIRL’: Modality-Invariant Representation Learning. ‘*Intra Baseline*’: Intra-modality baseline. ‘*Global Baseline*’: Two-stage baseline with vanilla global clustering (Shi et al. 2024b).

vanilla global clustering with our bias-mitigated global clustering, and gains a 10% Rank-1 accuracy improvement in *All* search mode. The significant accuracy boost validates the effectiveness of the improved global clustering. In addition, a comparison between  $\mathcal{M}5$  and  $\mathcal{M}6$  further proves the robustness of our global clustering when combined with other components, and highlights the importance of mitigating the modality bias for reliable cross-modality association.

**Effectiveness of modality-invariant representation learning.** By comparing  $\mathcal{M}4$  with  $\mathcal{M}6$ , we observe that adding the modality-invariant representation learning brings further accuracy boost, improving the Rank-1 accuracy by 2.3% in *All* search. The results indicate that improving the cross-modality representation learning is indeed an aspect worth exploring. Our method captures the modality-induced cluster variance through a modality-invariant contrastive loss, so that better representation can be learned.

**Effectiveness of subset clustering.** By sampling on the visible modality, subset clustering mitigates the over-clustering as well as modality imbalance. From the comparison of  $\mathcal{M}3$  and  $\mathcal{M}4$ , we see that utilizing the subset sampling strategy leads to clear improvement on the *Indoor* search accuracy. This indicates that incorporating subset clustering is helpful for the recognition of basic and less challenging cases.

**Effectiveness of the full model.** Integrating all the proposed components,  $\mathcal{M}6$  produces the best performance. Specifically, the accuracy boost compared to  $\mathcal{M}1$  shows the overall effectiveness of our method for cross-modality learning. Compared to  $\mathcal{M}3$ - $\mathcal{M}5$ , the full model’s enhanced performance also proves the synergy of the proposed components.

### Parameter Analysis

In this section, we delve into the performance of the proposed method by analyzing two aspects: clustering accuracy and feature visualization.

**Analysis on clustering accuracy.** The clustering accuracy is a direct reflection of association quality. To investigate how well our proposed global clustering performs, we present in Figure 4 the cluster accuracy obtained by our method against other state-of-the-art methods. It can be observed that our method obtains a much higher ARI accuracy on both intra-modality and global clustering. The results prove our model is able to predict reliable pseudo labels for model learning.

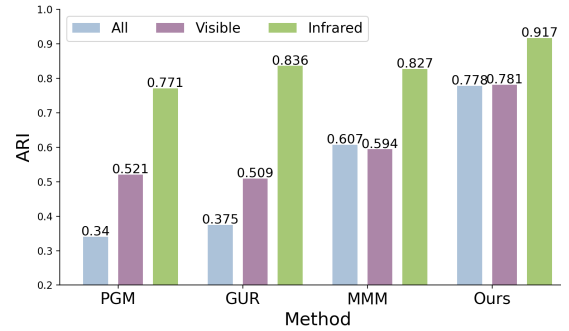


Figure 4: Comparison of clustering accuracy ARI (Adjusted Rand Index) on different methods.

**Feature visualization.** Figure 5 presents the feature visualization of our method compared to the global baseline. First, we observe that the baseline model produces multiple modality-specific clusters for the same identity, such as the purple and the green ones. This reflects that the baseline model is relatively weak at aggregating cross-modality images. As a comparison, our model generates compact clusters in most cases. With our model, samples of the purple identity are now gathered into the same cluster, so are the samples of the green identity. The visualization indicates that our model is able to enhance the intra-identity compactness while maintain the inter-identity separation.




Figure 5: T-SNE visualization of features from 10 randomly selected identities in SYSU-MM01 training set. Different color represents different identity. Circle and triangle denote visible and infrared image, respectively.

### Conclusion

In this paper, we propose a modality-aware bias mitigation and invariance learning method for USVI-ReID. To mitigate the bias in cross-modality distance measure, we propose modality-aware Jaccard distance so that reliable cross-modality association is obtained. To enhance the robustness to modality discrepancy, we design modality-specific global prototypes so that modality-invariant representation can be learned through a unified multi-positive contrastive loss. Despite our method’s simplicity, it achieves the state-of-art performance on benchmark datasets, validating its effectiveness and showing the potential for further improvements in cross-modality learning.

## Acknowledgements

This work was supported by the Zhejiang Province Pioneer Research and Development Project “Research on Multi-modal Traffic Accident Holographic Restoration and Scene Database Construction Based on Vehicle-cloud Intersection” (No. 2024C01017), General Project of Basic Science Research in Higher Education Institutions of Jiangsu Province (No. 25KJD520008).

## References

- Alehdaghi, M.; Josi, A.; Cruz, R. M. O.; and Granger, E. 2022. Visible-Infrared Person Re-Identification Using Privileged Intermediate Information. In *Proceedings of the European conference on computer vision*, 720–737.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- Chen, Y.; Fan, Z.; Chen, Z.; and Zhu, Y. 2024. CA-Jaccard: Camera-aware Jaccard Distance for Person Re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17532–17541.
- Chen, Z.; Zhang, Z.; Tan, X.; Qu, Y.; and Xie, Y. 2023. Unveiling the Power of CLIP in Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3667–3675.
- Cheng, D.; He, L.; Wang, N.; Zhang, S.; Wang, Z.; and Gao, X. 2023a. Efficient bilateral cross-modality cluster matching for unsupervised visible-infrared person reid. In *Proceedings of the 31st ACM international conference on multimedia*, 1325–1333.
- Cheng, D.; Huang, X.; Wang, N.; He, L.; Li, Z.; and Gao, X. 2023b. Unsupervised visible-infrared person reid by collaborative learning with neighbor-guided label refinement. In *Proceedings of the 31st ACM international conference on multimedia*, 7085–7093.
- Cho, Y.; Kim, W. J.; Hong, S.; and Yoon, S.-E. 2022. Part-based Pseudo Label Refinement for Unsupervised Person Re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7298–7308.
- Cuturi, M. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26.
- Dai, Z.; Wang, G.; Yuan, W.; Zhu, S.; and Tan, P. 2022. Cluster Contrast for Unsupervised Person Re-identification. In *ACCV*, 319–337.
- Fan, H.; Zheng, L.; Yan, C.; and Yang, Y. 2018. Unsupervised person re-identification: Clustering and fine-tuning. *ACM Trans. Multimedia Comput. Commun. Appl.*, 14(4).
- Ge, Y.; Chen, D.; Zhu, F.; Zhao, R.; and Li, H. 2020. Self-paced Contrastive Learning with Hybrid Memory for Domain Adaptive Object Re-ID. In *Advances in Neural Information Processing Systems*.
- Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*.
- Hong, Y.; and Yang, E. 2021. Unbiased classification through bias-contrastive and bias-balanced learning. *Advances in Neural Information Processing Systems*, 34: 26449–26461.
- Hu, Z.; Yang, B.; and Ye, M. 2024. Empowering Visible-Infrared Person Re-Identification with Large Foundation Models. In *Advances in Neural Information Processing Systems*.
- Jin, X.; He, T.; Shen, X.; Liu, T.; Wang, X.; Huang, J.; Chen, Z.; and Hua, X.-S. 2022. Meta clustering learning for large-scale unsupervised person re-identification. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2163–2172.
- Kim, M.; Kim, S.; Park, J.; Park, S.; and Sohn, K. 2023. Part-Mix: Regularization Strategy to Learn Part Discovery for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18621–18632.
- Kuhn, R. W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*.
- Lei, T.; Zhang, Y.; Han, K.; Dai, P.; Zhang, Y.; Wu, Y.; and Ji, R. 2024. RLE: A Unified Perspective of Data Augmentation for Cross-Spectral Re-identification. In *Advances in Neural Information Processing Systems*.
- Liang, W.; Wang, G.; Lai, J.; and Xie, X. 2021. Homogeneous-to-Heterogeneous: Unsupervised Learning for RGB-Infrared Person Re-Identification. *IEEE Trans. Image Process.*, 30: 6392–6407.
- Liu, J.; Sun, Y.; Zhu, F.; Pei, H.; Yang, Y.; and Li, W. 2022. Learning Memory-Augmented Unidirectional Metrics for Cross-modality Person Re-identification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19344–19353.
- Nguyen, D. T.; Hong, H. G.; Kim, K.; and Park, K. R. 2017. Person Recognition System Based on a Combination of Body Images from Visible Light and Thermal Cameras. *Sensors*, 17(3): 605.
- Shi, J.; Yin, X.; Chen, Y.; Zhang, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024a. Multi-Memory Matching for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the European conference on computer vision*.
- Shi, J.; Yin, X.; Zhang, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024b. Learning Commonality, Divergence and Variety for Unsupervised Visible-Infrared Person Re-identification. In *Advances in Neural Information Processing Systems*.
- Shi, J.; Zhang, Y.; Yin, X.; Xie, Y.; Zhang, Z.; Fan, J.; Shi, Z.; and Qu, Y. 2023. Dual Pseudo-Labels Interactive Self-Training for Semi-Supervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11218–11228.
- Tartaglione, E.; Barbano, C. A.; and Grangetto, M. 2021. End: Entangling and disentangling deep representations for bias correction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13508–13517.
- Teng, X.; Lan, L.; Chen, D.; Xu, K.; and Yin, N. 2025. Relieving Universal Label Noise for Unsupervised Visible-Infrared Person Re-Identification by Inferring from Neighbors. In *AAAI Conference on Artificial Intelligence*.
- Wad, T.; Sun, Q.; Pranata, S.; Jayashree, K.; and Zhang, H. 2022. Equivariance and invariance inductive bias for learning from insufficient data. In *European Conference on Computer Vision*, 241–258. Springer.
- Wang, J.; Zhang, Z.; Chen, M.; Zhang, Y.; Wang, C.; Sheng, B.; Qu, Y.; and Xie, Y. 2022. Optimal Transport for Label-Efficient Visible-Infrared Person Re-Identification. In *Proceedings of the European conference on computer vision*, 93–109.
- Wang, M.; Lai, B.; Huang, J.; Gong, X.; and Hua, X.-S. 2021a. Camera-aware Proxies for Unsupervised Person Re-Identification. In *AAAI Conference on Artificial Intelligence*, volume 35, 2764–2772.
- Wang, T.; Zhou, C.; Sun, Q.; and Zhang, H. 2021b. Causal attention for unbiased visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3091–3100.

- Wang, Z.; Wang, Z.; Zheng, Y.; Chuang, Y.-Y.; and Satoh, S. 2019. Learning to Reduce Dual-Level Discrepancy for Infrared-Visible Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 618–626.
- Wu, A.; Zheng, W.-S.; Yu, H.-X.; Gong, S.; and Lai, J. 2017. RGB-Infrared Cross-Modality Person Re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5390–5399.
- Wu, Q.; Dai, P.; Chen, J.; Lin, C.; Wu, Y.; Huang, F.; Zhong, B.; and Ji, R. 2021. Discover Cross-Modality Nuances for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4330–4339.
- Wu, Y.; Huang, T.; Yao, H.; Zhang, C.; Shao, Y.; Han, C.; Gao, C.; and Sang, N. 2022. Multi-centroid representation network for domain adaptive person re-id. In *AAAI Conference on Artificial Intelligence*, 2750–2758.
- Wu, Z.; and Ye, M. 2023. Unsupervised Visible-Infrared Person Re-Identification via Progressive Graph Matching and Alternate Learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9548–9558.
- Yang, B.; Chen, J.; Ma, X.; and Ye, M. 2023. Translation, Association and Augmentation: Learning Cross-Modality Re-Identification From Single-Modality Annotation. *IEEE TIP*, 32: 5099–5113.
- Yang, B.; Chen, J.; and Ye, M. 2023a. Towards Grand Unified Representation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11069–11079.
- Yang, B.; Chen, J.; and Ye, M. 2023b. Towards Grand Unified Representation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11069–11079.
- Yang, B.; Chen, J.; and Ye, M. 2024. Shallow-Deep Collaborative Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16870–16879.
- Yang, B.; Ye, M.; Chen, J.; and Wu, Z. 2022a. Augmented Dual-Contrastive Aggregation Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2843–2851.
- Yang, M.; Huang, Z.; Hu, P.; Li, T.; Lv, J.; and Peng, X. 2022b. Learning with Twin Noisy Labels for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14288–14297.
- Yang, Y.; Hu, W.; and Hu, H. 2025. Progressive Cross-Modal Association Learning for Unsupervised Visible-Infrared Person Re-Identification. *IEEE Transactions on Information Forensics and Security*, 20: 1290–1304.
- Ye, M.; Lan, X.; Leng, Q.; and Shen, J. 2020a. Cross-Modality Person Re-Identification via Modality-Aware Collaborative Ensemble Learning. *IEEE Transactions on Image Processing*, 29: 9387–9399.
- Ye, M.; Ruan, W.; Du, B.; and Shou, M. Z. 2021. Channel Augmented Joint Learning for Visible-Infrared Recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 13547–13556.
- Ye, M.; Shen, J.; Crandall, D. J.; Shao, L.; and Luo, J. 2020b. Dynamic Dual-Attentive Aggregation Learning for Visible-Infrared Person Re-identification. In *Proceedings of the European conference on computer vision*, 229–247.
- Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; and Hoi, S. C. H. 2022. Deep Learning for Person Re-Identification: A Survey and Outlook. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2872–2893.
- Ye, M.; Wu, Z.; Chen, C.; and Du, B. 2024. Channel Augmentation for Visible-Infrared Re-Identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 2299–2315.
- Yin, X.; Shi, J.; Zhang, Y.; Lu, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024. Robust Pseudo-label Learning with Neighbor Relation for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the 32th ACM International Conference on Multimedia*, 2242–2251.
- Yu, H.; Cheng, X.; Peng, W.; Liu, W.; and Zhao, G. 2023a. Modality Unifying Network for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11151–11161.
- Yu, H.; Cheng, X.; Peng, W.; Liu, W.; and Zhao, G. 2023b. Modality Unifying Network for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11185–11195.
- Yu, X.; Dong, N.; Zhu, L.; Peng, H.; and Tao, D. 2025. Clip-driven semantic discovery network for visible-infrared person re-identification. *IEEE Transactions on Multimedia*, 1–13.
- Zhang, Q.; Lai, C.; Liu, J.; Huang, N.; and Han, J. 2022. FMCNet: Feature-Level Modality Compensation for Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7339–7348.
- Zhang, Y.; and Wang, H. 2023. Diverse Embedding Expansion Network and Low-Light Cross-Modality Benchmark for Visible-Infrared Person Re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2153–2162.
- Zhong, Z.; Zheng, L.; and Li, S. 2017. Re-ranking Person Re-identification with k-Reciprocal Encoding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1318–1327.
- Zou, C.; Chen, Z.; Cui, Z.; Liu, Y.; and Zhang, C. 2023. Discrepant and Multi-instance Proxies for Unsupervised Person Re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11058–11068.