WRIM-Net: Wide-Ranging Information Mining Network for Visible-Infrared Person Re-Identification

Yonggan Wu^{1,2,3}, Ling-Chao Meng³, Yuan Zichao³, Sixian Chan^{2,4}, and Hong-Qiang Wang²*

 ¹ University of Science and Technology of China, Hefei 230026, China braverywu@mail.ustc.edu.cn
 ² Institute of Intelligent Machines/Zhongqi AI Joint Lab., HIPS, CAS, Hefei, China hqwang1260126.com
 ³ QiXinMingZhi Technology, Hefei 230088, China
 ⁴ Zhejiang University of Technology, Hangzhou 310014, China

Abstract. For the visible-infrared person re-identification (VI-ReID) task, one of the primary challenges lies in significant cross-modality discrepancy. Existing methods struggle to conduct modality-invariant information mining. They often focus solely on mining singular dimensions like spatial or channel, and overlook the extraction of specificmodality multi-dimension information. To fully mine modality-invariant information across a wide range, we introduce the Wide-Ranging Information Mining Network (WRIM-Net), which mainly comprises a Multidimension Interactive Information Mining (MIIM) module and an Auxiliary-Information-based Contrastive Learning (AICL) approach. Empowered by the proposed Global Region Interaction (GRI), MIIM comprehensively mines non-local spatial and channel information through intradimension interaction. Moreover, Thanks to the low computational complexity design, separate MIIM can be positioned in shallow layers, enabling the network to better mine specific-modality multi-dimension information. AICL, by introducing the novel Cross-Modality Key-Instance Contrastive (CMKIC) loss, effectively guides the network in extracting modality-invariant information. We conduct extensive experiments not only on the well-known SYSU-MM01 and RegDB datasets but also on the latest large-scale cross-modality LLCM dataset. The results demonstrate WRIM-Net's superiority over state-of-the-art methods.

 $\textbf{Keywords:} \ \text{Person Re-identification} \cdot \text{Wide-Ranging} \cdot \text{Cross-modality}$

1 Introduction

Person re-identification aims to match images of a person of interest in a query set with those in a gallery set, which may have been captured by different cameras [48]. Existing methods [1, 2, 18, 25, 30, 48] have achieved remarkable performance in visible-light person re-identification. However, visible-light cameras face

^{*} Corresponding author.

2 Y.G Wu et al.



Fig. 1: Example of noticeable modality discrepancy between cross-modality images. The motivation behind the proposed WRIM-Net is to mine modality-invariant information across a wide range (e.g. non-local spatial interaction, channel interaction, specific-modality, shared-modality) and to guide the network in better mining invariant information through a novel cross-modality loss.

challenges in capturing clear images under poor lighting conditions. To address this deficiency, infrared cameras are widely employed in modern surveillance systems. Based on this, the VI-ReID task has garnered increasing attention. The VI-ReID task is to match individuals in a gallery set who share the same ID but belong to the opposite modality (infrared or visible).

VI-ReID is an immensely challenging task, primarily attributed to significant modality discrepancies, as shown in Figure 1. Existing methods can be broadly categorized into two main types: feature-level methods and image-level methods. Feature-level methods aim to map cross-modality features into a common space [5,38,43,48], while image-level methods [6,28,29,33] generate diverse-mode images to reduce the modality gap. These methods demonstrate outstanding performance in VI-ReID tasks. However, they still fall short in fully mining modality-invariant information. On one hand, the feature-level methods easily overlook specific-modality multi-dimension information. On the other hand, the image-level methods encounter difficulties in extracting sufficient information due to the scarcity of image pairs.

In Figure 1, we observe significant differences in images across different modalities. Drawing inspiration from the human visual system, we prioritize the overall shape and the connections between details without focusing on very subtle details per se. For example, in the spatial dimension, our primary focus lies in the non-local spatial relationships, such as contour, posture, and proportions between different body parts (head, arms, knees, etc.). In the channel dimension, due to significant differences in color between modalities, our focus is on information such as texture, which reflects variations in color. All of these constitute modality-invariant information, an aspect where the methods mentioned above are less proficient at mining. Specifically, these methods struggle to simultane-



Fig. 2: Framework of WRIM-Net. Two separate MIIMs are inserted after each of the first two blocks of the network and a shared MIIM is inserted after each of the last two blocks of the network. Apart from separate MIIM, all other network parameters are shared. AICL uses the traditional ID loss after Block 3 of the network and CMKIC loss after Block 4 of the network.

ously explore information across multi-dimensions and multi-modalities (specific and shared modality). For example, AGW [48], DDAG [47], and MPANet [38] employed attention mechanism for extracting information on spatial or channel dimension. However, they applied attention mechanisms in either spatial or channel dimensions, neglecting multi-dimension interaction. Furthermore, they only consider deep layers, disregarding the significance of specific-modality information. While some methods like AGW [48] proposed a dual-stream architecture to extract specific-modality information, they merely segregate different modalities without specific-modality multi-dimension information mining.

Building upon the above issues and aiming to mine modality-invariant information across a wide range, we propose the Multi-dimension Interactive Information Mining (MIIM) module, as illustrated in Figure 3. Enhanced by the well designed Global Region Interaction (GRI) module, MIIM establishes longrange dependencies through global spatial interactions, effectively extracting non-local spatial-related information such as posture and shape, and simultaneously extracts channel-related information like texture through channel interactions. Furthermore, by being placed in both shallow and deep layers, MIIM captures specific-modality and shared-modality multi-dimension information.

To address the significant modality gap in VI-ReID tasks, guiding the network in better mining modality-invariant information is crucial. [27,31] suggest that contrastive loss aids the network in mining invariant information by attracting positive pairs and separating negative samples. Building upon the aforementioned considerations, We introduce an Auxiliary-Information-based Contrastive Learning (AICL) approach, depicted in the right portion of Figure 2, which devises a Cross-Modality Key-Instance Contrastive (CMKIC) loss. The core of AICL lies in employing the CMKIC loss to guide the network in mining modalityinvariant information. Additionally, an extra layer of auxiliary information is utilized to further enhance the mining of modality-invariant information.

In summary, our main contributions are summarized below:

- 4 Y.G Wu et al.
- We present the Wide-Ranging Information Mining Network (WRIM-Net). It enables deep information mining across a broad range, encompassing two aspects: (1) Multiple dimensions: spatial and channel-dimension. (2) Multiple modalities: specific-modality and shared-modality.
- A plug-and-play Multi-dimension Interactive Information Mining (MIIM) module is designed. It effectively establishes non-local spatial and channel interactions to better mine modality-invariant information.
- An Auxiliary-Information-based Contrastive Learning (AICL) approach is devised, employing the novel Cross-Modality Key-Instance Contrastive (CM-KIC) loss. AICL leverages auxiliary information and CMKIC loss to more effectively capture modality-invariant information.
- Extensive experiments conducted on three benchmark datasets demonstrate the superiority of our model over previous methods. WRIM-Net achieved the best performance of almost all the metrics on all three benchmarks.

2 Related Work

2.1 Visible-Infrared Person Re-Identification

VI-ReID is a challenging task due to the modality differences between infrared and visible images. Existing image-level methods fuse modalities at the data level to alleviate the modality discrepancy [7,28,29], while feature-level methods map visible (VIS) and infrared (IR) features into a common space [5, 38, 43, 48]. Often, image-level methods suffer from a lack of high-quality VIS-IR image pairs, leading to noise disruption. Feature-level methods are limited due to insufficient modality-invariant information mining. To better capture specific-modality and shared-modality information, the dual-stream network architecture is commonly employed for VI-ReID tasks [16, 44, 45, 48], where the parameters are separated in the shallow layers to obtain specific-modality information and shared in the deep layers to obtain shared-modality information. However, they merely segregate images from different modalities without multi-dimension information interaction. Our proposed WRIM-Net leverages non-local spatial and channel interactions to extract richer long-dependencies information. Additionally, by strategically placing MIIM at different positions, it effectively explores both specific-modality and shared-modality multi-dimension information.

2.2 Attention Mechanisms

Attention mechanisms have found widespread application in visual tasks to promote neural network performance by improving visual representation [12, 26, 32, 35]. In the field of person re-identification, attention mechanisms have been incorporated as demonstrated in [2, 32, 40, 48]. To capture long-range spatial relationships, [40] introduced second-order non-local attention modules; [48] introduced the non-local [32] mechanism for attentional feature extraction within the last two blocks of the network. DDAG [47] extracted potent features for cross-modality person re-identification by employing spatial attention modeling on the local features of pedestrians. However, most of the methods primarily employed attention mechanisms in the deep layers of the network, leading to the neglect of mining specific-modality multi-dimension information. The MIIM we propose here effectively mitigates this issue. MIIM achieves non-local spatial interactions through spatial compression and Global Region Interaction (GRI), and channel interactions via channel compression and expansion. The spatial and channel compression of MIIM effectively reduces computational complexity, enabling MIIM to be positioned in shallower layers for specific-modality multidimension information mining.

2.3 Contrastive Learning

Contrastive learning plays an important role in self-supervised learning [3, 10, 39.49, allowing the neural network to extract invariant features efficiently. Contrastive learning also has a significant impact on image-text multi-modalities learning [15, 22, 51]. [22] used contrastive learning to perform modality alignment of images and text for cross-modality retrieval. [15,51] employed contrastive learning initially to achieve modality alignment and shared networks for crossmodality feature learning. [41] employed joint contrastive learning to acquire color-invariant features in unsupervised VI-ReID tasks. [14] introduced a contrastive regularization loss to regularize the model, where positive samples for this contrastive loss involve mixing part descriptors with the same identity. Unlike the previous methods, our CMKIC loss neither employs additional methods to create data-augmented positive samples nor contrasts a single positive for each anchor. Instead, our CMKIC loss introduces a novel approach for selecting positive samples, specifically choosing top-K samples with the same ID as the anchor, different modalities, and the least similarity. This approach increases task difficulty, enabling the network to mine modality-invariant information.

3 Methodology

In this section, we provide a detailed description of the proposed method for the VI-ReID task. Our network uses a pre-trained single-stream network (ResNet50 [11]) to extract visible and infrared features. As shown in Figure 2, the proposed method consists of two main components. (1) **MIIM module**. MIIM enhances its non-local spatial interaction capability by incorporating the spatial compression and GRI module. (2) **AICL approach**. AICL utilizes the CMKIC loss to effectively guide the network in learning modality-invariant information and incorporates auxiliary information to further enhance its extraction.

3.1 Multi-dimension Interactive Information Mining Module

MIIM improves information extraction in several ways: (1) Employing global spatial interactions and establishing long-range dependencies to capture non-local spatial-related information, such as shape, pose, and object proportions.





Fig. 3: Diagram of MIIM module. The input features first pass through a standard batch normalization layer and then pass through the Spatial-Channel Compress (SCC) component to compress the size. Subsequently, the feature is passed to the Global Region Interaction (GRI) component, which employs Multi Head Attention (MHA). Finally, the feature weights are restored to the same size as the input features through the Spatial-Channel Restore (SCR) component. \otimes denotes element-wise multiplication.

(2) Employing channel interactions to capture channel-related information, such as texture. (3) Employing the separate MIIM in the shallow layers to extract specific-modality multi-dimension information (often overlooked by many existing methods) and the shared MIIM in the deep layers to extract shared-modality multi-dimension information. The MIIM module primarily comprises three components: Spatial-Channel Compression (SCC), Global Region Interaction (GRI), and Spatial-Channel Restore (SCR). The input features first pass through a standard Batch Normalization (BN) layer and then proceed to the SCC. In SCC, features undergo two compressions. The first compression is done with a convolutional layer, and the resulting features serve as the Query (Q) for subsequent GRI. The second compression is executed by employing average pooling on the aforementioned features. Then, the resulting features serve as the Key (K) and Value (V) components of GRI. This approach not only achieves non-local spatial interaction but also further reduces computational complexity.

Specifically, the input features $\mathbf{F}_1 \in \mathbb{R}^{H_1 \times W_1 \times C_1}$ pass through a convolutional layer with a kernel size and stride of r_s , resulting in $\mathbf{F}_2 \in \mathbb{R}^{H_2 \times W_2 \times C_2}$. \mathbf{F}_2 undergoes average pooling with kernel size and stride of k_s , yielding $\mathbf{F}_3 \in \mathbb{R}^{H_3 \times W_3 \times C_2}$. The resulting \mathbf{F}_2 and \mathbf{F}_3 features are flattened and then fed into the Global Region Interaction (GRI) component, which employs Multi Head Attention (MHA) [26]. In this setup, the flattened \mathbf{F}_2 functions as the Query (Q), while flattened \mathbf{F}_3 serves as both the Key (K) and Value (V). The resultant features, denoted as \mathbf{F}_4 , are obtained from the GRI module's operation, as indicated by the following formula:

$$\boldsymbol{F}_{4} = \text{GRI}(Q = \boldsymbol{F}_{2}, K = \boldsymbol{F}_{3}, V = \boldsymbol{F}_{3}).$$
(1)

The specific implementation of MHA is detailed in [26], where the number of heads is set to 8. To better leverage the positional information, we conduct sin-cos positional embedding [26]. In the processing of the GRI component, we utilize higher-resolution F_2 as the Query (Q), while lower-resolution F_3 serves as the Key (K) and Value (V), thereby shifting the focus of spatial interaction from interactions among all spatial points to a greater emphasis on long-range spatial regions, enabling global spatial interaction. After passing through the GRI component, the spatial resolution of F_4 remains the same as that of F_2 . This design further reduces the computational complexity, enabling the placement of this module in the shallower layers of networks.

Next, C_2 is restored to the original number of C_1 channels using a linear layer W, similar to Senet [12], resulting in feature F_5 . Subsequently, a sigmoid activation is applied to the F_5 to generate weights M_1 :

$$M_1 = \text{sigmoid}(W(F_5)). \tag{2}$$

 M_1 is then reshaped and upsampled to recover the original spatial scale, resulting in the formation of weights $M_2 \in \mathbb{R}^{H_1 \times W_1 \times C_1}$. Finally, the input features F_1 are element-wise multiplied with M_2 and then passed through the ReLU and BN layers to produce the output feature F_6 :

$$\boldsymbol{F_6} = \mathrm{BN}(\mathrm{ReLU}(\boldsymbol{F_1} \otimes \boldsymbol{M_2})). \tag{3}$$

In our implementation, we significantly reduce the complexity of the module by compressing the channels and spatial dimensions. In the shallow layers of the network, the features have a larger spatial size, so we primarily focus on spatial dimension compression. Through judicious compression, we not only establish more effective long-range spatial dependencies but also enable MIIM to be conveniently implemented on shallow specific modalities. In the deeper layers of the network, characterized by longer channel lengths, our primary focus is on channel compression. Compressing both ensures that the network maintains a lightweight level of complexity during the following GRI operations.

To enable MIIM to better assist the network in information mining, we employ separate MIIM for each modality after Block1 and Block2 to perform specific-modality multi-dimension information mining, as shown in Figure 2. Following Block3 and Block4, a shared MIIM is applied, which allows the network to perform the mining of shared-modality multi-dimension information.

3.2 Auxiliary-Information-based Contrastive Learning Approach

In this subsection, we introduce the proposed AICL approach, as illustrated in Figure 2. The approach involves employing our designed CMKIC loss behind the block4 layer of the network to guide the learning of modality-invariant features. Subsequently, auxiliary information of the block3 layer is leveraged to enhance the mining of modality-invariant information. Our CMKIC loss is different from common contrastive loss. It doesn't solely use one augmentation of the anchor as the positive sample. Instead, it selects the top-K least similar samples (Key-Instance) with the same ID but different modalities as positive samples. As the anchor and positive samples are from different modalities, this approach facilitates the network's exploration of crucial modality-invariant information.

8 Y.G Wu et al.

Following the approach used in PCB [25] and MGN [30], we partition P_5 , the features extracted by Block4, followed by a global average pooling layer, to obtain local features R_1 to R_N . Also, another global average pooling layer is directly applied to P_5 to obtain the global feature R_g . R_g is then passed through a MLP head [3,4] to obtain Z_5 , which is formulated as:

$$\mathbf{Z}_5 = \mathbf{W}_2(\text{ReLU}(\text{BN}(\mathbf{W}_1(\mathbf{R}_g)))), \tag{4}$$

where W_1 and W_2 represent a linear layer. Then, CMKIC loss is applied to \mathbb{Z}_5 .

CMKIC loss serves as the core component enabling AICL to fulfill its mission. The design of our CMKIC loss for WRIM-Net is as follows:

$$\mathcal{L}_{CMKIC_P_5_VI} = \sum_{i \in I_{vis}} \frac{-1}{\text{top-K}} \sum_{p \in P_{infra}(i)} \log \frac{\exp(z_i \cdot z_p/\tau)}{\sum_{a \in A(i)} \exp(z_i \cdot z_a/\tau)}, \quad (5)$$

The VI in $\mathcal{L}_{CMKIC_P_5_VI}$ indicates that the anchor represents the Visible feature, while the other samples involved in the loss correspond to the Infrared light features. I_{vis} denotes the set of visible samples in the current training batch. $P_{infra}(i)$ represents the set of the top-K least similar infrared samples with the same ID as visible sample *i*. z is the feature obtained after applying L_2 normalization to the feature z_5 from Figure 2. z_p refers to the feature zcorresponding to a sample extracted from the set $P_{infra}(i)$, while z_a denotes the feature z corresponding to a sample extracted from the set A(i). Here, A(i)represents the set of all samples in the infrared modality with different ID from i, as well as all samples in $P_{infra}(i)$. Similarly, for the case where the anchor is the infrared feature, the corresponding CMKIC loss is $\mathcal{L}_{CMKIC_P_5_IV}$ (See supplementary materials for the formula). And the final $\mathcal{L}_{CMKIC_P_5}$ is:

$$\mathcal{L}_{CMKIC_P_5} = \frac{1}{2} (\mathcal{L}_{CMKIC_P_5_IV} + \mathcal{L}_{CMKIC_P_5_VI}).$$
(6)

In addition, the global and local features are also jointly optimized using a cross-entropy classification loss as:

$$\mathcal{L}_{cls_P_5} = \frac{\left(\mathrm{E}(-\log p(\boldsymbol{R}_{\boldsymbol{g}})) + \sum_{i=1}^{N} \mathrm{E}(-\log p(\boldsymbol{R}_{i})) \right)}{N+1}, \tag{7}$$

where p() is the probability that the feature is correctly predicted by the classifier and E represents the expectation.

On the other hand, to fully leverage the auxiliary information from the additional layer, we use the ID loss to guide the features generated by Block3, P_4 . Similar to those for P_5 , we partition P_4 and obtain global features denoted as Q_g and local features denoted as Q_1 to Q_M . We apply the triplet loss [23] to the global features Q_g and cross-entropy loss for both the global features Q_g and the local features. Finally, the ID loss is written as:

$$\mathcal{L}_{ID_P_4} = \mathcal{L}_{cls_P_4} + \mathcal{L}_{tri_P_4}.$$
(8)

The loss function $\mathcal{L}_{cls_P_4}$ is similar to $\mathcal{L}_{cls_P_5}$. The detailed information for $\mathcal{L}_{cls_P_4}$ and $\mathcal{L}_{tri_P_4}$ can be found in supplementary materials.

3.3 Training and Inference

During training, we consider the three losses, CMKIC, the cross-entropy classification loss, and the ID loss, and have the total loss as a weighted sum:

$$\mathcal{L}_{total} = \mathcal{L}_{cls \ P_5} + \lambda_1 \mathcal{L}_{CMKIC \ P_5} + \lambda_2 \mathcal{L}_{ID \ P_4}, \tag{9}$$

where the hyper-parameters λ_1 and λ_2 are used to balance the contribution of each loss function. During the testing phase, we combine R_g , R_1 to R_N , Q_g and Q_1 to Q_M to form the final features for inference.

4 Experiments

4.1 Datasets and Evaluation Setting

SYSU-MM01 [36] dataset comprises 491 identities captured by four visible (VIS) cameras and two infrared (IR) cameras, including two modes: All-Search and Indoor-Search. The training set comprises 22,258 visible images and 11,909 infrared images of 395 unique individuals. The test set contains 96 IDs and uses 3,803 infrared images for the query set.

RegDB [20] dataset consists of 412 pedestrians, each of which has 10 visiblelight images and 10 infrared images. It randomly divides images of 412 individuals into two equal parts to create the training and test sets. RegDB has two test modes, including the VIS2IR setting and the IR2VIS setting.

LLCM [53] dataset is the latest VI-ReID benchmark, which comprises 30,921 images from 713 unique identities. The testing set includes 13, 909 images from 351 different identities. Both VIS2IR and IR2VIS modes are employed for evaluating the performance of the VI-ReID. Compared to RegDB and SYSU-MM01, LLCM poses a greater challenge, primarily in three aspects: (1) more images from various perspectives. (2) complex low-light conditions. (3) a longer time span. These collectively make the VI-ReID task considerably more challenging.

Evaluation metrics: To make fair comparisons, all experiments in this study use two commonly used metrics to evaluate the performance: Rank-1 accuracy and mean average precision (mAP).

4.2 Implementation details

We chose the pre-trained ResNet50 as the backbone. The BNNeck [18] was used in the classification head, and the features of the last layer were split according to the ideas of PCB [25] and MGN [30]. For SYSU-MM01 and LLCM, the number of split local features is N = 2 and M = 0, respectively. For RegDB, N = 6 and M = 2. The image size is resized to 384×144 . The remaining implementation details can be found in the supplementary materials.

10 Y.G Wu et al.

Table 1: Comparison of Rank-1 (%) and mAP (%) performances with the stateof-the-art methods on SYSU-MM01 and RegDB. R1 is the abbreviation of Rank-1. Red denotes the optimal value under the current metric, while blue represents the suboptimal value.

	SYSU-MM01							Reg	DB			
	All-Search			Indoor-Search			VICOID		IBOMIC			
Methods	Singl	e-shot	Mult	i-shot	Singl	e-shot	Mult	i-shot	VIL	52110	1112	V10
	R1	mAP	R1	mAP	R1	mAP	R1	mAP	R1	mAP	R1	mAP
LbA [21]	55.4	57.4	57.4	59.1	-	-	-	-	74.2	67.6	72.4	65.5
HCT [16]	61.7	57.5	-	-	63.4	68.2	76.5	65.1	91.1	83.3	89.3	81.5
CM-NAS	62.0	60.0	68.7	53.5	67.0	73.0	-	-	84.5	80.3	82.6	78.3
MCLNet [9]	65.4	62.0	-	-	72.6	76.6	79.6	66.6	80.3	73.1	75.9	69.5
SMCL [34]	67.4	61.8	72.2	54.9	68.8	75.6			83.9	79.8	83.1	78.6
MPANet [38]	70.6	68.2	75.6	62.9	76.7	81.0	84.2	75.1	83.7	80.9	82.8	80.7
MAUM [17]	71.7	68.8	-	-	77.0	81.9	-	-	87.9	85.1	87.0	84.3
CMT [13]	71.9	68.6	80.2	63.1	76.9	79.9	84.9	74.1	95.2	87.3	92.0	84.5
MSCLNet [52]	77.0	71.6	-	-	78.5	81.2	-	-	84.2	81.0	83.9	78.3
SGIEL [8]	77.1	72.3	-	-	82.1	83.0	-	-	92.2	86.6	91.1	85.2
PartMix [14]	77.8	74.6	80.5	69.8	81.5	84.4	88.0	80.0	85.7	82.3	84.9	82.5
DEEN [53]	74.7	71.8	-	-	80.3	83.3	-	-	91.1	85.1	89.5	83.4
MUN [50]	76.2	73.8	-	-	79.4	82.1	-	-	95.2	87.2	91.9	85.0
CAL [37]	74.7	71.7	77.1	64.9	79.7	83.7	87.0	78.5	94.5	88.7	93.6	87.6
WRIM-Net	77.4	75.4	83.2	71.1	86.2	88.1	92.1	84.6	94.5	90.5	93.7	89.7

4.3 Comparison with state-of-the-art methods

SYSU-MM01 and RegDB: Table 1 compares our method with SOTA methods. From Table 1, it can be observed that our model achieved the best results by almost all the metrics on both datasets, with only a slight gap in the Rank-1 metric for one mode in RegDB and SYSU-MM01. Specifically, on the SYSU-MM01 dataset, WRIM-Net achieves substantial improvement in the Indoor-Search scenario and Multi-shot mode.

LLCM: Table 2 compares WRIM-Net and previous methods on the LLCM dataset. The results show that our model consistently outperforms the state-of-the-art, with all metrics significantly exceeding it.

4.4 Ablation Study

The WRIM-Net consists two components: the MIIM module and the AICL approach. In this section, we perform a detailed ablation study to evaluate the role of each component. The experiment was done on the LLCM dataset. The results are shown in Table 3.

Baseline. The baseline method uses ResNet50 as the backbone network, followed by the BNNeck [18], and finally a fully connected layer as the classifier. The baseline employs both global and local features, guiding P_5 solely through cross-entropy loss, without auxiliary information of P_4 .

Table	2:	Compa	rison	of	Rank-1	(%)	and
mAP (%),	with the	state-	of-	the-art	on LL	CM.

		LLCM				
Method	IR2	2VIS	VIS2IR			
	R1	mAP	R1	mAP		
DDAG [47]	41.0	49.6	48.5	53.0		
LbA [21]	44.6	53.8	50.8	55.9		
AGW [48]	46.4	54.8	56.0	59.1		
CAJ [46]	48.8	56.6	56.5	59.8		
DART [42]	52.2	59.8	60.4	63.2		
MMN [54]	52.5	58.9	59.9	62.7		
DEEN [53]	54.9	62.9	62.5	65.8		
WRIM-Net	58.4	64.8	67.0	69.2		

Table 3: Evaluation of different components of the proposed method on LLCM in terms of Rank-1 (%) and mAP (%). The mode is IR2VIS.

Baseline	MIIM	AICL	Rank-1	mAP
\checkmark	×	×	50.42	57.05
\checkmark	\checkmark	×	55.31	61.56
\checkmark	×	\checkmark	54.83	61.27
\checkmark	\checkmark	\checkmark	58.38	64.75

Effectiveness of MIIM and AICL. As shown in Table 3, the results are significantly improved with the addition of MIIM compared to the baseline. This implies that the MIIM module effectively mines modality-invariant information. Table 3 also reveals significant improvements in the model's capabilities with the inclusion of AICL. Specifically, introducing MIIM alone results in a 4.89% increase in Rank-1 and a 4.51% increase in mAP. Introducing AICL alone leads to a Rank-1 increase of 4.41% and an mAP increase of 4.22%. When both MIIM and AICL are incorporated into the model, the Rank-1 further rises by 7.96%, and the mAP increases by 7.7%. For further details on the ablation of fusion modes within the MIIM's internal SCR module and auxiliary information in AICL, please refer to the supplementary materials.

Impact of Different Configurations of MIIM. The results of different configurations of MIIM are presented in Table 4. From Table 4, it can be observed that when employing separate MIIM after Block 1 and 2 (Block1/2) alone, the Rank-1 accuracy increases by 2.99%, and the mAP increases by 2.76% compared to the configuration without MIIM. Furthermore, utilizing only shared MIIM after Block 3 and 4 (Block3/4) yields a 1.33% increase in Rank-1 and a 1.45% improvement in mAP. It is evident that placing MIIM in shallower layers of the network for specific-modality multi-dimension information mining, an aspect overlooked by existing methods, can lead to more significant performance improvements. This is attributed to the innovative and low-complexity design of MIIM. For more details on the ablation studies of MIIM configurations, please refer to the supplementary materials.

Comparison with triplet loss and common contrastive loss. To validate the superiority of the CMKIC loss, we compare it with the triplet loss and common contrastive loss (randomly selecting one positive sample). The experiments were conducted under complete configuration conditions, and the results are presented in Table 5. From the table, we can observe that our CMKIC loss significantly outperforms traditional triplet loss and contrastive loss.

Comparison with other attention mechanisms. To demonstrate the superiority of MIIM over other attention methods, we compare MIIM with

12 Y.G Wu et al.

Table 4: Impact of different placement of MIIMs on LLCM. The MIIM placed after Block1 and 2 is the separate MIIM, while the MIIM placed after Block3 and 4 is the shared MIIM.

Block1/2	Block3/4	Rank-1	mAP
×	×	54.83	61.27
\checkmark	×	57.82	64.03
×	\checkmark	56.16	62.72
\checkmark	\checkmark	58.38	64.75

Table 5: Comparison with tripletloss and common contrastive loss onLLCM. "triplet" and "contrastive"respectively representWRIM-NetreplacingCMKICloss with tripletloss and common contrastive loss.

	Rank-1	mAP
triplet	56.53	63.19
contrastive	56.79	63.39
WRIM-Net	58.38	64.75

Table 6: Comparison with other attention mechanisms on LLCM. The mode is IR2VIS.

	flops(B)	$\operatorname{params}(M)$	Rank-1	mAP
CBAM [35]	6.25	26.03	50.71	57.83
Non-Local [32]	7.27	27.20	50.66	57.25
MIIM	7.19	28.99	55.31	61.56

CBAM [35] and Non-Local [32]. CBAM conducts spatial and channel attentions within each block of the backbone network, while Non-Local focuses solely on spatial attention within specific blocks. From Table 6, we can observe that MIIM significantly outperforms CBAM and Non-Local. Here, MIIM refers to the network with the baseline augmented by MIIM. Non-Local and CBAM configurations are applied to the backbone network based on [32,35]. Our Rank-1 accuracy is 4.60% higher, and mAP is 3.73% higher compared to CBAM. Furthermore, although MIIM demonstrates comparable complexity to Non-Local, it notably outperforms Non-Local.

4.5 Parameters analysis

Analysis of parameters λ_1 and λ_2 . Setting λ_1 and λ_2 to 0.5 and 0.1 respectively gives the best results. Details can be found in the supplementary materials.

Table 7: Comparison of the results of different spatial compression ratios at each stage on LLCM. The channel compression ratios is set to 2/2/4/4.

Table 8: Comparison of the results of different channel compression ratios at each stage on LLCM. The spatial compression ratio is set to 4/2/1/1.

$r_{s1}/r_{s2}/r_{s3}/r_{s4}$	8/4/2/2	4/2/1/1	2/1/1/1	$r_{c1}/r_{c2}/r_{c3}/r_{c4}$	4/4/8/8	2/2/4/4	1/1/2/2
Rank-1	57.87	58.00	57.71	Rank-1	57.51	58.00	57.45
mAP	64.17	64.32	64.04	mAP	63.79	64.32	63.77

Analysis of spatial and channel compression ratio parameters. In this section, to enable size compression at various ratios, we configure the input size to be 384×128 . We compare and analyze the results of different spatial and channel compression ratios at each stage on the LLCM dataset, as shown in Tables 7 and 8. We denote $r_{s1}/r_{s2}/r_{s3}/r_{s4}$ as the spatial compression ratio and $r_{c1}/r_{c2}/r_{c3}/r_{c4}$ as the channel compression ratio after Block1, Block2, Block3, and Block4 in the MIIM, respectively. From Table 7 and Table 8, it is evident that the most optimal performance is attained when the spatial compression ratio is configured as 4/2/1/1 and the channel compression ratio is set to 2/2/4/4. In both spatial and channel dimensions, performance initially improves before declining as the compression ratio increases. In our view, excessively high compression ratios lead to the loss of critical information, while excessively low ratios impede effective information interaction. In terms of spatial dimensions, appropriate spatial compression enables the network to focus on global relationships, such as the overall shape of the human body.

Table 9: Comparison of results with different k_s values in the SCC on LLCM. The mode is IR2VIS.

Table 10: Comparison of results with different top-K values in CMKIC loss on LLCM, where 'W' indicates top-K utilization and 'W/O' denotes random selection. The mode is IR2VIS.

k_s	1	2	3	4
Rank-1	57.77	58.01	58.38	57.98
mAP	64.11	64.49	64.75	64.23

top-K	1 W	2 W	4 W	4 W/O	6 W	8 W
Rank-1	57.33	58.04	58.38	56.46	57.84	57.74
mAP	63.84	64.35	64.75	63.14	64.36	64.00

Analysis of the k_s in SCC. To determine the optimal k_s parameter, we conducted parameter experiments. The results, shown in Table 9, indicate that setting k_s to 3 yields the best performance. Specifically, setting k_s to 3 leads to a 0.61% increase in Rank-1 and a 0.64% increase in mAP compared to setting k_s to 1 (equivalent to no compression). This compressed design enables GRI to further facilitate global information interaction, enhancing network performance while maintaining low complexity.

Analysis of top-K in CMKIC Loss. To validate the effectiveness of top-K and determine the optimal value for CMKIC loss, experiments were conducted as shown in Table 10. The results indicate that the best performance is achieved when top-K is set to 4. In CMKIC loss, if we randomly select top-K positive samples, there is a significant decrease of 1.92% in Rank-1 (as shown in the fifth column). This indicates that our top-K Key-Instance design is effective in assisting the network in mining modality-invariant information.

4.6 Visualization Analysis

To observe the MIIM module, we utilize Grad-Cam [24] to visually analyze the features extracted by both the baseline network and the baseline network added with MIIM. Figure 4 illustrates the visualizations obtained from an infrared query image and its corresponding visible image from the gallery set. From



Fig. 4: Grad-Cam feature visualization analysis of MIIM. The second column shows the Visualization Heat Map from the baseline network, while the third column displays the Visualization Heat Map with the MIIM module integrated.



Fig. 5: T-SNE feature visualization. Each color represents an ID and the circles and triangles represent different modalities. As can be seen, AICL better alleviates the modality discrepancy and improves the discriminability.

Figure 4, it can be observed that after adding MIIM, the network effectively focuses on pedestrians themselves, whether in the visible modality or the infrared modality, while paying more attention to global information. We attribute this to MIIM's improved non-local spatial and channel interactions, along with its specific-modality multidimensional information mining.

To observe the AICL, we use T-SNE [19] to visualize and compare the features extracted by the baseline and AICL. As shown in Figure 5. After using AICL, the feature distance of different modalities of the same ID is pulled closer while the feature distance between different IDs is pushed farther. This allows the network to more effectively mine modality-invariant information.

5 Conclusion

We have introduced the WRIM-Net for VI-ReID, which emphasizes modalityinvariant information mining across a wide range. To achieve this goal, we developed the MIIM module for extracting specific-modality and shared-modality features through multi-dimension interactions. And as a plug-and-play module, it is placed at different positions of WRIM-Net to expand the range of mining information. We also designed the AICL approach, which guides the network to better explore modality-invariant information by utilizing Cross-Modality Key-Instance Contrastive loss. Extensive experiments on three standard datasets demonstrated the superiority of WRIM-Net with the best performance of almost all the metrics on all three benchmarks.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (Grant numbers 61973295, 81872276), Key Research and Development Project of Anhui Province (Grant number 201904a07020092), and Zhejiang Provincial Natural Science Foundation of China (Grant No. LY23F020023)

References

- Bai, X., Yang, M., Huang, T., Dou, Z., Yu, R., Xu, Y.: Deep-person: Learning discriminative deep features for person re-identification. Pattern Recognition 98, 107036 (2020)
- Chen, T., Ding, S., Xie, J., Yuan, Y., Chen, W., Yang, Y., Ren, Z., Wang, Z.: Abd-net: Attentive but diverse person re-identification. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8351–8361 (2019)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
- Chen, Y., Wan, L., Li, Z., Jing, Q., Sun, Z.: Neural feature search for rgb-infrared person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 587–597 (2021)
- Choi, S., Lee, S., Kim, Y., Kim, T., Kim, C.: Hi-cmd: Hierarchical cross-modality disentanglement for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10257– 10266 (2020)
- 7. Dai, P., Ji, R., Wang, H., Wu, Q., Huang, Y.: Cross-modality person reidentification with generative adversarial training. In: IJCAI. vol. 1, p. 6 (2018)
- Feng, J., Wu, A., Zheng, W.S.: Shape-erased feature learning for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22752–22761 (2023)
- Hao, X., Zhao, S., Ye, M., Shen, J.: Cross-modality person re-identification via modality confusion and center aggregation. In: Proceedings of the IEEE/CVF International conference on computer vision. pp. 16403–16412 (2021)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7132–7141 (2018)
- Jiang, K., Zhang, T., Liu, X., Qian, B., Zhang, Y., Wu, F.: Cross-modality transformer for visible-infrared person re-identification. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV. pp. 480–496. Springer (2022)

- 16 Y.G Wu et al.
- Kim, M., Kim, S., Park, J., Park, S., Sohn, K.: Partmix: 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. pp. 18621–18632 (2023)
- Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. Advances in neural information processing systems 34, 9694–9705 (2021)
- Liu, H., Tan, X., Zhou, X.: Parameter sharing exploration and hetero-center triplet loss for visible-thermal person re-identification. IEEE Transactions on Multimedia 23, 4414–4425 (2020)
- Liu, J., Sun, Y., Zhu, F., Pei, H., Yang, Y., Li, W.: Learning memory-augmented unidirectional metrics for cross-modality person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19366–19375 (2022)
- Luo, H., Jiang, W., Gu, Y., Liu, F., Liao, X., Lai, S., Gu, J.: A strong baseline and batch normalization neck for deep person re-identification. IEEE Transactions on Multimedia 22(10), 2597–2609 (2019)
- Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(11) (2008)
- Nguyen, D.T., Hong, H.G., Kim, K.W., Park, K.R.: Person recognition system based on a combination of body images from visible light and thermal cameras. Sensors 17(3), 605 (2017)
- Park, H., Lee, S., Lee, J., Ham, B.: Learning by aligning: Visible-infrared person re-identification using cross-modal correspondences. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12046–12055 (2021)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
- Sun, Y., Zheng, L., Yang, Y., Tian, Q., Wang, S.: Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In: Proceedings of the European conference on computer vision (ECCV). pp. 480–496 (2018)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- Wang, F., Liu, H.: Understanding the behaviour of contrastive loss. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2495–2504 (2021)
- Wang, G.A., Zhang, T., Yang, Y., Cheng, J., Chang, J., Liang, X., Hou, Z.G.: Cross-modality paired-images generation for rgb-infrared person re-identification. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 12144– 12151 (2020)

17

- Wang, G., Zhang, T., Cheng, J., Liu, S., Yang, Y., Hou, Z.: Rgb-infrared crossmodality person re-identification via joint pixel and feature alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3623– 3632 (2019)
- 30. Wang, G., Yuan, Y., Chen, X., Li, J., Zhou, X.: Learning discriminative features with multiple granularities for person re-identification. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 274–282 (2018)
- Wang, T., Isola, P.: Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: International Conference on Machine Learning, pp. 9929–9939. PMLR (2020)
- Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7794–7803 (2018)
- Wang, Z., Wang, Z., Zheng, Y., Chuang, Y.Y., Satoh, S.: Learning to reduce duallevel discrepancy for infrared-visible person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 618–626 (2019)
- Wei, Z., Yang, X., Wang, N., Gao, X.: Syncretic modality collaborative learning for visible infrared person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 225–234 (2021)
- Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018)
- Wu, A., Zheng, W.S., Yu, H.X., Gong, S., Lai, J.: Rgb-infrared cross-modality person re-identification. In: Proceedings of the IEEE international conference on computer vision. pp. 5380–5389 (2017)
- 37. Wu, J., Liu, H., Su, Y., Shi, W., Tang, H.: Learning concordant attention via targetaware alignment for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11122–11131 (2023)
- Wu, Q., Dai, P., Chen, J., Lin, C.W., Wu, Y., Huang, F., Zhong, B., Ji, R.: Discover cross-modality nuances for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4330–4339 (2021)
- 39. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via nonparametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018)
- Xia, B.N., Gong, Y., Zhang, Y., Poellabauer, C.: Second-order non-local attention networks for person re-identification. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3760–3769 (2019)
- Yang, B., Ye, M., Chen, J., Wu, Z.: Augmented dual-contrastive aggregation learning for unsupervised visible-infrared person re-identification. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 2843–2851 (2022)
- 42. Yang, M., Huang, Z., Hu, P., Li, T., Lv, J., Peng, X.: Learning with twin noisy labels for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14308–14317 (2022)
- Ye, M., Lan, X., Leng, Q., Shen, J.: Cross-modality person re-identification via modality-aware collaborative ensemble learning. IEEE Transactions on Image Processing 29, 9387–9399 (2020)
- 44. Ye, M., Lan, X., Li, J., Yuen, P.: Hierarchical discriminative learning for visible thermal person re-identification. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)

- 18 Y.G Wu et al.
- Ye, M., Lan, X., Wang, Z., Yuen, P.C.: Bi-directional center-constrained topranking for visible thermal person re-identification. IEEE Transactions on Information Forensics and Security 15, 407–419 (2019)
- Ye, M., Ruan, W., Du, B., Shou, M.Z.: Channel augmented joint learning for visible-infrared recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13567–13576 (2021)
- 47. Ye, M., Shen, J., J. Crandall, D., Shao, L., Luo, J.: Dynamic dual-attentive aggregation learning for visible-infrared person re-identification. In: Computer Vision– ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16. pp. 229–247. Springer (2020)
- Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., Hoi, S.C.: Deep learning for person re-identification: A survey and outlook. IEEE transactions on pattern analysis and machine intelligence 44(6), 2872–2893 (2021)
- 49. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6210–6219 (2019)
- Yu, H., Cheng, X., Peng, W., Liu, W., Zhao, G.: Modality unifying network for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11185–11195 (2023)
- Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917 (2022)
- 52. Zhang, Y., Zhao, S., Kang, Y., Shen, J.: Modality synergy complement learning with cascaded aggregation for visible-infrared person re-identification. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV. pp. 462–479. Springer (2022)
- 53. Zhang, Y., Wang, H.: Diverse embedding expansion network and low-light crossmodality benchmark for visible-infrared person re-identification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2153–2162 (2023)
- Zhang, Y., Yan, Y., Lu, Y., Wang, H.: Towards a unified middle modality learning for visible-infrared person re-identification. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 788–796 (2021)