

Improving Multispectral Pedestrian Detection by Addressing Modality Imbalance Problems

Kailai Zhou¹, Linsen Chen¹, and Xun Cao¹

Nanjing University, Nanjing, China
{calayzhou,linsen}@smail.nju.edu.cn {caoxun}@nju.edu.cn

Abstract. Multispectral pedestrian detection is capable of adapting to insufficient illumination conditions by leveraging color-thermal modalities. On the other hand, it is still lacking of in-depth insights on how to fuse the two modalities effectively. Compared with traditional pedestrian detection, we find multispectral pedestrian detection suffers from modality imbalance problems which will hinder the optimization process of dual-modality network and depress the performance of detector. Inspired by this observation, we propose Modality Balance Network (MBNet) which facilitates the optimization process in a much more flexible and balanced manner. Firstly, we design a novel Differential Modality Aware Fusion (DMAF) module to make the two modalities complement each other. Secondly, an illumination aware feature alignment module selects complementary features according to the illumination conditions and aligns the two modality features adaptively. Extensive experimental results demonstrate MBNet outperforms the state-of-the-arts on both the challenging KAIST and CVC-14 multispectral pedestrian datasets in terms of the accuracy and the computational efficiency. Code is available at <https://github.com/CalayZhou/MBNet>.

Keywords: Multispectral pedestrian detection · Modality imbalance problems · Multimodal feature fusion

1 Introduction

Recent years have witnessed increasing researches towards object detection among vision community by taking the advantages of multi-modal inputs, such as RGB + thermal, RGB + depth, RGB + LiDAR and so on [19,17,15,1]. Compared than traditional single-modal RGB images, which present great challenges at complex scenarios (e.g. dim environment, face spoofing detection [47], autonomous driving [24,39], etc), the introducing of another modality dramatically benefits the tasks of object detection. For instances, spectral images are able to detect the optical radiation of matter and reveal the essential color properties of target object, to avoid the metamerism ambiguity. Thermal images can be captured based on the heat radiation difference of the object, which does not rely on external light sources. Time-of-flight (TOF) or LiDAR sensors provide additional depth information of the target scene, which has been widely used




Fig. 1. The modality imbalance problems which consist of two parts: the illumination modality imbalance problem and the feature modality imbalance problem.

as data representation for many vision applications. Even with these remarkable benefits, however, how to effectively fuse multi-modal information in the context of advanced algorithms, like convolutional neural network, still remains much to be studied.

As for ordinary optimization process of object detection from multi-modality inputs, the imbalance problems [33] are crucial. The most known imbalance problem is the foreground-to-background imbalance [26]. This drawback is caused by an extremely inequality between the number of positive examples and negative ones. Nevertheless, the imbalance problems are not limited to the class imbalance. For instance, in multi-task losses minimization, the imbalance problems exist since the norms of gradients are different and the ranges of loss functions vary [14]. The common solution is to add coefficients upon each loss function to guide a balanced optimization process. Similarly, the modality imbalance issue in multispectral detection has a substantial influence on the algorithm performance.

The traditional Caltech [10] and CityPersons [45] pedestrian detection datasets only have RGB modality images captured during the day, so as shown in Fig. 1, modality imbalance problems existing in multispectral pedestrian detection datasets can be divided into two categories: the illumination modality imbalance and the feature modality imbalance. Illumination modality imbalance means the difference of illumination conditions between the daytime and the night images. Intuitively, pedestrians in RGB images have clearer texture features than thermal images in daytime. Comparatively, thermal images can provide more distinct pedestrian shapes than RGB images during night time. The RGB modality branch and the thermal modality branch tend to obtain different confidence scores and have uneven contributions to the object losses under diverse illumination conditions. It is expected that the RGB modality branch and the thermal modality branch should be optimized adaptively according to illumination conditions [4,23].

Feature modality imbalance problem signifies that the misalignment and inadequate integration of different modalities can lead to an uneven contribution and representation of the features. On the one hand, as the visualization results shown in Fig. 1, it is obvious that the RGB and thermal modality features are diverse in terms of pedestrian morphology, texture and properties in the two independent backbone networks. In RGB modality, the complexion and hair of the pedestrian can be some important hints of the pedestrian characteristics [6], but none of the thermal images has such cues. It is necessary to sufficiently incorporate the cross-modality complementarity to generate robust features. On the other hand, the misalignment between the RGB and thermal modalities will cause unbalanced modality feature representation in the fixed receptive fields of a convolution kernel. Both the balance and the integration of different modalities are the cornerstone we should consider in multispectral pedestrian detection. Unfortunately, existing RGB-Thermal based detection methods simply fuse the RGB and the thermal input/features by concatenation [23,42,22,37]. The inherent complementary is not fully exploited yet between different modalities.

To address the modality imbalance problems above, we investigate the impact and explore solutions in this paper. First, we construct the Modality Balance Network (MBNet) based on SSD [29] to extract the characteristics of two modalities separately. Then for the purpose of fully fusing features at different scales in the network, Differential Modality Aware Fusion (DMAF) module is proposed to tap the difference between RGB and thermal feature maps which brings more complementary information at each channel. Finally we design an illumination aware feature alignment module to align two modality features and induce the network to be optimized adaptively according to illumination conditions.

The main contributions of this paper are as follows: (1) We present modality imbalance problems specific to multispectral pedestrian detection, and analyse that modality imbalance problems will affect the performance of the detector due to the modality inconsistency in the optimization of the network; (2) We propose a one-stage detector named Modality Balance Network (MBNet) which consists of Differential Modality Aware Fusion (DMAF) module and illumination aware feature alignment (IAFA) module to address the modality imbalance problems. With DMAF module and IAFA module, the contribution of each feature map from two modalities will be explicitly integrated and balanced. In addition, The MBNet backbone (ResNet embedded with DMAF) may also do a favor to other computer vision communities; (3) MBNet achieves state-of-the-art results on both the challenging KAIST and CVC-14 multispectral pedestrian datasets in terms of the accuracy while maintaining the fastest speed.

2 Relate Work

2.1 Multispectral Pedestrian Detection

CNN-based pedestrian detection has achieved notable progress in recent years with methods of occlusion handling [46,32,38], cascaded detection systems [30,2], semantic attention [49,3], anchor-free approach [31], etc. Nevertheless, current

pedestrian detectors using single RGB modality may fail under the insufficient illumination condition. The KAIST multispectral pedestrian detection dataset [17] provides a new way to solve this problem by combining RGB modality and thermal modality. The initial baseline F + T + THOG is extended from Aggregated Channel Features (ACF) [9] with the thermal channel added. As the popularization of deep learning, the CNN-based methods [37,7,34,40] greatly reduce the miss rate of multispectral pedestrian detection. Inspired by [41], Boosted Decision Trees classifier [21] is built on high-resolution RPN feature maps to reduce potential false positive detections. MSDS RCNN [22] is learned by jointly optimizing pedestrian detection and semantic segmentation tasks.

How to fuse the information of two modalities is the common concerned problem in multispectral pedestrian detection. Liu et al. [27] design four distinct fusion architectures that integrate two modality branches on different DNNs stages and reveal the Halfway Fusion model provides the best performance. GFD-SSD [48] proposes two variations of novel Gated Fusion Units (GFU) that learn the combination of feature maps generated by the two SSD middle layers. Zhang et al. [42] explore the cross-modality disparity problem in multispectral pedestrian detection and propose a novel region feature alignment module to solve this problem. CIAN [43] makes the middle-level feature maps of two streams converge to a unified one under the guidance of cross-modality interactive attention and adopts the context enhancement blocks (CEBs) to further augment contextual information. Illumination-aware Faster R-CNN [23] adaptively merges color and thermal sub-networks to obtain the final confidence scores via a gate function defined over the illumination value. As the most popular solution, the two-stream architecture with concatenating RGB-Thermal feature maps has achieved significant improvements. Nevertheless, direct concatenation will inevitably introduce redundant features and a selection module is required to unveil the relation of modality complementary features.

2.2 Imbalance Problems In Object Detection

Oksuz et al. [33] present a comprehensive review of the imbalance problems in object detection and group these problems in a taxonomic tree with four main types: spatial imbalance, objective imbalance, class imbalance and scale imbalance. Spatial imbalance and objective imbalance focus on spatial properties of the bounding boxes and multiple loss functions respectively. Class imbalance occurs due to the significant inequality among different classes of training data. RetinaNet [26] addresses class imbalance by means of reshaping the standard cross entropy loss to prevent the vast number of easy negatives from overwhelming the detector. AP-Loss [5] and DR Loss [35] also provide ideas of designing loss function to solve the class imbalance problem. Scale imbalance occurs when certain sizes of the object bounding boxes are over-represented in the network. For instances, SSD [29] makes independent predictions from features at different layers. Since abstractness of information varies among different layers, it is unreliable to make predictions directly from different layers of the backbone network. Feature Pyramid Network [25] exploits an additional top-down pathway in order




Fig. 2. Overview framework of the Modality Balance Network (MBNet). The MBNet consists of three parts: feature extraction module, illumination aware feature alignment module and illumination mechanism. The feature extraction module adopts ResNet-50 [16] as the backbone network and embeds DMAF module to supplement modality information. Illumination mechanism is designed to acquire illumination values which will assign weights to two modality streams. Illumination aware feature alignment module plays the role of adapting the model to different illumination conditions and aligning the two modality features in the region proposal stage.

to have a balanced mixed of features from different scales. FPN can be further enhanced [28] by integrating and refining pyramidal feature maps.

In addition to the integration balance of different level, we argue that the integration of different modality features should also be balanced in the two-stream network. In other words, different modality features should be fully integrated and represented in order to have a balanced modality optimization in the training.

3 Approach

The overall architecture of the proposed method is shown in Fig. 2. The MBNet extends the framework of SSD [29] and it consists of three parts: feature extraction module, illumination aware feature alignment module and illumination mechanism. Details of DMAF module are introduced in Sec. 3.1. The design of illumination aware feature alignment module is introduced in Sec. 3.2.

3.1 Differential Modality Aware Fusion Module

To address feature modality imbalance problem, we propose to enhance the one modality from another modality with differential modality information. Previous RGB-T fusion models [37,22,43,42] based on deep convolutional networks typically employ a two-stream architecture, in which the RGB and thermal modalities are learned independently. The most straightforward method is to concatenate the features at different levels, e.g., early fusion, halfway fusion as well as late fusion [27,21,23]. However, it is ambiguous to capture the cross-modality complementary information by traditional direct concatenation scheme. Both modalities have their own characteristic representations which are mixed with useful hints and noises. While simple fusion strategies such as linear combination or concatenation are lacking in clarity to extract cross-modality complementary. In our view, the inherent difference between the two modalities can be exploited with an explicit and simple mechanism named Differential Modality Aware Fusion (DMAF) module.

We are inspired by differential amplifier circuits in which the common-mode signals are suppressed and the differential-mode signals are amplified. Our DMAF module retains the original features and compensates according to differential features. The RGB convolution feature map F_R and the thermal convolution feature map F_T can be represented with common modality part and differential modality part at each channel as follows:

$$\begin{aligned} F_T &= \frac{F_T + F_T}{2} + \frac{F_R - F_R}{2} = \frac{F_R + F_T}{2} + \frac{F_T - F_R}{2} \\ F_R &= \frac{F_R + F_R}{2} + \frac{F_T - F_T}{2} = \frac{F_R + F_T}{2} + \frac{F_R - F_T}{2} \end{aligned} \quad (1)$$

where the common modality part reflects the common features and the differential modality part reflects the unique features captured by two modalities. Eq. 1 illustrates the principle of splitting which is same behind differential amplifier circuits and DMAF module. The key idea of our DMAF module is acquiring complementary features from another modality with channel-wise differential weighting. We expect the learning of complementary features to be enhanced by explicitly modeling modality interdependencies, so that the network sensitivity to informative features from another modality can be increased.

In order to make sufficient use of cross-modality complements, the DMAF module is densely inserted in each ResNet block. As the top right corner of Fig. 2 shows, we obtain the differential feature F_D by direct subtraction of two modalities first. Then we squeeze global spatial information F_D into a global differential vector which contains channel-wise differential statistics with global average pooling. The global differential vector can be interpreted as a channel descriptor whose statistics are expressive for the discrepancy between RGB and thermal modality. The tanh activation function ranging from -1 to 1 is applied for the global differential vector to obtain the fusion weight vector V_w . The two modality features F_T and F_R are recalibrated by the fusion weight vector V_w with channel-wise multiplication. The recalibration results F_{RD}, F_{TD} will be




Fig. 3. Feature map visualization of one channel in stage3 (shown in Fig. 2) before and after DMAF module. The two modality feature maps are remedied with the differential information from each other.

added to the original modality path as complementary information. After the enhancement from another modality with DMAF module, the more informative and robust features are generated and sent to the next ResNet block in the following step. The whole procedure of DMAF module can be formulated as:

$$\begin{aligned} F'_T &= F_T + \mathcal{F}(F_T \oplus F_{RD}) = F_T + \mathcal{F}(F_T \oplus (\sigma(GAP(F_D)) \odot F_R)) \\ F'_R &= F_R + \mathcal{F}(F_R \oplus F_{TD}) = F_R + \mathcal{F}(F_R \oplus (\sigma(GAP(F_D)) \odot F_T)) \end{aligned} \quad (2)$$

where $\mathcal{F}(\mathcal{X})$ is considered as the residual function. σ refers to the tanh function, GAP refers to Global Average Pooling, and \oplus, \odot represent element-wise sum and element-wise multiplication respectively. It is noteworthy that the F_{RD}, F_{TD} are added to the residual branch which formulates the complementary feature learning as residual learning inspired by RFBNet [8]. With residual mapping, the complementary feature would not directly impact the modality-specific stream. The DMAF module acts as a part of residual function in the ResNet block.

The visualization result of DMAF module is illustrated in Fig. 3. Due to the differences in the characteristics of two modalities, thermal and RGB modalities have certain limitations respectively in capturing pedestrian and background features. As the CNN goes deeper, pedestrian features gradually become salient and background features are re-integrated. The integration of background features means useful background information is refined and noisy background information is eliminated as much as possible. The DMAF module which effectively combine modality features can contribute to the integration of background information and make pedestrian features prominent from low level to high level. In our opinion, the DMAF module facilitates modality interaction in the network which reduces the learning of redundancy and conveys more information (refer the detailed analysis in appendix). In terms of no extra parameters and low computational complexity, the MBNet backbone (ResNet embedded with DMAF) may do a favor to other computer vision communities such as RGB-Depth tasks, stereo image SR, RGB-LiDAR tasks, etc.




Fig. 4. The structure of illumination aware feature alignment module. Anchor Proposed (AP) stage generates an approximate location and Illumination Aware Feature Complement (IAFC) stage predicts based on the results of AP stage with the illumination aware balance of the two modality features. Modality Alignment (MA) module fixes the misalignments between the RGB modality and the thermal modality.

3.2 Illumination Aware Feature Alignment Module

Illumination Aware Feature Alignment module plays the role of adapting the model to different illumination conditions and aligning the two modality features in the region proposal stage. As the top of Fig. 2 shows, we design a tiny neural network to capture the illumination values in which only the RGB images are used because the thermal images are difficult to reflect the environment illumination condition. In order to reduce computational complexity, the RGB images are resized to 56×56 and sent into the illumination aware module which consists of two convolutional layers and three fully-connected layers. The ReLU activation function and a 2×2 maxpooling layer are followed after the convolutional layer to compress and extract features. The network is optimized by minimizing cross entropy loss function between the predicted illumination values and the true labels. The illumination loss L_I is formulated as:

$$L_I = -\hat{w}_d \cdot \log(w_d) - \hat{w}_n \cdot \log(w_n) \quad (3)$$

$$w_r = \left(\frac{w_d - w_n}{2}\right) \cdot (\alpha_w \cdot |w| + \gamma_w) + \frac{1}{2} \quad w_t = 1 - w_r$$

where w_d and w_n are the softmax output of full connection layers. \hat{w}_d and \hat{w}_n represent the true labels of the day and night. To be self-adaptable in the network, w_d, w_n are readjusted in the illumination mechanism in which $|w| \in [0, 1]$ is the independent prediction of the bias from 0.5 and α_w, γ_w are two learnable parameters initialized with 1, 0. Then the re-scaled results w_r, w_t are embedded into the network to have a balanced optimization during different illumination conditions. We tailor an illumination gate to control the weight of thermal modality stream and RGB modality stream before the Anchor Proposal (AP) stage. By element-wise multiplying with illumination value, the feature maps from two modalities have different scales after reweighting, and we use L2-normalization to rescale their norms to 10.

Considering that RGB and thermal cameras are not always captured at the same time, there are slight misalignments between RGB and thermal modality as shown in Fig. 4. In the fixed receptive field of a convolution kernel, the modality misalignments will cause the unbalanced feature representations and contributions of the two modalities. We contrive a Modality Alignment (MA) module which predicts offsets (dx, dy) for every pixel (x, y) of each modality. Channel dimensions are compressed because the rearrangement of two modality features according to the learned offsets is time-consuming. Since (dx, dy) is the float type, we adopt bilinear interpolation to obtain the final pixel value (x+dx, y+dy) from the four neighborhood pixels.

Due to the vague and sparse pedestrian distribution, we employ a cascade architecture inspired by [20,18,30]. Fig. 4 shows the cascade region proposal module which consists of two stages, Anchor Propose (AP) stage and Illumination Aware Feature Complement (IAFC) stage. First, the reweighted RGB and thermal feature maps are combined to generate an approximate location estimation by AP stage. The predicted regression offsets t_0 are used to propose the deformable anchors as the basic reference for position prediction in the next IAFC stage. Then the deformable anchors and confidence scores are further fine-tuned through the IAFC stage. The confidence scores s_r, s_t predicted by RGB and thermal feature maps are reweighted according to illumination values. The final confidence scores s_{final} and regression offsets t_{final} are computed as follows:

$$s_{final} = s_0 \times s_1 = s_0 \times (w_r \cdot s_r + w_t \cdot s_t) \quad t_{final} = t_0 + t_1 \quad (4)$$

The multiplication in confidence scores is to encourage the final score only if two-stage scores s_0, s_1 are both high. While for regression offsets, summation is adopted to progressively approach the pedestrian bounding boxes. Inspired by [26], we append the focal weight in classification loss L_{cls} to address the positive-negative imbalance. The L_{cls} is formulated as:

$$L_{cls} = -\alpha \sum_{i \in S_+} (1 - s_i)^\gamma \log(s_i) - (1 - \alpha) \sum_{i \in S_-} s_i^\gamma \log(1 - s_i) \quad (5)$$

S_+, S_- are the positive and negative anchor boxes. As suggested in [26], we experimentally set $\alpha = 0.25$ and $\gamma = 2$. s_i is the positive probability of samples i . The total loss is the sum of illumination loss L_I , classification loss L_{cls} and regression loss L_{reg} , where the regression loss L_{reg} is the smooth L1 loss raised by Faster-RCNN [36]. The total loss function L is as follows:

$$L = L_I + L_{cls0} + L_{cls1} + [y = 1]L_{reg0} + [y = 1]L_{reg1} \quad (6)$$

With the progressive detection of AP stage and IAFC stage, more positive cases are generated to benefit bounding box regression in the second IAFC stage. The adaptive illumination aware feature alignment of RGB modality and thermal modality provides a solution to feature modality imbalance problems by aligning two modality features, meanwhile, it also makes the detector more robust to the illumination variation.

Table 1. Comparisons with the state-of-the-art methods on the KAIST reasonable subset in terms of MR^{-2} [17] with different thresholds of IoU. In addition, Comparisons of running time are also provided.

Methods	MR^{-2} (IoU = 0.5)			MR^{-2} (IoU = 0.75)			Plateform	Speed(s)
	All	Day	Night	All	Day	Night		
ACF [17]	47.32	42.57	56.17	88.79	87.70	91.22	MATLAB	2.73
Halfway Fusion[27]	25.75	24.88	26.59	81.29	78.43	86.80	TITAN X	0.43
Fusion RPN+BF [21]	18.29	19.57	16.27	72.97	68.14	81.35	MATLAB	0.80
IAF R-CNN [23]	15.73	14.55	18.26	75.50	72.34	81.12	TITAN X	0.21
IATDNN + IASS[13]	14.95	14.67	15.72	76.69	76.46	77.05	TITAN X	0.25
RFA[42]	14.61	16.78	10.21	-	-	-	TITAN X	0.08
CIAN [43]	14.12	14.77	11.13	74.45	71.42	80.16	1080 Ti	0.07
MSDS-RCNN [22]	11.34	10.53	12.94	70.57	67.36	79.25	TITAN X	0.22
AR-CNN [44]	9.34	9.94	8.38	64.22	57.87	76.82	1080 Ti	0.12
MBNet(ours)	8.13	8.28	7.86	60.12	54.90	68.34	1080 Ti	0.07

4 Experiments

In this section, we first introduce the KAIST dataset [17] and CVC-14 dataset [12]. Then we show implementation details and experiment results to compare logMR and runtime of the proposed MBNet with the state-of-the-art methods. The evaluation is based on the reasonable setup (55 pixel or taller under partial or no occlusion) unless otherwise mentioned. Finally, we will carry out ablation studies for the proposed method on the KAIST dataset.

4.1 Datasets

Our approach is evaluated on the KAIST dataset [17] and CVC-14 dataset [12].

KAIST. The KAIST dataset [17] contains 95,328 aligned color-thermal image pairs, with a total of 103,128 bounding boxes covering 1,182 unique pedestrians. Due to the problematic annotations in original training data, we adopt the annotations improved by Zhang et al. [44] for training. The test set consists of 2, 252 frames sampled every 20th frame from video, among which 1,455 images are captured during daytime and the rest 797 images are during nighttime. The evaluation metric follows the standard KAIST evaluation [17]: log-average Miss Rate over False Positive Per Image (FPPI) range of $[10^{-2}, 10^0]$ (denoted as MR^{-2}). We evaluate the detection performance on the KAIST test set with annotations improved by Liu et al. [27] and report the runtime of the proposed MBNet using a single NVIDIA GTX 1080Ti GPU for fair comparison with the state-of-the-art methods before.

CVC-14. The CVC-14 dataset [12] contains visible (grayscale) and thermal paired images. It was recorded in various scenes at day and night by on-board color and thermal cameras at 10 Hz. The training and testing set contains 7, 085

Table 2. Evaluations on the KAIST dataset under all nine test subsets and ablation experiments are also provided. The lowest MR are highlighted with bold font.

Methods	Rea.	Day	Night	Near	Medium	Far	None	Partial	Heavy
ACF [17]	47.32	42.57	56.17	28.74	53.67	88.20	62.94	81.40	88.08
Halfway Fusion [27]	25.75	24.88	26.59	8.13	30.34	75.70	43.13	65.21	74.36
FusionRPN+BF [21]	18.29	19.57	16.27	0.04	30.87	88.86	47.45	56.10	72.20
IAF R-CNN [23]	15.73	14.55	18.26	0.96	25.54	77.84	40.17	48.40	69.76
IATDNN+IASS [13]	14.95	14.67	15.72	0.04	28.55	83.42	45.43	46.25	64.57
CIAN [43]	14.12	14.77	11.13	3.71	19.04	55.82	30.31	41.57	62.48
MSDS-RCNN [22]	11.63	10.60	13.73	1.29	16.19	63.73	29.86	38.71	63.37
AR-CNN [44]	9.34	9.94	8.38	0.00	16.08	69.00	31.40	38.63	55.73
MBNet (ours)									
IAFC DMAF Aligned									
	11.93	12.51	10.86	0.00	20.08	63.70	34.16	44.10	65.11
✓	10.96	11.13	10.48	0.00	20.33	61.25	33.50	39.80	62.68
✓	10.53	11.00	9.75	0.00	16.50	58.47	29.39	40.25	59.13
✓	9.36	9.72	8.63	0.00	16.18	54.66	28.02	38.19	60.70
✓	8.13	8.28	7.86	0.00	16.07	55.99	27.74	35.43	59.14

and 1, 433 frames, respectively. Annotations are individually provided in each modality since the cameras are not well calibrated.

4.2 Implementation Details

Our MBNet detector uses ResNet-50 [16] as the backbone network, which is pretrained on ImageNet unless otherwise stated. The training IoU of the AP stage is set to $\{0.3, 0.5\}$ and the IAFC stage is set to $\{0.5, 0.7\}$. The Xavier method [11] is used to randomly initialize other convolutional layers. In the training, we crop a patch with the size of $[0.3, 1]$ of the input image and resize it to 640×512 , then each image is randomly color distorted and horizontal flipped with a probability of 0.5 to increase the diversity. The whole network is trained by adam optimizer for 7 epoches with a learning rate of 0.0001 and a batch size of 10. Followed by [22], the width of initial anchors are set to $[25.84, 29.39]$, $[33.81, 38.99]$, $[44.47, 52.54]$, $[65.80, 131.40]$ for stage3 to stage6 (shown in Fig. 2) with a single anchor ratio of 0.41.

4.3 Evaluation on the KAIST Dataset

We show the superiority of our method from both aspects of miss rate (MR) and speed.

Miss Rate. Our proposed approach achieves 8.28 MR, 7.86 MR, and 8.13 MR on the reasonable day, night and all-day subset respectively under the IoU threshold of 0.5, all of them are lower than the previous best competitor AR-CNN [44]. In the case of a stricter IoU threshold of 0.75, Tab. 1 shows our proposed

method achieves about 4.10% lower on MR^{-2} which implies that the MBNet has a substantially better localization accuracy compared with AR-CNN. The larger the IoU threshold is set, the harder the predicted bounding boxes are considered to be True Positives (TP) in the evaluation. In order to have a comprehensive understanding of detector performance, we also make an evaluation under all nine subsets including the pedestrian distances and the occlusion levels. Tab. 2 shows MBNet outperforms other methods under most subsets with no extra treatment to handle the small and occlusion pedestrians, especially on the none subset (27.74 vs. 29.86 MR) and partial subset (35.43 vs. 38.63 MR).




Fig. 5. (a) Log-average miss rate versus the running time of each detector. (b) Performance comparisons with the state-of-the-art methods on the KAIST dataset under reasonable subset.

Speed. We also compare the running time of MBNet with state-of-the-art methods. MBNet directly takes 640×512 multispectral images as input without image up-scaling. Since modality alignment module is time-consuming, we draw the Speed vs. FPPI results of MBNet and MBNet without MA module respectively in Fig. 5 (a). The MBNet without MA module reaches the fastest speed of 20 fps and has a comparable performance with AR-CNN [44]. MBNet achieves the state-of-art performance on the test annotations improved by Liu et al. [27] while maintaining high computational efficiency.

Overall, the Speed vs. FPPI and performance comparisons under reasonable setting are shown in Fig. 5 (a) and (b). The result indicates that MBNet is an attractive multispectral pedestrian detector in both accuracy and speed.

4.4 Evaluation on the CVC-14 Dataset

We fine-tune from the KAIST pretrained model in the training of CVC-14 dataset. We follow the protocol in [34] to conduct the evaluation experiments and adopt the strategy in [44]. Specifically, there existing serious misalignments

Table 3. Evaluation results on the CVC-14 dataset. The first column refers to input modalities of the approach.

	Methods	MR^{-2}			Visible+ Thermal	Methods	MR^{-2}		
		Day	Night	All			Day	Night	All
Visible	SVM [12]	37.6	76.9	-		MACF [34]	61.3	48.2	60.1
	DPM [12]	25.2	76.4	-		Choi et al. [7]	49.3	43.8	47.3
	Random Forest [12]	26.6	81.2	-		Halfway Fusion [34]	38.1	34.4	37.0
	ACF [34]	65.0	83.2	71.3		Park et al. [34]	31.8	30.8	31.4
	Faster R-CNN [34]	43.2	71.4	51.9		AR-CNN [44]	24.7	18.1	22.1
					MBNet (ours)	24.7	13.5	21.1	

between thermal and RGB modalities, so we consider pedestrians in the RGB modality as the training target, and pedestrians in the thermal modality act as a reference. It can be observed from Tab. 3. that MBNet can still achieve good results even in the case of serious modality misalignments, which demonstrates that modality balance strategy improves the robustness to the position shift problem.

4.5 Ablation Study

Ablation experiments are performed on the KAIST dataset for a detailed analysis in this section. The baseline is initialized from two-stream SSD [29] and adopts simple concatenation to fuse two modalities as most previous methods do. We show how to construct the MBNet with the principle of modality balance.

Differential Modality Aware Fusion. Simple concatenation modulates the complementary part implicitly and has a large computational cost. It encourages us to seek a better integration of modalities. DMAF modules are inserted in each ResNet block and act as a part of residual function to be learned which will impel a deep integration of RGB and thermal modalities. From the ablation experiment results in Tab. 2, we could see that the MBNet generates more accurate detection results with DMAF module added. The DMAF module captures the complementary modality information in a more explicit way with no extra parameters and very little computation, so that the MBNet can maintain a fast speed. The feature representation can be more balanced after the two modalities are fully integrated.

Modality Alignment. It could be seen from the Tab. 4 that the MBNet with modality alignment module has a much lower miss rate on the reasonable subset. It indicates that the MA module with bilinear interpolation used can locate pedestrians more precisely especially in the deeper network which has a smaller resolution of the feature maps. After the MA module, the RGB and thermal modality features are rearranged and the aligned features can make a balanced contribution at one position. In addition, it can be observed that MA module narrows the gap of miss rates between day and night subsets. The experiment results demonstrate that the illumination modality imbalance problem

Table 4. Ablation experiments of illumination aware feature alignment module evaluated on KAIST reasonable set.

Component		Choice				
RGB IAFC		✓		✓	✓	✓
Thermal IAFC			✓	✓	✓	✓
Illumination Gate		✓	✓		✓	✓
Modality Alignment						✓
MR(%)	Day	9.80	10.05	10.74	9.72	8.28
	Night	10.17	9.76	9.54	8.63	7.86
	All	9.86	9.89	10.27	9.36	8.13

will be mitigated with the alignment of feature modalities. The illumination and feature modality imbalance problems exist side by side and play a part together.

Illumination Aware Feature Alignment. To have a deep insight into the effectiveness of illumination aware feature alignment module, we investigate the performance of different design choices in Tab. 4. RGB IAFC and thermal IAFC represent complementing the prediction results of AP stage with RGB and thermal modality features according to the illumination conditions. It is observed that performance gains can generally be achieved by illumination aware feature complement. RGB modality is beneficial to the pedestrian detection during the day while thermal modality is beneficial to the night. By introducing the illumination gate which applies weights to the RGB and thermal stream, MBNet has a more balanced performance under different illumination conditions. This demonstrates that the detection performance can be further improved by illumination aware mechanism, since it helps the network mitigate illumination modality imbalance problem.

5 Conclusion

In this work, we explore a one-stage detector named MBNet to alleviate the modality imbalance problems in multispectral pedestrian detection. Specifically, the DMAF module is densely inserted in the ResNet block to fully integrate features and the MA module aligns two modalities so that the RGB and thermal features can have an equal contribution and representation. Meanwhile, the illumination gate embedded in the backbone network and the adaptive illumination aware feature complement in the region proposal stage make the detector robust to the variant illumination. We argue the modality imbalance problems are not limited to multispectral pedestrian detection. They are widespread in multimodal computer vision task to which the balance and integration of different modality features should be paid attention. We will further study how to reconcile the balance and reduce the learning of redundancy between different modalities in other computer vision task in the future.

Acknowledgment. This work was supported in part by the National Natural Science Foundation of China (Nos. 61627804).

References

1. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV). vol. 3 (2019)
2. Brazil, G., Liu, X.: Pedestrian detection with autoregressive network phases. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7231–7240 (2019)
3. Brazil, G., Yin, X., Liu, X.: Illuminating pedestrians via simultaneous detection & segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4950–4959 (2017)
4. Cao, Y., Guan, D., Wu, Y., Yang, J., Cao, Y., Yang, M.Y.: Box-level segmentation supervised deep neural networks for accurate and real-time multispectral pedestrian detection. ISPRS journal of photogrammetry and remote sensing **150**, 70–79 (2019)
5. Chen, K., Li, J., Lin, W., See, J., Wang, J., Duan, L., Chen, Z., He, C., Zou, J.: Towards accurate one-stage object detection with ap-loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5119–5127 (2019)
6. Chi, C., Zhang, S., Xing, J., Lei, Z., Li, S.Z., Zou, X.: Relational learning for joint head and human detection. arXiv preprint arXiv:1909.10674 (2019)
7. Choi, H., Kim, S., Park, K., Sohn, K.: Multi-spectral pedestrian detection based on accumulated object proposal with fully convolutional networks. In: 2016 23rd International Conference on Pattern Recognition (ICPR). pp. 621–626. IEEE (2016)
8. Deng, L., Yang, M., Li, T., He, Y., Wang, C.: Rfbnet: Deep multimodal networks with residual fusion blocks for rgb-d semantic segmentation. arXiv preprint arXiv:1907.00135 (2019)
9. Dollár, P., Appel, R., Belongie, S., Perona, P.: Fast feature pyramids for object detection. IEEE transactions on pattern analysis and machine intelligence **36**(8), 1532–1545 (2014)
10. Dollar, P., Wojek, C., Schiele, B., Perona, P.: Pedestrian detection: An evaluation of the state of the art. IEEE transactions on pattern analysis and machine intelligence **34**(4), 743–761 (2011)
11. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the thirteenth international conference on artificial intelligence and statistics. pp. 249–256 (2010)
12. González, A., Fang, Z., Socarras, Y., Serrat, J., Vázquez, D., Xu, J., López, A.M.: Pedestrian detection at day/night time with visible and fir cameras: A comparison. Sensors **16**(6), 820 (2016)
13. Guan, D., Cao, Y., Yang, J., Cao, Y., Yang, M.Y.: Fusion of multispectral data through illumination-aware deep neural networks for pedestrian detection. Information Fusion **50**, 148–157 (2019)
14. Guo, M., Haque, A., Huang, D.A., Yeung, S., Fei-Fei, L.: Dynamic task prioritization for multitask learning. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 270–287 (2018)
15. Ha, Q., Watanabe, K., Karasawa, T., Ushiku, Y., Harada, T.: Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5108–5115. IEEE (2017)

16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
17. Hwang, S., Park, J., Kim, N., Choi, Y., So Kweon, I.: Multispectral pedestrian detection: Benchmark dataset and baseline. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1037–1045 (2015)
18. Jang, H.D., Woo, S., Benz, P., Park, J., Kweon, I.S.: Propose-and-attend single shot detector. arXiv preprint arXiv:1907.12736 (2019)
19. Ju, R., Ge, L., Geng, W., Ren, T., Wu, G.: Depth saliency based on anisotropic center-surround difference. In: 2014 IEEE International Conference on Image Processing (ICIP). pp. 1115–1119. IEEE (2014)
20. Kong, T., Sun, F., Liu, H., Jiang, Y., Shi, J.: Consistent optimization for single-shot object detection. arXiv preprint arXiv:1901.06563 (2019)
21. Konig, D., Adam, M., Jarvers, C., Layher, G., Neumann, H., Teutsch, M.: Fully convolutional region proposal networks for multispectral person detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 49–56 (2017)
22. Li, C., Song, D., Tong, R., Tang, M.: Multispectral pedestrian detection via simultaneous detection and segmentation. arXiv preprint arXiv:1808.04818 (2018)
23. Li, C., Song, D., Tong, R., Tang, M.: Illumination-aware faster r-cnn for robust multispectral pedestrian detection. *Pattern Recognition* **85**, 161–171 (2019)
24. Li, X., Li, L., Flohr, F., Wang, J., Xiong, H., Bernhard, M., Pan, S., Gavrilu, D.M., Li, K.: A unified framework for concurrent pedestrian and cyclist detection. *IEEE transactions on intelligent transportation systems* **18**(2), 269–281 (2016)
25. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125 (2017)
26. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
27. Liu, J., Zhang, S., Wang, S., Metaxas, D.N.: Multispectral deep neural networks for pedestrian detection. arXiv preprint arXiv:1611.02644 (2016)
28. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8759–8768 (2018)
29. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: European conference on computer vision. pp. 21–37. Springer (2016)
30. Liu, W., Liao, S., Hu, W., Liang, X., Chen, X.: Learning efficient single-stage pedestrian detectors by asymptotic localization fitting. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 618–634 (2018)
31. Liu, W., Liao, S., Ren, W., Hu, W., Yu, Y.: High-level semantic feature detection: A new perspective for pedestrian detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5187–5196 (2019)
32. Noh, J., Lee, S., Kim, B., Kim, G.: Improving occlusion and hard negative handling for single-stage pedestrian detectors. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 966–974 (2018)
33. Oksuz, K., Cam, B.C., Kalkan, S., Akbas, E.: Imbalance problems in object detection: A review. arXiv preprint arXiv:1909.00169 (2019)
34. Park, K., Kim, S., Sohn, K.: Unified multi-spectral pedestrian detection based on probabilistic fusion networks. *Pattern Recognition* **80**, 143–155 (2018)

35. Qian, Q., Chen, L., Li, H., Jin, R.: Dr loss: Improving object detection by distributional ranking. arXiv preprint arXiv:1907.10156 (2019)
36. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015)
37. Wagner, J., Fischer, V., Herman, M., Behnke, S.: Multispectral pedestrian detection using deep fusion convolutional neural networks. In: ESANN (2016)
38. Wang, X., Xiao, T., Jiang, Y., Shao, S., Sun, J., Shen, C.: Repulsion loss: Detecting pedestrians in a crowd. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7774–7783 (2018)
39. Wu, B., Iandola, F., Jin, P.H., Keutzer, K.: Squeezedet: Unified, small, low power fully convolutional neural networks for real-time object detection for autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 129–137 (2017)
40. Xu, D., Ouyang, W., Ricci, E., Wang, X., Sebe, N.: Learning cross-modal deep representations for robust pedestrian detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5363–5371 (2017)
41. Zhang, L., Lin, L., Liang, X., He, K.: Is faster r-cnn doing well for pedestrian detection? In: European conference on computer vision. pp. 443–457. Springer (2016)
42. Zhang, L., Liu, Z., Chen, X., Yang, X.: The cross-modality disparity problem in multispectral pedestrian detection. arXiv preprint arXiv:1901.02645 (2019)
43. Zhang, L., Liu, Z., Zhang, S., Yang, X., Qiao, H., Huang, K., Hussain, A.: Cross-modality interactive attention network for multispectral pedestrian detection. *Information Fusion* **50**, 20–29 (2019)
44. Zhang, L., Zhu, X., Chen, X., Yang, X., Lei, Z., Liu, Z.: Weakly aligned cross-modal learning for multispectral pedestrian detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5127–5137 (2019)
45. Zhang, S., Benenson, R., Schiele, B.: Citypersons: A diverse dataset for pedestrian detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3213–3221 (2017)
46. Zhang, S., Yang, J., Schiele, B.: Occluded pedestrian detection through guided attention in cnns. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6995–7003 (2018)
47. Zhang, S., Liu, A., Wan, J., Liang, Y., Guo, G., Escalera, S., Escalante, H.J., Li, S.Z.: Casia-surf: A large-scale multi-modal benchmark for face anti-spoofing. arXiv preprint arXiv:1908.10654 (2019)
48. Zheng, Y., Izzat, I.H., Ziaee, S.: Gfd-ssd: Gated fusion double ssd for multispectral pedestrian detection. arXiv preprint arXiv:1903.06999 (2019)
49. Zhou, C., Wu, M., Lam, S.K.: Ssa-cnn: Semantic self-attention cnn for pedestrian detection. arXiv preprint arXiv:1902.09080 (2019)