

# Multi-Modal Pedestrian Detection with Large Misalignment Based on Modal-Wise Regression and Multi-Modal IoU

Napat Wanchaitanawong<sup>1,\*</sup>, Masayuki Tanaka<sup>2,\*</sup>, Takashi Shibata<sup>3,\*\*</sup>, Masatoshi Okutomi<sup>4,\*</sup>  
Tokyo Institute of Technology\*, NTT Corporation\*\*  
Tokyo 152-8550, Japan\*, Kanagawa 243-0198, Japan\*\*  
wnapat@ok.sc.e.titech.ac.jp<sup>1</sup>, {mtanaka<sup>2</sup>, mxo<sup>4</sup>}@sc.e.titech.ac.jp, t.shibata@ieee.org<sup>3</sup>

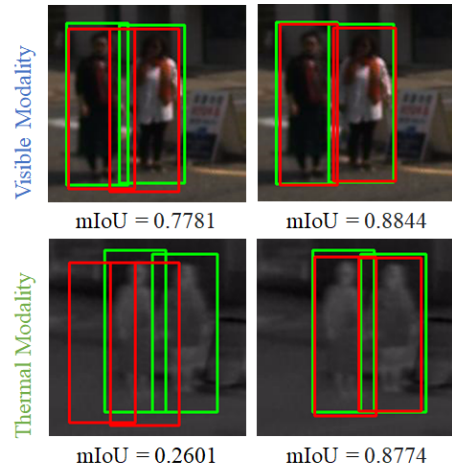
## Abstract

The combined use of multiple modalities enables accurate pedestrian detection under poor lighting conditions by using the high visibility areas from these modalities together. The vital assumption for the combination use is that there is no or only a weak misalignment between the two modalities. In general, however, this assumption often breaks in actual situations. Due to this assumption's breakdown, the position of the bounding boxes does not match between the two modalities, resulting in a significant decrease in detection accuracy, especially in regions where the amount of misalignment is large. In this paper, we propose a multi-modal Faster-RCNN that is robust against large misalignment. The keys are 1) modal-wise regression and 2) multi-modal IoU for mini-batch sampling. To deal with large misalignment, we perform bounding box regression for both the RPN and detection-head with both modalities. We also propose a new sampling strategy called “multi-modal mini-batch sampling” that integrates the IoU for both modalities. We demonstrate that the proposed method’s performance is much better than that of the state-of-the-art methods for data with large misalignment through actual image experiments.

## 1 Introduction

Pedestrian detection is still an important issue in machine vision and its applications. In practical situations, detection accuracy is significantly degraded under poor lighting conditions when only visible images are used [30, 32, 20, 18, 35]. To achieve a robust pedestrian detection in the poor lightning condition, various approaches have been proposed to combine multiple modalities (e.g., visible and far-infrared) [11, 7]. The critical assumption for the fusion is that there is no or only weakly misalignment between the two modalities. In general, however, these assumptions often break down due to lack of time synchronization, inaccurate calibration, or the effects of disparity for stereo [27, 1].

Recently, to address this issue, several methods that are robust to weak misalignment have been proposed. For example, L. Zhang et al. [34] proposed incorporating a module inside the Faster-RCNN that predicts and then aligns the travel distance between modalities



(a) Existing method (b) Proposed method

Figure 1. Visualization examples of ground truth annotations by [34] (boxes in green), detection results (boxes in red), and mIoU between them. Image patches are cropped from visible-thermal image pairs on the same position.

ties for each region. In general, however, this existing method assumes weak misalignment and is very sensitive to large misalignment. As shown in Fig. 1 (a), the position of bounding boxes detected by the existing algorithms are identical in both modalities, resulting in the poor mean intersection over union (mIoU) in thermal images, both pedestrians in thermal modality would be evaluated as false negatives. Thus, the existing methods often fail to detect in either (or both) modal when the misalignment is large. In summary, multi-modal image detection with large misalignment is an unsolved problem, even though it is typical for machine vision with multi-modal information.

In this paper, we propose a multi-modal Faster-RCNN that is robust against large misalignment. To the best of our knowledge, this paper is the first work to address the problem of “object detection from multi-modal images with large misalignment”. The keys are 1) a new sampling strategy called “mini-batch sampling based on the amount of misalignment” by introducing a new metric called multi-modal IoU, and 2) modal-wise regression: bounding-box regression for




Figure 2. **Comparison between multi-modal pedestrian detection frameworks.** Our proposed method produces pairs of bounding boxes for both modalities as output. Blue and green blocks/paths represent properties of visible and thermal modality respectively.

each modal to deal with large misalignment. Using the proposed method, the correct bounding boxes detect objects in both visible and thermal images with high mIoU, as shown in Fig. 1 (b). Real image experiments show that the proposed method’s performance significantly outperforms that of state-of-the-art methods for data containing large misalignment.

This paper’s contributions are as follows: 1) a new problem: the detection from multi-modal images with large misalignment, 2) modal-wise regression to deal with the large misalignment, and 3) multi-modal IoU and mini-batch sampling strategy for training for multi-modal inputs.

## 2 Related Work

**Multi-modal pedestrian detection.** KAIST Multi-spectral Pedestrian Detection (KAIST) dataset [11] has been widely used in the research field of multi-modal pedestrian detection. Despite non-CNN-based approach such as Aggregate Channel Features (ACF) [4] in the early days, the CNN-based approach is mainstream in this field currently [12, 28, 9, 15, 29, 22, 16, 8, 17, 33, 31]. The main challenge in the early days was how to combine and make use of information from both modalities as with other computer vision applications [19, 24, 25]. Most importantly, most of the existing methods strictly assume that visible-thermal image pairs are geometrically aligned. Those methods merely fuse both modalities’ features in corresponding pixel position directly, as shown in Fig. 2 (a). Al-

though many geometric calibration and image alignment methods for multi-modal cameras have been proposed [21, 13, 5], accurate and dense alignment for each pixel is still open problem. As a result, their detectors suffer dramatically worse performance in poorly aligned regions.

**Weak misalignment.** AR-CNN [34] is the first work that immensely tackles the misalignment issue in multi-modal CNN-based pedestrian detection. They also provided a novel KAIST-Paired annotation. Their method predicts shift distance between modalities for each Region of Interest (RoI), relocates visible region into the thermal area, then proceeds to align them together, as illustrated in Fig. 2 (b). MBNet [36] also proposes a method that takes Modality Imbalance into account. However, those methods assume that the misalignment is weak, which leads to inaccurate detection of bounding boxes in one (or both) modality in large misalignment. To tackle this problem, we introduce the modal-wise regressor to detect each object in a pair of bounding boxes with different coordinates in each modality, as shown in Fig. 2 (c), resulting in more accurate object localization in both modalities.

## 3 Proposed Method

We adopt Faster R-CNN [23] architecture and extend it into two-stream network for multi-modal imaging, which consists of multi-modal RPN and multi-modal detector. Overview of our network structure is shown in Fig. 3. Moreover, multi-modal IoU ( $\text{IoU}^M$ ) and our mini-batch sampling strategy are introduced.

### 3.1 Multi-modal RPN

The proposed multi-modal RPN has a regressor for each modality. This will enable proposals from each modality to adjust their sizes and positions independently. After receiving channel-wise concatenated features from backbone networks, the proposed multi-modal RPN will generate proposal pairs as its output, via classifier predicting each proposal pair a confidence score. To keep paired relations of proposals after applying NMS, we use thermal modality proposals as a reference, if any of them are suspended, they also suspend their corresponding pairs in visible modality. All remaining proposals will be applied with RoIAlign [10] operation before returning to channel-wise concatenate with their corresponding pairs, resulting in well-aligned RoI for the detector. We employ the loss function of RPN from [23] and add one more regression loss to optimize precision of both modals, which is defined as:

$$L(\{p_i\}, \{t_i^V\}, \{t_i^T\}) = \frac{1}{N_{cls}} \sum_i L_{cls}(p_i, p_i^*) \quad (1)$$

$$+ \lambda \frac{1}{N_{reg}} \left[ \sum_i p_i^* L_{reg}^V(t_i^V, t_i^{V*}) + \sum_i p_i^* L_{reg}^T(t_i^T, t_i^{T*}) \right],$$

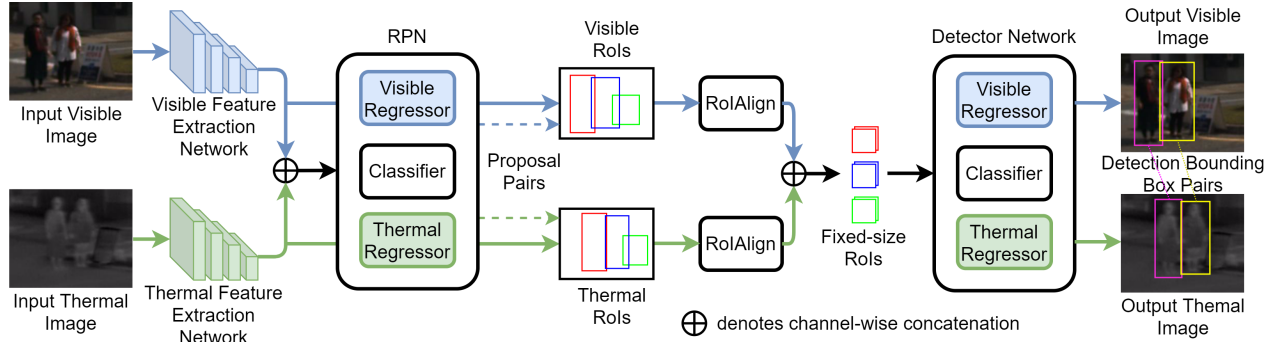


Figure 3. **The overall architecture of our network.** We extend Faster R-CNN into two-stream network to take visible-thermal image pairs as input, then return pairs of detection bounding boxes as output for both modalities. Blue and green blocks/paths represent properties of visible and thermal modality respectively. RoIs and bounding boxes with the same color represent their paired relations.

where  $i$  is the index of the anchor,  $p_i$  is the predicted probability of anchor  $i$  being an object.  $p_i^*$  is ground truth label, which equals 1 if anchor  $i$  is positive, and equals 0 if anchor  $i$  is negative.  $t_i^V, t_i^T$  are vector representing coordinates of predicted bounding box pair in visible and thermal modalities respectively.  $t_i^{V*}, t_i^{T*}$  are ground truth bounding box pair that associate with anchor  $i$ .  $L_{cls}$  is a cross entropy over object and not object classes. Regression losses  $L_{reg}^V, L_{reg}^T$  are smooth  $L_1$  loss defined in [6] for visible and thermal modality respectively.  $N_{cls}$  is mini-batch size and  $N_{cls}$  is number of anchor locations. We set  $\lambda = 1$  for all experiments.

### 3.2 Multi-modal detector

Similar to RPN, the proposed multi-modal detector network has one regressor for each modality to adjust bounding boxes' positions independently, and one classifier to predict each bounding box pair a confidence score. NMS also works the same way as RPN's. In the end, we will have detection result as pairs of bounding boxes for both modalities, which have different sizes and positions in different modalities, results in detection bounding boxes that are precise for both modalities and also keep their paired relations. We adopt loss function of detector from [6] and add one more regression loss, which is defined as:

$$L(p, t^V, t^T) = L_{cls}(p, u) + \lambda |u| [L_{loc}^V(t^V, v^V) + L_{loc}^T(t^T, v^T)], \quad (2)$$

where  $L_{cls}$  is a cross entropy for class probability  $p$  and true class  $u$ . Regression losses  $L_{loc}^V, L_{loc}^T$  are smooth  $L_1$  loss over predicted regression offsets  $t^V, t^T$  and regression targets  $v^V, v^T$  for visible and thermal modality respectively.  $|u|$  is one-hot encoding vector, equals 1 when  $u$  is in object classes and 0 otherwise. We set  $\lambda = 1$  for all experiments.

### 3.3 Multi-modal IoU

Traditionally, we use Intersection-over-Union (IoU) to classify prediction results into true/false positives and negatives categories in evaluation, defined as:

$$IoU = \frac{GT \cap DT}{GT \cup DT}, \quad (3)$$

where  $GT, DT$  denote ground truth bounding boxes and detection bounding boxes respectively.  $GT \cap DT$  represents the area of intersection of ground truth and detection bounding boxes,  $GT \cup DT$  represents the area of union of ground truth and detection bounding boxes. However, when there is misalignment between modalities, the coordinates of each object in both modalities are not the same. If we only concern about precision of one modality, another modality will have poor precision. In order to measure the ability to handle with both modalities, especially when level of misalignment is high, we introduce a new evaluation metric, which we call "multi-modal IoU (IoU<sup>M</sup>)" defined as:

$$IoU^M = \frac{(GT^V \cap DT^V) + (GT^T \cap DT^T)}{(GT^V \cup DT^V) + (GT^T \cup DT^T)}, \quad (4)$$

where  $GT^V, GT^T$  denote paired ground truth bounding boxes referring to the same object from visible and thermal modality respectively.  $DT^V, DT^T$  denote paired detection bounding boxes referring to the same object from visible and thermal modality respectively. IoU<sup>M</sup> can be used to determine the precision of detection bounding boxes in both modalities. Moreover, in order to thoroughly evaluate each modality, we define visible IoU (IoU<sup>V</sup>) as IoU in visible modality, and thermal IoU (IoU<sup>T</sup>) as IoU in thermal modality.

**Mini-batch sampling.** We follow sampling strategies from [23] and [6]. But since our approach has one regressor for each modality exclusively and we need to keep paired relations for all proposals and RoIs, we select training samples as anchor pairs and RoI pairs.

Table 1. Comparison with state-of-the-art methods on KAIST dataset, with simulated disparity of misalignment, by MR<sup>M</sup>.

Thermal Shift Distance	MR <sup>M</sup>							
	IoU <sup>M</sup> threshold: 0.5				IoU <sup>M</sup> threshold: 0.7			
	MSDS [16]	AR-CNN [34]	MBNet [36]	Ours	MSDS [16]	AR-CNN [34]	MBNet [36]	Ours
-20	84.15	82.05	82.13	<b>62.24</b>	98.27	98.12	97.06	<b>76.46</b>
-15	59.67	59.80	58.19	<b>37.51</b>	93.41	94.88	92.65	<b>62.12</b>
-10	27.06	21.61	23.13	<b>16.71</b>	85.06	89.47	82.98	<b>49.52</b>
-5	13.77	<b>9.65</b>	10.01	11.15	63.15	66.04	55.96	<b>41.66</b>
0	11.09	8.79	<b>7.76</b>	9.67	51.50	42.14	38.95	<b>38.65</b>
5	13.73	<b>10.35</b>	<b>10.35</b>	11.14	62.10	56.75	55.41	<b>39.55</b>
10	27.48	19.84	21.51	<b>17.10</b>	87.49	87.22	84.52	<b>45.46</b>
15	60.08	52.85	55.49	<b>35.44</b>	96.37	96.43	95.10	<b>58.45</b>
20	86.23	82.10	84.15	<b>59.93</b>	98.88	99.14	98.55	<b>74.35</b>

For this purpose, we use IoU<sup>M</sup> as selection criteria instead of IoU. For RPN, we assign positive and negative labels to anchor pairs that have IoU<sup>M</sup> overlap higher than 0.63 with any ground truth bounding box pair and lower than 0.3 for all ground truth bounding box pairs respectively. For detector, we assign positive and negative labels to RoI pairs that have IoU<sup>M</sup> overlap with any ground truth bounding box pair higher than 0.5 and lower than 0.5 but higher than 0.1 respectively.

## 4 Experiments

Detection performance was measured by log-average miss rate (MR) suggested by [3]. MR is defined by geometrical mean of miss rates at nine specific false positives per image (FPPI), at which we divide them by evenly spaced FPPI in range of  $[10^{-2}, 10^0]$ . Since number of false negatives and false positives are required to calculate miss rate and FPPI, IoU is used to determine objects and detection results into those categories with threshold of 0.5 and 0.7. In order to evaluate precision of detection results in both modalities, visible MR (MR<sup>V</sup>) representing MR based on IoU<sup>V</sup>, thermal MR (MR<sup>T</sup>) representing MR based on IoU<sup>T</sup>, and multi-modal MR (MR<sup>M</sup>) representing MR based on IoU<sup>M</sup>, were used in our experiments. Furthermore, to evaluate the effectiveness of our method against misalignment, we simulated disparity of misalignment between modalities by shifting thermal images horizontally from 0 to 20 pixels in both directions. All experiments were performed under reasonable configuration [11], i.e., only pedestrians taller than 55 pixel under partial or no occlusion are considered. For other methods that do not have both DT<sup>V</sup> and DT<sup>T</sup>, we substituted both with their detection bounding boxes.

**Dataset.** We used KAIST dataset [11] in our experiments. It was recorded in both day and night to consider changes in light conditions. Since we focus on misalignment, we adopted annotations provided by L. Zhang et al. [34], which localize objects in each modality independently and keep all of their paired relations, as ground truth. Only 2,252 frames from test set were used in performance test as traditional.

**Implementation details.** We adopt VGG-16 [26] pre-trained on ImageNet [2, 14] as our two-stream

Table 2. Comparison with state-of-the-art methods on KAIST dataset, without simulated disparity of misalignment, by MR<sup>V</sup>, MR<sup>T</sup>, MR<sup>M</sup>.

Method		MSDS [16]	AR-CNN [34]	MBNet [36]	Ours
MR <sup>V</sup>	IoU <sup>V</sup> threshold: 0.5	11.28	9.86	<b>7.89</b>	10.69
	IoU <sup>V</sup> threshold: 0.7	50.01	43.50	41.90	<b>41.45</b>
MR <sup>T</sup>	IoU <sup>T</sup> threshold: 0.5	12.51	8.26	<b>8.12</b>	9.24
	IoU <sup>T</sup> threshold: 0.7	55.56	43.55	41.27	<b>39.14</b>
MR <sup>M</sup>	IoU <sup>M</sup> threshold: 0.5	11.09	8.79	<b>7.76</b>	9.67
	IoU <sup>M</sup> threshold: 0.7	51.50	42.14	38.95	<b>38.65</b>

backbone networks as in AR-CNN [34]. We train the network for 3 epochs with learning rate of 0.005 and 1 additional epoch with learning rate of 0.0005 by Stochastic Gradient Descent (SGD) optimizer with 0.9 momentum and 0.0005 weight decay. We select 8,892 images containing informative pedestrians for the training. Image resolution is fixed to 640×512. All images are horizontally flipped for data augmentation.

**Comparison with state-of-the-art methods.** We selected three state-of-the-art methods for our experiments, MSDS (MSDS-RCNN) [16] is representative for methods without misalignment consideration, AR-CNN [34] and MBNet [36] are methods that consider misalignment. From Table 1, when IoU<sup>M</sup> threshold is 0.5, we only achieve the lowest MR when the misalignment is larger than 10 pixels. However, when IoU<sup>M</sup> threshold is 0.7, we achieve the lowest MR at all shift distances, which indicates our proposed method’s robustness to large misalignment.

From Table 2, our method’s performance is comparable to state-of-the-art methods. Still, it is not the best on any evaluation metrics when IoU thresholds are 0.5. However, when IoU thresholds increase to 0.7, i.e., requirement of precision is higher, our method achieves the best performance among available competitors in all evaluation metrics, which demonstrates our superior precision of detection bounding boxes in both modalities. This can benefit applications that require high and reliable precision of detection, such as autonomous vehicle, precise location of pedestrians is crucial.

## 5 Conclusion

We have proposed the novel multi-modal detection method based on modal-wise regression and multi-modal IoU, the proposed method is robust to large misalignment and also keeps paired relations of all detection bounding boxes between both modalities. To the best of our knowledge, this paper is the first work to tackle the problem: detection from multi-modal images with large misalignment. Our experiments showed that when the precision requirement of the bounding box or the level of misalignment is high, our proposed method achieves the best performance, demonstrating our robustness to misalignment and superior precision of detection bounding boxes in both modalities.

## References

- [1] Yukyung Choi, Namil Kim, Soonmin Hwang, Kibaek Park, Jae Shin Yoon, Kyoungwan An, and In So Kweon. Kaist multi-spectral day/night data set for autonomous and assisted driving. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):934–948, 2018.
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [3] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 34(4):743–761, 2012.
- [4] Piotr Dollár, Ron Appel, Serge Belongie, and Pietro Perona. Fast feature pyramids for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 36(8):1532–1545, 2014.
- [5] Jing Dong, Byron Boots, Frank Dellaert, Ranveer Chandra, and Sudipta Sinha. Learning to align images using weak geometric supervision. In *International Conference on 3D Vision (3DV)*, pages 700–709. IEEE, 2018.
- [6] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015.
- [7] Alejandro González, Zhijie Fang, Yainuvis Socarras, Joan Serrat, David Vázquez, Jiaolong Xu, and Antonio Manuel López. Pedestrian detection at day/night time with visible and fir cameras: A comparison. *Sensors*, 16(6), 2016.
- [8] Dayan Guan, Yanpeng Cao, Jiangxin Yang, Yanlong Cao, and Michael Ying Yang. Fusion of multispectral data through illumination-aware deep neural networks for pedestrian detection. *Information Fusion*, 50:148–157, 2019.
- [9] Hangil Choi, Seungryong Kim, Kihong Park, and Kwanghoon Sohn. Multi-spectral pedestrian detection based on accumulated object proposal with fully convolutional networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 621–626, 2016.
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017.
- [11] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection: Benchmark dataset and baseline. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1037–1045, 2015.
- [12] Shu Wang, Jingjing Liu, Shaoting Zhang and Dimitris Metaxas. Multispectral deep neural networks for pedestrian detection. In *British Machine Vision Conference (BMVC)*, pages 73.1–73.13, 2016.
- [13] Seungryong Kim, Dongbo Min, Bumsub Ham, Seungchul Ryu, Minh N Do, and Kwanghoon Sohn. Dasc: Dense adaptive self-correlation descriptor for multi-modal and multi-spectral correspondence. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2103–2112, 2015.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Everest Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105. 2012.
- [15] Daniel König, Michael Adam, Christian Jarvers, Georg Layher, Heiko Neumann, and Michael Teutsch. Fully convolutional region proposal networks for multispectral person detection. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 243–250, 2017.
- [16] Chengyang Li, Dan Song, Ruofeng Tong, and Min Tang. Multispectral pedestrian detection via simultaneous detection and segmentation. In *British Machine Vision Conference (BMVC)*, 2018.
- [17] Chengyang Li, Dan Song, Ruofeng Tong, and Min Tang. Illumination-aware faster r-cnn for robust multispectral pedestrian detection. *Pattern Recognition*, 85:161–171, 2019.
- [18] Jianan Li, Xiaodan Liang, Shengmei Shen, Tingfa Xu, Jiashi Feng, and Shuicheng Yan. Scale-aware fast r-cnn for pedestrian detection. *IEEE Transactions on Multimedia*, 20(4):985–996, 2018.
- [19] Shutao Li, Xudong Kang, and Jianwen Hu. Image fusion with guided filtering. *IEEE Transactions on Image Processing (TIP)*, 22(7):2864–2875, 2013.
- [20] Jiayuan Mao, Tete Xiao, Yuning Jiang, and Zhimin Cao. What can help pedestrian detection? In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6034–6043, 2017.
- [21] Yuka Ogino, Takashi Shibata, Masayuki Tanaka, and Masatoshi Okutomi. Coaxial visible and fir camera system with accurate geometric calibration. In *Thermosense: Thermal Infrared Applications XXXIX*, volume 10214, page 1021415. International Society for Optics and Photonics, 2017.
- [22] Kihong Park, Seungryong Kim, and Kwanghoon Sohn. Unified multi-spectral pedestrian detection based on probabilistic fusion networks. *Pattern Recognition*, 80:143–155, 2018.
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015.

- [24] Thapanapong Rukkanchanunt, Masayuki Tanaka, and Masatoshi Okutomi. Full thermal panorama from a long wavelength infrared and visible camera system. *Journal of Electronic Imaging*, 28(3):1 – 10, 2019.
- [25] Takashi Shibata, Masayuki Tanaka, and Masatoshi Okutomi. Misalignment-robust joint filter for cross-modal image pairs. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3295–3304, 2017.
- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR), San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [27] Wayne Treible, Philip Saponaro, Scott Sorensen, Abhishek Kolagunda, Michael O’Neal, Brian Phelan, Kelly Sherbondy, and Chandra Kambhampettu. Cats: A color and thermal stereo benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [28] Jörg Wagner, Volker Fischer, Michael Herman, and Sven Behnke. Multispectral pedestrian detection using deep fusion convolutional neural networks. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2016.
- [29] Dan Xu, Wanli Ouyang, Elisa Ricci, Xiaogang Wang, and Nicu Sebe. Learning cross-modal deep representations for robust pedestrian detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4236–4244, 2017.
- [30] Bin Yang, Junjie Yan, Zhen Lei, and Stan Z. Li. Convolutional channel features. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [31] Heng Zhang, Elisa Fromont, Sebastien Lefevre, and Bruno Avignon. Guided attentive feature fusion for multispectral pedestrian detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 72–80, 2021.
- [32] Liliang Zhang, Liang Lin, Xiaodan Liang, and Kaiming He. Is faster r-cnn doing well for pedestrian detection? In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 443–457, 2016.
- [33] Lu Zhang, Zhiyong Liu, Shifeng Zhang, Xu Yang, Hong Qiao, Kaizhu Huang, and Amir Hussain. Cross-modality interactive attention network for multispectral pedestrian detection. *Information Fusion*, 50:20–29, 2019.
- [34] Lu Zhang, Xiangyu Zhu, Xiangyu Chen, Xu Yang, Zhen Lei, and Zhiyong Liu. Weakly aligned cross-modal learning for multispectral pedestrian detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [35] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. Occlusion-aware r-cnn: Detecting pedestrians in a crowd. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [36] Kailai Zhou, Linsen Chen, and Xun Cao. Improving multispectral pedestrian detection by addressing modality imbalance problems. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 787–803, 2020.