

Robust Pedestrian Detection with Uncertain Modality

Qian Bie^{1*}, Xiao Wang^{2*}, Bin Yang^{3†}, Zhixi Yu¹, Jun Chen³, Xin Xu^{1†}

¹ School of Computer Science and Technology, Wuhan University of Science and Technology, Wuhan 430065 China

² Hubei Province Key Laboratory of Intelligent Information Processing and Real-time Industrial System, Wuhan University of Science and Technology, China

³ National Engineering Research Center for Multimedia Software, School of Computer Science, Wuhan University, China
bieqian@wust.edu.cn, wangxiao2021@wust.edu.cn, yangbin_cv@whu.edu.cn,
yuzhixi@wust.edu.cn, chenjun@whu.edu.cn, xuxin@wust.edu.cn

Abstract

Existing cross-modal pedestrian detection (CMPD) employs complementary information from RGB and thermal-infrared (TIR) modalities to detect pedestrians in 24h-surveillance systems. RGB captures rich pedestrian details under daylight, while TIR excels at night. However, TIR focuses primarily on the person’s silhouette, neglecting critical texture details essential for detection. While the near-infrared (NIR) captures texture under low-light conditions, which effectively alleviates performance issues of RGB and detail loss in TIR, thereby reducing missed detections. To this end, we construct a new Triplet RGB–NIR–TIR (TRNT) dataset, comprising 8,281 pixel-aligned image triplets, establishing a comprehensive foundation for algorithmic research. However, due to the variable nature of real-world scenarios, imaging devices may not always capture all three modalities simultaneously. This results in input data with unpredictable combinations of modal types, which challenge existing CMPD methods that fail to extract robust pedestrian information under arbitrary input combinations, leading to significant performance degradation. To address these challenges, we propose the Adaptive Uncertainty-aware Network (AUNet) for accurately discriminating modal availability and fully utilizing the available information under uncertain inputs. Specifically, we introduce Unified Modality Validation Refinement (UMVR), which includes an uncertainty-aware router to validate modal availability and a semantic refinement to ensure the reliability of information within the modality. Furthermore, we design a Modality-Aware Interaction (MAI) module to adaptively activate or deactivate its internal interaction mechanisms per UMVR output, enabling effective complementary information fusion from available modalities. AUNet enables accurate modality validation and robust inference without fixed modality pairings, facilitating the effective fusion of RGB, NIR, and TIR information across diverse inputs.

Introduction

Robust pedestrian detection is essential for round-the-clock applications such as autonomous driving (Kim et al. 2021; Chen et al. 2022; Zhou et al. 2025) and intelligent surveillance (Cao et al. 2023; Chen et al. 2025; Xu et al. 2025;

*These authors contributed equally.

†Corresponding author.

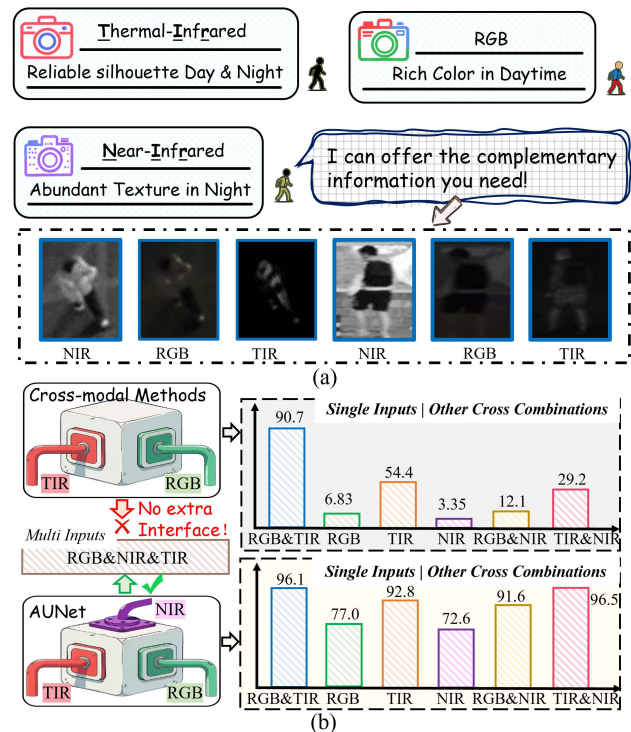


Figure 1: Illustration of our idea. (a) demonstrates the complementary benefits of the NIR for RGB and TIR. (b) illustrates the performance degradation of existing CMPD methods under single-modal and other cross-modal inputs. we conduct evaluation using ICAFusion (Shen et al. 2024) and our AUNet trained on our DRNT.

Liu, Ye, and Du 2024), where cross-modal pedestrian detection (CMPD) leveraging complementary information from RGB and thermal-infrared (TIR) data is a primary approach. Over the past decade, various CMPD methods have been proposed for RGB and TIR information fusion, ranging from direct feature concatenation (Liu et al. 2016; Konig et al. 2017), illumination-guided weighting (Li et al. 2019; Xie et al. 2023), gated mechanisms (Zheng, Izzat, and Ziaee 2019), differential fusion (Zhou, Chen, and Cao 2020),

attention-based fusion (Zhang et al. 2021; Shen et al. 2024), frequency-domain fusion (Li et al. 2025), to causal inference mechanisms (Kim et al. 2024). However, TIR imaging primarily captures the coarse pedestrian’s silhouette based on the principle of heat radiation, neglecting the vital texture details necessary for accurate detection, as depicted in Figure 1 (a). In contrast, the near-infrared (NIR) sensor is capable of capturing abundant person texture information, particularly under active nighttime illumination. Therefore, incorporating the complementary NIR (abundant texture at night) information for RGB (rich color and texture in daytime) and TIR (reliable silhouette day & night) can effectively alleviate detail loss in TIR and target blurring in RGB images under low-light conditions, thereby reducing the rate of missed pedestrian detections.

To this end, we construct a novel comprehensive **Triplet RGB–NIR–TIR** (TRNT) dataset, captured by three collaborative sensors across diverse real-world scenarios. As illustrated in Figure 2 (a), it contains various detection-relevant factors, including variations in illumination (day → night), occlusion (partial → heavy), weather conditions (winter → summer), and complex backgrounds. Furthermore, TRNT consists of 7,429 and 852 image triplets captured from the ground view and Unmanned Aerial Vehicle (UAV) perspective, respectively, accounting for dynamic camera angles in real-world deployment scenarios. TRNT provides a comprehensive and reliable experimental foundation for exploring the potential of these three modalities.

However, due to the variable factors in real-world scenarios, imaging devices may not always capture RGB, NIR, and TIR images simultaneously. This results in input data with various and unpredictable modal types and quantities. Existing CMPD methods are typically designed for fixed RGB–TIR input pairs, thus failing to extract reliable pedestrian information under single-modal and other cross-modal inputs, which dramatically decreases their detection performance, as demonstrated in Figure 1 (b). Besides, they are limited to two branches and cannot provide additional interfaces for the NIR modality. Although completing missing modality data is an intuitive solution, it requires additional generative modules and heavily relies on the availability of fully complete three-modality data during training.

To address these challenges, we propose a novel **Adaptive Uncertainty-aware Network** (AUNet), designed for robust pedestrian detection under arbitrary input combinations of RGB, NIR, and TIR data. Specifically, to accurately validate the availability of input modalities, we implement a **Unified Modality Validation Refinement** (UMVR) with two components: a lightweight uncertainty-aware router net to validate the availability of each modality, with its output yielding a binary decision value (0 or 1). Motivated by the success of CLIP’s powerful global semantic extraction capabilities, we design a CLIP-driven semantic refinement that allows semantic priors extracted by CLIP’s image encoders to serve as guiding signals, discriminating critical pedestrian features from cluttered backgrounds to ensure the reliability of information within each modality. To fully leverage the information of the available modalities, we design a **Modality-Aware Interaction** (MAI) module that adaptively acquires the fu-

sion strategy and dynamically activates or deactivates its internal interaction mechanisms according to the availability value obtained by UMVR, effectively integrating complementary information from different modalities. For thorough examination, we conduct extensive experiments on both our TRNT and the LLVIP (Jia et al. 2021) to demonstrate the effectiveness and robustness of our AUNet under various uncertain input combinations. In summary, we launch a new benchmark for multi-modality pedestrian detection and a new baseline for uncertain-modal pedestrian detection. The main contributions can be summarized as follows:

- We construct a multi-scenario, multi-modality **Triplet RGB–NIR–TIR** dataset called TRNT, comprising 8281 precisely aligned RGB–NIR–TIR image triplets simultaneously captured by triple co-aligned sensors in diverse real-world environments.
- We propose an **Adaptive Uncertainty-aware Network** (AUNet) designed to tackle the challenge of uncertain modal inputs, which incorporates a **Unified Modality Validation Refinement** (UMVR) and a **Modality-Aware Interaction** (MAI) module to accurately validate modal availability and effectively exploit complementary information from available modalities, respectively.
- We conduct a comprehensive evaluation of various state-of-the-art methods, along with detailed experimental analysis on both our TRNT and the LLVIP cross-modal dataset to demonstrate the effectiveness and robustness of our AUNet under arbitrary input combinations.

Related Work

Cross-modal Pedestrian Detection Datasets

To overcome the illumination constraints inherent in RGB-based pedestrian detection, Hwang et al. (Hwang et al. 2015) pioneer cross-modal pedestrian detection by introducing the KAIST dataset, which aims to leverage complementary information from RGB and TIR data for accurate detection. CVC-14 (González et al. 2016) and FLIR (F. A. Group. 2018) are constructed to complement KAIST by motion-blurred scenarios relevant to autonomous driving. Jia et al. (Jia et al. 2021) introduce LLVIP, a high-resolution RGB–TIR dataset designed for pedestrian detection in low-light conditions. Liu et al. (Liu et al. 2022) propose the M3FD dataset, which incorporates adverse weather conditions such as rain and fog. However, these datasets overlook the prevalent NIR modality used in night vision systems, which can provide rich pedestrian texture details complementary to both RGB and TIR, thereby significantly limiting the potential for multi-modal synergy.

Cross-modal Pedestrian Detection Methods

Liu et al. (Liu et al. 2016) craft four sophisticated convolutional neural network architectures and demonstrate that the halfway fusion model exhibits superior detection synergy. Illumination-aware Faster R-CNN (Li et al. 2019) dynamically weights RGB and TIR sub-networks via illumination-based gating to optimize multispectral fusion. Yang et al. (Yang et al. 2022) design a BAANet with bi-directional

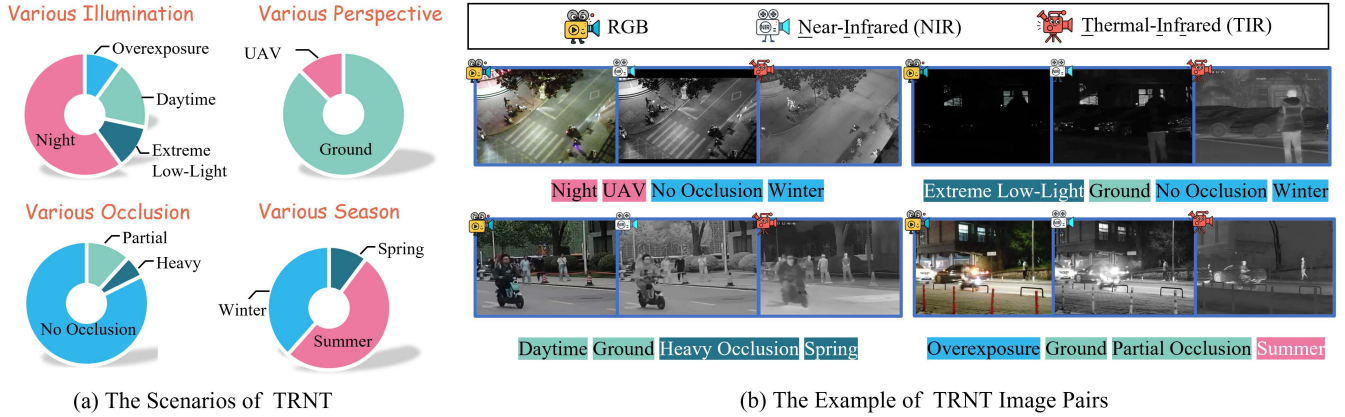


Figure 2: Data Diversity. The scenarios and examples of our benchmark TRNT. TRNT covers diverse scenarios with various illumination, perspective, occlusion, and season.

adaptive attention gates, progressively distilling modality-specific noise and recalibrating features via illumination-guided interactions. Shen et al. (Shen et al. 2024) introduce a framework with dual cross-attention transformers to overcome the limitations of CNNs in local-range feature interaction. Fusion-mamba (Dong et al. 2024) introduces Mamba for cross-modal fusion, designs state space channel swapping and dual state space fusion modules to reduce modality disparities. FD²Net (Li et al. 2025) extracts RGB and TIR information through high-frequency units and low-frequency units, respectively, and combines them with a parameter-independent fusion strategy for further enhancing feature representation. However, these methods primarily focus on exploring the complementary correlations between fixed TIR and RGB images, neglecting variable input modality types and counts, which significantly limits their applicability and generalizability.

TRNT Dataset

To effectively utilize the complementary advantages of near-infrared (NIR), we construct a novel dataset, **Triplet RGB-NIR-TIR (TRNT)**, encompassing all three modalities and supporting seven possible input combinations.

Data Acquisition Data collection employs two camera platforms: a customized triplicated Pan-Tilt-Zoom (PTZ) to capture images from the viewpoint of a shot on the ground, and a purchased Matrice 350 Real-Time Kinematic (RTK) drone device to capture images from the UAV viewpoint. Both platforms are equipped with synchronized RGB, NIR, and TIR sensors to ensure simultaneous tri-modal image pair captures. According to the KAIST (Hwang et al. 2015) dataset protocol, raw TIR images are converted to grayscale to obtain the final TIR images.

Data Description TRNT comprises 8281 rigorously aligned RGB-NIR-TIR image triplets captured from different camera angles with diverse illumination conditions and various occlusions. The division of the TRNT dataset into training and testing follows CVC-14, comprising 7026 image pairs in the training subset and 1255 image pairs in the

test set. The bounding boxes are annotated accurately with a resolution of 640×480 for each image. Crucially, every image contains at least one pedestrian. As shown in Table 1, compared with the existing cross-modal pedestrian detection datasets, TRNT has the following main advantages:

- **Novel Modality Integration.** TRNT pioneers the inclusion of near-infrared (NIR) alongside RGB and thermal-infrared (TIR) in pedestrian detection. NIR sensors are widely available in modern surveillance systems but overlooked in prior cross-modal pedestrian detection research. The inclusion of all three common modalities significantly expands the diversity of possible input combinations compared to cross-modal datasets.
- **Dual-Perspective Coverage.** TRNT is the first multi-modal dataset that integrates both ground-views (mimicking vehicle and robotic perspectives) and high-altitude UAV perspectives ($\geq 50\text{m}$ altitude), particularly the UAV perspective that has never been present in any of the previous cross-modal pedestrian detection datasets.
- **Longitudinal Illumination and Seasonal transitions.** TRNT images span 14 continuous hours (8:00–22:00) covering diverse lighting conditions: natural illumination (day, night), extremely low illumination (dark environments without urban/car lights). Besides, TRNT images span three distinct seasons (winter, spring, and summer), introducing dynamic modality reliability shifts. For example, thermal radiation attenuation from pedestrians’ upper bodies due to insulation by heavy winter clothing (please refer to the Figure 2 (b)).

Methods

To dynamically handle arbitrary input combinations (single, dual, or triple modalities), we propose AUNet consisting of a Unified Modality Validation Refinement (UMVR) and a Modality-Aware Interaction (MAI) module to discriminate modal availability and exploit the information of the available modalities, respectively, as shown in Figure 3.

Shared-weight Feature Extraction Formally, let $I_R \in \mathbb{R}^{H \times W \times 3}$, $I_T \in \mathbb{R}^{H \times W \times 3}$, and $I_N \in \mathbb{R}^{H \times W \times 3}$ be the in-

Datasets	Number of images	Modality			Camera angle	Illumination condition		Pedestrian
		RGB	TIR	NIR		Day	Night	
TNO	261 × 2	✓	✓	×	Shot on the ground	few	✓	few
INO	2100 × 2	✓	✓	×	surveillance	✓	✓	few
OSU	285 × 2	✓	✓	×	surveillance	×	✓	✓
KAIST	4750 × 2	✓	✓	×	driving	✓	✓	✓
CVC-14	849 × 2	✓	✓	×	driving	✓	✓	✓
FILR	5258 × 2	✓	✓	×	driving	✓	✓	✓
LLVIP	15488 × 2	✓	✓	×	surveillance	✓	few	✓
DRNT(ours)	8281 × 3	✓	✓	✓	Shot on the ground, UAV	✓	✓	✓

Table 1: A comparative analysis of the TRNT dataset with existing cross-modal pedestrian detection datasets. The number of images represents the dataset’s selection rate of one frame per second. TRNT comprises 8281 RGB-NIR-TIR image triplets (24,843 images). Moreover, the TRNT dataset ensures that each image contains at least one or more pedestrians.

put RGB, TIR, and NIR images, respectively. Unlike previous cross-modal works (Li et al. 2019; Chen et al. 2023) that employ three independent backbones, we utilize a single weight-shared backbone $\mathcal{F}_\theta(\cdot)$ with shared parameters θ across all modalities, supporting flexible input configurations (single-, cross-, or multi-modal) without requiring architectural changes or retraining. If modality-specific backbones or partially shared backbones are employed instead, the AUNet architecture must be adjusted whenever the number of input modalities increases or decreases. Compared to separate backbone networks for each modality, a weight-shared backbone significantly decreases computational costs. The whole process can be represented as:

$$F_X = \mathcal{F}_\theta(I_X), \quad X \in R, T, N, \quad (1)$$

where RGB, TIR and NIR features are denoted as $F_R \in \mathbb{R}^{h \times w \times c}$, $F_T \in \mathbb{R}^{h \times w \times c}$, and $F_N \in \mathbb{R}^{h \times w \times c}$, respectively. Here, $\mathcal{F}_\theta(\cdot)$ can be any convolutional neural network.

Unified Modality Validation Refinement To validate the availability of each input modality, we propose a lightweight uncertainty-aware router (UAR) to validate the availability of the input RGB, NIR, and TIR images. Specifically, UAR employs a parameterized multilayer perceptron (MLP) with sigmoid activations, with its output yielding three binary decision values $\mathbf{v} \in \{0, 1\}^3$. For example, when RGB, NIR, and TIR images are all available, the output decision of the router is $v = [1, 1, 1]$. To enforce consistency between UAR decisions and realistic availability $v' \in \{0, 1\}^3$ of input images, we introduce availability validation loss \mathcal{L}_V applied per modality:

$$\mathcal{L}_V = \sum_{x \in X} [-v'_x \log v_x - (1 - v'_x) \log(1 - v_x)] \quad (2)$$

While UAR can determine the availability of each modality, it doesn’t guarantee that all its information will benefit detection. For instance, as shown in Figure 3, given nighttime RGB inputs, unrefined feature maps exhibit misguided attention toward artifacts, such as reflective pavement surfaces or luminous streetlights, rather than pedestrians. To mitigate this interference of intrinsic modality noise with the aggregation of available modal information (UAR output value is 1), we introduce a CLIP-driven specific refinement (CSR) module. The CSR module leverages semantic

priors extracted by CLIP’s image encoders as guiding signals to discriminate critical pedestrian features from background noise across three modalities. Specifically, CLIP-derived weighting maps M are generated to suppress irrelevant background regions while enhancing pedestrian-aware features. Since the weights of CLIP are fixed, M is fine-tuned to adapt to the pedestrian detection task. The process for RGB is formulated as:

$$M_R = \delta(\text{conv}(\text{conv}(\mathcal{E}_{\text{CLIP}}(I_R)))), \quad (3)$$

where δ denotes the sigmoid activation function, conv represents a 1×1 convolution operation, and $\mathcal{E}_{\text{CLIP}}(\cdot)$ denotes CLIP’s image encoder. The refined feature F_R^{csr} is then obtained via:

$$F_R^{csr} = F_R \oplus F_R \odot M_R, \quad (4)$$

where \oplus and \odot represent element-wise sum and element-wise multiplication, respectively. Following the same methodology, we can derive the output refined maps for TIR and NIR, denoted as F_T^{csr} and F_N^{csr} , respectively. To ensure the maps (M_R, M_N, M_T) generated by CSR precisely localize regions where features need to be enhanced or diminished, ground-truth pedestrian distribution maps are utilized for supervision. The process for acquiring realistic pedestrian distribution maps is formulated as:

$$F_{truth}^{csr} = \frac{(x - cx)^2}{a^2} + \frac{(y - cy)^2}{b^2} < 1 + \epsilon, \quad (5)$$

where x and y are acquired by the mesh grid function. (cx, cy) , a , and b correspond to the coordinates of the center point and half the width and height of a bounding box, respectively. ϵ is set to 0.0001. The CSR is optimized by minimizing the contrastive loss between the F^{csr} and the ground-truth pedestrian distribution maps. Taking the contrastive loss in the CSR module of RGB as an illustrative example, the contrastive loss \mathcal{L}_{CR} is formulated as:

$$\begin{aligned} \mathcal{L}_{CR} &= \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{2|F_R^{csr} \cap F_{truth}^{csr}| + 1}{|F_R^{csr}| + |F_{truth}^{csr}| + 1} \right) \\ &= \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{\left(2 \times \sum_{p=1}^{h \times w} F_{R,p}^{csr} F_{truth,p}^{csr} \right) + 1}{\left(\sum_{p=1}^{h \times w} F_{R,p}^{csr} \right) + \left(\sum_{p=1}^{h \times w} F_{truth,p}^{csr} \right) + 1} \right), \end{aligned} \quad (6)$$

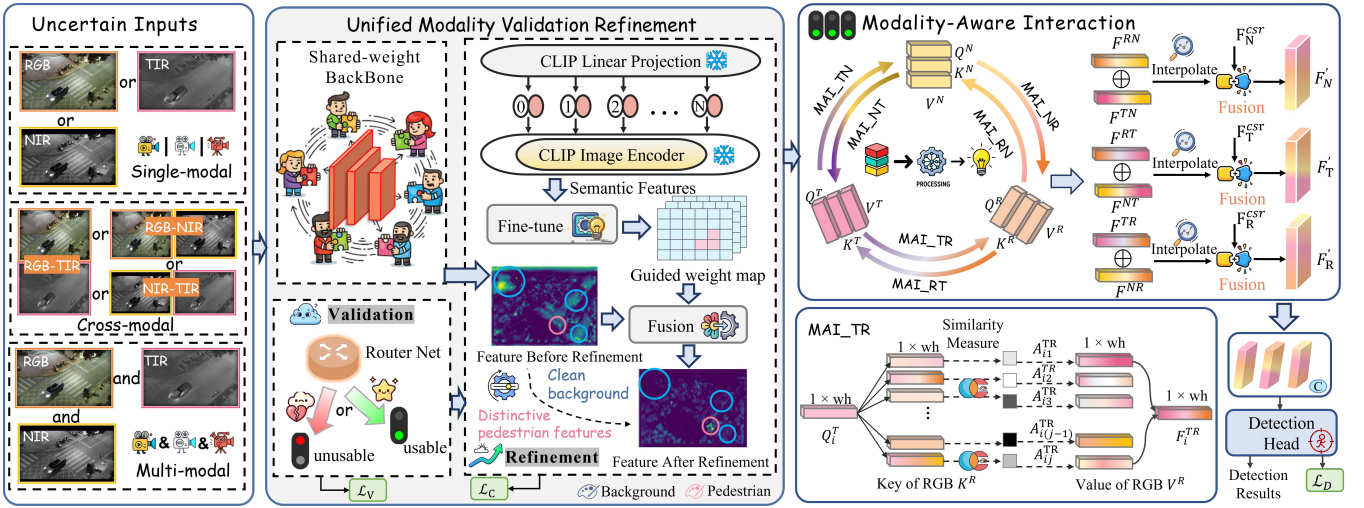


Figure 3: Illustration of our AUNet, including two main components: 1) Unified Modality Validation Refinement (UMVR), which integrates an Uncertainty-aware Router (UAR) validation and CLIP-driven Semantic Refinement (CSR); 2) Modality-Aware Interaction (MAI), which effectively integrates complementary information from the available modalities. This figure depicts a scenario where all three modalities are available, corresponding to the UMVR output of [1, 1, 1].

where $F_{truth,p}^{csr}$ and $F_{R,p}^{csr}$ denote the elements of F_{truth}^{csr} and F_R^{csr} at the location p , respectively. N denotes the number of samples in a mini-batch. The total contrastive loss \mathcal{L}_C is formulated as:

$$\mathcal{L}_C = W_R \mathcal{L}_{CR} + W_T \mathcal{L}_{CT} + W_N \mathcal{L}_{CN}, \quad (7)$$

where W_R , W_T , and W_N represent the learnable fusion weight parameters, initially set to 1.0.

Modality-Aware Interaction To adaptively acquire the fusion strategy per UMVR output and effectively integrate complementary information from available modalities, we introduce the Modality-Aware Interaction (MAI) module with a diversified information interaction mechanism. As illustrated at the right of Figure 3, the MAI module contains six predefined and independent cross-modal information interaction mechanisms called MAI_{TN}, MAI_{NT}, MAI_{TR}, MAI_{RT}, MAI_{NR}, and MAI_{RN}, respectively. When RGB, TIR, and NIR are all available, MAI activates all cross-modal interaction mechanisms to capture complementary information among the three modalities. When only two modalities are available, for example, RGB and TIR, MAI activates the MAI_{TR} and MAI_{RT}. To elaborate on the process of the interaction mechanism, we take the MAI_{TR} interaction as an example without loss of generality. As shown in the bottom right corner of Figure 3, the feature from F_T^{csr} is passed into a pooling layer, a linear, and layernorm transformation to generate a query matrix $Q^T = \{Q_1^T, \dots, Q_i^T, \dots, Q_L^T\}$. Similarly, the feature from F_R^{csr} is passed into two pooling layers, two linear, and layernorm transformations to generate a key matrix $K^R = \{K_1^R, \dots, K_j^R, \dots, K_L^R\}$ and a value matrix $V^R = \{V_1^R, \dots, V_j^R, \dots, V_L^R\}$. L denotes the number of queries, keys, and values in MAI_{TR}. The i -th element of

F_i^{TR} is calculated as follows:

$$F_i^{TR} = \sum_{j=1}^L V_j^R \cdot A_{ij}^{TR}, \quad (8)$$

$$A_{ij}^{TR} = \frac{\exp(s_{ij})}{\sum_{j=1}^L \exp(s_{ij})}, \quad s_{ij} = \frac{Q_i^T (K_j^R)^T}{\sqrt{c}}, \quad (9)$$

where $(\cdot)^T$ means the matrix transposition. Here, $A_i^{TR} = \{A_{i1}^{TR}, \dots, A_{ij}^{TR}, \dots, A_{iL}^{TR}\} \in R^{1 \times L}$ represents the relevancy of the correlation between Q_i^T and K^R . $F^{TR} = \{F_i^{TR}\}_{i=1}^L$. We can acquire the output F^{RN} , F^{TN} , F^{NR} , F^{TR} , F^{RT} and F^{NT} by six independent information interactions. Taking F^{RN} and F^{TN} as examples, they are the complement features of RGB to NIR and TIR to NIR, respectively. F^{RN} and F^{TN} are summed with learnable weights with F_N^{csr} . The process for acquiring final NIR maps is expressed as:

$$F'_N = \mathcal{I}(W_{RN} F^{RN} \oplus W_{TN} F^{TN}) \oplus F_N^{csr}, \quad (10)$$

where $\mathcal{I}(\cdot)$ means interpolation. The output feature maps of MAI are reshaped to output feature maps of CSR. Here, W_{RN} and W_{TN} are two different learnable weights. Following this analogy, we can derive F'_T and F'_R . Finally, the F'_N , F'_R , and F'_T are concatenated by two 1×1 convolution blocks, followed by batch normalizations and RELU activations, to acquire the final fused map called F'_m .

Objective Function As illustrated in Figure 3, we utilize a multi-scale detection head to detect pedestrians of different sizes. Our objective function comprises three components: loss for UAR (\mathcal{L}_V), loss for CSR (\mathcal{L}_C), and loss for the detection module (\mathcal{L}_D) that contains classification and localization loss of our one-stage AUNet.

$$\mathcal{L} = \mathcal{L}_D + \lambda_1 \mathcal{L}_V + \lambda_2 \mathcal{L}_C, \quad \mathcal{L}_D = \mathcal{L}_{cls} + \mathcal{L}_{loc}, \quad (11)$$

Methods	VGG16			ResNet50			CSPDarknet53		
	AP ₅₀ ↑	AP ₇₅ ↑	mAP ↑	AP ₅₀ ↑	AP ₇₅ ↑	mAP ↑	AP ₅₀ ↑	AP ₇₅ ↑	mAP ↑
RGB	73.91	23.70	31.38	78.03	26.23	34.33	77.00	22.03	32.00
NIR	68.94	18.89	27.99	72.89	19.79	29.81	72.61	20.21	29.83
TIR	90.78	46.48	47.80	92.38	53.60	51.93	92.81	52.22	51.36
(RGB-NIR)_AUNet	91.73	45.79	46.97	91.48	49.97	50.06	91.63	49.64	50.29
(NIR-TIR)_AUNet	95.24	62.27	57.20	96.50	64.85	57.48	96.51	64.92	57.74
(RGB-TIR)_AUNet	95.93	63.24	57.63	96.86	66.33	59.12	96.10	66.27	58.32
(RGB-NIR-TIR)_AUNet	96.63	65.49	58.24	97.03	67.30	59.82	97.07	67.78	60.08

Table 2: Experimental comparison of the different modalities and backbones on TRNT (in %). It is obvious that AUNet with CSPDarknet53 achieves the best performance compared to the other backbones.

Methods	TRNT			FPS(Hz) ↑
	AP ₅₀ ↑	AP ₇₅ ↑	mAP ↑	
ProbEn	86.29	50.64	41.23	6.97
TINet	96.60	64.10	56.47	10.40
INSANet	50.40	17.50	13.80	42.36
ICAFusion	95.32	64.19	56.68	32.62
DE-YOLO	95.73	65.26	57.00	38.64
AUNet	97.07	67.78	60.08	60.99

Table 3: Experimental results of AUNet on TRNT comparing with state-of-the-art methods while handling multi-modality pedestrian detection data.

where λ_1 and λ_2 denote the learnable weight parameters, initially set to 0.5 in the experiments.

Experiments

Impact of Different Input Combinations

To ascertain the contribution of NIR information and validate the robustness of AUNet under seven possible modality inputs, we conduct a comprehensive evaluation of our method across diverse modal combinations on three different backbones, including VGG16 (Simonyan and Zisserman 2014), ResNet50 (He et al. 2016), CSPDarkNet53 (Bochkovskiy, Wang, and Liao 2020). As shown in Table 2, this outcome ascertains the information contribution for RGB and TIR. Specifically, the AP₇₅ achieved by RGB-NIR (49.64%) with our AUNet is higher than the single RGB (22.03%) and NIR (20.21%) added together when utilizing CSPDarkNet53 as the backbone. Even if we change the backbone, the proposed AUNet with RGB-NIR-TIR inputs still outperforms other methods. Furthermore, it demonstrates AUNet’s capability to dynamically harness arbitrary subsets of RGB-NIR-TIR modalities, mitigating the negative impact of uncertainty inputs through Unified Modality Validation Refinement and Modality-Aware Interaction, which greatly increases its applicability and generalizability. In light of its ranking performance, particularly concerning the mean Average Precision (mAP) metric, we select CSPDarkNet53 as our default backbone architecture.

Comparison with State-of-the-art Methods

Since it is the first work of multi-modality pedestrian detection, we extend five state-of-the-art cross-modality pedestrian

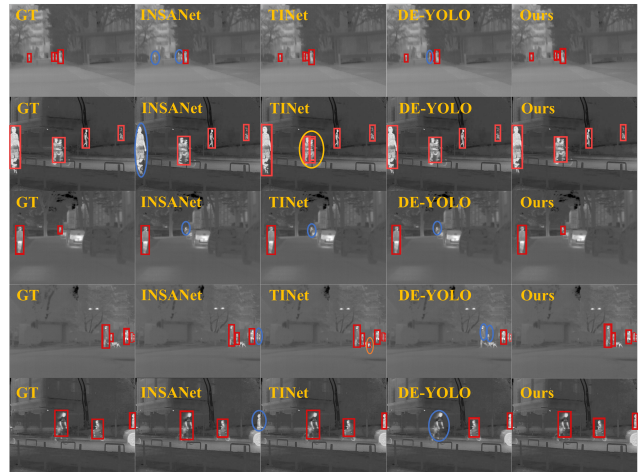


Figure 4: Qualitative Comparison of INSANet, TINet, DE-YOLO, and our AUNet on the TRNT dataset. Red boxes indicate detection results. Blue circles highlight pedestrians missed. Orange circles mark the location of an incorrect detection. Yellow circles represent overlapping multiple detection boxes for the same instance.

trian detection methods (ProbEn (Chen et al. 2022), TINet (Zhang et al. 2023), INSANet (Lee et al. 2024), ICA-Fusion (Shen et al. 2024), and DE-YOLO (Chen et al. 2025)) and fuse the deep features from the NIR modality with their RGB-TIR fused feature maps for comparison. As reported in Table 3, INSANet introduces an attention-based novel fusion network to capture global intra- and inter-information. TINet designs an Intra-modality and Inter-modality module and fuses the inter- and intra-modality features with illumination-guided feature weights for pedestrian detection. DE-YOLO aims to achieve the detection-center mutual enhancement of RGB-TIR by semantic-spatial cross-modality and bi-directional decoupled focus modules. ICA-Fusion and DE-YOLO achieve remarkable performance while handling uncertain-modality scenarios. However, they are still significantly inferior to our AUNet (mAP improves **3.08%** compared to DE-YOLO). The suboptimal performance of INSANet (mAP 13.80%) is due to its approach of fusing features processed through the INSA modules using

Methods	VGG16			ResNet50			CSPDarknet53		
	AP ₅₀ ↑	AP ₇₅ ↑	mAP↑	AP ₅₀ ↑	AP ₇₅ ↑	mAP↑	AP ₅₀ ↑	AP ₇₅ ↑	mAP↑
Baseline	95.68	61.73	55.33	96.02	62.10	56.29	96.10	62.16	56.06
+MAI	96.20	62.98	56.68	96.50	63.62	57.74	96.85	63.95	57.94
+MAI+CSR (AUNet)	96.63	65.49	58.24	97.03	67.30	59.82	97.07	67.78	60.08

Table 4: An ablation study is conducted on the CSR and the MAI within the TRNT framework, utilizing various backbones. The baseline model refers to the direct concatenation of RGB-NIR-TIR modality features without CSR or MAI. The progressive enhancement in detection performance is attributed to the sequential integration of the key modules, one at a time.

Methods	Venue	LLVIP		
		AP ₅₀ ↑	AP ₇₅ ↑	mAP↑
Halfway Fusion	BMVC'16	91.4	60.1	55.1
GAFF	WACV'21	94.0	60.2	55.8
ProEn	ECCV'22	93.4	50.2	51.5
CMDet	TCSVT'22	96.3	-	-
CAMF	TMM'23	89.0	-	55.6
MetaFusion	CVPR'23	91.0	-	56.9
CALNet	MM'23	-	-	63.9
CSAA	CVPR'23	94.3	66.6	59.2
ICAFusion	PR'24	96.3	71.7	62.3
Fusion-mamba	Arxiv'24	96.5	-	62.8
FD ² Net	AAAI'24	96.2	70.0	-
AUNet	—	97.2	73.5	64.8

Table 5: Experimental results of AUNet on LLVIP compared with state-of-the-art methods while handling RGB-TIR cross-modality pedestrian detection data. Where ‘-’ means the corresponding result is not available.

fixed static weights rather than employing adaptive, learnable weights. As depicted in Figure 4, INSA Net and DEYOLO exhibit more missed detections, while TINet tends to generate multiple high-confidence score boxes for one instance and is prone to false positives. Furthermore, we test the FPS of these methods and our AUNet on the RTX 3090 platform. Our AUNet runs at **60.99** Hz, which demonstrates that AUNet combines fast inference speed with robust detection performance.

Furthermore, we conduct a comprehensive evaluation of AUNet with state-of-the-art CMPD methods on LLVIP to evaluate the performance of AUNet in handling RGB-TIR cross-modal pedestrian detection, as shown in Table 5. These detectors include Halfway Fusion (Liu et al. 2016), GAFF (Zhang et al. 2021), CFT (Qingyun, Dapeng, and Zhaokui 2021), ProEn (Chen et al. 2022), CMDet (Sun et al. 2022), CAMF (Tang et al. 2023), MetaFusion (Zhao et al. 2023), CALNet (He et al. 2023), CSAA (Cao et al. 2023), ICA Fusion (Shen et al. 2024), Fusion-mamba (Dong et al. 2024) and FD²Net (Li et al. 2025). Through this rigorous comparison, we aim to highlight the robustness and accuracy of AUNet (marking a **2.0%** mAP enhancement compared to the FD²Net), demonstrating its superior performance in the cross-modal pedestrian detection task.

Ablation Studies

We conduct an ablation study on the CSR and MAI of AUNet on TRNT, as reported in Table 4. As expected, both MAI (mAP ↑ **+1.88%** vs. 56.06%) and UMVR (mAP ↑ **+2.14%** vs. 57.94%) enhance the performance of the baseline, which validates the effectiveness of each module in handling complex RGB-NIR-TIR multi-modal pedestrian detection scenarios. Our AUNet consistently outperforms the concatenation operation across all evaluated scenarios and metrics, appraising the superior effectiveness of the AUNet in multi-modal information fusion. Additionally, we provide a visual figure of the detection results for the ablation of modules.

Conclusion and Discussion

To explore the potential of the near-infrared spectrum, which provides critical pedestrian texture information under low-light conditions, we first establish a spatially and temporally aligned TRNT dataset, breaking the fixed RGB-TIR constraint in existing cross-modal datasets. TRNT provides 8,281 rigorously aligned RGB-NIR-TIR image triplets that account for key challenges in pedestrian detection, including variations in illumination (day → night), occlusion (partial → heavy), weather conditions (spring → winter), perspectives (ground → UAV), and backgrounds (clean → cluttered). To achieve robust detection in real-life applications, where the input modal types and quantities are uncertain, we propose a novel Adaptive Uncertainty-aware Network (AUNet) to dynamically handle arbitrary input combinations of RGB, NIR, and TIR data, which is more flexible than existing CMPD methods. AUNet integrates a Unified Modality Validation Refinement module for reliable modal availability validation and a Modality-Aware Interaction (MAI) module for effectively exploiting complementary information from available modalities. Extensive experiments confirm AUNet’s effectiveness and robust performance under uncertain inputs on both TRNT and LLVIP datasets. AUNet is the unique solution that supports arbitrary modality combinations, whereas existing methods require training a separate model for each input configuration.

Acknowledgments

This work was supported by the Natural Science Foundation of China (62302351, 62501428, and 62376201), Hubei Provincial Science & Technology Talent Enterprise Services Program (2025DJB059), and Hubei Provincial Special Fund for Central-Guided Local S&T Development

(2025CSA017).

References

- Bochkovskiy, A.; Wang, C.-Y.; and Liao, H.-Y. M. 2020. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
- Cao, Y.; Bin, J.; Hamari, J.; Blasch, E.; and Liu, Z. 2023. Multimodal object detection by channel switching and spatial attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 403–411.
- Chen, N.; Xie, J.; Nie, J.; Cao, J.; Shao, Z.; and Pang, Y. 2023. Attentive alignment network for multispectral pedestrian detection. In *Proceedings of the 31st ACM International Conference on Multimedia (MM)*, 3787–3795.
- Chen, Y.; Wang, B.; Guo, X.; Zhu, W.; He, J.; Liu, X.; and Yuan, J. 2025. DEYOLO: Dual-Feature-Enhancement YOLO for Cross-Modality Object Detection. In *International Conference on Pattern Recognition (ICPR)*, 236–252. Springer.
- Chen, Y.-T.; Shi, J.; Ye, Z.; Mertz, C.; Ramanan, D.; and Kong, S. 2022. Multimodal object detection via probabilistic ensembling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 139–158. Springer.
- Dong, W.; Zhu, H.; Lin, S.; Luo, X.; Shen, Y.; Liu, X.; Zhang, J.; Guo, G.; and Zhang, B. 2024. Fusion-mamba for cross-modality object detection. *arXiv preprint arXiv:2404.09146*.
- F. A. Group. 2018. Flir thermal dataset for algorithm training. <https://www.flir.com/oem/adas/adas-dataset-form>. Accessed: 2023-10-26.
- González, A.; Fang, Z.; Socarras, Y.; Serrat, J.; Vázquez, D.; Xu, J.; and López, A. M. 2016. Pedestrian detection at day/night time with visible and FIR cameras: A comparison. *Sensors*, 16(6): 820.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- He, X.; Tang, C.; Zou, X.; and Zhang, W. 2023. Multispectral object detection via cross-modal conflict-aware learning. In *Proceedings of the 31st ACM International Conference on Multimedia (MM)*, 1465–1474.
- Hwang, S.; Park, J.; Kim, N.; Choi, Y.; and So Kweon, I. 2015. Multispectral pedestrian detection: Benchmark dataset and baseline. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1037–1045.
- Jia, X.; Zhu, C.; Li, M.; Tang, W.; and Zhou, W. 2021. Llvip: A visible-infrared paired dataset for low-light vision. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 3496–3504.
- Kim, J.; Kim, H.; Kim, T.; Kim, N.; and Choi, Y. 2021. MLPD: Multi-label pedestrian detector in multispectral domain. *IEEE Robotics and Automation Letters (IEEE RAL)*, 6(4): 7846–7853.
- Kim, T.; Shin, S.; Yu, Y.; Kim, H. G.; and Ro, Y. M. 2024. Causal Mode Multiplexer: A Novel Framework for Unbiased Multispectral Pedestrian Detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 26784–26793.
- Konig, D.; Adam, M.; Jarvers, C.; Layher, G.; Neumann, H.; and Teutsch, M. 2017. Fully convolutional region proposal networks for multispectral person detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 49–56.
- Lee, S.; Kim, T.; Shin, J.; Kim, N.; and Choi, Y. 2024. IN-SANet: INtra-INter spectral attention network for effective feature fusion of multispectral pedestrian detection. *Sensors*, 24(4): 1168.
- Li, C.; Song, D.; Tong, R.; and Tang, M. 2019. Illumination-aware faster R-CNN for robust multispectral pedestrian detection. *Pattern Recognition (PR)*, 85: 161–171.
- Li, K.; Wang, D.; Hu, Z.; Li, S.; Ni, W.; Zhao, L.; and Wang, Q. 2025. Fd2-net: Frequency-driven feature decomposition network for infrared-visible object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 4797–4805.
- Liu, F.; Ye, M.; and Du, B. 2024. Learning a generalizable re-identification model from unlabelled data with domain-agnostic expert. *Visual Intelligence*, 2(1): 28.
- Liu, J.; Fan, X.; Huang, Z.; Wu, G.; Liu, R.; Zhong, W.; and Luo, Z. 2022. Target-aware dual adversarial learning and a multi-scenario multi-modality benchmark to fuse infrared and visible for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5802–5811.
- Liu, J.; Zhang, S.; Wang, S.; and Metaxas, D. N. 2016. Multispectral deep neural networks for pedestrian detection. *arXiv:1611.02644*.
- Qingyun, F.; Dapeng, H.; and Zhaokui, W. 2021. Cross-modality fusion transformer for multispectral object detection. *arXiv preprint arXiv:2111.00273*.
- Shen, J.; Chen, Y.; Liu, Y.; Zuo, X.; Fan, H.; and Yang, W. 2024. Icafusion: Iterative cross-attention guided feature fusion for multispectral object detection. *Pattern Recognition (PR)*, 145: 109913.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, Y.; Cao, B.; Zhu, P.; and Hu, Q. 2022. Drone-based RGB-infrared cross-modality vehicle detection via uncertainty-aware learning. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 32(10): 6700–6713.
- Tang, L.; Chen, Z.; Huang, J.; and Ma, J. 2023. CAMF: An interpretable infrared and visible image fusion network based on class activation mapping. *IEEE Transactions on Multimedia*, 26: 4776–4791.
- Xie, Q.; Cheng, T.-Y.; Dai, Z.; Tran, V.; Trigoni, N.; and Markham, A. 2023. Illumination-Aware Hallucination-Based Domain Adaptation for Thermal Pedestrian Detec-

tion. *IEEE Transactions on Intelligent Transportation Systems (TITS)*.

Xu, Y.; Wu, M.; Guo, Z.; Cao, M.; Ye, M.; and Laaksonen, J. 2025. Efficient text-to-video retrieval via multi-modal multi-tagger derived pre-screening. *Visual Intelligence*, 3(1): 1–13.

Yang, X.; Qian, Y.; Zhu, H.; Wang, C.; and Yang, M. 2022. Baanet: learning bi-directional adaptive attention gates for multispectral pedestrian detection. In *2022 International Conference on Robotics and Automation (ICRA)*, 2920–2926. IEEE.

Zhang, H.; Fromont, E.; Lefèvre, S.; and Avignon, B. 2021. Guided attentive feature fusion for multispectral pedestrian detection. In *Proceedings of the IEEE winter conference on applications of computer vision (WACV)*, 72–80.

Zhang, Y.; Yu, H.; He, Y.; Wang, X.; and Yang, W. 2023. Illumination-guided RGBT object detection with inter-and intra-modality fusion. *IEEE Transactions on Instrumentation and Measurement (TIM)*, 72: 1–13.

Zhao, W.; Xie, S.; Zhao, F.; He, Y.; and Lu, H. 2023. Meta-fusion: Infrared and visible image fusion via meta-feature embedding from object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13955–13965.

Zheng, Y.; Izzat, I. H.; and Ziaee, S. 2019. GFD-SSD: Gated fusion double SSD for multispectral pedestrian detection. *arXiv preprint arXiv:1903.06999*.

Zhou, K.; Chen, L.; and Cao, X. 2020. Improving multispectral pedestrian detection by addressing modality imbalance problems. In *European conference on computer vision*, 787–803. Springer.

Zhou, Z.; Yang, B.; Huang, W.; Chen, J.; and Ye, M. 2025. Unbiased Prototype Consistency Learning for Multi-Modal and Multi-Task Object Re-Identification. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.