

Multispectral Pedestrian Detection with Sparsely Annotated Label

Chan Lee*, Seungho Shin*, Gyeong-Moon Park†, Jung Uk Kim†

Kyung Hee University, Yong-in, South Korea
{cksdlakstp12, ssh9918, gmpark, ju.kim}@khu.ac.kr

Abstract

Although existing Sparsely Annotated Object Detection (SAOD) approaches have made progress in handling sparsely annotated environments in multispectral domain, where only some pedestrians are annotated, they still have the following limitations: (i) they lack considerations for improving the quality of pseudo-labels for missing annotations, and (ii) they rely on fixed ground truth annotations, which leads to learning only a limited range of pedestrian visual appearances in the multispectral domain. To address these issues, we propose a novel framework called Sparsely Annotated Multispectral Pedestrian Detection (SAMPD). For limitation (i), we introduce Multispectral Pedestrian-aware Adaptive Weight (MPAW) and Positive Pseudo-label Enhancement (PPE) module. Utilizing multispectral knowledge, these modules ensure the generation of high-quality pseudo-labels and enable effective learning by increasing weights for high-quality pseudo-labels based on modality characteristics. To address limitation (ii), we propose an Adaptive Pedestrian Retrieval Augmentation (APRA) module, which adaptively incorporates pedestrian patches from ground-truth and dynamically integrates high-quality pseudo-labels with the ground-truth, facilitating a more diverse learning pool of pedestrians. Extensive experimental results demonstrate that our SAMPD significantly enhances performance in sparsely annotated environments within the multispectral domain.

Introduction

Recently, multispectral pedestrian detection has gained attention in computer vision (Chen et al. 2022, 2023; Kim et al. 2024). Unlike single-modal detectors that use only one modality (e.g., visible or thermal), multispectral pedestrian detection combines visible and thermal images (Kim, Park, and Ro 2021a; Kim and Ro 2023). Visible images capture texture and color, while thermal images provide heat signatures. Combining the two modalities enhances detection robustness across various conditions, including low-light environments (Jia et al. 2021; Dasgupta et al. 2022) and adverse weather conditions (Hwang et al. 2015).

However, multispectral pedestrian detection faces challenges due to frequently occurring sparse annotations, of-

*Equal contribution

†Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

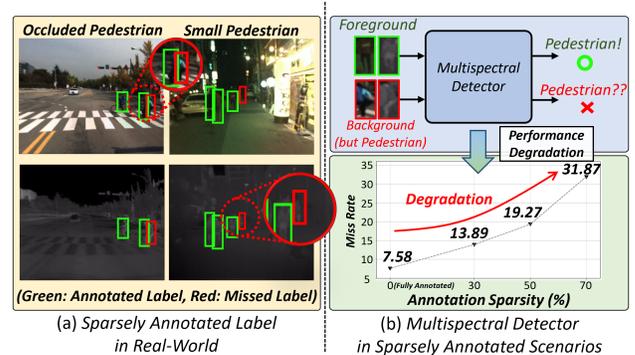


Figure 1: (a) Examples of sparsely annotated labels in multispectral domain. (b) In sparse annotation scenarios, multispectral pedestrian detectors struggle to learn pedestrians.

ten caused by human errors in annotating small or occluded pedestrians, even when they are present (Li et al. 2019). In sparsely annotated environments, a real pedestrian might be annotated as a pedestrian in some instances but not in others (Figure 1(a)). This inconsistency causes the network to struggle with effectively learning knowledge from both modalities, resulting in a significant decline in pedestrian detection performance (Figure 1(b)). For this reason, while some existing multispectral datasets (Hwang et al. 2015; Jia et al. 2021) have reached near-perfect annotation at high labor and time costs, they may still have gaps, making robust research in sparsely annotated environments essential.

To address the above-mentioned issue, Sparsely Annotated Object Detection (SAOD) task has been introduced (Niitani et al. 2019; Zhang et al. 2020; Wang et al. 2021, 2023; Suri et al. 2023). The task aims to generate pseudo-labels to address missing annotations. Current methods include selecting high-confidence boxes from model predictions (Niitani et al. 2019; Wang et al. 2023), adjusting loss functions to incorporate both original and augmented images (Wang et al. 2021), and applying self-supervised loss to avoid negative gradient propagation (Suri et al. 2023). However, these methods often assume pseudo-labels are reliable, even when they are inaccurate or do not fully capture pedestrian appearance. Furthermore, reliance on a fixed set of ground-truth annotations makes it challenging to incorporate valuable missing annotations, limiting the ability to

learn about diverse pedestrian visual appearances.

In this paper, we present a novel method, Sparsely Annotated Multispectral Pedestrian Detection (SAMPD), to tackle the challenges of sparse annotation in the multispectral domain. Our approach considers two main challenges: (i) how to effectively learn multispectral pedestrian information from pseudo-labels and enhance their quality, and (ii) how to integrate the identified missing annotations during training to enable more comprehensive learning.

To address the challenge (i), we first adopt a teacher-student structure, as used in existing SAOD works. However, we newly introduce two modules: Multispectral Pedestrian-aware Adaptive Weight (MPAW) and Positive Pseudo-label Enhancement (PPE). The MPAW module aims to help the student model learn multispectral modalities by assigning higher weights to high-quality pseudo-labels based on modality characteristics (single and multispectral). The PPE module aligns feature representations to make high-quality pseudo-labels more similar to each other, while distancing them from low-quality pseudo-labels. We also consider the uncertain pseudo-labels to prevent them from misguiding the model. This approach allows the student model to learn from more reliable pseudo-labels, thereby improving robustness in environments with sparse annotations.

For the challenge (ii), we propose an Adaptive Pedestrian Retrieval Augmentation (APRA) module. In sparsely annotated environments, finding missing annotations by leveraging diverse visual representations of pedestrians is crucial. The APRA module adaptively attaches ground-truth pedestrian patches that best match the input image based on lighting conditions. Additionally, high-quality pseudo-labels are dynamically integrated with the ground-truth. This dynamic augmentation approach enriches the visual representation of pedestrians, allowing the teacher model to generate more reliable pseudo-labels and enhancing the robustness of the student model, which is our final detector during inference.

As a result, SAMPD effectively addresses sparsely annotated settings by generating more reliable pseudo-labels. It also improves performance in fully annotated scenarios, as real-world images may still have missing annotations.

The major contributions can be summarized as follows:

- We propose the Multispectral Pedestrian-aware Adaptive Weight (MPAW) module to adjust the learning of each modality for improved pseudo-label learning.
- We introduce the Positive Pseudo-label Enhancement (PPE) module to help our SAMPD generate higher-quality pseudo-labels.
- We develop the Adaptive Pedestrian Retrieval Augmentation (APRA) module to enrich pedestrian knowledge by adaptively augmenting images and integrating high-quality pseudo-labels into ground-truth labels.

Related Work

Multispectral Pedestrian Detection

Multispectral pedestrian detection has gained attention for its impressive performance (Park et al. 2022; Hu, Zhang, and Weng 2023; Xu et al. 2023; Zhu et al. 2023). AR-CNN

tackles modality alignment issues (Zhang et al. 2019), while IATDNN+IASS introduces illumination-aware mechanisms for robustness (Guan et al. 2019). MBNet and AANet address modal discrepancy (Zhou, Chen, and Cao 2020; Chen et al. 2023), and MLPD proposes multi-label and non-paired augmentation (Kim et al. 2021). ProbEn adopts late-fusion with probability ensembling (Chen et al. 2022), and DCM-Net uses local and non-local aggregation for contextual information (Xie et al. 2022). Liu et al. (Liu et al. 2024) improves performance by utilizing illumination and temperature information. Beyond Fusion (Xie et al. 2024) introduces a hallucination branch to map thermal to visible domains.

While these studies have shown promising detection performances, they assume perfect bounding box annotations (Kim, Park, and Ro 2021b,a; Kim and Ro 2023), which may not be realistic in real-world scenarios. Challenges such as differences between visible and thermal images, small or occluded pedestrians, and human error can lead to imperfect (*i.e.*, sparse) annotations. Existing methods struggle in such scenarios. To address this, we introduce Sparse Annotation Multispectral Pedestrian Detection (SAMPD), which fully leverages multispectral knowledge to effectively generate and enhance pseudo-labels even with sparse annotations.

Sparsely Annotated Object Detection

When pedestrian annotations are incomplete, it can lead to inconsistencies where pedestrians are misclassified as either foreground or background. To tackle this issue, Sparsely Annotated Object Detection (SAOD) has introduced. Pseudo label methods (Niitani et al. 2019) establish logical connections between object co-occurrence and pseudo-label application. BRL (Zhang et al. 2020) introduces automatic adjustments for areas prone to mislabeling. Co-mining (Wang et al. 2021) jointly trains models using predictions from both original and augmented images. SparseDet (Suri et al. 2023) introduces a self-supervised loss to prevent negative gradient propagation. Calibrated Teacher (Wang et al. 2023) validates pseudo-label quality using a calibrator that distinguishes positive from negative labels. However, existing methods rely on the pseudo-labels without considering their quality. In contrast, our SAMPD enhances pseudo-label quality by integrating multispectral knowledge and incorporating filtered high-quality pseudo-labels through the Adaptive Pedestrian Retrieval Augmentation (APRA) module.

Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) has proven to enhance the accuracy and reliability of generative AI, addressing the challenge of limited access to the latest information after training phase. Initially developed for natural language processing (NLP) (Lewis et al. 2021), RAG has proven effective in various applications, including image captioning (Ramos, Elliott, and Martins 2023; Sarto et al. 2022), multimodal learning. Traditionally, it compensates for the lack of information in high-accuracy modalities by storing features of that modality in memory. We propose a novel RAG method, *i.e.*, APRA module, for effective multispectral pedestrian detection. By leveraging the unique characteristics of multispectral data, our APRA module enables

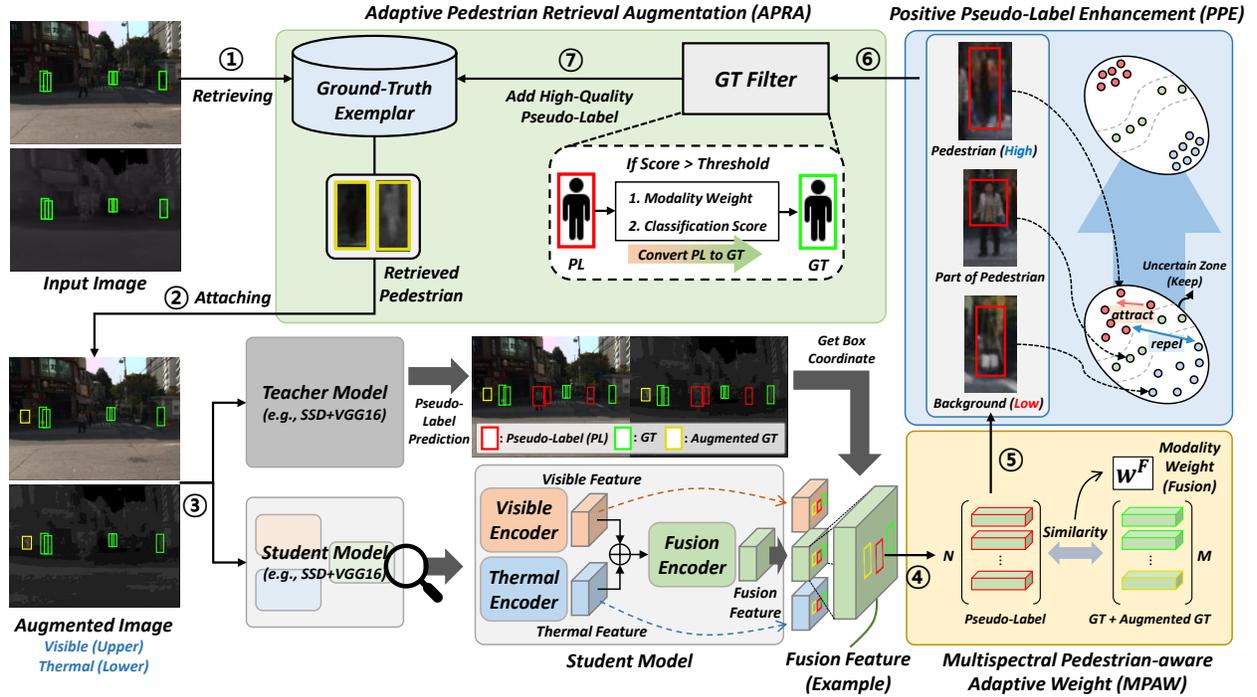


Figure 2: Overall architecture of our SAMPD. Given the multispectral input image, the APRA module generates an augmented image, which is then used as the input for the training phase. \oplus denotes the concatenation. During inference phase, only student model is used for detection, and only the original input image (without augmentation) is utilized.

the learning of diverse pedestrian information. It helps create a teacher model that generates more reliable pseudo-labels, improving the overall accuracy of pedestrian detection.

Proposed Method

Figure 2 shows the overall architecture of our SAMPD in the training phase. The APRA module retrieves m pedestrian patches to create augmented image pair. The teacher and student model receive an augmented image pair and encode through each modality backbone. Then, the teacher model generates pseudo-labels and passes them to the student model. The student model employs the MPAW module that assigning higher weights to high-quality pseudo-labels. The PPE module guides the pseudo-labels in the feature space to distinguish the high and low quality pseudo-labels. Finally, the APRA enriches pedestrian knowledge by adaptively augmenting images and integrating high-quality pseudo-labels into the ground-truth annotation.

Multispectral Pedestrian-aware Adaptive Weight

As shown in Figure 2, we adopt a teacher-student structure commonly used in SAOD methods. The student model, our final detector, learns from pseudo-labels generated by the teacher model. The multispectral data comprises visible (V), thermal (T), and fusion (F) modalities. To generate effective pseudo-labels in sparsely annotated scenarios, both the teacher and student models are constructed with 3-way encoding paths, fully exploiting single-modal (V and T) and multispectral (F) knowledge during training.

However, some pseudo-labels from the teacher model may be incorrect, capturing only parts of pedestrians or including background elements. These errors can negatively impact the performance of the student model and cause confusion during training. To address this, we propose a multispectral pedestrian-aware adaptive weight (MPAW) module to increase the influence of reliable pseudo-labels and reduce the impact of unreliable ones for each modality. This approach emphasizes high-quality pseudo-labels and reliable modalities. For each modality $k = \{V, T, F\}$, we extract N feature maps from the pseudo-label boxes of the student model $\mathbf{f}_{PL}^{k(s)} = \{f_{PL_i}^{k(s)}\}_{i=1}^N$ and M feature maps from the ground-truth labels $\mathbf{f}_{GT}^{k(s)} = \{f_{GT_i}^{k(s)}\}_{i=1}^M$. We then apply global average pooling (GAP) to obtain latent vectors $\mathbf{l}_{PL}^{k(s)} = \{l_{PL_i}^{k(s)}\}_{i=1}^N$ and $\mathbf{l}_{GT}^{k(s)} = \{l_{GT_i}^{k(s)}\}_{i=1}^M$. We compute the modality weight for k modality, w^k , as defined by:

$$w^k = \frac{1}{N} \sum_{i=1}^N \max_{j \in \{1, \dots, M\}} d(l_{PL_i}^{k(s)}, l_{GT_j}^{k(s)}), \quad (1)$$

$$d(\alpha, \beta) = \frac{\alpha \cdot \beta}{\|\alpha\| \|\beta\|}, \quad (2)$$

where N and M denote the number of pseudo-labels and ground-truth labels, respectively, and $d(\cdot, \cdot)$ is the cosine similarity. We compute the maximum cosine similarity between $l_{PL_i}^{k(s)}$ and the closest vector in $l_{GT}^{k(s)}$. A high w^k indicates that the pseudo-labels for modality k closely match the ground-truth labels, reflecting reliable quality, while a low w^k signifies less reliable pseudo-labels.

Using w^k , the detection losses \mathcal{L}_{det}^{sum} from the 3-way encoding paths are computed as follows:

$$\mathcal{L}_{det}^{sum} = w^V \mathcal{L}_{det}^{V(s)} + w^T \mathcal{L}_{det}^{T(s)} + w^F \mathcal{L}_{det}^{F(s)}, \quad (3)$$

where $\mathcal{L}_{det}^{k(s)}$ represents the detection loss (Kim et al. 2021; Kim, Park, and Ro 2021a) of the student model, including the classification loss $\mathcal{L}_{cls}^{k(s)}$ and the localization loss $\mathcal{L}_{loc}^{k(s)}$ for k modality. By applying Eq. (3), the teacher model focuses on reliable pseudo-labels and modalities, enabling more stable learning in sparsely annotated scenarios.

Positive Pseudo-label Enhancement

Through the MPAW module, the student model can learn from the pseudo-labels. In addition, we introduce a Positive Pseudo-label Enhancement (PPE) module to enable the teacher model to generate higher-quality pseudo-labels. With $\mathbf{I}_{PL}^{k(s)}$ and $\mathbf{I}_{GT}^{k(s)}$ from student model of k modality ($k = V, T, F$), we calculate the cosine similarity $d(\cdot, \cdot)$ to construct the *positive* (foreground) and *negative* (background) pseudo-labels, which can be represented as:

$$a = \arg \max_{j \in \{1, \dots, M\}} d(l_{PL_i}^{k(s)}, l_{GT_j}^{k(s)}). \quad (4)$$

We classify the i -th pseudo-label as *positive* pseudo-label, if the similarity between the i -th pseudo-label and the most similar ground-truth label exceeds the τ_1 . Conversely, if the similarity is below τ_2 , i -th pseudo-label is classified as *negative* pseudo-label. Also, the pseudo-labels with similarity scores between τ_1 and τ_2 fall into an *uncertain* pseudo-label, acting as a buffer to avoid bias towards either *positive* or *negative*, thereby improving the reliability of the learning process. We set $\tau_1 = 0.9$ and $\tau_2 = 0.7$, respectively.

After the distinction, we devise a positive pseudo-label guiding (PG) loss for the k modality \mathcal{L}_{PG}^k , which can be represented as:

$$p_{PL}(p_i) = \sum_{j=1}^{N_p} \exp(d(l_{PL_{p_i}}^{k(s)}, l_{PL_{p_j}}^{k(s)})) / \tau, \quad (5)$$

$$n_{PL}(p_i) = \sum_{j=1}^{N_n} \exp(d(l_{PL_{p_i}}^{k(s)}, l_{PL_{n_j}}^{k(s)})) / \tau, \quad (6)$$

$$\mathcal{L}_{PG}^k = -\frac{1}{N_p} \sum_{i=1}^{N_p} \log \frac{p_{PL}(p_i)}{p_{PL}(p_i) + n_{PL}(p_i)}, \quad (7)$$

where p and n indicates *positive* and *negative* pseudo-labels, N_p and N_n is the number of *positive* and *negative* pseudo-labels, and τ is the temperature parameter.

The purpose of the PG loss is to use the differentiated *positive* and *negative* labels from the student model to guide the feature representations of *positive* pseudo-labels, promoting those similar to the ground-truth labels provided by the teacher model to move closer together. This process helps train the teacher model to generate higher-quality pseudo-labels. Since the PG loss is applied across visible, thermal, and fusion modalities, it leverages diverse multimodal knowledge, enhancing learning stability.

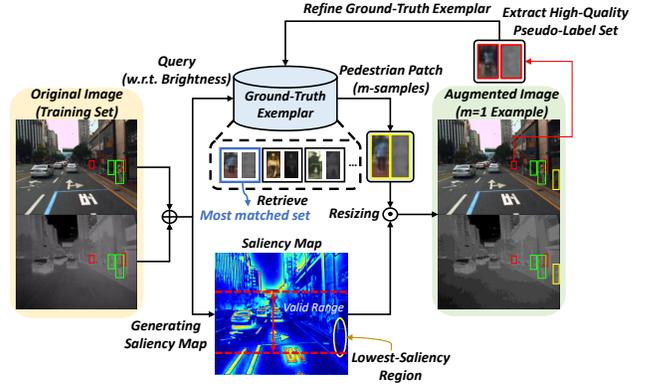


Figure 3: Given an original image, our adaptive pedestrian retrieval augmentation (APRA) module finds m pedestrian patches from a ground-truth exemplar with similar brightness to the original image ($m = 1$ example). The APRA module uses the saliency map of the original image to locate the region with the lowest saliency value and attaches the resized m pedestrian patches, according to the ground-truth bounding-box size, to this region. In addition, if the pseudo-labels are considered reliable, the corresponding pedestrian patch is included in the ground-truth exemplar and used as ground-truth in the subsequent training process. \oplus denotes the concatenation, while \odot indicates the attaching process.

Adaptive Pedestrian Retrieval Augmentation

While MPAW and PPE modules effectively utilize pseudo-labels, there remains a challenge due to the limited visual diversity in ground-truth annotations within sparsely annotated environments. To address this, we propose an Adaptive Pedestrian Retrieval Augmentation (APRA) module that adaptively integrates additional pedestrian patches into the original image to enhance the ground-truth information.

Augmenting Images with Pedestrian Patches. Figure 3 shows the operation of the APRA module. A ground-truth exemplar retrieves relevant pedestrian patches by selecting m patches with the smallest brightness difference from the original image. These patches are resized to the average annotation size for each image before being integrated into the image. If no annotations are present, the patches are resized to the average size of all annotations before being attached.

To ensure effective integration of these patches, we carefully select regions within the image, avoiding areas with a low likelihood of containing pedestrians, such as the sky or directly in front of vehicles. We define valid pedestrian locations using the mean, variance, and standard deviation of the y -coordinates from existing annotations within a 90% confidence interval. A saliency map identifies regions with a low probability of containing pedestrians, targeting these areas for patch integration. By focusing on low-saliency regions, the APRA module enhances the ground-truth representation, leading to improved performance.

Dynamic Ground-Truth Refinement. Additionally, we re-

Methods	30%			50%			70%		
	All	Day	Night	All	Day	Night	All	Day	Night
Supervised	13.89	15.50	10.88	19.27	21.16	15.66	31.87	33.73	27.81
Pseudo label (CVPR'19)	11.95	14.09	8.14	18.09	21.00	13.33	29.43	31.24	26.00
BRL (ICASSP'20)	11.90	13.61	8.69	17.93	19.62	14.42	28.71	32.10	21.48
Co-mining (AAAI'21)	11.66	12.68	9.80	17.90	19.43	14.91	28.80	28.48	29.41
SparseDet (ICCV'23)	10.92	11.99	8.78	18.07	<u>19.30</u>	15.40	28.20	28.80	26.54
Calibrated Teacher (AAAI'23)	<u>10.47</u>	<u>11.81</u>	<u>7.82</u>	<u>17.67</u>	<u>19.77</u>	<u>13.19</u>	<u>25.48</u>	<u>28.43</u>	<u>19.28</u>
SAMPD (Ours)	8.56	10.55	5.62	15.27	17.28	11.15	23.52	26.15	17.87

Table 1: Sparsely annotated detection results (MR) on the KAIST dataset by varying removal percentages (30%, 50%, and 70%) for sparsely annotated scenarios. We compared our method with state-of-the-art SAOD methods that address sparsely annotated scenario. **Bold/underlined** fonts indicate the best/second-best results.

fine ground-truth annotations by dynamically converting pseudo-labels into ground-truth labels. This process uses w^k ($k = V, T, F$) from Eq. (1) and the classification score of the pseudo-label. If both values exceed the threshold τ_1 (indicating high-quality pseudo-labels in the PPE module), the pseudo-label is converted into a ground-truth label. These converted labels are then integrated into the ground-truth annotations, ensuring consistent recognition of reliable pseudo-labels and improving model stability. Through the APRA module, our framework achieves more reliable annotations, reducing gaps in ground-truth information within sparsely annotated environments.

Total Loss

The total loss function of our SAMPD is represented as:

$$\mathcal{L}_{Total} = \lambda_1 \mathcal{L}_{det}^{sum} + \lambda_2 (\mathcal{L}_{PG}^V + \mathcal{L}_{PG}^T + \mathcal{L}_{PG}^F), \quad (8)$$

where \mathcal{L}_{det}^{sum} denotes the detection loss using. λ_1 and λ_2 denote the balancing hyper-parameters. Using the \mathcal{L}_{Total} , our SAMPD shows the robust detection performance even in the sparsely annotation scenarios.

Experiments

Dataset and Evaluation Metric

KAIST Dataset. The KAIST dataset (Hwang et al. 2015) comprises 95,328 pairs of visible and thermal images, enriched with 103,128 bounding box annotations to identify pedestrians. We use a test set of 2,252 images to evaluate performance. Following (Suri et al. 2023), we simulated a sparsely annotated scenario by increasing the probability of removing bounding-box annotations with smaller widths from the training set. More details are in the supplementary document. Finally, we removed 30%, 50%, and 70% of the bounding-box annotations among the total annotations, consistent with the ratios described in (Niitani et al. 2019).

LLVIP Dataset. The LLVIP dataset (Jia et al. 2021) contains visible-thermal paired dataset for low-light vision. It consists 15,488 visible-thermal image pairs. Following the same protocol as the KAIST dataset, we simulated a sparsely annotated scenario for the LLVIP dataset by removing 30%, 50%, and 70% of the bounding-box annotations from the total annotations.

Methods	MR			AP ₅₀		
	30%	50%	70%	30%	50%	70%
Supervised	11.87	14.25	17.77	92.31	89.78	87.11
Pseudo label (CVPR'19)	10.91	13.21	17.45	93.85	92.41	89.49
BRL (ICASSP'20)	10.93	12.74	17.24	94.46	93.37	90.89
Co-mining (AAAI'21)	10.88	12.89	16.51	93.37	91.97	88.92
SparseDet (ICCV'23)	10.06	12.60	15.82	94.34	91.80	88.97
Calibrated Teacher (AAAI'23)	<u>9.41</u>	<u>12.10</u>	<u>15.57</u>	<u>95.27</u>	<u>93.75</u>	<u>91.51</u>
SAMPD (Ours)	7.65	9.03	11.71	95.38	94.39	92.08

Table 2: Detection results (MR and AP₅₀) on the LLVIP dataset by varying removal percentages (30%, 50%, and 70%) for sparsely annotated scenarios. **Bold/underlined** fonts indicate the best/second-best results.

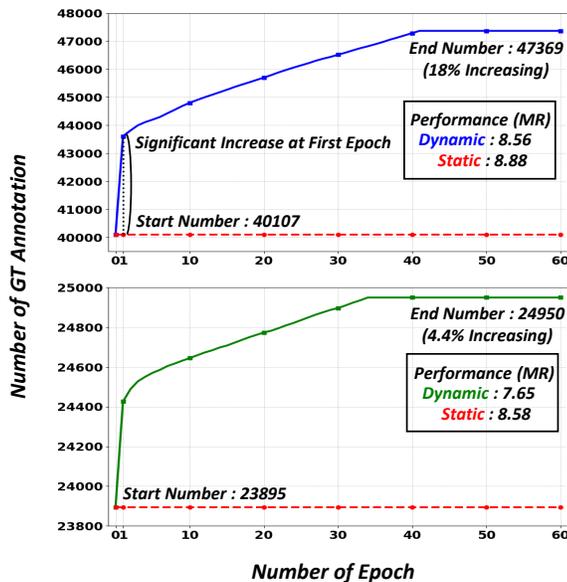
MPAW	PPE	APRA	30%		
			All	Day	Night
-	-	-	13.89	15.50	10.88
✓	-	-	10.14	11.18	7.82
✓	✓	-	9.65	10.89	7.00
✓	-	✓	9.23	10.76	6.26
✓	✓	✓	8.56	10.55	5.62

Table 3: Ablation studies of our multispectral pedestrian-aware adaptive weight (MPAW), positive pseudo-label enhancement (PPE), and adaptive pedestrian retrieval augmentation (APRA) on the KAIST dataset (30% removal).

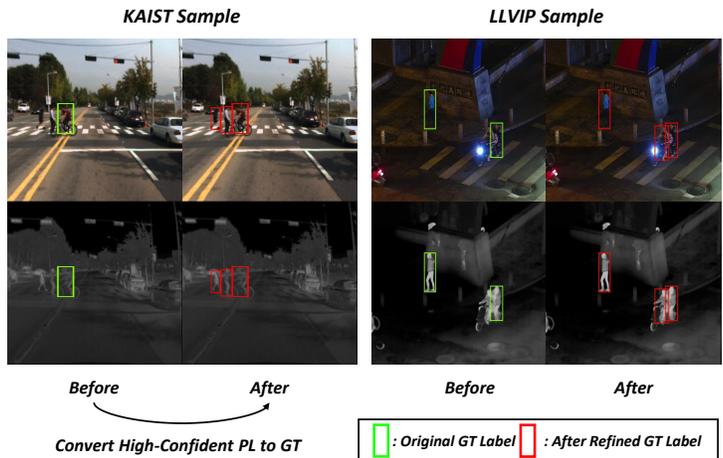
Evaluation Metric. Following (Kim, Park, and Ro 2022; Kim et al. 2021), we measure the performance by adopting miss rate (MR), which is averaged over the range of false positives per image (FPPI), spanning from $[10^{-2}, 10^0]$. The detection performance improves as the MR decreases. Following (Kim et al. 2021), evaluations were conducted across three distinct settings: ‘All’, ‘Day’, and ‘Night’. For the LLVIP dataset, we also use Average Precision (AP) with an Intersection over Union (IoU) threshold of 0.5 (AP₅₀).

Implementation Details

We deploy the SAOD methodology leveraging an SSD (Liu et al. 2016) structure combined with a VGG16 (Simonyan and Zisserman 2015) backbone. We optimize our framework using Stochastic Gradient Descent (SGD) (Kiefer and Wolfowitz 1952), coordinating the process across two GTX 3090 GPUs and processing 6 images in each mini-batch. The



(a) Comparison of GT Annotation (Static vs. Dynamic)



(b) Dynamic Refined Ground-Truth Examples

Figure 4: (a) Line graph of the changes in ground-truth annotations, and (b) visualization examples of the refined ground-truth. It shows that our APRA module (‘Dynamic’) effectively captures missed annotations in sparsely annotated environments.

augmentation include random horizontal flipping, color jittering, and cutout. Parameters are set with $m = 1$, $\tau = 0.1$ and $\lambda_1 = \lambda_2 = 1$. We train our detector for 80 epochs with 0.0001 learning rate. All experimental procedures are performed utilizing the Pytorch framework (Paszke et al. 2017).

Comparisons

Results on the KAIST Dataset. Table 1 shows the performance of our SAMPD with the state-of-the-art SAOD methods (Niitani et al. 2019; Zhang et al. 2020; Wang et al. 2021; Suri et al. 2023; Wang et al. 2023) on the KAIST dataset. As removal percentages decrease from 70% to 30%, existing methods showed some improvement over the baseline (‘Supervised’), which is trained only on sparsely annotated bounding boxes. However, SAMPD outperforms them across ‘All’, ‘Day’, and ‘Night’ settings, with notable advantages at a 70% removal rate. It highlights the effectiveness of our method in commonly encountered real-world scenarios with sparse or missing annotations.

Results on the LLVIP Dataset. We also conduct experiments on the LLVIP dataset by varying removal percentages (30%, 50%, and 70%). As shown in Table 2, Calibrated Teacher (Wang et al. 2023) achieves the highest improvements among the existing methods. In contrast, our approach outperforms the Calibrated Teacher. By incorporating the proposed components during the training phase, our SAMPD effectively addresses sparsely annotated scenarios.

Ablation Study

We conduct the ablation study to explore the effect of the proposed modules, *i.e.*, multispectral pedestrian-aware

Augmentation Methods	MPAW	PPE	30%		
			All	Day	Night
Supervised	-	-	13.89	15.50	10.88
Robust Teacher (CVIU’23)	✓	-	11.25	12.85	8.09
Ours (Static)	✓	-	9.35	10.86	6.51
Ours (Dynamic)	✓	-	9.23	10.76	6.26
Robust Teacher (CVIU’23)	✓	✓	10.45	12.44	6.82
Ours (Static)	✓	✓	8.88	10.57	6.29
Ours (Dynamic)	✓	✓	8.56	10.55	5.62

Table 4: Effect of pedestrian augmentation method. ‘Static’ refers to our APRA method without a dynamic adding mechanism and ‘Dynamic’ refers dynamic adding mechanism.

adaptive weight (MPAW) module, positive pseudo-label enhancement (PPE) module, and adaptive pedestrian retrieval augmentation (APRA) module. We conducted the ablation study with 30% sparsely annotated scenario on the KAIST dataset. As shown in Table 3, with the consideration of each module, the performances of our method are consistently improved. When all our modules are considered, we achieve the highest performance. In summary, our SAMPD effectively learns pedestrian representations in sparse annotation scenarios by: (1) reducing the impact of low-quality pseudo-labels with the multispectral pedestrian-aware adaptive weight W_{AL} in MPAW module, (2) using PPE module to enhance the teacher model in generating higher-quality pseudo-labels, and (3) incorporating a wide range of pedestrian patches through APRA module to enrich pedestrian knowledge during training.

Dataset	\mathcal{R}	All	Day	Night	Dataset	\mathcal{R}	All
KAIST	70%	23.52	26.15	17.87	LLVIP	70%	11.71
	50%	15.27	17.28	11.15		50%	9.73
	30%	8.56	10.55	5.62		30%	7.65
	0%	7.58	7.95	6.95		0%	6.01

Table 5: Comparison of performance (MR) between our SAMPD at different removal percentages (\mathcal{R} : 30%, 50%, 70%) fully annotated baseline (0% annotation removal scenario) on the KAIST and LLVIP datasets.

Methods	0% (Fully Labeled)		
	All	Day	Night
Supervised	7.58	7.95	6.95
Calibrated Teacher (AAAI'23)	9.87 (-2.29)	11.25 (-3.30)	7.24 (-0.29)
SAMPD (Ours)	6.50 (+1.08)	6.85 (+1.10)	5.99 (+0.96)

Table 6: Detection results (MR) of our SAMPD and Calibrated Teacher in the fully annotated setting (0% annotation removal scenario) on the KAIST dataset.

Discussions

APRA Module (Static vs. Dynamic (Ours)). Our APRA module has a dynamic property that uses high-quality pseudo-labels as new ground-truths. Figure 4 illustrates the changes over epochs during training on the KAIST and LLVIP datasets. As shown in Figure 4(a), the number of annotations increases progressively with each epoch. The dynamic mechanism of the APRA module learns various visual appearances of potentially missing pedestrians and shows improved performance compared to the static approach (*i.e.*, without dynamic refining), with improvements such as 8.88 to 8.56 (KAIST dataset) and 8.58 to 7.65 (LLVIP dataset). In addition, Figure 4(b) demonstrates that the dynamic mechanism significantly reduces annotation sparsity, successfully filling in many previously missing annotations.

Pedestrian Augmentation. To see the effectiveness of our APRA module, we compared it with the recent pedestrian augmentation approach, Robust Teacher (Li et al. 2023). As shown in Table 4, while the Robust Teacher improved performance over the baseline ('Supervised'), our dynamic approach of APRA module outperformed it. In fact, since Robust Teacher does not account for scene-specific characteristics (*e.g.*, scene brightness, patch sizes, etc.) leads to boundary bias and visual distortion from augmentation that do not match the lighting conditions of the scene. As a result, it performs worse compared to the static approach (w/o dynamic refinement).

SAMPD in Sparsely-/Fully-Annotated Scenarios. Table 5 compares the performance our SAMPD at various removal percentages (30%, 50%, and 70%) with the fully annotated baseline (0% removal) on the KAIST and LLVIP datasets. Our method at 30% removal percentage shows comparable performance to the fully annotated scenarios of the baseline. Interestingly, despite the 30% removal percentage of anno-



Figure 5: Visualization of APRA module on (a) KAIST dataset and (b) LLVIP dataset (left: visible, right: thermal). Yellow boxes are pedestrian patches from APRA module.

tations in the KAIST dataset, our method shows an improvement performance (1.33 MR) on 'Night' setting.

Also, Table 6 shows that applying our method in the 0% removal scenario improves performance across all cases. This demonstrates that our SAMPD, by diversifying pedestrian sample knowledge through the proposed modules, is effective in both sparsely annotated environment and fully annotated environment, *where potential incompleteness may exist*.

Limitation. In SAMPD, while various results show the effectiveness of the APRA module, it currently offers only image-level guidance and samples pedestrians based on overall image brightness rather than individual characteristics. Thus, exploring a method to effectively guide at both the image-level and feature-level, considering individual pedestrian characteristics in addition to the entire image, can be a promising avenue for our future work.

Visualization Results

Effect of Adaptive Pedestrian Retrieval Augmentation. The APRA module enriches pedestrian information by attaching patches with similar brightness to the original image, ensuring smooth integration. Figure 5 shows examples of this augmentation of APRA module on the KAIST and LLVIP datasets. The APRA module integrates these patches without disrupting original image by placing them only in low-saliency regions, preserving visual appearance of pedestrians. More examples are in the supplementary document.

Conclusion

We present a new framework for multispectral pedestrian detection in sparsely annotated scenarios. Our method includes three key modules: MPAW module to increase the transfer weight of reliable pseudo-label knowledge, PPE module to guide the teacher model in generating better pseudo-labels, and APRA module to adaptively refine ground-truth annotations with pedestrian patches. Experimental results demonstrate the effectiveness of our approach in handling sparse annotations. We believe our method provides valuable insights for various sparsely annotated scenarios.

Acknowledgements

This work was supported by the NRF grant funded by the Korea government (MSIT) (No. RS-2023-00252391), and by IITP grant funded by the Korea government (MSIT) (No. RS-2022-00155911: Artificial Intelligence Convergence Innovation Human Resources Development (Kyung Hee University), IITP-2022-II220078: Explainable Logical Reasoning for Medical Knowledge Generation, No. RS-2024-00509257: Global AI Frontier Lab), and by the MSIT (Ministry of Science and ICT), Korea, under the National Program for Excellence in SW (2023-0-00042) supervised by the IITP in 2025, and conducted by CARAI grant funded by DAPA and ADD (UD230017TD).

References

- 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*, 3787–3795.
- Chen, Y.-T.; Shi, J.; Ye, Z.; Mertz, C.; Ramanan, D.; and Kong, S. 2022. Multimodal object detection via probabilistic ensembling. In *European Conference on Computer Vision*, 139–158. Springer.
- Dasgupta, K.; Das, A.; Das, S.; Bhattacharya, U.; and Yogamani, S. 2022. Spatio-Contextual Deep Network-Based Multimodal Pedestrian Detection for Autonomous Driving. *IEEE Transactions on Intelligent Transportation Systems*, 23(9): 15940–15950.
- Guan, D.; Cao, Y.; Yang, J.; Cao, Y.; and Yang, M. Y. 2019. Fusion of multispectral data through illumination-aware deep neural networks for pedestrian detection. *Information Fusion*, 50: 148–157.
- Hu, Y.; Zhang, N.; and Weng, L. 2023. Retrieve the Visible Feature to Improve Thermal Pedestrian Detection Using Discrepancy Preserving Memory Network. In *2023 IEEE International Conference on Image Processing (ICIP)*, 1125–1129.
- 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*, 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/CVF international conference on computer vision*, 3496–3504.
- Kiefer, J.; and Wolfowitz, J. 1952. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 462–466.
- 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*, 6(4): 7846–7853.
- Kim, J. U.; Park, S.; and Ro, Y. M. 2021a. Robust small-scale pedestrian detection with cued recall via memory learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3050–3059.
- Kim, J. U.; Park, S.; and Ro, Y. M. 2021b. Uncertainty-guided cross-modal learning for robust multispectral pedestrian detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(3): 1510–1523.
- Kim, J. U.; Park, S.; and Ro, Y. M. 2022. Towards versatile pedestrian detector with multisensory-matching and multispectral recalling memory. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36(1), 1157–1165.
- Kim, J. U.; and Ro, Y. M. 2023. Similarity Relation Preserving Cross-Modal Learning for Multispectral Pedestrian Detection Against Adversarial Attacks. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- 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/CVF Conference on Computer Vision and Pattern Recognition*, 26784–26793.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv:2005.11401.
- Li, C.; Song, D.; Tong, R.; and Tang, M. 2019. Illumination-aware faster R-CNN for robust multispectral pedestrian detection. *Pattern Recognition*, 85: 161–171.
- Li, S.; Liu, J.; Shen, W.; Sun, J.; and Tan, C. 2023. Robust Teacher: Self-correcting pseudo-label-guided semi-supervised learning for object detection. *Computer Vision and Image Understanding*, 235: 103788.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, 21–37. Springer.
- Liu, Y.; Hu, C.; Zhao, B.; Huang, Y.; and Zhang, X. 2024. Region-Based Illumination-Temperature Awareness and Cross-Modality Enhancement for Multispectral Pedestrian Detection. *IEEE Transactions on Intelligent Vehicles*.
- Niitani, Y.; Akiba, T.; Kerola, T.; Ogawa, T.; Sano, S.; and Suzuki, S. 2019. Sampling Techniques for Large-Scale Object Detection from Sparsely Annotated Objects. arXiv:1811.10862.
- Park, S.; Choi, D. H.; Kim, J. U.; and Ro, Y. M. 2022. Robust thermal infrared pedestrian detection by associating visible pedestrian knowledge. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4468–4472. IEEE.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. *NeurIPS-W*.
- Ramos, R.; Elliott, D.; and Martins, B. 2023. Retrieval-augmented Image Captioning. arXiv:2302.08268.
- Sarto, S.; Cornia, M.; Baraldi, L.; and Cucchiara, R. 2022. Retrieval-Augmented Transformer for Image Captioning. arXiv:2207.13162.

Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556.

Suri, S.; Rambhatla, S. S.; Chellappa, R.; and Shrivastava, A. 2023. SparseDet: Improving Sparsely Annotated Object Detection with Pseudo-positive Mining. arXiv:2201.04620.

Wang, H.; Liu, L.; Zhang, B.; Zhang, J.; Zhang, W.; Gan, Z.; Wang, Y.; Wang, C.; and Wang, H. 2023. Calibrated Teacher for Sparsely Annotated Object Detection. arXiv:2303.07582.

Wang, T.; Yang, T.; Cao, J.; and Zhang, X. 2021. Co-mining: Self-Supervised Learning for Sparsely Annotated Object Detection. arXiv:2012.01950.

Xie, J.; Anwer, R. M.; Cholakkal, H.; Nie, J.; Cao, J.; Laaksonen, J.; and Khan, F. S. 2022. Learning a dynamic cross-modal network for multispectral pedestrian detection. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4043–4052.

Xie, Q.; Cheng, T.-Y.; Zhong, J.-X.; Zhou, K.; Markham, A.; and Trigoni, N. 2024. Beyond Fusion: Modality Hallucination-based Multispectral Fusion for Pedestrian Detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 655–664.

Xu, X.; Zhan, W.; Zhu, D.; Jiang, Y.; Chen, Y.; and Guo, J. 2023. Contour Information-Guided Multi-Scale Feature Detection Method for Visible-Infrared Pedestrian Detection. *Entropy*, 25(7).

Zhang, H.; Chen, F.; Shen, Z.; Hao, Q.; Zhu, C.; and Savvides, M. 2020. Solving Missing-Annotation Object Detection with Background Recalibration Loss. arXiv:2002.05274.

Zhang, L.; Zhu, X.; Chen, X.; Yang, X.; Lei, Z.; and Liu, Z. 2019. Weakly aligned cross-modal learning for multispectral pedestrian detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5127–5137.

Zhou, K.; Chen, L.; and Cao, X. 2020. Improving multispectral pedestrian detection by addressing modality imbalance problems. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, 787–803. Springer.

Zhu, H.; Wu, H.; Wang, X.; He, D.; Liu, Z.; and Pan, X. 2023. DPACFuse: Dual-Branch Progressive Learning for Infrared and Visible Image Fusion with Complementary Self-Attention and Convolution. *Sensors*, 23(16): 7205.