

RGB-X Object Detection via Scene-Specific Fusion Modules

Sri Aditya Deevi^{1*}

Connor Lee^{1*}

Lu Gan^{1*}

Sushruth Nagesh²

Gaurav Pandey²

Soon-Jo Chung¹

¹California Institute of Technology ²Ford Motor Company

{sdeevi, clee, ganlu, sjchung}@caltech.edu {snagesh1, gpandey2}@ford.com

Abstract

Multimodal deep sensor fusion has the potential to enable autonomous vehicles to visually understand their surrounding environments in all weather conditions. However, existing deep sensor fusion methods usually employ convoluted architectures with intermingled multimodal features, requiring large coregistered multimodal datasets for training. In this work, we present an efficient and modular RGB-X fusion network that can leverage and fuse pretrained single-modal models via scene-specific fusion modules, thereby enabling joint input-adaptive network architectures to be created using small, coregistered multimodal datasets. Our experiments demonstrate the superiority of our method compared to existing works on RGB-thermal and RGB-gated datasets, performing fusion using only a small amount of additional parameters. Our code is available at <https://github.com/dsriaditya999/RGBXFusion>.

1. Introduction

Autonomous vehicles rely on object detection algorithms to understand and interact with their surrounding environments. In order to be robust against different driving conditions, these algorithms operate on data from various sensor modalities ranging from optical cameras to LiDAR, each with their own advantages and disadvantages. Because no single sensor modality is robust to all possible conditions that may be encountered during driving, multiple sensor modalities are often used in conjunction via *deep sensor fusion (DSF)* to boost performance during normal driving operations, as well as to ensure segmentation and object detection reliability in adverse weather conditions [9].

Unlike traditional sensor fusion which merges processed sensor data outputs coming from independent pipelines, current works in DSF generally require joint end-to-end training of multi-branch sensor networks on large multi-

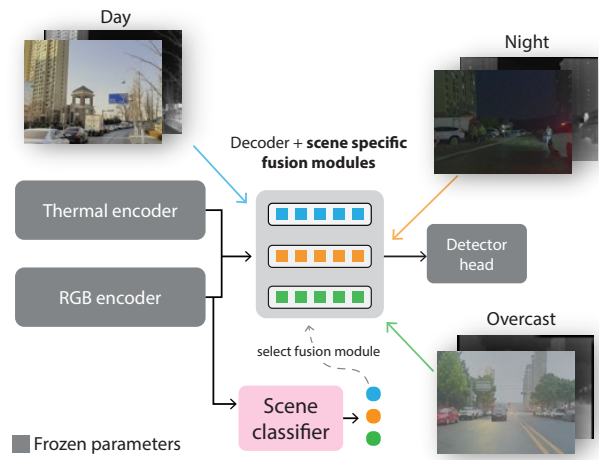


Figure 1. Our multimodal object detection approach combines RGB and thermal pretrained networks using lightweight, scene-specific fusion modules. Fusion modules are trained using categorized scene images and are used adaptively during inference with an auxiliary scene classifier.

modal datasets [1, 3, 16, 25, 39, 40] such as NuScenes [2], Berkeley Deep Drive [38], and Waymo [28] prior to deployment in the wild [9]. This means that fusion architectures must undergo time-consuming and potentially expensive re-training (in cost and carbon emissions) anytime a sensor modality is removed or added [24], and that they fail to take full advantage of state-of-the-art RGB pretrained networks.

In this paper, we propose the use of existing, well-known attention blocks as lightweight, scene-specific attention modules in order to easily fuse pretrained networks and to better adapt to common weather disturbances. We demonstrate our approach (Fig. 1) for object detection applications, training RGB-thermal and RGB-gated fusion models on RGB, thermal, and gated imagery collected in adverse driving conditions such as night, fog, snow, and rain [1, 7, 19]. We also leverage the attention modules as a method to visually interpret the contributions of each sensor modality. Compared to prior works, our approach takes

[†]This work was funded by the Ford University Research Program.

*Equal contribution

us another step closer to enabling a modular, *drag-and-drop* design for deep sensor fusion that absolves the need for extensive and expensive retraining while delivering on-par or better performance. **Our contributions are as follows:** **1.** A lightweight, modular RGB-X fusion network for object detection that leverages pretrained single-modality networks. **2.** A scene-adaptive fusion approach that selectively uses different fusion modules for different scene/weather conditions. **3.** Extensive experiments on publicly available RGB-X datasets that demonstrate the superiority of our approach in terms of detection performance and computational efficiency.

2. Related Work

Object Detection: Most modern methods for detecting objects utilize convolutional neural networks (CNN) or transformers. CNN object detectors include two-stage and single-shot detectors [21, 26, 27, 31]. A two-stage detector has an additional region proposal step while a single-shot detector relies only on a feature extractor and a detection head that directly predicts bounding boxes and classes, resulting in faster inference [43]. To deploy on mobile devices, neural architecture search (NAS) has been used to develop faster and lighter networks and detection architectures [14, 31]. In this work, we adopt the EfficientDet [31] detection architecture to target self-driving car applications that operate on mobile computing devices. Recent large vision transformer models have achieved state-of-the-art object detection results, but are not suitable for real-time use on robotic platforms [4, 22].

Deep Sensor Fusion: Robotic perception applications, notably for self-driving cars, rely on DSF to add sensor redundancy and to increase perception robustness and performance in both common and adverse operating scenarios. Current DSF algorithms consume multimodal data using deep networks and are trained end-to-end, combining different features at various points throughout a network depending on their particular fusion policy [8, 9]. Early fusion policies aggregate raw inputs or features extracted early on in the network [20, 32] while mid-fusion approaches [17, 35] operate on deeper, intermediate representations. Late fusion methods operate directly on bounding box outputs and can be used directly with pretrained detectors, but are subject to the performance of pretrained models [5]. In our work, we opt for a mid-fusion approach in order to take full advantage of the different feature modalities at various stages.

Regardless of fusion policies, current DSF algorithms and datasets for self-driving cars mainly focus on incorporating sensors like LiDAR and radar with RGB cameras [2, 12, 28, 35, 38]. In our work, we are interested in supplementing RGB with 2D image data from thermal and gated cameras due to the rich semantic information they provide and their robustness to fog and lighting in driving

scenarios [9].

RGB-Thermal Object Detection: Current RGB-thermal (RGB-T) object detection methods typically operate on aligned RGB-thermal image pairs and utilize some form of attention-based modules to perform mid-fusion on RGB and thermal image features. [40] utilizes intra-modality and inter-modality spatial attention modules to enhance and adaptively fuse intermediate features, respectively, prior to passing downstream to a detection head. Recently, [3] proposed mid-fusion modules that utilize channel attention to dynamically swap RGB and thermal feature channels. This helps to maximize feature usefulness before enhancing local features via parameter-free spatial attention. Other works including [10, 25, 42] fuse multi-modal data in a similar fashion but instead leverage transformer-based attention modules that increase model size and computational cost. [1] does not use thermal images, but similarly fuses RGB, gated, and projected LiDAR and radar data using local entropy masks in lieu of attention. In our work, we demonstrate that pretrained, single-modality detectors can be fused using simple, scene-specific channel and spatial attention modules to achieve strong RGB-T object detection performance.

3. Approach

We propose a modular RGB-X fusion network for object detection that is built upon pretrained single-modal detection architecture and multi-stage convolutional block attention modules (CBAM) [34] for cross-modal feature fusion. This modularity separates the training of single-modal backbones that contain the majority of network parameters and the training of a small fusion module, mitigating the requirement of large-scale multi-modal training data. The overall architecture for RGB-X fusion is shown in Fig. 2 using RGB-T as an example. We have an individual EfficientDet [31] for each image modality consisting of an EfficientNet [30] backbone network, a bidirectional feature pyramid network (BiFPN) and a detector head. While we choose to use EfficientDet to demonstrate our approach, we note that this architecture can be built using any single-modal detection network.

We employ CBAM to fuse the RGB and thermal features output from the respective BiFPN at various stages. Each CBAM fuses features at the same scale, resulting in 5 CBAM fusion modules. During training, only CBAM parameters are updated while pretrained object detector weights are frozen. CBAM modules are trained per scene category and are selected for use during inference time using an auxiliary scene classifier. In the rest of this section, we go over the details of our fusion mechanism and the auxiliary scene classifier, before describing the overall scene-adaptive fusion algorithm for RGB-X object detection.

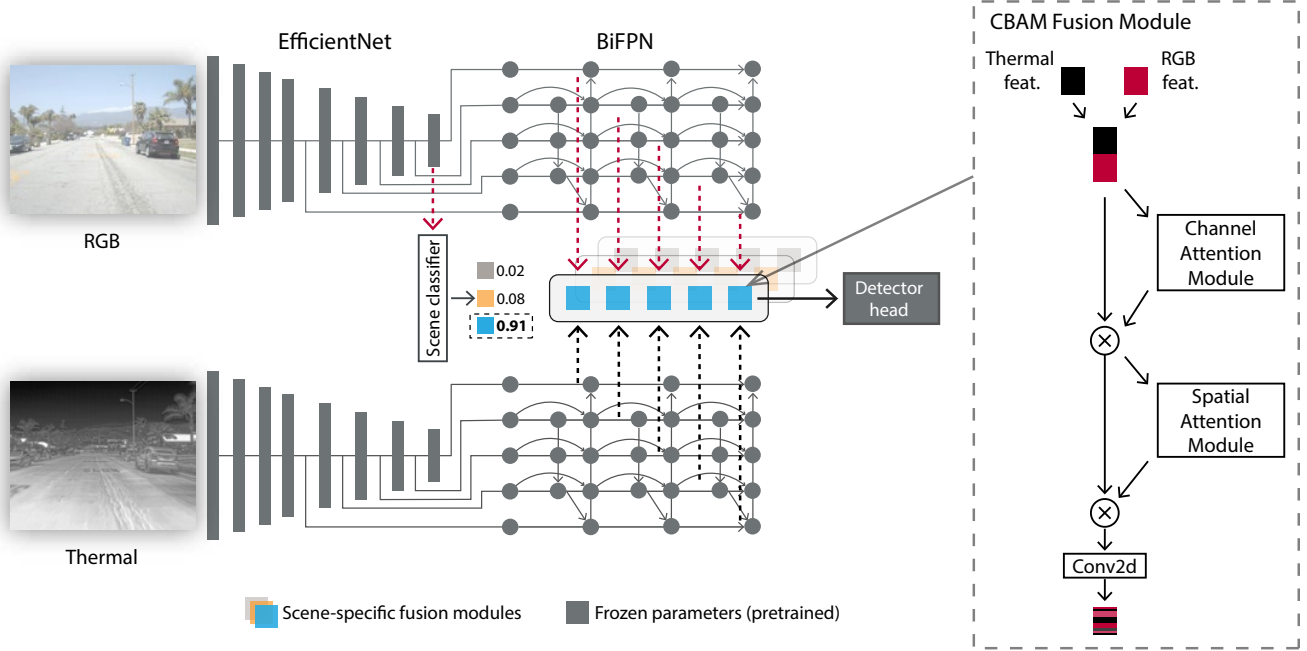


Figure 2. Overall framework of our scene-adaptive CBAM model for RGB-X fusion illustrated by RGB-T fusion. RGB and thermal images are processed by separate EfficientNet backbones, followed by BiFPNs. The features from BiFPNs are used for cross-modal feature fusion using modules selected by the scene classifier. The detector head utilizes these fused features to obtain the final detection results. The right side of the figure illustrates the CBAM fusion module, consisting of channel and spatial attention blocks, for feature fusion.

3.1. Convolutional Block Attention Fusion

We use CBAM to fuse RGB and thermal (or gated) CNN feature maps \mathbf{F}_{rgb} and \mathbf{F}_{x} , respectively. We concatenate features from both modalities across the channel dimension to create an input feature map \mathbf{F} for CBAM:

$$\mathbf{F} = [\mathbf{F}_{\text{rgb}}; \mathbf{F}_{\text{x}}] \in \mathbf{R}^{B \times C \times H \times W}, \quad (1)$$

where B denotes the batch size, and C, H, W denote the channel and spatial dimensions of the feature, respectively. Following the notation in [34], a CBAM module takes the feature map \mathbf{F} and masks it using channel and spatial attention operators M_c, M_s such that

$$\mathbf{F}' = M_c(\mathbf{F}) \otimes \mathbf{F}, \quad (2)$$

$$\mathbf{F}'' = M_s(\mathbf{F}') \otimes \mathbf{F}', \quad (3)$$

where \otimes denotes element-wise multiplication. We further convolve \mathbf{F}'' with $C/2$ kernels resulting in $C/2$ channels which is the original feature dimension.

Channel attention operator M_c is computed via

$$M_c(\mathbf{F}) = \sigma(\mathbf{W}_1 \mathbf{W}_0 \mathbf{F}_{\text{avg}}^c + \mathbf{W}_1 \mathbf{W}_0 \mathbf{F}_{\text{max}}^c), \quad (4)$$

where $\sigma, \mathbf{W}, \mathbf{F}_{\text{avg}}^c, \mathbf{F}_{\text{max}}^c$ denotes the sigmoid function, linear layer weights, the global average and max pooled features, respectively. Spatial attention is computed via

$$M_s(\mathbf{F}) = \sigma(f^{7 \times 7}([\mathbf{F}_{\text{avg}}^c; \mathbf{F}_{\text{max}}^c])), \quad (5)$$

where $\mathbf{F}_{\text{avg}}^c, \mathbf{F}_{\text{max}}^c$ are computed via channel-wise mean and max operations and $f^{7 \times 7}$ denotes convolution with a kernel size of 7.

3.2. Auxiliary Scene Classification

We utilize a simple scene classifier during inference time to adaptively select the most suitable set of fusion modules for the current setting, based on the intuition that the fusion module should attend different modalities under different scene/weather conditions. The scene classifier consists of a 2D adaptive average pooling operator followed by a fully connected layer, taking in the features from the RGB object detector encoder and outputs probabilities of possible scene categories. We choose RGB features as the input for scene classification due to their high variance in different scenes.

3.3. Scene-Specific Fusion

We train different CBAM fusion modules for various scenes by considering scene-specific dataset splits (Tab. 2). The number of parameters in different parts of the proposed fusion model is shown in Tab. 1. The total number of trainable parameters per scene is significantly less than the total number of parameters, making our approach expeditious. During inference of scene-adaptive fusion, we use the CBAM fusion modules trained on the scene with the highest probability, as indicated by the scene classifier.

Table 1. Parameter statistics of the proposed RGB-X fusion model.

Network Part	# Parameters
Backbones (RGB + X)	24.8 M
BiFPNs (RGB + X)	0.12 M
Detection Head	1.60 M
Fusion Modules (one per fusion level)	0.21 M
Total	26.7 M
Total Trainable (per scene)	0.21 M

Table 2. Dataset scene and train/val/test splits in our experiments.

Split	Scene Condition					
	Clear		Fog		Snow	
	Day	Night	Day	Night	Day	Night
Train	2147	1572	712	438	1365	1455
Val	537	393	438	110	342	364
Test	895	655	297	183	570	607

(a) Seeing Through Fog

Split	Scene Condition	
	Day	Night
Train	3476	653
Val	—	—
Test	702	311

(b) FLIR Aligned [39]

Split	Scene Condition			
	Day	Night	Overcast	Challenge
Train	992	488	746	484
Val	216	108	190	122
Test	323	140	205	156

(c) M³FD [19]

4. Results

4.1. Implementation and Training Details

Our code is written in PyTorch and based on the EfficientDet¹ repository. Pretrained RGB detectors on COCO dataset [18] were taken from the same repository. All other networks were trained using the Adam optimizer, a batch size of 8, an initial learning rate of $1e^{-3}$ with an exponential learning rate schedule, and a L_2 weight decay of $1e^{-3}$. The maximum number of epochs is set to 300 and 50 for pretraining single modality networks and fine-tuning RGB-X fusion networks, respectively. The scene classifier is trained for 50 epochs while the RGB backbone remains frozen. Networks were trained using an Nvidia P100 GPU.

4.2. Datasets

We use the following RGB-X datasets to validate our method and compare against state-of-the-art baselines. The train/val/test split statistics we use for various datasets and scene conditions are shown in Tab. 2.

FLIR Aligned: The FLIR Aligned dataset [39] consists of 5,142 aligned RGB-thermal image pairs from the original FLIR ADAS object detection dataset [7]. This derived dataset consists of bounding box annotations for *person*, *bicycle* and *car* classes. The provided train and test splits contain 4,129 and 1,013 image pairs, respectively. We manually divided them into *day* and *night* scene categories based on the appearance.

¹<https://github.com/rwightman/efficientdet-pytorch>



(a) Scene-Agnostic CBAM (b) Scene-Adaptive CBAM

Figure 3. Qualitative detection results on M³FD dataset. Zoomed-in images (yellow rectangle) are shown on the right of the original images for better visualization.

M³FD: The M³FD object detection dataset consists of 4,200 coregistered, time-synchronized RGB-thermal image pairs [19]. Bounding box annotations for *people*, *car*, *bus*, *motorcycle*, *truck* and *lamp* classes are provided. The data is also split into four scene categories (*day*, *night*, *overcast*, *challenge*) in [19] according to environment characteristics. We use the train/val/test splits provided by [16] due to the unavailability in the original dataset.

Seeing Through Fog: The Seeing Through Fog (STF) multispectral object detection dataset [1] consists of synchronized RGB/gated/LiDAR/radar/unaligned thermal data for a variety of weather conditions. The dataset also provides bounding box annotations for *pedestrian*, *truck*, *car*, *cyclist*, and *dontcare* classes. For training our scene-adaptive model, we considered the scene splits in Tab. 2a due to overlaps in original splits. For evaluation, we follow the original scene splits including *clear*, *light fog*, *dense fog*, and *snow/rain*. We use pairs of aligned 12-bit RGB and 10-bit gated images throughout this work.

4.3. Performance Evaluation

In this section, we validate the proposed method on the three datasets for RGB-X object detection. Auxiliary scene classification is employed to adaptively select suitable fusion modules per input image.

Auxiliary Scene Classification: We train our scene classifiers using ground truth scene labels provided in Tab. 2 by minimizing the standard cross-entropy loss for image classification. Top-1 accuracy of the scene classification is reported in Tab. 3, where the classifier attains high accuracy for categorizing various scenes in M³FD, FLIR and STF-Clear (the subset of STF dataset consists of *clear-day* and *clear-night*) datasets. The classifier does not perform as high for STF-Full, possibly because a large portion of *snow* images are also foggy and confused the classifier.

Scene-Adaptive Object Detection: This subsection reports quantitative and qualitative object detection results

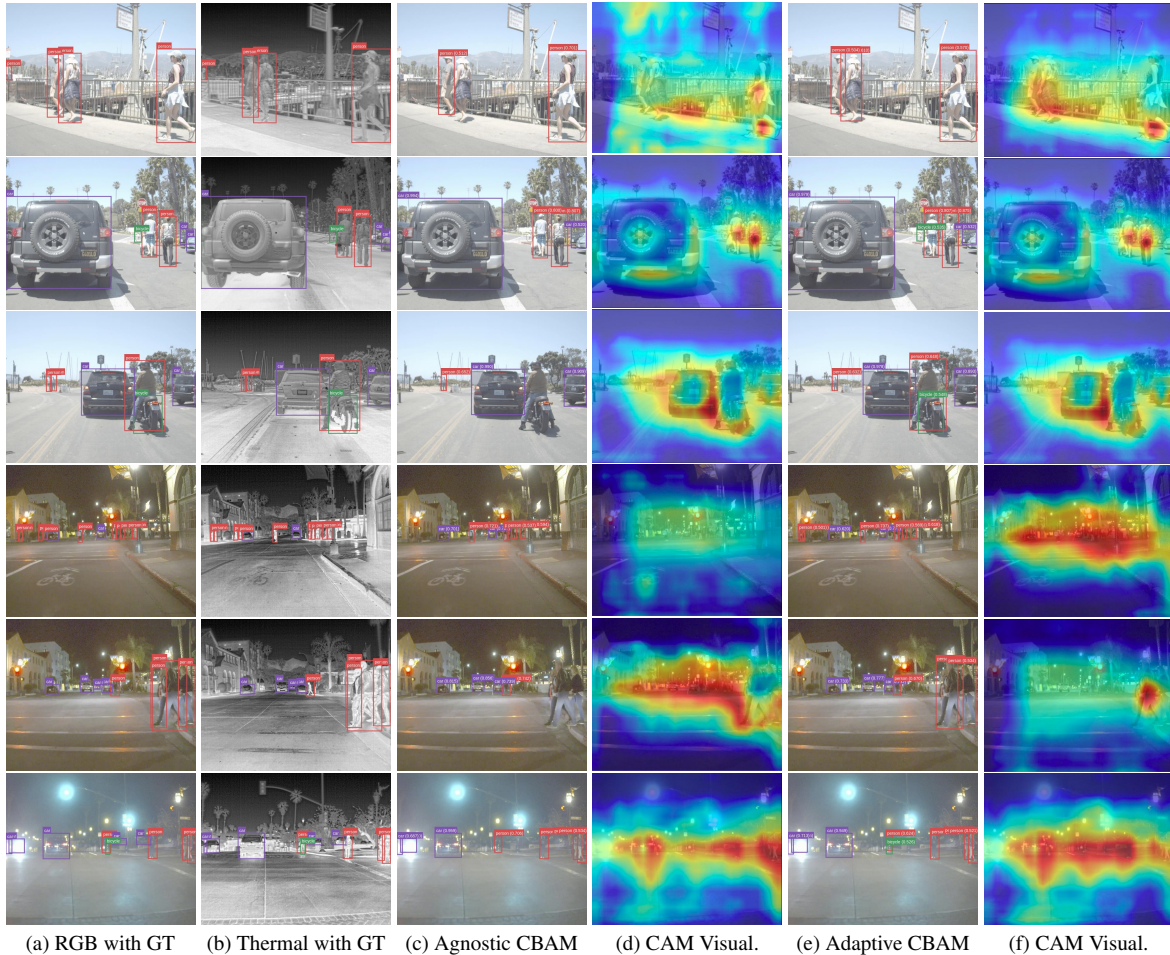


Figure 4. Qualitative detection results on FLIR Aligned dataset with *day* examples in the upper rows and *night* examples in the lower rows. The input RGB and thermal images are overlaid with ground truth (GT) bounding boxes. For each fusion model, we plot the detected bounding boxes and Eigen-CAM [23] visualizations of the CBAM fusion module. (d) and (f) are visualizations of (c) and (e), respectively.

Table 3. Top-1 Accuracy (%) of our scene classifier on the test set of three datasets.

Dataset	M ³ FD	FLIR	Seeing Through Fog	
			Clear	Full
Accuracy	91.42	96.35	96.01	77.02

of our proposed methods, compared with existing works. From Tab. 4, our scene-adaptive CBAM model outperforms existing methods on the M³FD dataset using the mean Average Precision IoU = 0.5 (mAP@0.5) metric used in [16, 19]. On the *full* test set, it outperforms EAEFNet [16] by 1.4% and the scene-agnostic CBAM model (in which only one set of CBAM fusion modules are trained using all training images) by 1%. A comparison of qualitative detection results on M³FD dataset between the scene-agnostic and scene-adaptive models is shown in Fig. 3. From the zoomed-in area of the figures, we can see that the scene-

Table 4. Object detection results (mAP@0.5) and speed (s) on M³FD dataset. Due to the difference in scene splits between baselines and our models, only results on the *full* test set are comparable across all methods.

Method	Test Scene					Inference Speed (s)
	Day	Night	Overcast	Challenge	Full	
RGB only	71.59	91.06	81.55	80.03	77.79	0.016
Thermal only	65.68	89.17	79.66	76.39	74.64	0.016
U2F [36]	73.80	86.8	73.10	97.6	77.5	0.129 [†]
TarDAL [19]	74.50	89.30	74.10	98.30	77.80	0.047 [†]
EAEFNet [16]	78.30	89.50	78.60	97.90	80.10	—
Scene-Agnostic CBAM (ours)	74.53	93.09	84.11	81.06	80.46	0.028
Scene-Adaptive CBAM (ours)	75.92	92.55	85.14	82.72	81.46	0.032

[†] Includes image fusion and object detection inference time.

adaptive model detects some occluded, blurred objects that the scene-agnostic model fails to detect. Note that, the single-modality models used in this experiment are pre-trained on COCO and further fine-tuned on the M³FD train-

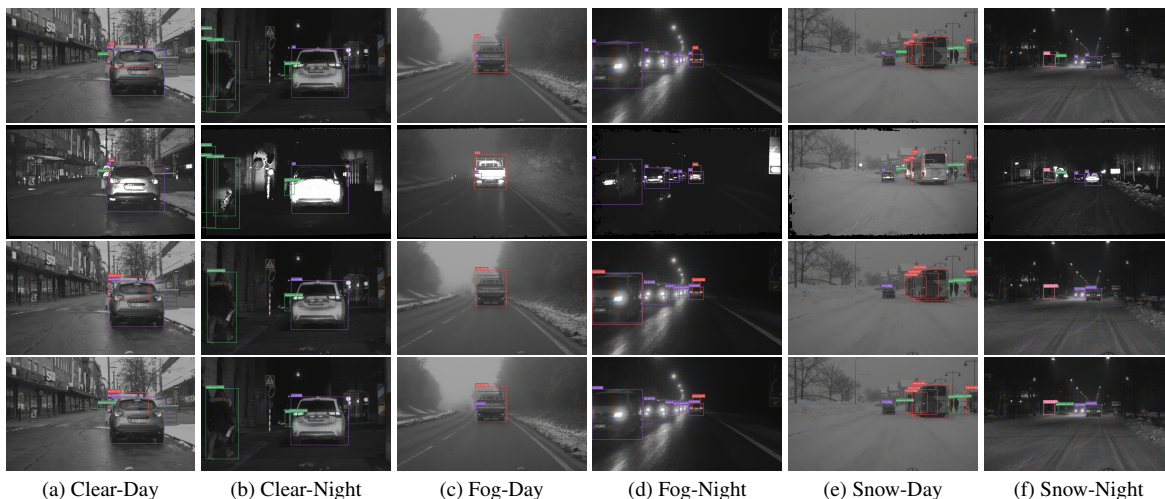


Figure 5. Object detection results on STF dataset in various scene conditions. From top to bottom: RGB images, gated images, scene-agnostic CBAM detections, and scene-adaptive CBAM detections. RGB and gated images are overlaid with ground truth bounding boxes.

Table 5. Object detection results and speed (s) on FLIR Aligned dataset. AP@0.5 for each object category is reported.

Method	Person	Bicycle	Car	mAP@0.5	mAP@0.75	mAP [†]	Inference Speed (s)
RGB only	60.79	37.25	73.94	57.32	17.6	24.7	0.016
Thermal only	82.86	50.80	82.83	72.16	33.4	37.0	0.016
GAFF [40]	76.60	59.40	85.50	72.9	32.9	37.5	0.061
CFR-3 [39]	74.49	57.77	84.91	72.93	—	—	0.050
RetinaNet + MFPT [42]	78.1	65.0	87.3	76.80	—	—	0.050
UA-CMDet [29]	83.20	64.30	88.40	78.60	—	—	—
CFT [25]	—	—	—	78.7	35.5	40.2	0.026
CSAA [3]	—	—	—	79.20	37.4	41.3	0.031
FasterRCNN + MFPT [42]	83.2	67.7	89.0	80.00	—	—	0.080
LRAF-Net [10]	—	—	—	80.50	—	42.8	—
Scene-agnostic CBAM (ours)	88.26	77.43	90.68	85.45	43.3	46.8	0.028
Scene-adaptive CBAM (ours)	88.92	78.61	90.94	86.16	43.0	47.1	0.032

[†] mAP refers to mAP@0.5:0.95

ing set for better performance. We also show some failure cases on M³FD in Fig. 6 where both fusion models struggled with distant small objects in *overcast* and *night* scenes, and cluttered objects under daylight.

For the FLIR Aligned dataset, we evaluate fusion networks built from an RGB network pretrained on COCO and a thermal network trained on the unaligned FLIR thermal training set. In general, both our scene-agnostic and scene-adaptive fusion models outperform the baselines by a large margin (Tab. 5), due to the increase in data the thermal and RGB networks had access to. Some qualitative detection results on FLIR test images along with attention visualizations are given in Fig. 4. We observe that scene-adaptive model tends to detect bicycles more successfully than scene-agnostic model, especially when the bicycle is rode by a person (see row 2 and 6 in Fig. 4). The higher margin of AP@0.5 for *bicycle* in Tab. 5 also aligns with this observation. In order to exam the effects of scene-adaptive CBAM, we visualize the CBAM using class activation map (CAM) [23] where the spatial attention is shown by a heat

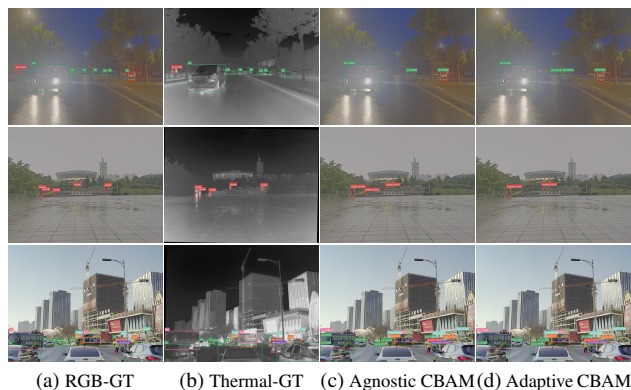


Figure 6. Example of failure cases on M³FD dataset. Both models struggled with distant small objects in *night* and *overcast* images and cluttered objects in *day* images.

map. From the visualization, we can see there is generally no difference between scene-agnostic CBAM and scene-adaptive CBAM for day images. However, the spatial attention in scene-adaptive CBAM attend more on small areas.

We visualize the channel attention of the scene-specific fusion module by plotting the normalized attention weights of thermal (black) and RGB (crimson) features for various scenes in Fig. 7. Higher value implies CBAM attends more on that feature channel. We find that scene-agnostic CBAM exhibits similar channel attention patterns across all scenes, while scene-adaptive CBAM shows tailored attention patterns per scene. Moreover, we observe attention weight increases on thermal features compared with RGB features from day to overcast to night images, likely as RGB images contain less information under lower illumination.

For the STF dataset, we first follow [1] and train our fu-

Table 6. Quantitative detection AP on the *clear* scene and unseen scenes for *car* following the KITTI evaluation [13] used in [1]. Models are all trained on the training set of the *clear* scene. Our scene-adaptive CBAM model is trained on *clear-day* and *clear-night* splits.

Method	Test Scene											
	Clear			Light Fog			Dense Fog			Snow/Rain		
	easy	mod.	hard	easy	mod.	hard	easy	mod.	hard	easy	mod.	hard
RGB only	90.14	87.56	80.87	91.19	88.47	82.02	90.43	85.59	80.79	89.44	82.87	77.81
Gated only	88.51	80.09	74.65	87.98	78.92	73.59	80.52	75.86	70.42	80.58	75.59	69.52
Fusion SSD [1]	87.73	78.02	69.49	88.33	78.65	76.54	74.07	68.46	63.23	85.49	75.28	67.48
Deep Fusion [1]	90.07	80.31	77.82	90.60	81.08	79.63	86.77	77.28	73.93	89.25	79.09	70.51
Deep Entropy Fusion [1]	89.84	85.57	79.46	90.54	87.99	84.90	87.68	81.49	76.69	88.99	83.71	77.85
Scene-agnostic CBAM (ours)	90.33	88.53	81.16	91.43	89.05	84.94	90.75	88.66	82.07	89.99	86.57	79.79
Scene-adaptive CBAM (ours)	90.29	88.53	81.07	91.13	89.13	84.20	90.77	88.37	81.68	89.96	86.30	79.74

Table 7. Quantitative detection AP on all scenes for *pedestrian*, *truck*, *car*, and *cyclist* following the KITTI evaluation [13] used in [1]. Models are trained on the training set of all scenes. The last column shows mAP@0.5 for all objects on all test images.

Method	Test Scene															
	Clear			Light Fog			Dense Fog			Snow/Rain			Full			
	easy	mod.	hard	easy	mod.	hard	easy	mod.	hard	easy	mod.	hard	easy	mod.	hard	all
RGB only	87.05	83.88	82.93	89.68	88.88	87.99	88.61	88.28	87.90	88.92	86.01	83.73	84.22	79.94	76.30	80.85
Gated only	81.69	76.19	74.57	85.63	84.01	80.19	83.40	82.00	79.88	84.03	79.54	77.38	80.70	73.58	70.13	75.15
Scene-agnostic CBAM (ours)	88.65	85.12	84.25	90.30	89.68	88.95	89.78	89.18	88.82	89.25	87.01	85.77	86.11	81.84	78.52	83.01
Scene-adaptive CBAM (ours)	88.60	85.24	84.22	90.53	89.39	88.89	89.79	89.33	89.03	89.37	87.46	85.69	86.13	81.85	78.48	83.11

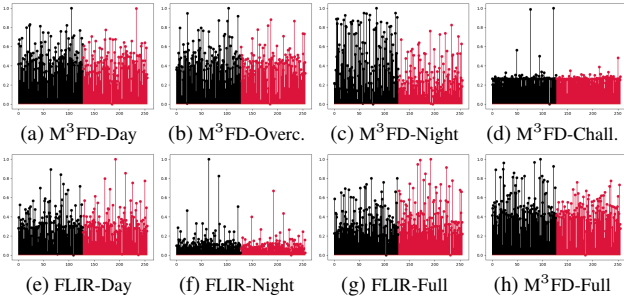


Figure 7. Normalized attention weights for 256 feature channels in CBAM fusion module trained on different scenes. Thermal channels are in black, and RGB channels in crimson. The fusion module trained on the entire dataset (g-h) exhibits similar attention patterns across all scene/weather conditions, whereas from *day* to *overcast* to *night*, the scene-specific fusion module (a-f) attends increasingly on thermal features.

sion modules only on *clear-day* and *clear-night* RGB-gated image pairs for fair comparison. As shown in Tab. 6, the scene-agnostic and scene-adaptive CBAM models achieve similar performance on different scenes and outperform the baseline models using even more modalities than RGB-gated images [1]. When training on all scenes in Tab. 7, we can see that our scene-adaptive model outperforms the scene-agnostic model by 0.1% on mAP@0.5. Single-modality models used for this experiment are also further trained on STF training data, due to their use of 10 and 12 bit gated and RGB imagery. Fig. 5 presents a few examples of the qualitative detection results in various scenes.

Computational Benchmarks: We compiled our CBAM

Table 8. Ablation study on different fusion modules. Object detection results (mAP@0.5) on M³FD dataset are reported.

Fusion Module	Train/Search Scene	Test Scene				
		Day	Night	Overcast	Challenge	Full
RGB only		71.59	91.06	81.55	80.03	77.79
Thermal only		65.68	89.17	79.66	76.39	74.63
ECAAttn (Tr)		73.38	93.39	83.55	82.28	80.17
ECAAttn (RH)		72.25	92.83	81.98	80.53	78.81
ECAAttn (TH)		74.02	93.38	84.25	81.48	80.32
ShuffleAttn (Tr)	Full	73.47	94.56	84.61	80.91	80.17
ShuffleAttn (RH)		72.78	92.63	83.61	80.37	79.28
ShuffleAttn (TH)		74.07	93.21	84.19	81.43	80.34
CBAM (Tr)		73.11	93.01	83.11	80.17	79.33
CBAM (RH)		72.85	92.46	83.54	80.73	79.21
CBAM (TH)		74.53	93.09	84.11	81.06	80.46
ECAAttn (TH)	Day	74.75	94.51	84.16	81.09	80.65
	Night	72.00	91.84	83.56	79.74	78.74
	Overcast	71.96	92.67	84.44	80.09	79.18
	Challenge	73.25	93.11	83.78	81.88	80.14
ShuffleAttn (TH)	Day	75.28	94.64	84.72	81.85	81.04
	Night	71.79	92.21	83.30	78.95	78.21
	Overcast	73.57	92.42	84.32	80.95	80.00
	Challenge	72.42	92.90	84.21	81.27	79.62
CBAM (TH)	Day	76.04	94.07	84.89	80.78	81.07
	Night	72.68	92.55	83.20	78.77	78.62
	Overcast	73.30	92.53	85.15	80.67	79.94
	Challenge	74.10	94.28	82.70	82.61	80.93
DSF-NAS	Day	75.68	94.25	84.35	81.85	81.03
	Night	72.32	91.94	83.85	80.51	79.12
	Overcast	73.15	93.46	83.79	80.60	79.52
	Challenge	72.90	93.44	83.29	81.59	79.81
	Full	74.68	92.65	83.90	81.67	80.56

Tr – Trained head TH – Thermal head RH – RGB head

fusion models using TorchInductor and conducted benchmarks on a Titan RTX. The inference time for the scene-adaptive fusion model is 0.032 seconds per individual image pair, while the scene-agnostic variant clocks in at 0.028 seconds. These times are comparable with other recent mul-

Table 9. Object detection results (mAP@0.5) of our scene-adaptive CBAM model trained using decreasing amounts of data.

Dataset	% of Original Training Set			
	100%	50%	25%	1%
FLIR	86.16	85.70	84.60	75.72
M ³ FD	81.46	78.34	77.65	41.94
STF-Clear	80.65	80.73	80.10	73.76
STF-Full	83.11	83.06	82.99	75.64

timodal object detection approaches (Tab. 4, 5) and meet the speed requirements for real-time autonomous driving applications.

4.4. Ablation Studies

Fusion Module Design: We conduct an ablation study using the M³FD dataset to explore the effects of different fusion modules and architectures (Tab. 8). We compare our CBAM-based RGB-X fusion approach against two other attention modules: ECAAttn [6, 33] and ShuffleAttn [41]. Furthermore, we also compare against custom fusion modules (DSF-NAS) designed purposely for this fusion task via neural architecture search. In particular, we use Bilevel Multimodal Neural Architecture Search [37] (BM-NAS) to automate this design as its gradient-based optimization approach makes it faster compared to other NAS methods based on reinforcement learning and genetic algorithms. Specifically, we allow BM-NAS to optimize over sequential applications of two operations chosen from sum, spatial attention, channel attentions from CBAM and ECAAttn, and 2D convolution of concatenated features.

We first train for fusion using scene-agnostic CBAM, ECAAttn, and ShuffleAttn modules along with either a trainable, frozen thermal, or frozen RGB detector head. We find that training with a frozen detector head initialized with thermal weights performs the best in Tab. 8, possibly due to the lower variance of thermal data across different scenes. We repeat the study under the scene-adaptive regime, with the previous three attention modules and frozen thermal detection heads, along with DSF-NAS fusion modules. Overall, we find similar performance between DSF-NAS and CBAM-based fusion networks. However, CBAM fusion models exhibit better performance on scene-specific data verifying its use in our proposed modular framework.

Effect of Training Dataset Size on Fusion: As our proposed fusion method looks to fuse pretrained networks with lightweight fusion modules, the fusion process should still be effective and be able to generalize even when done with limited amounts of training data. To determine the extent of this, we perform fusion using 100%, 50%, 25%, and 1% of the original datasets in Tab. 9. Overall, we find that competitive results can still be achieved using only 25% of the

Table 10. Object detection results (mAP@0.5) of our scene-adaptive CBAM model on unknown scenes in M³FD dataset.

Test	Excluded Training Scene			
	Day	Night	Overcast	Challenge
Excluded Scene	73.17	93.50	84.18	80.74
All Scenes	80.48	81.30	81.27	80.98

original training data results, with the exception of M³FD which decays quicker than the rest.

Performance on Unknown Scenes: As our proposed method requires scenes to be known during training, we further investigate the performance of our method on unknown/unexpected scenes. In this experiment, our scene-specific CBAM fusion modules and scene classifiers are trained with one scene data excluded, and tested on that excluded scene and all test images. We observed minor regression in overall performance (row 2 in Tab. 10) compared with our scene-adaptive model trained on all scenes (81.46 in Tab. 4), which is expected as there is no fusion module trained specifically for that unknown scene. However, the overall mAP@0.5 in all cases is still higher than scene-agnostic model trained on all scenes (80.46 in Tab. 4). Specifically, in the case of *night* or *overcast* scene excluded, the object detection performance on the unknown scene (row 1 in Tab. 10) is higher than the scene-agnostic model. This is possibly because our scene classifier tends to select a fusion module trained on a similar scene, for instance, classifying *night* image as *overcast* and vice versa.

5. Limitations and Future Work

Our method requires aligned RGB-X data, which is not always available. The scene-specific modules require scenes to be known during training, and the approach is not expected to work as well in unexpected weather conditions. Future work looks to incorporate unsupervised [11] and on-line learning [15] to adapt to unexpected conditions.

6. Conclusion

We presented a novel RGB-X object detection model that improves autonomous vehicle perception in different weather and lighting conditions. We showed that our method is superior compared to existing works on two RGB-T and one RGB-gated object detection benchmarks, demonstrating the robustness of our scene-adaptive models and generalizability to different modalities. Furthermore, our use of lightweight fusion modules brings us closer to achieving a more modular design for deep sensor fusion. For future work, we look to train and leverage larger pretrained models for both RGB and thermal modalities via multitask learning and to incorporate into an online learning framework to adapt to unexpected weather patterns.

References

- [1] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 2, 4, 6, 7
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenet: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 1, 2
- [3] Yue Cao, Junchi Bin, Jozsef Hamari, Erik Blasch, and Zheng Liu. Multimodal object detection by channel switching and spatial attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 403–411, 2023. 1, 2, 6
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. 2
- [5] Yi-Ting Chen, Jinghao Shi, Zelin Ye, Christoph Mertz, Deva Ramanan, and Shu Kong. Multimodal object detection via probabilistic ensembling. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, pages 139–158. Springer, 2022. 2
- [6] Sri Aditya Deevi and Deepak Mishra. Expeditious object pose estimation for autonomous robotic grasping. In *International Conference on Computer Vision and Image Processing*, pages 15–30. Springer, 2022. 8
- [7] Flir thermal adas dataset, 2019. 1, 4
- [8] Jamil Fayyad, Mohammad A. Jaradat, Dominique Gruyer, and Homayoun Najjaran. Deep Learning Sensor Fusion for Autonomous Vehicle Perception and Localization: A Review. *Sensors*, 20(15):4220, July 2020. 2
- [9] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, Mar 2021. 1, 2
- [10] Haolong Fu, Shixun Wang, Puhong Duan, Changyan Xiao, Renwei Dian, Shutao Li, and Zhiyong Li. Lraf-net: Long-range attention fusion network for visible–infrared object detection. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 2, 6
- [11] Lu Gan, Connor Lee, and Soon-Jo Chung. Unsupervised RGB-to-thermal domain adaptation via multi-domain attention network. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6014–6020. IEEE, 2023. 8
- [12] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 2
- [13] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 7
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1314–1324, 2019. 2
- [15] Connor Lee, Jonathan Gustafsson Frennert, Lu Gan, Matthew Anderson, and Soon-Jo Chung. Online self-supervised thermal water segmentation for aerial vehicles. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Accepted, 2023. 8
- [16] Mingjian Liang, Junjie Hu, Chenyu Bao, Hua Feng, Fuqin Deng, and Tin Lun Lam. Explicit attention-enhanced fusion for rgb-thermal perception tasks. *IEEE Robotics and Automation Letters*, 2023. 1, 4, 5
- [17] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 641–656, 2018. 2
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing. 4
- [19] Jinyuan Liu, Xin Fan, Zhanbo Huang, Guanyao Wu, Risheng Liu, Wei Zhong, and Zhongxuan Luo. 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 5802–5811, 2022. 1, 4, 5
- [20] Jingjing Liu, Shaoting Zhang, Shu Wang, and Dimitris Metaxas. Multispectral Deep Neural Networks for Pedestrian Detection. In *Proceedings of the British Machine Vision Conference 2016*, pages 73.1–73.13, York, UK, 2016. British Machine Vision Association. 2
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016. 2
- [22] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2
- [23] Mohammed Bany Muhammad and Mohammed Yeasin. Eigen-cam: Class activation map using principal compo-

- nents. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2020. 5, 6
- [24] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021. 1
- [25] Fang Qingyun, Han Dapeng, and Wang Zhaokui. Cross-modality fusion transformer for multispectral object detection. *arXiv preprint arXiv:2111.00273*, 2021. 1, 2, 6
- [26] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 2
- [28] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 1, 2
- [29] Yiming Sun, Bing Cao, Pengfei Zhu, and Qinghua Hu. Drone-based rgb-infrared cross-modality vehicle detection via uncertainty-aware learning. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(10):6700–6713, 2022. 6
- [30] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 2
- [31] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020. 2
- [32] Jörg Wagner, Volker Fischer, Michael Herman, and Sven Behnke. Multispectral Pedestrian Detection using Deep Fusion Convolutional Neural Networks. In *ESANN*, 2016. 2
- [33] Qilong Wang, Banggu Wu, Pengfei Zhu, Peihua Li, Wangmeng Zuo, and Qinghua Hu. Eca-net: Efficient channel attention for deep convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11534–11542, 2020. 8
- [34] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 2, 3
- [35] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2018. 2
- [36] Han Xu, Jiayi Ma, Junjun Jiang, Xiaojie Guo, and Haibin Ling. U2fusion: A unified unsupervised image fusion network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):502–518, 2020. 5
- [37] Yihang Yin, Siyu Huang, and Xiang Zhang. Bm-nas: Bilevel multimodal neural architecture search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8901–8909, 2022. 8
- [38] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 2
- [39] Heng Zhang, Elisa Fromont, Sébastien Lefevre, and Bruno Avignon. Multispectral fusion for object detection with cyclic fuse-and-refine blocks. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 276–280. IEEE, 2020. 1, 4, 6
- [40] Heng Zhang, Elisa Fromont, Sébastien Lefèvre, and Bruno Avignon. Guided attentive feature fusion for multispectral pedestrian detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 72–80, 2021. 1, 2, 6
- [41] Qing-Long Zhang and Yu-Bin Yang. Sa-net: Shuffle attention for deep convolutional neural networks. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2235–2239. IEEE, 2021. 8
- [42] Yaohui Zhu, Xiaoyu Sun, Miao Wang, and Hua Huang. Multi-modal feature pyramid transformer for rgb-infrared object detection. *IEEE Transactions on Intelligent Transportation Systems*, 2023. 2, 6
- [43] Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey. *Proceedings of the IEEE*, 2023. 2