

SAGA: Semantic-Aware Gray color Augmentation for Visible-to-Thermal Domain Adaptation across Multi-View Drone and Ground-Based Vision Systems

Manjunath D¹, Aniruddh Sikdar², Prajwal Gurunath¹, Sumanth Udupa¹, Suresh Sundaram^{1,2}

¹Department of Aerospace Engineering, Indian Institute of Science, Bengaluru, India

²Robert Bosch Centre for Cyber Physical Systems, Indian Institute of Science, Bengaluru, India

{manjunathd1, aniruddhss, vssuresh} @iisc.ac.in

Abstract

Domain-adaptive thermal object detection plays a key role in facilitating visible (RGB)-to-thermal (IR) adaptation by reducing the need for co-registered image pairs and minimizing reliance on large annotated IR datasets. However, inherent limitations of IR images, such as the lack of color and texture cues, pose challenges for RGB-trained models, leading to increased false positives and poor-quality pseudo-labels. To address this, we propose Semantic-Aware Gray color Augmentation (SAGA), a novel strategy for mitigating color bias and bridging the domain gap by extracting object-level features relevant to IR images. Additionally, to validate the proposed SAGA for drone imagery, we introduce the IndraEye, a multi-sensor (RGB-IR) dataset designed for diverse applications. The dataset contains 5,612 images with 145,666 instances, captured from diverse angles, altitudes, backgrounds, and times of day, offering valuable opportunities for multimodal learning, domain adaptation for object detection and segmentation, and exploration of sensor-specific strengths and weaknesses. IndraEye aims to enhance the development of more robust and accurate aerial perception systems, especially in challenging environments. Experimental results show that SAGA significantly improves RGB-to-IR adaptation for autonomous driving and IndraEye dataset, achieving consistent performance gains of +0.4% to +7.6% (mAP) when integrated with state-of-the-art domain adaptation techniques. The dataset and codes are available at <https://github.com/airl-iisc/IndraEye>

1. Introduction

Deep learning has advanced real-world applications like autonomous driving [57] and surveillance [46]. Robust object detection for aerial perception is essential for Unmanned Aerial Vehicles (UAVs) to operate effectively in challenging, low-light environments, facilitating critical tasks such

as infrastructure inspections [43], environmental monitoring [29] and surveillance [42, 44]. Most vision models focus on visible-spectrum (RGB) cameras due to the availability of extensive datasets, however, their robustness diminishes in low-light conditions [20, 45]. Increasing research in robotics and computer vision highlights thermal infrared (IR) imaging for its effectiveness in harsh weather and low visibility, as IR cameras capture unique spectral data and penetrate dust and smoke. While IR cameras are valuable, their images generally contain less semantic information than RGB images, leading to performance drops in deep learning models for downstream tasks [48, 49]. With affordable IR sensors, deep multimodal fusion uses RGB-IR integration to outperform unimodal methods [32, 59]. However, multimodal techniques face three key challenges: (i) *Limited Annotated Thermal Datasets*: The scarcity of labeled IR data limits training high-performance detection models. (ii) *Dependency on Co-Registered RGB-IR Image Pairs*: Accurate fusion relies on precisely aligned RGB and IR images, but achieving co-registration is challenging due to hardware constraints and sensor-specific variations. These challenges frequently result in alignment errors, which can substantially impact performance, especially in drone applications [2, 50], and (iii) *Domain Shift Between Modalities*: IR images differ significantly from RGB, causing domain shift and reducing model performance on IR data. Conventional Unsupervised Domain Adaptation (UDA) for object detection has been predominantly studied in the context of RGB images. In contrast, domain-adaptive thermal object detection [13] aims to enhance detection performance in the infrared (IR) spectrum by leveraging UDA techniques. Furthermore, UDA can be extended for RGB-to-IR domain adaptation to address the inherent challenges of deep multimodal fusion models. By minimizing the domain discrepancy between the source (i.e., RGB images) and target (i.e., IR images) domains, it alleviates the constraints of co-registered image pairs and mitigates the reliance on extensive annotated IR datasets, thereby enhancing the adaptability of object detection models across modalities.

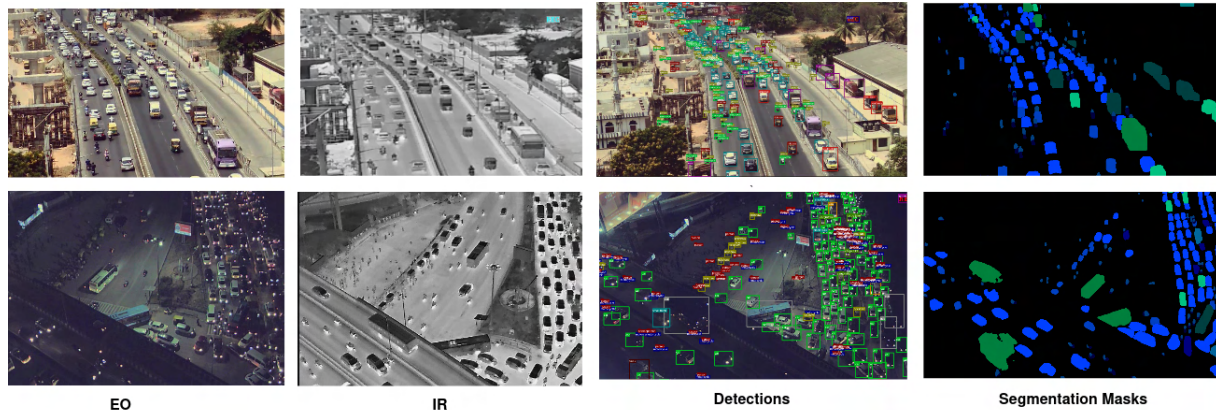


Figure 1. Snapshots from the IndraEye dataset showing different modalities RGB, IR and complete semantic annotations for detection & segmentation tasks taken from different slant angles

Domain-adaptive thermal object detectors use the Mean Teacher (MT) framework [55], where the teacher model generates pseudo-labels from the target domain to guide the student. The teacher’s weights are updated via an exponential moving average of the student’s weights for improved cross-domain adaptation. Object detection requires fine-grained, localized representations, but RGB-trained models often over-rely on color cues and texture patterns, which are absent in monochromatic IR images [58, 67]. Additionally, IR imaging struggles to capture fine edges visible in RGB due to reduced edge clarity [38, 56]. These inherent limitations in color and structural representation pose challenges for domain adaptation, leading to poor-quality pseudo-labels. Additionally, domain-adaptive thermal object detection for drone imagery is still in its early stages and presents significant challenges in the RGB-IR domain due to the distinct nature of aerial perspectives. The adaptation process is further challenged by dynamic environmental conditions, diverse viewing angles, and factors like weather fluctuations, and inconsistencies in acquisition parameters. Perhaps the most critical limitation is the lack of large-scale annotated datasets, which impedes the advancement of robust adaptation techniques for drone-based applications [33, 65].

To address this challenge, we introduce Semantic-Aware Gray color Augmentation (SAGA), a novel augmentation strategy applied to RGB images to mitigate color bias and bridge the domain gap between RGB and IR. The domain gap often results in poor-quality pseudo-labels from the teacher model [62], which SAGA alleviates by extracting object-level features more relevant to IR images. SAGA selectively converts only object instances to grayscale while preserving the background’s color information, creating a hybrid RGB-grayscale image that maintains the structural integrity of objects. As a simple instance-level augmentation strategy, SAGA is applied exclusively to the source domain using RGB images. By training the teacher model to perceive object instances in a manner similar to their IR

counterparts, SAGA facilitates precise adaptation of localized features in the target domain, consistently enhancing performance on IR imagery in a domain adaptation setting.

To validate SAGA for RGB-IR drone imagery, we introduce IndraEye, a comprehensive dataset tailored for aerial perception. It includes diverse object categories like road vehicles and pedestrians, captured across varied angles, backgrounds, and scales, spanning urban areas to highways with different population densities and environmental conditions. On average, each RGB image contains 35 annotated instances, with a strong emphasis on dense traffic scenarios. Figure 1 showcases examples of RGB and IR images captured from different slant angles. IndraEye supports multi-task learning, including object detection and semantic segmentation, making it a crucial benchmark for evaluating and improving model robustness across multiple modalities and tasks. Additionally, we introduce a structured benchmark for RGB-to-IR domain adaptation, making it one of the first for drone imagery. The main contributions of this paper are:

1. We propose Semantic-Aware Gray color Augmentation (SAGA), a novel instance-level augmentation strategy. SAGA creates hybrid RGB-grayscale images that preserve object structure, mitigating color bias and narrowing the domain gap for RGB-to-IR domain adaptation.
2. We introduce the IndraEye dataset, that includes diverse scenes captured from seven different locations with varying slant angles and heights, resulting in significant object scale variations. The dataset includes 145,666 dense instances across 13 classes, covering various road vehicles and pedestrians in both RGB and IR modalities, supporting object detection and segmentation.
3. SAGA enhances RGB-to-IR adaptation for autonomous driving and IndraEye drone imagery. Integrated with state-of-the-art domain adaptation techniques, it achieves consistent performance gains of +0.4% to +7.6% (mAP) across different scenarios.

Scene	Altitude (metres)	Dynamic Angle range (degrees)	RGB Instances daytime	RGB Instances nighttime	IR Instances daytime	IR Instances nighttime	Scale-variability
A	30	10-25	1982	-	9247	-	Mid
B	30	10-25	14149	-	8215	-	Mid
C	60	5-50	37752	5312	7559	1285	High
D	12	20-40	5369	-	-	-	Low
E	12	20-40	5394	3774	-	6124	Low
F	12	10-30	4728	-	8759	-	Mid
G	7	20-40	2234	-	1936	2024	Low
H	7	10-30	3033	2971	3731	7413	High

Table 1. IndraEye dataset description: with altitude of imaged scene, dynamic angle ranges, scene-wise instances, and scale-variability of each scene.

2. Related works

Thermal Object Detection Object detection using thermal images has become a critical research area in computer vision, with Aerial Thermal Object Detection (ATOD) posing significant challenges due to variations in object sizes and backgrounds [12]. To address these challenges, GANs have been explored [11, 47] for IR generation to reduce manual annotation and enhance object detection performance. Recent studies [39, 50] have explored object detection with misaligned RGB-IR images, revealing the limitations of co-registered datasets. These works emphasize the necessity of advanced techniques such as domain adaptation to address this challenge effectively.

Domain Adaptive Thermal Object Detection UDA for RGB-to-IR adaptation reduces domain gaps between a labeled source and an unlabeled target. It minimizes discrepancies by aligning features or styles, and self-training. Methods like adversarial training [7, 23] and style transfer [6, 25, 30] aim to bridge domain gaps, but they struggle to balance feature discriminability. Self-training methods enhance performance on the target domain by utilizing source domain knowledge without relying on labels. For example, studies like [9], [34], and [3] employ a mean-teacher framework, where a teacher network generates pseudo-labels to guide the student network in the target domain. Recently, [13] proposed a two-teacher, single-student architecture with zig-zag learning to reduce domain confusion. However, frequent teacher switching causes loss values to diverge, leading to model collapse on complex datasets. In contrast, our proposed SAGA augmentation, integrated with different adaptation algorithms, consistently enhances performance across multiple datasets.

Multi-modal Object Detection dataset Datasets like VEDAI [40] introduced RGB-IR imagery for object detection, featuring 12,000 bird’s-eye-view drone images across nine categories. The KAIST dataset [24], released the same year, focused on pedestrian detection for autonomous vehicles and gained significant traction despite challenges like misalignment caused by mechanical vibrations and calibration errors. While KAIST employed beam-splitter

co-registration for alignment, other datasets like CVC14 [19] faced persistent alignment issues due to hardware limitations. In subsequent years, more extensive datasets have emerged, including FLIR [51], LLVIP [26], DroneVehicle [52], and M3FD [35]. Drone-based detection presents distinct challenges, such as object occlusion and scale variations, which datasets like VEDAI and DroneVehicle aim to address. The LLVIP dataset focuses on low-light surveillance for pedestrian detection, while M3FD facilitates multimodal fusion for detection using high-resolution imagery. However, LLVIP, DroneVehicle, and M3FD depend on manual RGB-IR image co-registration, a labor-intensive process unsuitable for real-time drone operations[50]. Hardware-based alignment methods, such as calibration devices or sensors, often struggle due to environmental factors like temperature fluctuations, mechanical instability, and limitations in feature-matching algorithms[2]. Additionally, these datasets contain fewer samples, limited classes, multi-scale variations, and diverse slant angles. To tackle these challenges, we present the IndraEye dataset, specifically designed to accommodate real-world constraints like long-tail distributions, occlusion, and scale diversity. A qualitative comparison with existing datasets is presented in Table 2.

3. Semantic Aware Gray color Augmentation (SAGA) for Domain Adaptation

The following sections outline SAGA and IndraEye. Section 3.1 outlines the Mean Teacher framework for thermal object detection, Section 3.2 presents SAGA, and Section 3.3 introduces IndraEye.

3.1. Preliminaries on the Mean Teacher Framework for Domain Adaptation

The Mean Teacher (MT) framework [55] for domain-adaptive object detection leverages a teacher-student mutual learning framework to transfer knowledge from a labeled source domain to an unlabeled target domain [9, 10, 36]. Specifically, in the context of domain-adaptive thermal ob-

Datasets	Multi-sensory	Diverse Viewpoints	Diverse backgrounds	Diverse classes	Diverse illumination	Detection	Segmentation
DOTA [63]	✗	✗	✗	✓	✗	✓	✗
HIT-UAV [53]	✗	✓	✓	✗	✗	✓	✗
VisDrone [4]	✗	✓	✓	✓	✓	✓	✗
UAVDT [14]	✗	✓	✓	✓	✓	✓	✗
Vedai [40]	✓	✗	✓	✓	✓	✓	✗
M3FD [35]	✓	✗	✓	✓	✓	✓	✗
FLIR [51]	✓	✗	✓	✓	✓	✓	✗
MSRS [54]	✓	✗	✓	✓	✓	✗	✓
InfraParis [15]	✓	✗	✓	✓	✓	✓	✓
IndraEye (Ours)	✓	✓	✓	✓	✓	✓	✓

Table 2. Qualitative comparison of multiple aerial vehicle object detection datasets.

ject detection, the source domain comprises RGB images, while the target domain consists of infrared (IR) images. The teacher and student models are two object detection networks with identical architectures. This framework processes source and target data simultaneously, enabling mutual knowledge transfer. The teacher model, trained on the labeled source domain, generates pseudo labels for the unlabeled target domain data. The student model learns from these pseudo labels, and its updated weights are periodically transferred to the teacher model, enabling continuous improvement. The framework is optimized using the following loss function:

$$L = L_{src} + L_{tgt} \quad (1)$$

where, L_{src} denotes the loss associated with the source domain, while L_{tgt} represents the loss for the target domain. The teacher model’s weights are updated using the student model’s weights through the Exponential Moving Average (EMA) mechanism. Throughout the training process, the teacher model update effectively forms an ensemble of the student model’s weights, represented as:

$$\theta^T = \alpha\theta^T + (1 - \alpha)\theta^S \quad (2)$$

where, θ^T represents the teacher model’s weights, while θ^S denotes the student model’s weights. CMT [3] introduces a framework that integrates contrastive learning with the MT methodology, where the student network is trained using stronger augmentations, while the teacher model is trained with weaker augmentations. D3T [13], on the other hand, employs two separate teacher models followed by a zigzag training mechanism. For additional information, please refer to [3, 13].

3.2. Semantic Aware Gray color Augmentation (SAGA) for Domain Adaptation

Color and structure serve as fundamental elements of an image, both essential for object detection models, with the impact of color bias extensively studied through an emphasis on structural features [18, 22]. Models trained on RGB data often heavily depend on these color cues, as chromatic

diversity is crucial for distinguishing objects in RGB imagery [67]. However, infrared (IR) imaging modalities inherently lack the ability to capture rich color and structural details compared to RGB, posing significant challenges for cross-modal analysis and domain adaptation [8, 68]. In the target IR domain, the absence of chromatic diversity leads to a significant reduction in color information, which can contribute to an increased rate of false positives. Domain-adaptive thermal object detection strives to balance the dependence on RGB color bias while retaining essential structural features, in the target IR domain [62]. To mitigate this bias while preserving essential visual cues, we introduce Semantic-Aware Gray Color Augmentation (SAGA), a simple, effective augmentation easily integrated into any object detection algorithm.

When training an object detection model on the source domain using RGB images, let the dataset be represented as $S = \{x_i, y_i\}$, where x_i denotes the images and y_i represents the corresponding object labels. These labels are utilized to selectively convert object instances in x_i to grayscale while preserving the original color information in the background, rather than applying grayscale conversion to the entire image. SAGA selectively converts object instances to grayscale while retaining the background’s color information, ensuring semantic awareness.

Consider an image I containing n objects, represented as $I = \{o_1, o_2, \dots, o_n\}$, where $o_1, o_2,$ and o_n denote the individual objects within the image. Each object is extracted, converted to grayscale, and then reintegrated into the original image, ensuring that only the object instances appear in grayscale while the background retains its original colors. This process produces a hybrid RGB-grayscale image, as illustrated in Fig. 2, while preserving the structural integrity of the objects. SAGA can be represented as the following equation, which converts given image I by multiplying the red, green and blue channel of each pixels by the values [0.2989, 0.587, 0.114] respectively,

$$I = I \times [0.2989, 0.5870, 0.1140] \quad (3)$$

RGB-to-IR adaptation highlights that directly transferring knowledge from the RGB source domain to the IR target

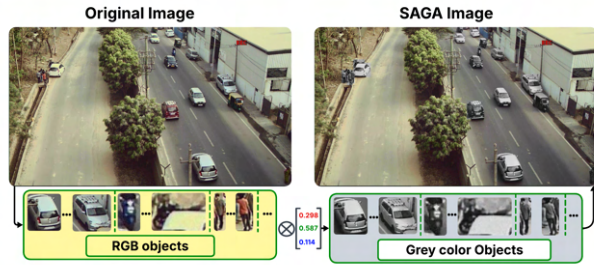


Figure 2. Illustration of SAGA augmentation. The process involves extracting objects from the image, converting them to grayscale, and reintegrating them into the original image while preserving background color information.

domain can lead to the transfer of irrelevant features via the EMA technique [13], as there is no explicit mechanism to selectively extract meaningful information. This challenge contributes to an increase in false positives within the pseudo labels generated by the teacher network. To mitigate this, SAGA is integrated into the mean-teacher framework, enabling effective knowledge transfer from labeled RGB images to the unlabeled IR target domain. SAGA is used for domain-adaptive thermal object detection, where source images undergo instance-level gray augmentation to enhance feature adaptation before being trained within their respective frameworks, as shown in Fig. 3.

Sensor	Resolution	Wavelength	FoV
RGB camera	1280x720	400-700nm	60°
IR Camera	640x480	8-14 μ m	32°

Table 3. Sensor specifications of the DragonEye 2 RGB-IR camera used for the IndraEye dataset.

3.3. Aerial Multi-Modal Perception Dataset

Acquisition process The IndraEye dataset was collected using a DJI M600 Pro drone equipped with a multi-sensor RGB-IR camera mounted on a gimbal, allowing for image capture from various slant angles by dynamically adjusting the pitch and yaw. The dataset also consists of data collected by mounting the RGB-IR camera on a 3 meter elevated tripod placed over places such as bridge and underpass, there by accounting for low range slant angle that is similar to that of a traffic pole mounted camera images. The information about the sensor can be found in Table 3. IndraEye is one of the very few mutlispatial and multitemporal slant angle dataset available at two modalities i.e. RGB and IR. The dataset was manually structured by annotating the potential objects of interest with bounding boxes in RGB and IR images. The video was shot for approximately 4 to 5 minutes in each of the

different locations as shown in Table 1. This generated around 20,000 frames per scene as the refresh rate of the RGB camera is 50Hz and that of IR is 10Hz. To avoid redundancy of objects we consider only 35th frame in the video to obtain diverse object rich frames.

Camera calibration Recent work [2] has shown that calibrating RGB and IR camera on a drone vehicle under motion is challenging and not feasible due to the variations in camera parameters such as focal length, field of view (FOV) and refresh rate during movement of the drone. This is a serious issue since most of the camera calibration methods requires computing homography matrix and running feature matching algorithm such as SIFT[37], SURF[1] on the device to align both RGB and IR images. This process limits edge-based machine learning due to computational constraints [2]. Stereo camera calibration is effective but challenged by differing RGB-IR FOVs, leading to depth-dependent parallax [17]. This occurs due to variations in sensor parameters and viewing angles, causing objects at different depths to appear in different image positions.

This misalignment becomes evident in the fused data when the vehicle with the camera is in motion [2]. Image alignment can lead to inaccuracies, particularly with long distance and small objects. As a result, the process of co-registration is avoided intentionally. While the RGB and IR images are captured with the same timestamp, they still exhibit minor misalignment due to change in refresh rate of the sensor. This limitation highlights the importance of multimodal approaches that eliminate the necessity for co-registered images, such as domain adaptation [13], [16]. **Statistics of the dataset** IndraEye comprises 5,612 images that includes multiple viewing angles, altitudes, backgrounds, and times of day. The RGB images are divided into 2,336 samples (2,026 for training, 60 for validation, and 250 for testing), while the IR images are divided into 3,276 samples (2,973 for training, 58 for validation, and 245 for testing). The dataset additionally includes day and night splits. The goal was to create a diverse and comprehensive dataset with both RGB and IR images, suitable for various conditions and contexts. IndraEye features 13 classes: backhoe loader, bicycle, bus, car, cargo truck, cargo trike (a medium-sized three-wheeled cargo vehicle), ignore, motorcycle, person, rickshaw (a small three-wheeled passenger vehicle), small truck, truck, tractor, and van. More information regarding the dataset is provided in Table 1. To improve object diversity and reduce redundancy, every 35th frame was extracted from the 20,000 frames captured per scene.

Annotation For the IndraEye dataset, we provide ground-truth labels for both object detection and pixel-level semantic segmentation across all classes using a two-stage




Figure 3. Domain-adaptive thermal object detection with RGB as the source domain and IR as the target domain. (a) Vanilla CMT on the IndraEye dataset. (b) CMT with SAGA on the IndraEye dataset.

approach.

- 1. Zero-Shot Annotations with Human in the Loop:** In this initial step, we generate zero-shot annotations by utilizing pre-trained models. These models are used to produce preliminary annotations with human oversight, ensuring that the annotations align closely with the true object locations and characteristics.
- 2. Manual Verification and Refinement:** After generating the initial annotations, each label undergoes thorough manual verification. This step involves a detailed review of the annotations to correct any discrepancies and refine the labels for higher accuracy.

This structured approach ensures that both object detection and pixel-level annotations are meticulously crafted and reliable. For the object detection task, RGB and IR images are manually annotated with bounding boxes using the X-AnyLabeling tool [61]. This tool facilitates efficient annotation by allowing users to load pre-trained models, such as those from the YOLO family of networks. After multiple iterations of correction and evaluation, highly precise annotations have been generated. For semantic segmentation, we use SAMRS [60], an extension of SAM [31], optimized for aerial imagery. SAMRS leverages remote sensing datasets to efficiently generate masks for large-scale segmentation tasks. This approach has streamlined our process for producing pixel-level annotations for IndraEye. The annotations for both tasks are carefully designed to match the established class schema.

Ethics and policy The proposed dataset undergoes a thorough manual review to ensure privacy protection. Faces of individuals and vehicle license plates that are clearly visible in the images are blurred to maintain confidentiality and safeguard personal rights. This process ensures that individuals’ privacy is upheld while preserving the dataset’s utility.

4. Experiments

In this section, we present the experimental results of the proposed SAGA on state-of-the-art unsupervised domain adaptation models, including D3T [13], CMT [3], and AT

[34]. These models are evaluated on two open-source datasets, FLIR [66] and LLVIP [26], along with the proposed IndraEye dataset. We also assess object detection models across various settings to establish a benchmark on the IndraEye dataset.

4.1. Datasets and Evaluation metric

Datasets Experimental evaluation is conducted using the FLIR, LLVIP, and IndraEye datasets. FLIR [66] is collected from an autonomous driving perspective. This dataset consists of 5,142 aligned RGB-IR image pairs, with 4,129 images allocated for training and 1,013 for testing, covering daytime and nighttime scenes. LLVIP [26] is a low-light vision dataset captured from a slant angle, roughly three stories above ground. It contains 15,488 RGB-IR image pairs, mainly featuring extremely dark scenes and focusing solely on the Pedestrian class. IndraEye is a slant-angle dataset with 2,336 RGB and 3,276 IR images. Captured from diverse backgrounds and angles, it covers 13 classes, including vehicles and pedestrians, with day and night scenes under varying illumination.

Evaluation metrics For experimental validation, we evaluate all the networks using the Mean Average Precision (mAP) score, calculated from precision and recall scores. An mAP threshold of 0.5 is used for all experiments.

Implementation Details For evaluating domain adaptation networks, we train the models using strategies from [13], [3], and [34]. During the burn-up stage, the student network trains on the source domain for 10,000 iterations. Adaptation to the target domain occurs in subsequent iterations, with a total of 50,000 iterations on the FLIR and LLVIP datasets. For the IndraEye dataset, we train all algorithms for 20,000 iterations, as the D3T [13] network diverges after this point, leading to poor results. We hypothesize this is due to the dataset’s complexity and the frequent alternation between the student and teacher networks during training (zig-zag training). All experiments in 5 used a batch size of 16 on four NVIDIA V100 GPUs with a VGG16 backbone and the same hyperparameters as the respective works.

Object detection models, including FasterRCNN, ReDet, and ORCNN, were benchmarked on the IndraEye dataset using four Nvidia V100 GPUs with a batch size of 6. We employed the SGD optimizer with a learning rate of 0.0025, weight decay of 0.0001, and momentum of 0.9, training for 12 epochs. The YOLOv8[28] model was trained with a learning rate of 0.00001, a batch size of 16, and an image size of 640 for a total of 100 epochs.

4.2. Experimental Results

4.2.1. Unsupervised Domain Adaptation

As observed in table 5, the proposed SAGA consistently enhances the performance of both CMT and D3T across all three datasets. Specifically, it improves CMT by 2.4 mAP

Models	Train on RGB	Test on RGB	Train on IR	Test on IR	mAP50
FasterRCNN [41]	✓	✓	✗	✗	47.6
ReDet [21]	✓	✓	✗	✗	43.9
ORCNN [64]	✓	✓	✗	✗	56.3
Faster RCNN [41]	✗	✗	✓	✓	57.3
ReDet [21]	✗	✗	✓	✓	43.8
ORCNN [64]	✗	✗	✓	✓	65.6

Table 4. Performance of object detection algorithms on the proposed IndraEye dataset.

on FLIR, 3.6 mAP on LLVIP, and 2.2 mAP on IndraEye respectively. Similarly, it enhances the performance of D3T by 4.29 mAP on FLIR. Similar trends are observed in other settings. This clearly demonstrates that transferring all information from the student to the teacher is not beneficial. Instead, mitigating color bias while preserving semantic information leads to consistently better performance. This makes the proposed augmentation technique highly effective for RGB-IR domain adaptation, as it consistently improves performance across different look angles, which is crucial for drone imagery.

Compatibility with existing techniques. Table 9 demonstrates that SAGA is compatible with various existing techniques along with domain adaptation approaches. For methods designed to enhance the representation capacity of encoders, such as Jigen[5], SAGA does not negatively impact training despite being an augmentation technique. Instead, it complements these methods, supporting their learning and improving the representation ability of models for RGB-IR adaptation.

Comparison with augmentation techniques: Although grayscale augmentation helps align RGB to Thermal for domain adaptation, we show that full grayscale augmentation reduces CMT model performance on the FLIR dataset, as observed in table 8. Converting the entire image to grayscale removes important local details, like fine boundaries, making it difficult for the model to identify key features. In contrast, instance-level grayscale augmentation retains these details, enhancing performance.

Qualitative results Qualitative predictions show improved performance for small-scale objects (Fig. 4), a common challenge in IR imagery. Unlike RGB, IR sensors lose information from certain viewpoints, especially when objects are viewed from different angles. Grayscale augmentation helps address this issue. We observe a significant reduction in false positives in the vanilla CMT model, and incorporating the proposed augmentation effectively mitigates this problem, while consistently improving performance across three benchmarked datasets, as shown in Table 5.

4.2.2. Benchmarking on IndraEye dataset

Object Detection Detection models like ReDet [21], ORCNN [64], and Faster-RCNN [41] show limited in-domain performance (training and testing on the same modality),

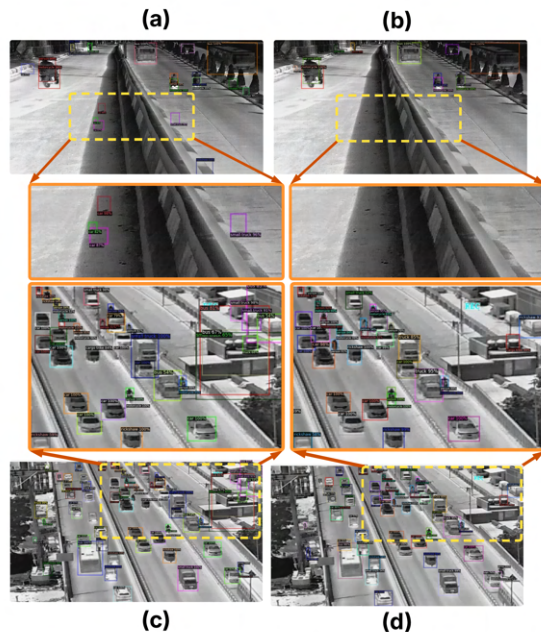


Figure 4. Output predictions to highlight the importance of the SAGA augmentation on CMT algorithm. (a) and (c) shows the increase in false positives while using vanilla CMT. Meanwhile (b) and (d) shows the reduction in false positives when using SAGA with CMT, showcasing its effectiveness.

Algorithm	FLIR	LLVIP	IndraEye
Source	37.8	82.8	50.5
Oracle	67.5	83	65.5
D3T	61.8	88.1	30.3
D3T + SAGA	66.1 ↑ 4.3	88.5 ↑ 0.4	32.9 ↑ 2.6
CMT	63.3	41.6	17.1
CMT + SAGA	65.7 ↑ 2.4	49.2 ↑ 7.6	19.3 ↑ 2.2

Table 5. Experimental Results on State-of-the-art algorithms on two standard datasets i.e, FLIR and LLVIP. Source refers to the accuracy obtained by training on RGB domain and testing on IR domain, while Oracle refers to training and testing on the same i.e. IR domain.

falling short of real-world, safety-critical application requirements, as shown in Table 4. This can be attributed to factors like scale variability, varying illumination and background conditions, and challenges specific to RGB sensors, which are common in drone imagery. This highlights the significance of the IndraEye dataset and the need for model architectures that are generalizable and adaptable to diverse scenarios.

Table 6 evaluates the performance of Yolov8x [27] on IndraEye on five different configurations. The four configurations include: first, training on the RGB modality with daytime images and testing on nighttime images under low il-

Train Setting	Test Configuration	mAP50
RGB-Day	RGB-Night	30.0
RGB-Day + RGB-Night	RGB-Night	52.0
IR-Day + IR-Night	IR-Night	73.6
RGB-Day + RGB-Night	RGB-Day	90.7
IR-Day + IR-Night	IR-Day	77.0

Table 6. Table illustrating the features of the two modalities complementing each other, utilizing a day-night split within the train-test split of the IndraEye dataset.

illumination conditions, highlighting the limitations of RGB imagery in low-light environments. The second configuration assesses the effectiveness of using only RGB sensors for aerial perception, where training incorporates both day and nighttime images, but inference is conducted solely under nighttime conditions. The third configuration is completely based on the capability of the model learning from the IR imagery. Precisely, in this setting IR modality with both day and nighttime images are used to train the DNN. As can be seen from the table 6, comparison from the first three test configurations, it is clear that the IR modality is better suited for low-illumination conditions. The fourth and fifth configurations are counterparts to the second and third, evaluated under well-illuminated (daytime) conditions. Results show that RGB sensors, capturing rich texture, perform better in these conditions, while IR sensors excel in low-light environments due to their ability to capture thermal emissivity.

To benchmark other aerial perception datasets against IndraEye in the RGB domain, we train YOLOv8x on the VisDrone dataset, which has more training image instances than IndraEye. For a fair comparison, we focus on the common classes (5 classes) between both datasets. Despite VisDrone’s larger training set, the cross-domain gap (between IndraEye to VisDrone and VisDrone to IndraEye) reveals that the IndraEye dataset helps the model generalize better, as shown in Table 7. We attribute this improvement to the greater diversity and higher quality annotations in the IndraEye dataset.

Train set	Test set	mAP50
VisDrone	VisDrone [69]	53.3
IndraEye	VisDrone	32.5 ↓ 20.8
IndraEye	IndraEye	83.1
VisDrone	IndraEye	52.1 ↓ 31.05

Table 7. Generalization of object detection model trained on Vis-drone RGB imagery datasets.

Model	Vanilla	Full gray	SAGA	mAP50
	✓			63.3
CMT	✗	✓		61.1
	✗	✗	✓	65.7

Table 8. Table illustrating the ablation study of SAGA augmentation and Jigen Pretraining effect on CMT architecture.

Model	mAP50
Oracle	63.3
Source	61.1
D3T	61.8
D3T+SAGA	66.1
D3T + Jigen	61.5
D3T + SAGA + Jigen	68.2 ↑ 6.4
CMT	65.7
CMT + SAGA	65.7
CMT + Jigen	64.7
CMT + SAGA + Jigen	68.3 ↑ 2.6

Table 9. SAGA on state-of-the-art domain adaptation methods, highlighting its compatibility with existing techniques like Jigen.

5. Conclusion

In this paper, we present Semantic-Aware Gray color Augmentation (SAGA), a novel technique that mitigates color bias and bridges the domain gap by extracting relevant object-level features for IR images. By integrating SAGA with the mean-teacher framework, we enable efficient knowledge transfer from labeled RGB images to unlabeled IR domains thereby achieving consistent performance gains of +0.4% to +7.6% when integrated with SOTA domain adaptation methods. To tackle the scarcity of RGB-IR drone imagery datasets, we present IndraEye, the first multisensor aerial perception dataset tailored for diverse classes and illumination conditions. The dataset includes diverse background environments and scene characteristics, featuring 2,336 RGB images and 3,276 IR images across different viewing angles, altitudes, and times of day. Our experimental results demonstrate that SAGA significantly enhances RGB-to-IR adaptation for autonomous driving and the IndraEye dataset.

6. Acknowledgments

This work was supported by The Bengaluru Traffic Police (BTP), The Ministry of Electronics and Information Technology (MeitY) & The Kotak IISc AI-ML Centre (KIAC).

References

- [1] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Computer Vision—ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006. Proceedings, Part I 9*, pages 404–417. Springer, 2006. 5
- [2] Martin Brenner, Napoleon H Reyes, Teo Susnjak, and Andre LC Barczak. Rgb-d and thermal sensor fusion: a systematic literature review. *IEEE Access*, 2023. 1, 3, 5
- [3] Shengcao Cao, Dhiraj Joshi, Liang-Yan Gui, and Yu-Xiong Wang. Contrastive mean teacher for domain adaptive object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23839–23848, 2023. 3, 4, 6
- [4] Yaru Cao, Zhijian He, Lujia Wang, Wenguan Wang, Yixuan Yuan, Dingwen Zhang, Jinglin Zhang, Pengfei Zhu, Luc Van Gool, Junwei Han, et al. Visdrone-det2021: The vision meets drone object detection challenge results. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 2847–2854, 2021. 4
- [5] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2229–2238, 2019. 7
- [6] Chaoqi Chen, Zebiao Zheng, Xinghao Ding, Yue Huang, and Qi Dou. Harmonizing transferability and discriminability for adapting object detectors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8869–8878, 2020. 3
- [7] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 3
- [8] Ke Cheng, Xuecheng Hua, Hu Lu, Juanjuan Tu, Yuanquan Wang, and Shitong Wang. Multi-scale semantic correlation mining for visible-infrared person re-identification. *arXiv preprint arXiv:2311.14395*, 2023. 4
- [9] Jinhong Deng, Wen Li, Yuhua Chen, and Lixin Duan. Unbiased mean teacher for cross-domain object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4091–4101, 2021. 3
- [10] Jinhong Deng, Dongli Xu, Wen Li, and Lixin Duan. Harmonious teacher for cross-domain object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23829–23838, 2023. 3
- [11] Chaitanya Devaguptapu, Ninad Akolekar, Manuj M Sharma, and Vineeth N Balasubramanian. Borrow from anywhere: Pseudo multi-modal object detection in thermal imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 3
- [12] Jian Ding, Nan Xue, Gui-Song Xia, Xiang Bai, Wen Yang, Michael Ying Yang, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, et al. Object detection in aerial images: A large-scale benchmark and challenges. *IEEE transactions on pattern analysis and machine intelligence*, 44(11): 7778–7796, 2021. 3
- [13] Dinh Phat Do, Taehoon Kim, Jaemin Na, Jiwon Kim, Keonho Lee, Kyunghwan Cho, and Wonjun Hwang. D3t: Distinctive dual-domain teacher zigzagging across rgb-thermal gap for domain-adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23313–23322, 2024. 1, 3, 4, 5, 6
- [14] Dawei Du, Yuankai Qi, Hongyang Yu, Yifan Yang, Kaiwen Duan, Guorong Li, Weigang Zhang, Qingming Huang, and Qi Tian. The unmanned aerial vehicle benchmark: Object detection and tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 370–386, 2018. 4
- [15] Gianni Franchi, Marwane Hariat, Xuanlong Yu, Nacim Belkhir, Antoine Manzanera, and David Filliat. Infraparis: A multi-modal and multi-task autonomous driving dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2973–2983, 2024. 4
- [16] 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. 5
- [17] Zeren Gao, Yue Gao, Yong Su, Yang Liu, Zheng Fang, Yaru Wang, and Qingchuan Zhang. Stereo camera calibration for large field of view digital image correlation using zoom lens. *Measurement*, 185:109999, 2021. 5
- [18] Yunpeng Gong, Jiaquan Li, Lifei Chen, and Min Jiang. Exploring color invariance through image-level ensemble learning. *arXiv preprint arXiv:2401.10512*, 2024. 4
- [19] Alejandro González, Zhijie Fang, Yainuvis Socarras, Joan Serrat, David Vázquez, Jiaolong Xu, and Antonio M López. Pedestrian detection at day/night time with visible and fir cameras: A comparison. *Sensors*, 16(6):820, 2016. 3
- [20] Qishen Ha, Kohei Watanabe, Takumi Karasawa, Yoshitaka Ushiku, and Tatsuya Harada. Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5108–5115. IEEE, 2017. 1
- [21] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2786–2795, 2021. 7
- [22] Yunzhong Hou, Liang Zheng, and Stephen Gould. Learning to structure an image with few colors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10116–10125, 2020. 4
- [23] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020. Proceedings, Part IX 16*, pages 733–748. Springer, 2020. 3
- [24] Soonmin Hwang, Jaesik Park, Namil Kim, Yukyung Choi, and In So Kweon. Multispectral pedestrian detection:

- Benchmark dataset and baselines. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 3
- [25] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5001–5009, 2018. 3
- [26] Xinyu Jia, Chuang Zhu, Minzhen Li, Wenqi Tang, and Wenli Zhou. Llvip: A visible-infrared paired dataset for low-light vision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3496–3504, 2021. 3, 6
- [27] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics YOLO, 2023. 7
- [28] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics yolov8, 2023. 6
- [29] Josy John, K Harikumar, J Senthilnath, and Suresh Sundaram. An efficient approach with dynamic multiswarm of uavs for forest firefighting. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(5):2860–2871, 2024. 1
- [30] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12456–12465, 2019. 3
- [31] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 6
- [32] Zülfiye Kütük and Görkem Algan. Semantic segmentation for thermal images: A comparative survey. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 286–295, 2022. 1
- [33] Gao Yu Lee, Jinkuan Chen, Tanmoy Dam, Md Meftahul Ferdous, Daniel Puiu Poenar, and Vu N Duong. Dehazing remote sensing and uav imagery: A review of deep learning, prior-based, and hybrid approaches. *arXiv preprint arXiv:2405.07520*, 2024. 2
- [34] Yu-Jhe Li, Xiaoliang Dai, Chih-Yao Ma, Yen-Cheng Liu, Kan Chen, Bichen Wu, Zijian He, Kris Kitani, and Peter Vajda. Cross-domain adaptive teacher for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, 6
- [35] 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. 3, 4
- [36] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. *arXiv preprint arXiv:2102.09480*, 2021. 3
- [37] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 5
- [38] Decao Ma, Juan Su, Shaopeng Li, and Yong Xian. Aerial-irgan: unpaired aerial visible-to-infrared image translation with dual-encoder structure. *Scientific Reports*, 14(1):22105, 2024. 2
- [39] You Ma, Lin Chai, Lizuo Jin, and Jun Yan. Hierarchical alignment network for domain adaptive object detection in aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 208:39–52, 2024. 3
- [40] Sebastien Razakarivony and Frederic Jurie. Vehicle detection in aerial imagery: A small target detection benchmark. *Journal of Visual Communication and Image Representation*, 34:187–203, 2016. 3, 4
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2015. 7
- [42] Prodip Kumar Sarker, Qingjie Zhao, and Md Kamal Uddin. Transformer-based person re-identification: A comprehensive review. *IEEE Transactions on Intelligent Vehicles*, 9(7):5222–5239, 2024. 1
- [43] J Senthilnath, K Harikumar, and Suresh Sundaram. Metacognitive decision-making framework for multi-uav target search without communication. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(5):3195–3206, 2024. 1
- [44] Aniruddh Sikdar, Sumanth Udupa, Suresh Sundaram, and Narasimhan Sundararajan. Fully complex-valued fully convolutional multi-feature fusion network (fc 2 mfn) for building segmentation of insar images. In *2022 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 581–587. IEEE, 2022. 1
- [45] Aniruddh Sikdar, Jayant Teotia, and Suresh Sundaram. Contrastive learning-based spectral knowledge distillation for multi-modality and missing modality scenarios in semantic segmentation. *arXiv preprint arXiv:2312.02240*, 2023. 1
- [46] Aniruddh Sikdar, Sumanth Udupa, Prajwal Gurunath, and Suresh Sundaram. Deepmao: Deep multi-scale aware overcomplete network for building segmentation in satellite imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 487–496, 2023. 1
- [47] Aniruddh Sikdar, Qiranul Saadiyeen, Prahlad Anand, and Suresh Sundaram. Ssl-rgb2ir: Semi-supervised rgb-to-ir image-to-image translation for enhancing visual task training in semantic segmentation and object detection. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5017–5023. IEEE, 2024. 3
- [48] Aniruddh Sikdar, Jayant Teotia, and Suresh Sundaram. Skd-net: Spectral-based knowledge distillation in low-light thermal imagery for robotic perception. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9041–9047. IEEE, 2024. 1
- [49] Aniruddh Sikdar, Jayant Teotia, and Suresh Sundaram. Ogp-net: Optical guidance meets pixel-level contrastive distillation for robust multi-modal and missing modality segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(7):6922–6930, 2025. 1

- [50] Kechen Song, Xiaotong Xue, Hongwei Wen, Yingying Ji, Yunhui Yan, and Qinggang Meng. Misaligned visible-thermal object detection: a drone-based benchmark and baseline. *IEEE Transactions on Intelligent Vehicles*, 2024. 1, 3
- [51] Open source FLIR dataset available online. Free teledyne flir thermal dataset for algorithm training. <https://www.flir.com/oem/adas/adas-dataset-form/>. 3, 4
- [52] 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*, pages 1–1, 2022. 3
- [53] Jiashun Suo, Tianyi Wang, Xingzhou Zhang, Haiyang Chen, Wei Zhou, and Weisong Shi. Hit-uav: A high-altitude infrared thermal dataset for unmanned aerial vehicle-based object detection. *Scientific Data*, 10(1):227, 2023. 4
- [54] Linfeng Tang, Jiteng Yuan, Hao Zhang, Xingyu Jiang, and Jiayi Ma. Piafusion: A progressive infrared and visible image fusion network based on illumination aware. *Information Fusion*, 83:79–92, 2022. 4
- [55] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [56] WE Tennant. Limits of infrared imaging. *International Journal of High Speed Electronics and Systems*, 20(03):529–539, 2011. 2
- [57] Sumanth Udupa, Prajwal Gurunath, Aniruddh Sikdar, and Suresh Sundaram. Mrfp: Learning generalizable semantic segmentation from sim-2-real with multi-resolution feature perturbation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5904–5914, 2024. 1
- [58] Berkcan Ustun, Ahmet Kagan Kaya, Ezgi Cakir Ayerden, and Fazil Altinel. Spectral transfer guided active domain adaptation for thermal imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 449–458, 2023. 2
- [59] Abhinav Valada, Rohit Mohan, and Wolfram Burgard. Self-supervised model adaptation for multimodal semantic segmentation. *International Journal of Computer Vision*, 128(5):1239–1285, 2020. 1
- [60] Di Wang, Jing Zhang, Bo Du, Minqiang Xu, Lin Liu, Dacheng Tao, and Liangpei Zhang. Samrs: Scaling-up remote sensing segmentation dataset with segment anything model. *Advances in Neural Information Processing Systems*, 36, 2024. 6
- [61] Wei Wang. Advanced auto labeling solution with added features. <https://github.com/CVHub520/X-AnyLabeling>, 2023. 6
- [62] Ancong Wu, Wei-Shi Zheng, Hong-Xing Yu, Shaogang Gong, and Jianhuang Lai. Rgb-infrared cross-modality person re-identification. In *Proceedings of the IEEE international conference on computer vision*, pages 5380–5389, 2017. 2, 4
- [63] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018. 4
- [64] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3520–3529, 2021. 7
- [65] Junying Zeng, Yajin Gu, Chuanbo Qin, Xudong Jia, Senyao Deng, Jiahua Xu, and Huiming Tian. Unsupervised domain adaptation for remote sensing semantic segmentation with the 2d discrete wavelet transform. *Scientific Reports*, 14(1):23552, 2024. 2
- [66] 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. 6
- [67] Qiang Zhang, Changzhou Lai, Jianan Liu, Nianchang Huang, and Jungong Han. Fmcnet: Feature-level modality compensation for visible-infrared person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7349–7358, 2022. 2, 4
- [68] Zhiwei Zhao, Bin Liu, Qi Chu, Yan Lu, and Nenghai Yu. Joint color-irrelevant consistency learning and identity-aware modality adaptation for visible-infrared cross modality person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, pages 3520–3528, 2021. 4
- [69] Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Heng Fan, Qinghua Hu, and Haibin Ling. Detection and tracking meet drones challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7380–7399, 2021. 8