IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

Domain Adaptive Detection of MAVs: A Benchmark and Noise Suppression Network

Yin Zhang, Jinhong Deng^(D), Peidong Liu^(D), Wen Li^(D), Member, IEEE, and Shiyu Zhao^(D), Member, IEEE

Abstract-Visual detection of Micro Air Vehicles (MAVs) has attracted increasing attention in recent years due to its important application in various tasks. The existing methods for MAV detection assume that the training set and testing set have the same distribution. As a result, when deployed in new domains, the detectors would have a significant performance degradation due to domain discrepancy. In this paper, we study the problem of cross-domain MAV detection. The contributions of this paper are threefold. 1) We propose a Multi-MAV-Multi-Domain (M3D) dataset consisting of both simulation and realistic images. Compared to other existing datasets, the proposed one is more comprehensive in the sense that it covers rich scenes, diverse MAV types, and various viewing angles. A new benchmark for cross-domain MAV detection is proposed based on the proposed dataset. 2) We propose a Noise Suppression Network (NSN) based on the framework of pseudo-labeling and a large-tosmall training procedure. To reduce the challenging pseudo-label noises, two novel modules are designed in this network. The first is a prior-based curriculum learning module for allocating adaptive thresholds for pseudo labels with different difficulties. The second is a masked copy-paste augmentation module for pasting truly-labeled MAVs on unlabeled target images and thus decreasing pseudo-label noises. 3) Extensive experimental results verify the superior performance of the proposed method compared to the state-of-the-art ones. In particular, it achieves mAP of 46.9%(+5.8%), 50.5%(+3.7%), and 61.5%(+11.3%) on the tasks of simulation-to-real adaptation, cross-scene adaptation, and cross-camera adaptation, respectively.

Note to Practitioners— To study the cross-domain MAV detection problem, this paper establishes a novel benchmark that consists of three domain adaptation tasks: simulation-to-real

Manuscript received 25 January 2024; accepted 16 February 2024. This article was recommended for publication by Associate Editor S. Wang and Editor D. Song upon evaluation of the reviewers' comments. This work was supported in part by the Hangzhou Key Technology Research and Development Program under Grant 20212013B09 and in part by the Research Center for Industries of the Future at Westlake University under Grant WU2022C027. (*Corresponding author: Shiyu Zhao.*)

Yin Zhang is with the College of Computer Science and Technology, Zhejiang University, Hangzhou 310027, China, and also with the School of Engineering, Westlake University, Hangzhou 310024, China (e-mail: zhangyin@westlake.edu.cn).

Jinhong Deng and Wen Li are with the School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China (e-mail: jhdeng1997@gmail.com; liwen@uestc.edu.cn).

Peidong Liu is with the School of Engineering, Westlake University, Hangzhou 310024, China (e-mail: liupeidong@westlake.edu.cn).

Shiyu Zhao is with the Research Center for Industries of the Future, School of Engineering, Westlake University, Hangzhou 310024, China, and also with the Westlake Institute for Advanced Study, Hangzhou 310024, China (e-mail: zhaoshiyu@westlake.edu.cn).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TASE.2024.3370147.

Digital Object Identifier 10.1109/TASE.2024.3370147

adaptation, cross-scene adaptation, and cross-camera adaptation, respectively. The benchmark is based on a novel MAV dataset called Multi-MAV-Multi-Domain (M3D), which is available at: https://github.com/WestlakeAerialRobotics/M3D. To reduce the noises caused by pseudo labels, a noise suppression network is proposed to overcome the error accumulation. Extensive experiments are conducted to prove the effectiveness of the proposed approach.

1

Index Terms—MAV detection, domain adaptation, MAV dataset, noise suppression.

I. INTRODUCTION

AV detection [1], [2], [3] has attracted increasing attention due to its important applications in various tasks such as multi-MAV swarming [4] and detection of malicious MAVs [5]. Although in recent years vision-based MAV detection has been studied under different setups [6], [7], [8], [9], MAV detection still faces critical challenges to be applied in practice [10]. The existing MAV detection methods aim to improve the performance of MAV detection under the assumption that the training set and testing set have the same distribution. As a result, deploying these detectors in new domains (e.g., different environments) would result in significant performance degradation due to the domain discrepancy. To this end, it is necessary to develop novel cross-domain MAV detection approaches.

Cross-domain object detection has been studied in recent years. The benchmarks in this area contain inverse weather adaptation, simulation-to-realistic adaptation, real-to-artistic adaptation, and day-to-night adaptation [11], [12], [13]. Due to the small amount of data in some target domains, a few works even exceed the supervised learning approaches [14]. Existing methods for cross-domain object detection can be classified into domain mapping, adversarial learning, consistency learning, and pseudo-labeling. For the works focusing on the style differences, domain-mapping methods (e.g., imageto-image translation) are adopted in the adaptation algorithms to reduce the appearance gap between the source domain and target domain [15], [16]. Adversarial learning has been widely explored to minimize domain discrepancy [14], [17]. Some works propose consistency-based methods to further improve the adaptation performance [18]. Pseudo-labeling has become popular due to its simplicity and effectiveness despite the performance being restricted by the pseudo-label noises [19].

While advancements have been made in cross-domain object detection, it is crucial to notice that the challenges

1545-5955 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See https://www.ieee.org/publications/rights/index.html for more information.

2

IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING



Fig. 1. Samples from the proposed Multi-MAV-Multi-Domain (M3D) dataset. The top to bottom shows examples from the M3D-Sim subset and M3D-Real subset, respectively. The backgrounds of MAVs contain many types: mountains, buildings, villages, rivers, deserts, farmlands, parks, and roads. MAVs are usually small objects in captured images. The enlarged MAVs containing diverse types are presented at the lower right of each picture.

specific to MAV detection require tailored solutions. First, the backgrounds of MAVs are more diverse than common objects like cars or pedestrians. As shown in Fig. 1, MAVs captured by cameras with multi-viewing angles (i.e., ground-to-air and air-to-air) may appear in diverse scenes. The diversity and complexity of the backgrounds cause severe challenges for MAV detection. Second, due to the requirement for long-range detection, the imaging sizes of MAVs are usually much smaller than that of the objects in natural images. For example, the average area of the objects in the real-to-artistic adaptation task occupies 12.7% of the entire image [20], while the MAVs in Drone-vs-Bird dataset only take 0.10% [21]. Small objects further increase the difficulty of detection. Finally, MAV detection requires real-time detection, which is an issue that is usually ignored by cross-domain object detection researchers cause most methods are based on two-stage detection networks with high accuracy but low efficiency. Therefore, it is necessary to develop novel cross-domain MAV detection approaches tailored to the unique characteristics and challenges of MAV detection.

In light of these challenges, this paper focuses on the crossdomain MAV detection problem. We aim to use the unlabeled target data to help the MAV detectors adapt to the target domain. The contributions of this paper are threefold.

1) We propose a Multi-MAV-Multi-Domain (M3D) dataset to study the cross-domain MAV detection problem. Compared to the state-of-the-art MAV datasets [10], [21], [22], our dataset contains different styles of images, more diverse environments, and more types of MAVs. Our dataset consists of two subsets: M3D-Sim and M3D-Real with 28,740 and 55,259 images, respectively. Some image examples are shown in Fig. 1. M3D-Sim is collected from a simulation platform called Unreal Engine [23] and includes 22 MAVs and 16 environments. M3D-Real is a realistic dataset that contains diverse scenes and 10 types of MAVs.

Based on the proposed M3D and other existing MAV datasets, we construct a novel benchmark that contains several representative tasks: *simulation-to-real adaptation*, *cross-scene adaptation*, and *cross-camera adaptation*, respectively. A comprehensive experiment is also conducted on the M3D dataset for universal MAV detection. To our best knowledge, we are among the first to benchmark the problem of cross-domain MAV detection and study this issue in this area.

2) We propose a novel Noise Suppression Network (NSN) to address the domain discrepancy through pseudo-labeling.

The inaccuracy of the pseudo labels will increase when meeting smaller MAVs under more diverse backgrounds. NSN can effectively overcome the limitations of existing unsupervised domain adaptation methods. NSN consists of two modules, a prior-guided curriculum learning module, and a masked copy-paste augmentation module, to fully utilize the domain-specific information of target data and reduce the pseudo-label noises. First, the prior-guided curriculum learning module is for generating more accurate pseudo labels. A novel difficulty matrix is adopted to evaluate the complexity degree of MAV samples. Then, curriculum learning is used to help the model adapt to target samples with different detection difficulties. Second, the masked copy-paste augmentation module is for decreasing the pseudo-label noises by adding more true labels. Cropped MAV images with true labels are segmented automatically first and then pasted on the unlabeled target images. The participation of true labels can help the detection model adapt to new backgrounds and reduce the influence of noises.

These two simple yet effective modules eliminate the error accumulation from self-training caused by pseudo-label noise, which impedes the performance improvements of MAV detectors on the target domain. Besides, to meet the requirement of real-time MAV detection, we also employ a large-to-small model training procedure to help the large model gradually transfer knowledge to the small model. Our NSN method is model-agnostic and can be combined with other detection networks.

3) We conduct extensive experiments on the proposed domain adaptive MAV detection benchmark. Before adaptation, the detection models face significant performance degradation and only achieve mAPs of 39.9%, 28.7%, and 31.9% on the target datasets. The experimental results show the effectiveness of our proposed method. In particular, our method achieves mAPs of 46.9%(+5.8%), 50.5%(+3.7%), and 61.5%(+11.3%) on the tasks of simulation-to-real adaptation, cross-scene adaptation, and cross-camera adaptation, respectively, which significantly outperforms the baseline algorithm [19]. Our method also achieves state-of-the-art results on the proposed adaptation tasks.

We also conduct ablation studies to show the functionalities of different modules in NSN. The ablation study contains three parts, the participation of different modules, the choice of cropped images for data augmentation, and the varying of different hyper-parameters. We also compare our augmentation module with other methods to show the improvement ability on the domain adaptation tasks. The qualitative and experimental analysis provides ample evidence for the feasibility and effectiveness of the proposed method.

II. RELATED WORK

A. Object Detection and MAV Detection

Object detection is a well-studied area and can be separated into one-stage-based and two-stage-based methods [24]. The main difference between these two types of methods is whether the network contains a region proposal network (RPN) module [25]. Generally, one-stage-based methods cost less time but have lower accuracy. YOLO [26], [27], RetinaNet [28], FCOS [29] and Centernet [30] are some representative detection models. Two-stage methods like R-FCN [31], Cascade-RCNN [32], and DETR [33] can produce results with higher accuracy but cost more time. Thus, some researchers aim to make the one-stage-based detection models inference fast with higher precision. One-stage-based methods are suitable for MAV detection because the detection models require to be deployed to embedded devices with limited computation ability.

MAV detection methods can be classified into two types: appearance-based methods [4], [34] and motion-based methods [22], [35]. In this paper, we mainly focus on improving the adaptation performance of appearance-based methods. Many MAV detectors draw lessons from the object detection area and are built on detection models like Faster-RCNN [25] and YOLO [27]. Due to the importance of MAV detection, there have been some works that designed deep-learning networks that specialized in detecting MAVs to solve this issue [36], [37]. [38] proposes MAVNet to achieve real-time detection of MAVs from a semantic segmentation network with fewer parameters. TIB-Net [9] utilizes a tiny iterative backbone to enhance the performance of small objects. FastUAVNet [39] is modified on YOLOv3 [26] by using the inception blocks to extract local and glocal features. However, the performance of those MAV detectors degrades sharply when the training and testing domains are not in the same distribution. To improve the detection performance of MAV detectors in testing environments, we propose the cross-domain MAV detection benchmark for further research.

B. Cross-Domain Object Detection

Cross-domain object detection methods aim to adapt to the unlabelled target domain from the labeled source domain. The benchmarks for cross-domain detection have several types, which are concluded in Table I. Some target datasets in these benchmarks have a small amount of data, making the domain adaptation methods exceed the supervised learning with the help of a large amount of labeled source data [14], [15]. There are also domain adaptation works that specialize in computer vision applications like pedestrian detection [40], face recognition [41], and X-ray object detection [42]. Many methods have been proposed to address the domain discrepancy in object detection. Most cross-domain object detection methods adopt Faster-RCNN as their detection model. There are a few works

TABLE I

The Comparison of Different Tasks in Traditional Domain Adaptation Benchmark. N_S and N_T Represent the Image Numbers in the Source Dataset and Target Dataset

Scenarios	Datasets	N_S	N_T
Simulation-to-real	SIM10K vs Cityscapes	10,000	5,000
Normal-to-foggy	Cityscapes vs Foggyscapes	5,000	5,000
Different cameras	KITTI vs Cityscapes	7,481	5,000
Real-to-artistic	PASCAL VOC vs Clipart	21,503	1,000

that pay attention to one-stage models to study the domain adaptation issue [18], [19], [43], [44], [45].

The cross-domain object detection methods [46] could be mainly divided into adversarial training [14], [17], [47], domain alignment [12], [47], [48], [49], and self-training [50], [51]. Adversarial learning-based methods are suitable to deal with domain adaptation tasks with huge differences. The NLTE method [52] proposes a data mining module and a graph relation module for learning the domain-invariant representations. When the target data doesn't have significant differences from the source data, adversarial training may be ineffective in learning domain-invariant representations. Domain alignment-based methods aim to reduce the distribution discrepancy by aligning features at the region level or instance level. The AsyFOD method [53] aims to reduce the domain gap by conducting the feature alignment between target-dissimilar source instances and augmented target instances to avoid over-adaptation. The MGADA method [54] tries to align the source and target domains from the pixel level, instance level, and category level. However, these two types of methods do not fully exploit the domain-specific information of the target data.

Data augmentation algorithms are widely utilized in cross-domain detection and have many types: GAN-based methods, strong-weak augmentation and mixup augmentation. Many works [55], [56], [57] use GAN-based methods (e.g., CycleGAN [16]) to reduce the gap between different domains. However, this type of method can have satisfactory results in tackling style-difference domains but has poor performance in solving real-to-real or scene-to-scene issues. Strong-weak augmentation methods usually appear with the teacher-student models to learn the appearance-invariant features [58]. Some works [49] use mixup augmentation [59] to generate domain intermediate images. The AcroFOD method [60] proposes a multi-level data augmentation algorithm and a new training strategy to reduce the domain shift. The above methods all attempt to reduce the shifts between source and target domains by making their appearances look more similar. However, the unlabeled target data is still not fully utilized in an unsupervised learning way.

C. Pseudo-Labeling in Domain Adaptation

Pseudo-labeling is a popular way to deal with semi-supervised learning and unsupervised learning tasks due to its simplicity and effectiveness. For domain adaptation tasks, the pseudo-labeling technique has been widely combined with several types of methods to reduce the noises TABLE II THE IMAGE NUMBERS OF DIFFERENT SUBSETS IN OUR M3D DATASET

Subsets	Training Set	Validation Set	Testing Set
M3D-Sim	22,992	5,748	-
M3D-Real	35,081	5,018	15,215
M3D-Sim-Normal	14,855	3,621	-
M3D-Sim-Texture	8,137	2,127	-
M3D-Real-Source	24,441	3,520	12,027
M3D-Real-Target	10,640	1,498	3,188

and obtain pseudo-labels with higher accuracy. The TDD method [61] is based on a unified teacher-student learning framework and a target proposal perceiver module to reduce the domain shift. The SimROD method [19] proposes to use the heavy-weight teacher model to generate more precise pseudo labels. However, when there exists a large domain shift, the teacher model still cannot generate reliable pseudo labels for training. Some works utilize curriculum learning to improve the accuracy of pseudo labels and obtain better performances [62], [63]. Compared to traditional deep learning, curriculum learning is a novel training paradigm by learning from easy to hard examples [64]. However, due to the particularity of the MAV detection problem, existing curriculum domain adaptation methods are less effective in reducing the domain gap. Some works refine pseudo labels by exploiting the intra-class similarity and aligning the source and target domains [65]. Due to the diversity of MAVs, backgrounds, and perspectives, the similarity of the MAV category is non-centralized and hard to utilize.

To avoid recurrent problems, we propose an MAV dataset containing a large amount of data and build a cross-domain MAV detection benchmark based on the dataset. We pay attention to the pseudo-labeling noise generated by self-training methods. The noise of pseudo labels increases when meeting small MAV targets and diverse backgrounds. Two modules are proposed to reduce the noise and help the detection model adapt to target domains better.

III. MULTI-MAV-MULTI-DOMAIN DATASET

We propose an MAV dataset called Multi-MAV-Multi-Domain (M3D) to study the domain adaptation problem. The M3D dataset contains two types of images: simulation images (M3D-Sim) and realistic images (M3D-Real). This dataset contains diverse MAVs, scenes, and viewing angles in both the simulation set and the real-world set. Some examples of these two sets are shown in Fig. 1. Our dataset only contains the MAV class. We do not label other objects so there is only the MAV class in the images. There are a total of 83,999 images with 86,415 boxes in the M3D dataset. Thus there is one MAV average per image. The data is collected under the condition that one MAV pursues another MAV in the air. The image numbers of these subsets are shown in Table II. Details of the subsets are introduced as follows.

A. Simulation Subset: M3D-Sim

M3D-Sim subset contains 28,740 simulation images. There are 22,992 images in the training set and 5,748 images in

the validation set, respectively. To reduce the time and cost caused by human labeling, we use the simulation platform to automatically collect data. Multiple MAVs and environments are included in M3D-Sim to improve the generalization ability of the simulation set. With the help of Unreal Engine and Airsim [68], we collect data on 22 types of MAVs and 16 different environments, whose examples are shown in Fig. 2. The images are captured from multiple perspectives (i.e., groundto-air, air-to-air, and air-to-ground) by the same camera fixed on another MAV. The images collected in the 16 environments establish a subset called M3D-Sim-Normal. Besides, the partial dataset is generated by texture randomization [69] to reduce the gap between simulation and reality. Therefore, the M3D-Sim dataset consists of two subsets: M3D-Sim-Normal and M3D-Sim-Texture, containing 18,476 and 10,264 images, respectively. The purpose of the M3D-Sim-Texture subset is to validate the feasibility of texture randomization methods. Details are shown in Section V-F3. Compared to the stateof-the-art dataset in [70], our dataset contains more types of MAVs and more diverse environments with a moving camera

IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

B. Realistic Subset: M3D-Real

and multi-viewing angles.

The M3D-Real subset contains 55,259 realistic images. Unlike common objects like cars or pedestrians, the collection of MAV images requires the preparation of multiple types of MAVs and a moving camera (e.g., another MAV with a camera). The motion of the camera can guarantee the diversity of shooting perspectives and changes in backgrounds. Thus the data collection of MAVs is more arduous and timeconsuming. In the M3D-Real subset, we purchased 10 types of MAVs, which are shown in Fig. 3(a). These MAVs have different sizes, shapes, and poses in the images. The images are collected in two cities with different natural styles. We collect data at multiple locations whose categories contain mountains, buildings, villages, rivers, deserts, farmlands, parks, and roads. Each category also has several different locations for diversity. Some examples are shown in Fig. 3(b). The trajectory of the MAVs at each location covers an area whose maximum value achieves one square kilometer. Thus the backgrounds of this real subset are much more diverse compared to existing MAV datasets collected by institutes. These characteristics can increase the generalization of the MAV detection network.

To study the impact of background changes on detection, all images are captured by the same camera (DJI H20T). Benefiting from the wide variety of environments, we separate the dataset into a source domain (M3D-Real-Source) and a target domain (M3D-Real-Target) to study the scene-adaptation problem. Fig. 3(b) shows some examples from the two sets for comparison. The M3D-Real-Source subset and the M3D-Real-Target subset contain 39,933 images and 15,326 images, respectively.

C. Comparison With Other MAV Datasets

The comparison of different MAV datasets is shown in Table III. The table contains six representing datasets. We compared various attributes of these datasets, such as the



(a) The 22 types of MAVs in M3D-Sim subset;



(b) The 16 types of simulation environments in M3D-Sim subset.

Fig. 2. The MAVs and environments in M3D-Sim subset.

TABLE III

THE COMPARISON OF DIFFERENT MAV DATASETS. "AIR-TO-AIR" AND "AIR-TO-GROUND" REPRESENT THAT THE IMAGES ARE CAPTURED BY A CAMERA FIXED ON AN MAV IN THE AIR

Attributes	Det-Fly[10]	Real-World[66]	Drone-vs-Bird [21]	MIDGARD [67]	USC-Drone [7]	YouTube [7]	Ours
Images	13,271	55,446	19,027	8,775	18,778	1,764	83,999
Boxes	13,271	55,539	23,335	10,696	1,264	1,675	86,415
Data Source	Institute	Internet	Institute	Institute	Institute	YouTube	Institute
MAV Types	1	-	7	1	1	-	22+10
Simulation	X	×	×	×	X	X	✓
Air-to-Air	1	×	×	1	X	X	1
Ground-to-Air	×	1	1	✓	1	1	✓
Air-to-Ground	1	×	×	×	X	X	1
Moving Camera	1	1	×	1	1	1	1

number of images, data sources, and shooting perspectives. Det-Fly dataset [10] involves multiple shooting angles but only one type of MAV. Real-World dataset [66] has multiple MAVs, but all the images are collected from the Internet. Besides, the images are mainly shot from the ground-to-air perspective. Most of the images in the Drone-vs-Bird dataset [21] have small MAVs, but the camera is fixed on the ground. MIDGARD dataset [67] contains indoor and outdoor environments with a single MAV and low-resolution images.

There is only one MAV (DJI Phantom) in the USC-Drone dataset [7] with a ground-to-air viewing angle. Besides, only one-fifteenth of the data in this dataset is labeled. YouTube dataset [7] only contains a small number of images collected from YouTube.

Our M3D dataset has several advantages. First, our dataset contains images from multi-domains. Both the simulation set and realistic set contain large amounts of images. Thus it is suitable for studying cross-domain MAV detection. Second, 6

IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING



(b) The 20 examples of environments in M3D-Real subset.

Fig. 3. The MAVs and environments in M3D-Real subset.

compared to other datasets, our dataset contains more types of MAVs and more diverse environments. Third, all the images are captured by a moving camera with multiple viewing angles. The diversity of viewing angles guarantees the generality of the data. Due to the diversity of our dataset, we conduct several types of domain adaptation tasks, which are introduced in Section V-A.

IV. PROPOSED METHOD

The motivation of this work is to design a cross-domain MAV detection approach that can reduce domain discrepancy so that the model can perform well in the target domain. Suppose that the source domain $\{S, \mathcal{Y}_s\} = \{(I_i^s, y_i)|_{i=1}^{N_s}\}$ has N_s labeled samples and the target domain $\mathcal{T} = \{I_i^t|_{i=1}^{N_t}\}$ has N_t unlabeled samples, where I_i is the *i*-th image and y_i is the corresponding labels consisting of bounding boxes and category information. The source and the target domains have different but related distributions. For the cross-domain MAV detection task, we aim to obtain an adaptive detection model $\mathcal{M}(S, \mathcal{T})$ by using both the labeled source domain and unlabeled target domain samples.

We use self-training strategies to address the domain discrepancy of cross-domain MAV detection. Here, pseudo-label noises are restrictions to improve the model performance. Thus, we propose a novel Noise Suppression Network (NSN) to suppress the noises from both internal and external ways. NSN is based on the framework of pseudo-labeling and a large-to-small training procedure. This network consists of two novel modules. Section IV-A presents the overall training procedure of our model. Then, we elaborate on two critical modules of our model, i.e., the Prior-guided Curriculum Leaning (PCL) in Section IV-B and Masked Copy-paste Augmentation (MCA) in Section IV-C.

A. Overall Approach

NSN contains a large-to-small model training procedure to transfer the knowledge of a heavyweight model to a lightweight model. Thus, we can obtain a small but accurate adaptive model capable of real-time detection. The training procedure contains three stages and is explained in Algorithm 1.

Stage 1: Denote the large model and the small model as \mathcal{M}_l and \mathcal{M}_s , respectively. We first train these two models on the source domain and obtain a source large model $\mathcal{M}_l(S)$ and a source small model $\mathcal{M}_s(S)$.



Fig. 4. The framework of each training stage in the noise suppression network. 1) The prediction of unlabeled target data is employed as pseudo labels and corrected by the prior-guided curriculum learning module; (see Section IV-B) 2) The unlabeled target data is augmented by the masked copy-paste augmentation module that can generate true labels; (see Section IV-C) 3) The detection model is trained by labeled source data and augmented target data for *E* epochs.

Algorithm 1 Noise Suppression Network (NSN)

Input: Labeled source dataset $\{S, \mathcal{Y}_S\}$, Unlabeled target images \mathcal{T}

Parameter: Training epochs E

Output: Adaptive model $\mathcal{M}_s(\mathcal{S}, \mathcal{T})$

- 1: Stage 1: Train source model $\mathcal{M}_l(\mathcal{S})$ and $\mathcal{M}_s(\mathcal{S})$.
- 2: **Stage 2.1:** Generate pseudo labels $\mathcal{Y}_{\mathcal{T}}^{p} = F_{\mathcal{M}_{l}}(\mathcal{T})$;
- 3: Use the PCL module to correct pseudo labels $\mathcal{Y}^{p}_{\mathcal{T}}$;
- 4: Use MCA to augment \mathcal{T}_A and add true labels $\mathcal{Y}_{\mathcal{T}}^t$;
- 5: Freeze $w_{\mathcal{M}_l}$ except for BatchNorm Layers;
- 6: Use $\{\mathcal{Y}_{\mathcal{S}}|\mathcal{S}, (\mathcal{Y}_{\mathcal{T}}^{t}, \mathcal{Y}_{\mathcal{T}}^{p})|\mathcal{T}_{A}\}$ to train $\mathcal{M}_{l}(\mathcal{S}, \mathcal{T}_{A})$ for *E*.
- Stage 2.2: Repeat lines 2-4 based on the last model and update \$\mathcal{Y}_T\$, \$\mathcal{Y}_T^p\$, \$\mathcal{T}_A\$;
- 8: Unfreeze the $w_{\mathcal{M}_l}$ of $\mathcal{M}_l(\mathcal{S}, \mathcal{T}_A)$;
- 9: Use $\{\mathcal{Y}_{\mathcal{S}}|\mathcal{S}, (\mathcal{Y}_{\mathcal{T}}^{t}, \mathcal{Y}_{\mathcal{T}}^{p})|\mathcal{T}_{A}\}$ to train $\mathcal{M}_{l}(\mathcal{S}, \mathcal{T}_{A})$ for *E*.
- 10: **Stage 3:** Repeat lines 2-4 based on the last model and update $\mathcal{Y}_{\tau}^{t}, \mathcal{Y}_{\tau}^{p}, \mathcal{T}_{A}$;
- 11: Freeze $w_{\mathcal{M}_s}$ except for BatchNorm Layers;
- 12: Use $\{\mathcal{Y}_{\mathcal{S}}|\mathcal{S}, (\mathcal{Y}_{\mathcal{T}}^{t}, \mathcal{Y}_{\mathcal{T}}^{p})|\mathcal{T}_{A}\}$ to train $\mathcal{M}_{s}(\mathcal{S}, \mathcal{T}_{A})$ for E.
- 13: return $\mathcal{M}_s(\mathcal{S}, \mathcal{T}_A)$

Stage 2: The source large model $\mathcal{M}_l(S)$ generates initial pseudo labels \mathcal{Y}_T^p of the unlabeled target data. To improve the accuracy of pseudo labels, we use the prior-based curriculum learning module (Section IV-B) to correct them. Moreover, the masked copy-paste module (Section IV-C) is adopted to reduce the pseudo-label noises by generating augmented target images \mathcal{T}_A . The adaptive model $\mathcal{M}_l(S, \mathcal{T}_A)$ is obtained by training the source data and target data. To make the best of the large model, we retrain the model $\mathcal{M}_l(S, \mathcal{T}_A)$ using a new set of data with updated pseudo-labels which are more precise than the last period. Stage 2 contains two periods. The weight $w_{\mathcal{M}_l}$ of the model is frozen except for BatchNorm Layers at the first training period. In this way, the adaptive large model can perform well on the target domain.

Stage 3: The adaptive large model generates more precise pseudo labels to train the small model. To transfer the knowl-

edge from the large model to the small model, we still follow the same training steps to train the lightweight model. The small model $\mathcal{M}_s(S, \mathcal{T}_A)$ is eventually obtained.

The implementation procedure is shown in Algorithm 1. Stage 2.1, Stage 2.2, and Stage 3 contain the same data preparation and training procedure. The training framework of these three periods is illustrated in Fig. 4. The training loss consists of a supervised loss from the labeled source data and an unsupervised loss from the unlabeled target data. The prior-guided curriculum learning module is utilized to generate more accurate pseudo labels. The masked copy-paste module provides a set of true labels to reduce the proportion of pseudo-labels in training loss. Instead of updating the pseudo-labels and augmented images at each training epoch, we choose to update data only once at the initialization of each period. The detection model can stably learn features of true labels in this way.

B. Prior-Guided Curriculum Learning

In this section, we propose a Prior-guided Curriculum Learning module (PCL) to improve the accuracy of pseudo labels. Since the prediction confidences of simple examples are much higher than that of hard examples, PCL corrects pseudo labels by assigning different difficulties of examples adaptive thresholds instead of a constant threshold. First, we use prior knowledge of MAV detection to partition pseudo labels of the target dataset into subsets of different difficulties. Second, we allocate adaptive thresholds to pseudo labels in different subsets. The details are as follows.

1) Dataset Partition: How to evaluate the detection difficulty of targets is the key issue for curriculum learning. We propose a *difficulty matrix* to cluster different categories of a MAV dataset. Many factors influence the detection of MAVs, such as the target size, local contrast, and background complexity. The calculation of the difficulty matrix is shown in the following.

Target size. The target size influences the MAV detection performance especially when MAVs are extremely small objects in images. The target area is a bounding box with a





Fig. 5. Schematic diagram of the target area A_t and local background area A_t . The yellow area and the blue area represent A_t and A_b , respectively.

height h_t and a width w_t . We select the target size m_{ts} as a difficulty metric:

$$m_{\rm ts} = h_{\rm t} w_{\rm t}.\tag{1}$$

Local contrast. The contrast between the local background and the target area is also another factor that influences the detection performance. The schematic diagram of the target area and local background area is shown in Fig. 5, where A_b and A_t represent the total pixels in the local background area and target area, respectively. We define N_b as the number of pixels in A_b , which can be represented as

$$N_{\rm b} = h_{\rm b} w_{\rm b} - h_{\rm t} w_{\rm t},\tag{2}$$

where h_b and w_b are the height and width of the local background area. We select the local contrast m_{lc} between A_t and A_b as another metric

$$m_{\rm lc} = \sqrt{\frac{1}{N_{\rm b}} \sum_{(i,j)\in A_{\rm b}} (I_{\rm b}^{(i,j)} - \overline{I}_{\rm t})^2},$$
 (3)

where \overline{I}_t denotes the average intensity of the grayscale target area, and $I_b^{(i,j)}$ represents the pixel I(i, j) in the local background area.

Background complexity. The scenes of MAVs can be diverse, such as skies, buildings, and mountains. The background complexity m_{bc} has a great influence on the detection results. m_{bc} is defined as

$$m_{\rm bc} = \sqrt{\frac{1}{N_{\rm b}} \sum_{(i,j) \in A_{\rm b}} (I_{\rm b}^{(i,j)} - \overline{I}_{\rm b})^2},\tag{4}$$

where \overline{I}_{b} represents the average intensity of the local background area.

We use the difficulty matrix to classify the MAV targets into 4 types: Small Targets (c_{st}), Low-Contrast images (c_{lc}), Complex Backgrounds (c_{cb}), and Simple Examples (c_{se}). The classification criteria are as follows:

$$c_{t} = \begin{cases} c_{\text{st}} &, m_{\text{ts}} \leq \tau_{ts}, \\ c_{\text{lc}} &, m_{\text{ts}} > \tau_{ts}, m_{\text{lc}} \leq \tau_{lc}, \\ c_{\text{se}} &, m_{\text{ts}} > \tau_{ts}, m_{\text{lc}} > \tau_{lc}, m_{\text{bc}} \leq \tau_{\text{bc}}, \\ c_{\text{cb}} &, m_{\text{ts}} > \tau_{ts}, m_{\text{lc}} > \tau_{lc}, m_{\text{bc}} > \tau_{\text{bc}}, \end{cases}$$
(5)

where c_t represents the difficulty category of a target, $\tau_{\rm ts}$, $\tau_{\rm lc}$ and $\tau_{\rm bc}$ are the thresholds for difficulty matrix $m_{\rm ts}$, $m_{\rm lc}$ and $m_{\rm bc}$, respectively. When the target has a small area, the target type is to be $c_{\rm st}$ without the consideration of the other two indexes. Thus there are four types of targets instead of eight types. The thresholds $\tau_{\rm ts}$, $\tau_{\rm lc}$ and $\tau_{\rm bc}$ are set to 16×16 , 10, and 10, respectively. Some examples of the classification results are shown in Fig. 6. The MAVs in the category of small targets are more ambitious than other examples because they are extremely small objects in images. Compared to other categories, the examples in the low-contrast category are not easily discernible from the local backgrounds. The category of simple targets is the easiest to be detected. The samples have the simplest backgrounds and the highest contrast between the foregrounds and the backgrounds. The backgrounds of examples in the complex background category are more complex than other categories, leading to the detection difficulty.

2) Adaptive Threshold Adjusting: Pseudo-labels of different categories should be assigned adaptive thresholds instead of constant thresholds. Inspired by the work FlexMatch [71] which is used for semi-supervised image classification, we propose an adaptive threshold adjusting algorithm to tackle the cross-domain MAV detection issue. A set of candidate pseudo labels with confidence threshold τ_{min} are separated into different difficulty categories. The relative difficulty $\sigma_t(c)$ for each category is defined as

$$\sigma_t(c) = \frac{1}{N_c} \sum_{n=1}^{N_c} \mathbb{1}(p_{y_n} > \tau_{\max}),$$
(6)

where p_{y_n} is the output probability, N_c represents the number of pseudo labels in category c and τ_{max} is the maximum confidence threshold. The adaptive threshold $\tau_t(c)$ is defined as

$$\tau_t(c) = \max\left\{\frac{\sigma_t(c)}{\max \sigma_t}\tau_{\max}, \tau_{\min}\right\}.$$
 (7)

After obtaining the threshold for each category, the pseudo labels will be corrected for training.

C. Masked Copy-Paste Augmentation

In this section, we propose a module called Masked Copy-paste Augmentation (MCA) to generate true labels on target images and reduce the noise caused by pseudo-labeling. Traditional copy-paste algorithms [72], [73] need to use segmentation masks as one of the inputs and mainly focus on supervised object detection. There have not been copy-pastebased methods for unsupervised cross-domain object detection yet. Our masked copy-paste algorithm only requires bounding boxes of the targets instead of segmentation masks.

Fig. 7 illustrates the masked copy-paste augmentation algorithm. Since MAV targets are salient in the local areas, we use saliency segmentation [74] to obtain the segmentation masks of cropped MAV images. To guarantee the style consistency of MAV cropped images and target images, we use Poisson Merging [75] to merge two images and generate harmonious augmented target images. Then the MAVs pasted on the target images have true labels. Instead of utilizing the cropped images in the source dataset, we collect a set of MAV images from the Internet, cause MAVs in these images are more salient to generate precise segmentation masks. The reason that we adopted the additional dataset instead of the source dataset is that the precise masks can help the network adapt to the target domain better. This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

ZHANG et al.: DOMAIN ADAPTIVE DETECTION OF MAVs: A BENCHMARK AND NOISE SUPPRESSION NETWORK



(d) Complex Background

Fig. 6. The four categories generated by the difficulty matrix from our M3D-Real dataset. Here, we use the grayscale image for the calculation of the four categories. Since the area of a small target is only 16×16 , the enlarged samples are not clear in the figure.



Fig. 7. Masked copy-paste augmentation algorithm.

The pipeline of the masked copy-paste augmentation module is shown in Algorithm 2. The cropped MAV images with ground truth are randomly pasted on the unlabeled target images. Then the true labels can guide the detectors to identify MAVs from backgrounds. Besides, the true labels keep the same sizes as the pseudo labels. Thus the augmentation algorithm can also help the detectors adapt to the size distributions of MAVs in the target domain.

As shown in Fig. 4, the training loss \mathcal{L} contains a supervised loss \mathcal{L}_s and an unsupervised loss \mathcal{L}_u . After data augmentation, the labels of target data contain not only the pseudo labels but also the true labels. Then the training loss contains a new supervised loss \mathcal{L}_t from truly-labeled MAVs on target images, which can be represented as

$$\mathcal{L} = \mathcal{L}_s + \alpha \mathcal{L}_u + \beta \mathcal{L}_t, \tag{8}$$

Algorithm 2 Masked Copy-Paste Augmentation Algorithm

9

Input: MAV cropped images \mathcal{I} and bounding boxes, Target images \mathcal{T} with pseudo labels $\mathcal{Y}^p_{\mathcal{T}}$

Parameter: Paste times J, the number N_T of target images with pseudo labels

Output: Augmented target images \mathcal{T}_A with true labels $\mathcal{Y}_{\mathcal{T}}^t$ and pseudo labels $\mathcal{Y}_{\mathcal{T}}^p$

- 1: Segment MAV cropped masks of cropped images.
- 2: for each $n \in [1, N_T]$ do
- 3: Randomly choose one pseudo label in $y_{\mathcal{T}_{r}}^{p}$;
- 4: Randomly choose J MAV cropped images;
- 5: for each $j \in [1, J]$ do
- 6: Keep the size of I_j consistent with $y_{\mathcal{I}_n}^p$;
- 7: Randomly choose the center of I_i in the image \mathcal{T}_n ;
- 8: Paste resized I_j with its mask on \mathcal{T}_n ;
- 9: Calculate the new bounding box y_i of I_i ;
- 10: Add the true label y_j into $y_{T_r}^t$;
- 11: end for
- 12: Obtain augmented image \mathcal{T}_{A_n} and its true labels $y_{\mathcal{T}_n}^t$;
- 13: end for 14: return augmented target images \mathcal{T}_A and true labels $y_{\mathcal{T}}^t$.

where α and β are trade-off parameters. By doing this, the masked copy-paste augmentation module can significantly eliminate the pseudo-label noises and help the model learn the features between MAVs and backgrounds.

V. EXPERIMENTS

A. Experimental Setup

To verify the effectiveness of our method, the experiments are conducted on three types of domain adaptation tasks: simulation-to-real adaptation, cross-scene adaptation, and cross-camera adaptation, respectively. The comparison of these three tasks is shown in Table IV. Benefiting from the diversity of data, our M3D dataset is utilized to build three types of tasks. All the target domains only contain realistic images to meet the requirements for real-world applications.

TABLE IV THE COMPARISON OF DIFFERENT TASKS IN OUR BENCHMARK

Scenarios	Datasets	Image Numbers
Simulation vs Real	M3D-Sim-Normal vs M3D-Real-Target	18,476 vs 15,326 images
Different scenes	M3D-Real-Source vs M3D-Real-Target	39,933 vs 15,326 images
Different cameras	M3D-Real vs Drone-vs-Bird	55,259 vs 19,027 images

1) Simulation-to-Real Adaptation Benchmark: The M3D-Sim-Normal subset and M3D-Real-Target subset are chosen to study the simulation-to-real adaptation problem. The huge style and scene discrepancies between the simulation and the realistic domain can effectively validate the adaptation methods. Both the source dataset and the target dataset contain a large amount of data and challenging scenes.

2) Cross-Scene Adaptation Benchmark: In this benchmark, we choose the M3D-Real-Source subset and the M3D-Real-Target subset as the source and the target domain. The source images and the target images are captured by the same camera. The only difference between these two sets is the backgrounds of MAVs. Thus this task can help to verify the cross-scene adaptation without the perturbation of the devices.

3) Cross-Camera Adaptation Benchmark: Drone-vs-Bird dataset [21] mainly focuses on small MAV detection but with a fixed camera. Due to the difference between the M3D dataset and the Drone-vs-Bird dataset, we choose these two datasets to study the cross-camera adaptation (M3D-Real \rightarrow Drone-vs-Bird) problem. The videos in this dataset are extracted to a set of images every 5 frames and randomly separated into the training set, validation set, and testing set, which contain 13,318, 1,902, and 3,807 images, respectively.

These three tasks are challenging and worthy of study.

B. Implementation Details

Following the work [19], [43], we choose YOLOv5 as the backbone model. We choose the YOLOv5x as the large model and the YOLOv5s as the small model. The images are resized to 640×640 before training. The learning rate for the source model is set at 0.01. During the whole adaptive training procedure, the learning rate is set to 0.002. Both the large and small source models are trained for 50 epochs with a batch size of 32 on 4 GPUs and 1 GPU, respectively. During the large-to-small model training process, the detection model is trained for 30 epochs at each training period. Each target image is pasted with truly-labeled MAVs 3 times (J = 3). The setting is similar to the result reported in [70]. The width w_b and the height h_b of the local background area are chosen to be 1.5 times the width w_t and the height h_t of the target area.

We use mAP(%) as the main evaluation metric when IOU = 0.5. We also compare the detection results with the Source Only way (trained on source domain) and the Oracle way (trained on target domain) to evaluate the improved performance of domain adaptation algorithms. Thus we adopt the Adaptation Gain ρ in [19] as the quantitative evaluation index. The calculation of ρ is represented as

$$\rho = \frac{\text{mAP}(\theta^a) - \text{mAP}(\theta^s)}{\text{mAP}(\theta^o) - \text{mAP}(\theta^s)} \times 100\%, \tag{9}$$

where θ^a , θ^s , and θ^o represent the adaptation model, Source Only model, and the Oracle model, respectively. The Source Only method means that the detection model is only trained on the source dataset. Oracle method represents the model trained on the target dataset.

C. Main Results

The main results of the three adaptation tasks are shown as follows. NSN *w/o teacher* represents the detection results of our NSN method without the large model.

Adaptation: 1) Simulation-to-Real We report the simulation-to-real adaptation results in Table V. Benefiting from the diversity of our simulation dataset, we can observe that the Source Only model achieves 39.9% mAP, which is much higher than the Source Only model trained by the M3D-Real-Source subset. However, the baseline method SimROD only improves the detection performance by 1.2% due to the huge domain gap between the simulation and real datasets. In particular, our NSN model exceeds SimROD by 5.8% on the task of simulation-to-real adaptation. This demonstrates the effectiveness of our prior-guided curriculum learning and masked copy-paste augmentation modules.

From Table V, we can find that most state-of-the-art methods perform even poorer than the Source Only method. This phenomenon does not happen in the other two adaptation tasks. The reason for degradation is as follows.

First, the self-training-based domain adaptation methods are easily affected by the pseudo-label noises. The inaccuracy of the pseudo labels will increase when meeting smaller MAVs under more diverse backgrounds. Due to the large gap between the simulation domain and the reality domain, the noises of the pseudo labels are more challenging in this scenario. If the noises exceed a certain level, the model will get worse and worse as the training continues. Therefore, the influence of noises makes the detection networks confused and causes poor detection performances. As shown in Table V, many methods obtain poorer performances than the Source Only method. Second, the SimROD method tries to solve this problem by introducing a teacher-student model to take advantage of the large model to reduce the noise. Thus, the SimROD method performs better than the SimROD w/o teacher method. However, the improvement ability is still limited because the pseudo-label noises are more complex when the targets are MAVs. Finally, the reason that we designed the noise suppression network is to solve this problem. We tackle this problem from two approaches. On the one hand, we use the masked copy-paste augmentation module to help the detection model fully adapt to the backgrounds in the target domain. On the other hand, we use the prior-guided curriculum learning module to improve the precision of the pseudo labels.

ZHANG et al.: DOMAIN ADAPTIVE DETECTION OF MAVs: A BENCHMARK AND NOISE SUPPRESSION NETWORK

TABLE V The simulation-to-real Adaptation Performance of Different Methods

Methods	Backbone	mAP (%)	ρ
Source Only	Yolov5s	39.9	0%
CSL [63]	Yolov5s	24.6	-30.9%
AsyFOD [53]	Yolov5x	35.4	-9.1%
AcroFOD [60]	Yolov5x	21.9	-36.3%
ConfMix [43]	Yolov5s	37.4	-5.0%
SimROD [19]	Yolov5s	41.1	2.4%
SimROD w/o teacher	Yolov5s	38.5	-2.8%
NSN w/o teacher	Yolov5s	41.2	2.6%
NSN (Ours)	Yolov5s	46.9	14.1%
Oracle	Yolov5s	89.5	100%

TABLE VI THE CROSS-SCENE ADAPTATION PERFORMANCE OF DIFFERENT METHODS

Methods	Backbone	mAP (%)	ρ
Source Only	Yolov5s	28.7	0%
CSL [63]	Yolov5s	34.4	9.3%
AsyFOD [53]	Yolov5x	47.9	31.6%
AcroFOD [60]	Yolov5x	30.4	2.8%
ConfMix [43]	Yolov5s	32.2	5.8%
SimROD [19]	Yolov5s	46.8	29.8%
SimROD w/o teacher	Yolov5s	39.1	17.1%
NSN w/o teacher	Yolov5s	43.1	23.7%
NSN (Ours)	Yolov5s	50.5	35.9%
Oracle	Yolov5s	89.5	100%

2) Cross-Scene Adaptation: The cross-scene adaptation results are shown in Table VI. The M3D-Real dataset contains plenty of scenes and is captured by a moving camera. Even if there is no cross-device problem to be solved, the cross-scene adaptation of the M3D dataset is still a challenging issue. The Source Only model only achieves 28.7% mAP on the target dataset, which indicates that the MAV detection model suffers from significant domain discrepancy. The detection improvement of our method is 21.8%. Among all the comparison methods, AsyFOD [53] has the best performance and achieves 47.9% mAP on this task, which is still lower than our proposed method. Our NSN method exceeds the baseline algorithm of 3.7% on the task of cross-scene adaptation. Furthermore, the simulation-to-real adaptation and cross-scene adaptation tasks contain the same target dataset M3D-Real-Target, thus we compare the detection results of these two tasks. We can observe that even though the Source Only model of cross-scene adaptation is lower, the final adaptation result and adaptation gain are much higher than the simulation-toreal adaptation. The reason is that the pseudo-label noises are more influenced by the huge style differences in the simulation-to-real task and make the detection model more confused. Our method can perform better on the simulationto-real task when combined with adversarial training modules. In this paper, we pay more attention to the real-to-real issue.

3) Cross-Camera Adaptation: There are huge differences in the target size and backgrounds of images between the M3D-Real dataset and the Drone-vs-Bird dataset. The cross-camera

TABLE VII THE CROSS-CAMERA ADAPTATION PERFORMANCE OF DIFFERENT METHODS

Methods	Backbone	mAP (%)	ρ
Source Only	Yolov5s	31.9	0%
CSL [63]	Yolov5s	33.9	3.5%
AsyFOD [53]	Yolov5x	39.8	13.8%
AcroFOD [60]	Yolov5x	35.7	6.6%
ConfMix [43]	Yolov5s	25.8	-10.7%
SimROD [19]	Yolov5s	50.2	32.0%
SimROD w/o teacher	Yolov5s	36.5	8.0%
NSN w/o teacher	Yolov5s	42.6	18.7%
NSN (Ours)	Yolov5s	61.5	51.8%
Oracle	Yolov5s	89.1	100%

adaptation performance of different methods is shown in Table VII. Thus the detection result is low when only using the source data, which is only 31.9%. Compared to the SimROD algorithm, our method achieves a huge improvement of 11.3% on the task of cross-camera adaptation. The performance of the small model is significantly improved by 18.9% with the help of the large model. Our method achieves higher adaptation gain than other tasks because some images in the Drone-vs-Bird dataset are shot under a still camera, thus the detection model can adapt to the new environments faster and perform better on the target domain.

After comparing the results of these three tasks, we can conclude that the pseudo-labeling technique is more suitable for dealing with scene-to-scene and real-to-real adaptation issues. Our method can significantly reduce the pseudo-label noises and increase the adaptation performance.

D. Qualitative Analysis

The qualitative comparison of different methods is shown in Fig 8. The detection results on target testing examples in the Drone-vs-Bird dataset are presented in the figure. The comparison consists of three methods, which are Source Only, SimROD, and NSN (Ours) methods. When only the Source Only method is adopted, the detection model cannot detect MAVs and wrongly classify other objects as MAVs. Even though the baseline method SimROD detects some MAVs in the images, there are still some tiny MAVs not detected. Our method significantly improves the detection performance of missed targets with the help of true labels generated by our masked copy-paste augmentation algorithm. Furthermore, our method can also help to reduce the wrong detection rate benefit of the PCL module.

E. Ablation Study

1) Ablation Study of Different Modules: The ablation study is mainly conducted on the task of cross-camera adaptation. Table VIII shows a comprehensive comparison of the contributions of different modules in our NSN method. First, with the participation of the large model, the improvement of the adaptive small detection model is further increased from 10.7% to 29.6%. Second, the MCA module improves

IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING



(a) Source Only

(b) SimROD

(c) NSN (Ours)

Fig. 8. The qualitative comparison of different methods on the cross-camera adaptation task.

TABLE VI	Ш	
----------	---	--

Ablation Study on the Task of Cross-Camera Adaptation. The LM, PCL, and MCA Represent the Large Model, Prior-Guided Curriculum Learning Module, and Masked Copy-Paste Augmentation Module, Respectively

Method	LM	PCL	MCA	mAP (%)	ρ
Source Only	-	-	-	31.9	0%
		✓		37.5	9.8%
			\checkmark	41.1	16.1%
Proposed		\checkmark	\checkmark	42.6	18.7%
	\checkmark			50.2	32.0%
	\checkmark		\checkmark	60.5	50.0%
Ours	\checkmark	\checkmark	\checkmark	61.5	51.8%

the performance of the small model by 9.2%. After the small model owns the knowledge of the large model, the MCA module still contributes the improvement of 10.3%. Finally, the PCL module is also indispensable. For example, if we skip this module, the performance of the purely small model drops by 1.5%, and our NSN method drops by 1.0%. The integration of these modules is more efficient than individual modules.

2) The Choice of Cropped Images for Data Augmentation: In the masked copy-paste augmentation module, we use an additional saliency dataset with 2,113 images as cropped images, which can be downloaded from our GitHub. Some examples are shown in Fig. 7. Compared with the examples in the source dataset (see Fig. 6), the MAVs in the saliency dataset are bigger and more salient enough for saliency segmentation. We conduct two experiments to show the design of this module.

First, we compare our method using the additional salient images with the cropped images whose widths exceed 50 pixels from the source dataset. To reduce the influence of other modules, we only adopt the MCA module for comparison. The detection results are shown in Table IX. When we use the cropped images directly from the source dataset, the adaptation achieves 40.6% mAP, which is smaller than the training with the saliency dataset. The reason is that the segmentation masks of small objects are not easily obtained automatically from the saliency segmentation algorithm. If we use all the images from the source dataset, the detection results will drop deeply due

TABLE IX THE ABLATION STUDY OF THE MASKED COPY-PASTE MODULE

Methods	Backbone	mAP (%)	ρ
Source Only	Yolov5s	31.9	0%
MCA (Saliency)	Yolov5s	41.1	16.1%
MCA (Source, $w_t \ge 50$ pixels)	Yolov5s	40.6	15.2%
SimROD w/o teacher	Yolov5s	36.5	8.0%
SimROD w/o teacher + Saliency	Yolov5s	35.5	6.3%
Oracle	Yolov5s	89.1	100%

to the inaccurate masks. In conclusion, when the MAVs in the source dataset are not big and salient enough for saliency segmentation, it is suggested to choose the additional dataset for utilization. Our proposed masked copy-paste module could work well whether the cropped images come from the source dataset or an additional dataset as long as we can obtain the precise masks of the targets.

Second, to guarantee fairness, we also show the experimental results of SimROD when expanding the source dataset with the additional salient dataset during training. As shown in Table IX, the detection performance drops to 35.5% when more salient images are added into the training period. The reason is that the MAVs in the salient images take up a large area. The distribution of the saliency dataset is different from the source dataset and the target dataset. The additional salient dataset can only play a significant role in our masked copypaste module.

3) Varying of Different Hyper-Parameters: The proposed modules contain a few hyper-parameters that could influence the detection results. We conduct two experiments to show the effect of different values of the J and τ_{max} . To ensure that the results are not disturbed by the participation of other modules, the experiments are independently conducted on the modules themselves. The detection results are shown in Table X.

When we change the times of the copy-paste module, the detection model achieves the highest result when J equals three. This phenomenon is consistent with the experiment results reported in [73]. The reason is that when the copy-paste time is too high, the detection is easily overfitting. We also conduct experiments of varying the value of τ_{max} . The detection results in Table X show that the PCL module can achieve

ZHANG et al.: DOMAIN ADAPTIVE DETECTION OF MAVs: A BENCHMARK AND NOISE SUPPRESSION NETWORK

TABLE X THE EXPERIMENTS OF VARYING THE VALUES OF DIFFERENT HYPER-PARAMETERS ON THE TASK OF CROSS-CAMERA ADAPTATION

Method	Hyper-parameters	mAP (%)	ρ
Source Only	-	31.9	0%
MCA	J = 1	37.4	9.6%
MCA	J = 3	41.1	16.1%
MCA	J = 5	39.8	13.8%
PCL	$\tau_{max} = 0.6$	37.2	9.3%
PCL	$\tau_{max} = 0.75$	37.5	9.8%
PCL	$\tau_{max} = 0.9$	36.7	8.4%
Oracle	-	89.1	100%



Fig. 9. The confidence distribution of pseudo labels from different categories during the training process.

the best performance when τ_{max} equals 0.75. As for another hyper-parameter τ_{min} , we follow the default setting 0.25 of the Yolov5 network for stable performances in different scenarios.

F. Experimental Analysis

1) The Influence of Adaptive Thresholds: In this section, we study the influence of adaptive thresholds on pseudo labels. Furthermore, we use Fig. 9 to prove the effectiveness of the PCL module and how it influences the choice of pseudo labels at different training stages. The confidences of pseudo labels from different difficulty categories show different distributions. The simple category tends to have pseudo labels of high confidence. As the progress of training, the confidences of pseudo labels become larger and all the distributions move towards the right side of the figure. The adaptive thresholds generated by our PCL module also gradually increase to choose more reliable pseudo labels.

2) The Effectiveness of the Data Augmentation: To show the effectiveness of our data augmentation module, we also compare our method with other data augmentation methods. To reduce the influence of other modules in our method, we only adopt the MCA module in the experiment. The



(a) CutMix

(b) MCA (Ours)

Fig. 10. The augmentation results of CutMix and MCA (Ours) on Drone-vs-Bird dataset. Original images only contain one MAV which is enclosed within a blue box.

TABLE XI THE CROSS-CAMERA ADAPTATION PERFORMANCE OF DIFFERENT AUGMENTATION METHODS

Methods	Backbone	mAP (%)	ρ
Source Only	Yolov5s	31.9	0%
MCA (Ours)	Yolov5s	41.1	16.1%
ConfMix [43]	Yolov5s	25.8	-10.7%
CutMix [76]	Yolov5s	33.0	1.9%
Oracle	Yolov5s	89.1	100%



Fig. 11. The images generated by texture randomization in M3D-Sim subset. The walls in Unreal Engine dynamically change their textures during data collection.

results are shown in Table XI. When we replace our MCA module with CutMix, the detection result drops from 41.1% to 33.0%, proving the significance of our data augmentation module. The reason is that the local background doesn't change when pasting the cropped MAV images on the target images directly. The differences between CutMix and MCA are shown in Fig. 10. Our augmentation module can generate more harmonious results. The detection model could not learn the local background features of target images when using CutMix augmentation.

3) Texture Randomization: M3D-Sim subset contains normal images and images generated by texture randomization. The visual differences between these two types of images can be seen in Fig. 2 and Fig. 11. Texture randomization is one way to reduce the gap between simulation and reality, thus we

IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

TABLE XII THE DETECTION PERFORMANCES TRAINED ON DIFFERENT SUBSETS IN THE M3D-SIM DATASET

Dataset	Backbone	mAP (%)
M3D-Sim-Normal	Yolov5s	39.9
M3D-Sim-Texture	Yolov5s	22.2
M3D-Sim	Yolov5s	43.5

TABLE XIII

THE DETECTION PERFORMANCES TRAINED ON DIFFERENT TRAINING SUBSETS OF THE M3D DATASET. THE M3D-REAL-S AND M3D-REAL-T PERFORMENT THE M3D REAL SOURCE AND M3D REAL TARCET

I KEPKESEN	I THE MOD-KEA	T-200	JRCE AND MIS	D-KEAL-I	ARGET
SUBSETS,	RESPECTIVELY.	The	DETECTION	RESULTS	Are
	REPORTED ON 7	гне Ті	ESTING SUBSE	ETS	

Training Subset	M3D-Real-S	M3D-Real-T	M3D-Real
M3D-Sim-Normal	53.5%	39.9%	50.6%
M3D-Sim-Texture	38.4%	22.2%	34.6%
M3D-Real-Source	86.2%	28.7%	73.4%
M3D-Real-Target	46.1%	89.1%	56.1%
M3D-Sim	57.0%	43.5%	54.2%
M3D-Real	88.2%	90.0%	88.6%
M3D	88.2%	89.0%	88.3%

also conduct experiments to verify the effectiveness of this approach. The results are shown in Table XII. We train the detection model by only using the normal data, randomized data, and all the data, respectively. The model is tested on the M3D-Real-Target dataset. The results of these three approaches are 39.9%, 22.2%, and 43.5%, respectively. It can be concluded that texture randomization can indeed improve detection performance in a supervised way, but its ability is limited. Besides, this type of method can improve the generalization ability of the detection model but cannot help the model adapt to new environments.

4) Computational Cost: The computational cost can be divided into the training period and the inference period. First, the training time of our method takes a few hours more than the baseline SimROD method. Take the task of cross-camera adaptation as an example. The training of the SimROD method takes 45 hours and our method costs 47 hours. The additional two hours are cost by the MCA and the PCL module. Second, the baseline SimROD method and our NSN method contain the same detection model, thus the inference time is the same as the Yolov5s model.

G. The Universal MAV Detection Model Based on M3D Dataset

As a dataset containing diverse images, the M3D dataset can be chosen as the common dataset for MAV detection. This section shows the experiment results when using different training subsets of M3D for training. All the experiments still choose the Yolov5s network as the MAV detection model. The models are tested on the testing set of M3D-Real-Source, M3D-Real-Target, and M3D-Real, respectively. When we use the whole training set of the M3D dataset as the training dataset, the detection model achieves 88.3% mAP on the testing set of M3D-Real. The results can serve as the precision of the Oracle model for study in the future. MAV detection is critical for real-world applications. Thus the results can also be utilized as the universal MAV detection model for the community.

VI. CONCLUSION

This paper benchmarks the cross-domain MAV detection problem. We first propose a Multi-MAV-Multi-Domain (M3D) dataset and construct a novel domain adaptive MAV detection benchmark consisting of three representative domain adaptation tasks, i.e., simulation-to-real adaptation, cross-scene adaptation, and cross-camera adaptation. Moreover, we propose a novel noise suppression network with a prior-guided curriculum learning module, a masked copy-paste augmentation module, and a large-to-small model training procedure. Our method significantly reduces the pseudo noises and achieves real-time MAV detection. The results of extensive experiments and ablation studies demonstrate the effectiveness of the proposed method. We have released the dataset, hoping our study could advance the cross-domain MAV detection research.

Due to the large domain shift of MAV detection, there is still a huge improving space to achieve the performance of supervised learning. We will consider the spatial and temporal information in future work.

REFERENCES

- J. Li, D. H. Ye, M. Kolsch, J. P. Wachs, and C. A. Bouman, "Fast and robust UAV to UAV detection and tracking from video," *IEEE Trans. Emerg. Topics Comput.*, vol. 10, no. 3, pp. 1519–1531, Jul. 2022.
- [2] C. Wang, J. Tian, J. Cao, and X. Wang, "Deep learning-based UAV detection in pulse-Doppler radar," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5105612.
- [3] B. K. S. Isaac-Medina, M. Poyser, D. Organisciak, C. G. Willcocks, T. P. Breckon, and H. P. H. Shum, "Unmanned aerial vehicle visual detection and tracking using deep neural networks: A performance benchmark," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshops* (*ICCVW*), Oct. 2021, pp. 1223–1232.
- [4] M. Pavliv, F. Schiano, C. Reardon, D. Floreano, and G. Loianno, "Tracking and relative localization of drone swarms with a visionbased headset," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 1455–1462, Apr. 2021.
- [5] Y. Zheng, C. Zheng, X. Zhang, F. Chen, Z. Chen, and S. Zhao, "Detection, localization, and tracking of multiple MAVs with panoramic stereo camera networks," *IEEE Trans. Autom. Sci. Eng.*, vol. 20, no. 2, pp. 1226–1243, Apr. 2023.
- [6] S. Jamil, M. Rahman, A. Ullah, S. Badnava, M. Forsat, and S. S. Mirjavadi, "Malicious UAV detection using integrated audio and visual features for public safety applications," *Sensors*, vol. 20, no. 14, p. 3923, Jul. 2020.
- [7] Y. Chen, P. Aggarwal, J. Choi, and C.-C. J. Kuo, "A deep learning approach to drone monitoring," in *Proc. Asia–Pacific Signal Inf. Process. Assoc. Annu. Summit Conf. (APSIPA ASC)*, 2017, pp. 686–691.
- [8] J. Zhao, J. Zhang, D. Li, and D. Wang, "Vision-based anti-UAV detection and tracking," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 12, pp. 25323–25334, Dec. 2022.
- [9] H. Sun, J. Yang, J. Shen, D. Liang, L. Ning-Zhong, and H. Zhou, "TIB-Net: Drone detection network with tiny iterative backbone," *IEEE Access*, vol. 8, pp. 130697–130707, 2020.
- [10] Y. Zheng, Z. Chen, D. Lv, Z. Li, Z. Lan, and S. Zhao, "Air-to-air visual detection of micro-UAVs: An experimental evaluation of deep learning," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 1020–1027, Apr. 2021.
- [11] K. Tian, C. Zhang, Y. Wang, S. Xiang, and C. Pan, "Knowledge mining and transferring for domain adaptive object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9113–9122.

- [12] L. Zhao and L. Wang, "Task-specific inconsistency alignment for domain adaptive object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 14197–14206.
- [13] J. Ye, C. Fu, G. Zheng, D. P. Paudel, and G. Chen, "Unsupervised domain adaptation for nighttime aerial tracking," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 8896–8905.
- [14] Y.-J. Li et al., "Cross-domain adaptive teacher for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 7581–7590.
- [15] J. Deng, W. Li, Y. Chen, and L. Duan, "Unbiased mean teacher for cross-domain object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 4091–4101.
- [16] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. ICCV*, Oct. 2017, pp. 2223–2232.
- [17] Y. Chen, W. Li, C. Sakaridis, D. Dai, and L. Van Gool, "Domain adaptive faster R-CNN for object detection in the wild," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2018, pp. 3339–3348.
- [18] J. Deng, D. Xu, W. Li, and L. Duan, "Harmonious teacher for crossdomain object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2023, pp. 23829–23838.
- [19] R. Ramamonjison, A. Banitalebi-Dehkordi, X. Kang, X. Bai, and Y. Zhang, "SimROD: A simple adaptation method for robust object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 3550–3559.
- [20] N. Inoue, R. Furuta, T. Yamasaki, and K. Aizawa, "Cross-domain weakly-supervised object detection through progressive domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5001–5009.
- [21] A. Coluccia et al., "Drone vs. bird detection: Deep learning algorithms and results from a grand challenge," *Sensors*, vol. 21, no. 8, p. 2824, Apr. 2021.
- [22] N. Jiang et al., "Anti-UAV: A large-scale benchmark for vision-based UAV tracking," *IEEE Trans. Multimedia*, vol. 25, pp. 486–500, 2023.
- [23] A. Sanders, An Introduction to Unreal Engine 4. Boca Raton, FL, USA: CRC Press, 2016.
- [24] X. Lu, Q. Li, B. Li, and J. Yan, "MimicDet: Bridging the gap between one-stage and two-stage object detection," in *Computer Vision—ECCV*, Glasgow, U.K. Berlin, Germany: Springer, 2020, pp. 541–557.
- [25] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards realtime object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 1–9.
- [26] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," 2018, arXiv:1804.02767.
- [27] A. Bochkovskiy, C.-Y. Wang, and H.-Y. Mark Liao, "YOLOv4: Optimal speed and accuracy of object detection," 2020, arXiv:2004.10934.
- [28] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 2980–2988.
- [29] Z. Tian, C. Shen, H. Chen, and T. He, "FCOS: Fully convolutional one-stage object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.* (*ICCV*), Oct. 2019, pp. 9627–9636.
- [30] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "CenterNet: Keypoint triplets for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6569–6578.
- [31] J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object detection via regionbased fully convolutional networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 1–9.
- [32] Z. Cai and N. Vasconcelos, "Cascade R-CNN: Delving into high quality object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6154–6162.
- [33] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 213–229.
- [34] M. Vrba and M. Saska, "Marker-less micro aerial vehicle detection and localization using convolutional neural networks," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2459–2466, Apr. 2020.
- [35] A. Rozantsev, V. Lepetit, and P. Fua, "Detecting flying objects using a single moving camera," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 5, pp. 879–892, May 2017.
- [36] M. W. Ashraf, W. Sultani, and M. Shah, "Dogfight: Detecting drones from drones videos," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 7063–7072.

- [37] Y. Li, D. Yuan, M. Sun, H. Wang, X. Liu, and J. Liu, "A global-local tracking framework driven by both motion and appearance for infrared anti-UAV," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2023, pp. 3025–3034.
- [38] T. Nguyen et al., "MAVNet: An effective semantic segmentation micronetwork for mav-based tasks," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3908–3915, Oct. 2019.
- [39] A. Yavariabdi, H. Kusetogullari, T. Celik, and H. Cicek, "FastUAV-net: A multi-UAV detection algorithm for embedded platforms," *Electronics*, vol. 10, no. 6, p. 724, 2021.
- [40] Y. Jiao, H. Yao, and C. Xu, "SAN: Selective alignment network for cross-domain pedestrian detection," *IEEE Trans. Image Process.*, vol. 30, pp. 2155–2167, 2021.
- [41] M. Faraki, X. Yu, Y.-H. Tsai, Y. Suh, and M. Chandraker, "Cross-domain similarity learning for face recognition in unseen domains," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 15287–15296.
- [42] R. Tao et al., "Exploring endogenous shift for cross-domain detection: A large-scale benchmark and perturbation suppression network," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 21157–21167.
- [43] G. Mattolin, L. Zanella, E. Ricci, and Y. Wang, "ConfMix: Unsupervised domain adaptation for object detection via confidence-based mixing," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2023, pp. 423–433.
- [44] W. Li, X. Liu, X. Yao, and Y. Yuan, "SCAN: Cross domain object detection with semantic conditioned adaptation," in *Proc. AAAI*, vol. 36, 2022, pp. 1421–1428.
- [45] W. Li, X. Liu, and Y. Yuan, "SIGMA: Semantic-complete graph matching for domain adaptive object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 5281–5290.
- [46] P. Oza, V. A. Sindagi, V. V. Sharmini, and V. M. Patel, "Unsupervised domain adaptation of object detectors: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Mar. 17, 2023, doi: 10.1109/TPAMI.2022.3217046.
- [47] K. Saito, Y. Ushiku, T. Harada, and K. Saenko, "Strong-weak distribution alignment for adaptive object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 6956–6965.
- [48] Q. Cai, Y. Pan, C.-W. Ngo, X. Tian, L. Duan, and T. Yao, "Exploring object relation in mean teacher for cross-domain detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 11457–11466.
- [49] H. Wang, S. Liao, and L. Shao, "AFAN: Augmented feature alignment network for cross-domain object detection," *IEEE Trans. Image Process.*, vol. 30, pp. 4046–4056, 2021.
- [50] A. RoyChowdhury et al., "Automatic adaptation of object detectors to new domains using self-training," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 780–790.
- [51] M. Khodabandeh, A. Vahdat, M. Ranjbar, and W. Macready, "A robust learning approach to domain adaptive object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 480–490.
- [52] X. Liu, W. Li, Q. Yang, B. Li, and Y. Yuan, "Towards robust adaptive object detection under noisy annotations," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 14187–14196.
- [53] Y. Gao, K.-Y. Lin, J. Yan, Y. Wang, and W.-S. Zheng, "AsyFOD: An asymmetric adaptation paradigm for few-shot domain adaptive object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2023, pp. 3261–3271.
- [54] W. Zhou, D. Du, L. Zhang, T. Luo, and Y. Wu, "Multi-granularity alignment domain adaptation for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 9581–9590.
- [55] T. Kim, M. Jeong, S. Kim, S. Choi, and C. Kim, "Diversify and match: A domain adaptive representation learning paradigm for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.* (CVPR), Jun. 2019, pp. 12448–12457.
- [56] R. Gong, W. Li, Y. Chen, and L. Van Gool, "DLOW: Domain flow for adaptation and generalization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2477–2486.
- [57] S.-W. Huang, C.-T. Lin, S.-P. Chen, Y.-Y. Wu, P.-H. Hsu, and S.-H. Lai, "AugGAN: Cross domain adaptation with GAN-based data augmentation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 718–731.
- [58] M. Xu et al., "End-to-end semi-supervised object detection with soft teacher," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Jun. 2021, pp. 3060–3069.

16

- [59] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "Mixup: Beyond empirical risk minimization," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–13.
- [60] Y. Gao, L. Yang, Y. Huang, S. Xie, S. Li, and W.-S. Zheng, "AcroFOD: An adaptive method for cross-domain few-shot object detection," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2022, pp. 673–690.
- [61] M. He et al., "Cross domain object detection by target-perceived dual branch distillation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 9570–9580.
- [62] Y. Zhang, P. David, and B. Gong, "Curriculum domain adaptation for semantic segmentation of urban scenes," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2020–2030.
- [63] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, "Curriculum selfpaced learning for cross-domain object detection," *Comput. Vis. Image Understand.*, vol. 204, Mar. 2021, Art. no. 103166.
- [64] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, "Curriculum learning: A survey," Int. J. Comput. Vis., vol. 130, pp. 1526–1565, Jun. 2022.
- [65] J. Wang and X.-L. Zhang, "Improving pseudo labels with intra-class similarity for unsupervised domain adaptation," *Pattern Recognit.*, vol. 138, Jun. 2023, Art. no. 109379.
- [66] M. L. Pawelczyk and M. Wojtyra, "Real world object detection dataset for quadcopter unmanned aerial vehicle detection," *IEEE Access*, vol. 8, pp. 174394–174409, 2020.
- [67] V. Walter, M. Vrba, and M. Saska, "On training datasets for machine learning-based visual relative localization of micro-scale UAVs," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 10674–10680.
- [68] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," in *Proc. Field Service Robot., Results 11th Int. Conf.* New York, NY, USA: Springer, 2018, pp. 621–635.
- [69] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: From simulation to reality with domain randomization," *IEEE Trans. Robot.*, vol. 36, no. 1, pp. 1–14, Feb. 2020.
- [70] C. Rui, G. Youwei, Z. Huafei, and J. Hongyu, "A comprehensive approach for UAV small object detection with simulation-based transfer learning and adaptive fusion," 2021, arXiv:2109.01800.
- [71] B. Zhang et al., "FlexMatch: Boosting semi-supervised learning with curriculum pseudo labeling," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 18408–18419.
- [72] G. Ghiasi et al., "Simple copy-paste is a strong data augmentation method for instance segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2918–2928.
- [73] M. Kisantal, Z. Wojna, J. Murawski, J. Naruniec, and K. Cho, "Augmentation for small object detection," in *Proc. CS IT Conf.*, vol. 9, 2019, pp. 1–15.
- [74] X. Qin, Z. Zhang, C. Huang, M. Dehghan, O. R. Zaiane, and M. Jagersand, "U2-Net: Going deeper with nested U-structure for salient object detection," *Pattern Recognit.*, vol. 106, Oct. 2020, Art. no. 107404.
- [75] P. Pérez, M. Gangnet, and A. Blake, "Poisson image editing," ACM Trans. Graph., vol. 22, no. 3, pp. 313–318, Jul. 2003.
- [76] S. Yun, D. Han, and S. J. Oh, "CutMix: Regularization strategy to train strong classifiers with localizable features," in *Proc. ICCV*, Oct. 2019, pp. 6023–6032.



Yin Zhang received the B.Sc. degree in measurement and control technology and instrumentation from Tianjin University, Tianjin, China, in 2017, and the M.Sc. degree in instrument science and technology from BUAA University, Beijing, China, in 2020. She is currently pursuing the Ph.D. degree with the Intelligent Unmanned Systems Laboratory, Westlake University, Hangzhou, China. Her research interests include domain adaptation, MAV detection, and depth estimation.



Jinhong Deng received the B.Eng. degree from Southwest Jiaotong University, Chengdu, China, in 2019. He is currently pursuing the Ph.D. degree with the School of Computer Science and Engineering, University of Electronic Science and Technology of China (UESTC). His research interests include transfer learning, label-efficient learning, and their applications in computer vision.



Peidong Liu received the B.E. and M.E. degrees from the National University of Singapore, in 2012 and 2015, respectively, and the Ph.D. degree in computer science from ETH Zürich in 2021. He is currently an Assistant Professor with the School of Engineering, Westlake University, Hangzhou, China. His research interests lie in 3-D computer vision and robotics.



Wen Li (Member, IEEE) received the B.E. and M.S. degrees from Beijing Normal University in 2007 and 2010, respectively, and the Ph.D. degree from Nanyang Technological University in 2015. He is currently a Full Professor with the School of Computer Science and Engineering, University of Electronic Science and Technology of China. His main research interests include transfer learning, multi-view learning, multiple kernel learning, and their applications in computer vision.



Shiyu Zhao (Member, IEEE) received the B.Eng. and M.Eng. degrees from Beijing University of Aeronautics and Astronautics, Beijing, China, in 2006 and 2009, respectively, and the Ph.D. degree from the National University of Singapore, Singapore, in 2014, all in electrical engineering.

From 2014 to 2016, he was a Post-Doctoral Researcher with the Technion–Israel Institute of Technology, Haifa, Israel, and the University of California at Riverside, Riverside, CA, USA. From 2016 to 2018, he was a Lecturer with the

Department of Automatic Control and Systems Engineering, The University of Sheffield, Sheffield, U.K. He is currently an Associate Professor with the School of Engineering, Westlake University, Hangzhou, China. His research interests lie in sensing, estimation, and control of robotic systems.