

Meta-UDA: Unsupervised Domain Adaptive Thermal Object Detection using Meta-Learning

Vibashan VS¹, Domenick Poster², Suya You³, Shuowen Hu³,
Vishal M. Patel¹

¹ Johns Hopkins University, MD, USA, ² West Virginia University, WV, USA

³ U.S. Army CCDC Army Research Laboratory, USA

{vvishnu2,vpatel136}@jhu.edu, dposter@mix.wvu.edu, {suya.you.civ,shuowen.hu.civ}@mail.mil

Abstract

Object detectors trained on large-scale RGB datasets are being extensively employed in real-world applications. However, these RGB-trained models suffer a performance drop under adverse illumination and lighting conditions. Infrared (IR) cameras are robust under such conditions and can be helpful in real-world applications. Though thermal cameras are widely used for military applications and increasingly for commercial applications, there is a lack of robust algorithms to robustly exploit the thermal imagery due to the limited availability of labeled thermal data. In this work, we aim to enhance the object detection performance in the thermal domain by leveraging the labeled visible domain data in an Unsupervised Domain Adaptation (UDA) setting. We propose an algorithm agnostic meta-learning framework to improve existing UDA methods instead of proposing a new UDA strategy. We achieve this by meta-learning the initial condition of the detector, which facilitates the adaptation process with fine updates without overfitting or getting stuck at local optima. However, meta-learning the initial condition for the detection scenario is computationally heavy due to long and intractable computation graphs. Therefore, we propose an online meta-learning paradigm which performs online updates resulting in a short and tractable computation graph. To this end, we demonstrate the superiority of our method over many baselines in the UDA setting, producing a state-of-the-art thermal detector for the KAIST and DSIAC datasets.

1. Introduction

Object detection is a well-known problem in computer vision which has been actively researched for over two decades. With recent developments in deep Convolutional Neural Networks (CNNs) [25], CNN-based object detectors produce state-of-the-art performance in many benchmark

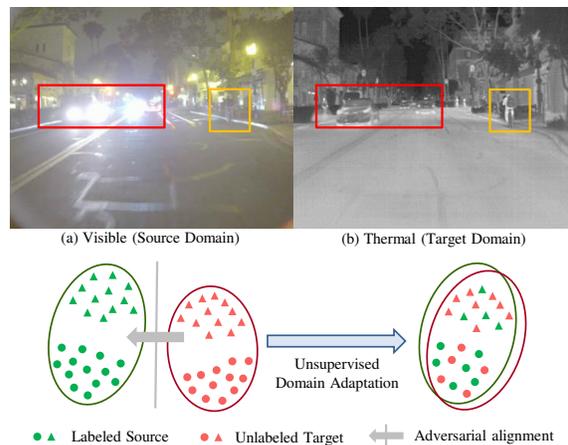


Figure 1. In the top row, a comparison between the visible and thermal domains is illustrated. The red box corresponds to the region with adverse illumination and the yellow box corresponds to the region with low lighting conditions. The top row comparison shows that thermal images are more robust to adverse illumination and lighting conditions, resulting in better image representation than visible images. The bottom row shows the unsupervised domain adaptation setting, where a domain shift is mitigated between labeled visible domain and unlabeled thermal domain by performing adversarial feature alignment.

datasets. These advancements have enabled object detection as a fundamental component in perception systems for many real-world applications such as autonomous driving, surveillance and human activity recognition [31]. These object detection models are generally trained on large-scale RGB datasets such as ImageNet [10], MS-COCO [30] and Pascal-VOC [12]. However, these RGB-trained models [32, 37, 38] fail due to the domain shift under adverse illumination, occlusion, and lighting conditions. A recent study by NTSB [2] showed that accidents caused by autonomous driving are due to a lack of sensory information regarding the surroundings and do not adequately detect pedestrians and vehicles. In addition, 75% out of 5,987 U.S. pedestrian

fatalities that happened in 2016 are during nighttime [39]. Hence, it is highly challenging for an autonomous system to decide solely based on visible sensory information, as visible sensors fail under such conditions (see Figure 1).

In contrast, the Infrared (IR) sensors are robust under adverse illumination and nighttime conditions and capture more information than visible sensors under those scenarios. Some of these thermal IR cameras are compact, low-cost and small in size. As a result, thermal IR cameras have become increasingly popular in many applications such as autonomous driving, drones and military surveillance for enhanced sensory information [14]. Hence, addressing the detection and classification of objects in thermal imagery plays an essential role in deploying such models in the aforementioned real-world applications. Nevertheless, directly deploying the detectors trained on large-scale visible domain (RGB) datasets fail to detect objects in the thermal domain. This is due to the shift between the visible and thermal domains resulting in performance drop. This is mainly due to the following reasons. First, the RGB-trained detectors lack generalizability and are not robust to unseen domains. Second, shortage of annotated large-scale datasets in the thermal domain. These two reasons hinder the performance level of the detectors from detecting and classifying objects in the thermal domain compared to the visible domain.

Domain shift can be tackled by solving the lack of annotated data issues or increasing the model generalizing capacity and robustness. However, including more annotated data while training the detector is not feasible, as annotating data is a labor-intensive and time-consuming task. This leaves us with the only viable solution to improve the model’s generalizing capability and make it robust by realizing the domain shift. Several domain adaptation settings [7, 41] and methods have been proposed to mitigate the domain shift between the source (i.e. visible) and target (i.e. thermal) domain. In this work, we explore an unsupervised domain adaptation (UDA) setting [16], where domain alignment is achieved solely by learning from the unlabeled target (thermal) domain.

Most UDA works try to mitigate the domain shift using adversarial domain training [7, 42, 48], pseudo-label self-training [24, 40] or image-to-image translation techniques [5, 21]. In this work, we tackle the domain shift problem by proposing an algorithm agnostic meta-learning strategy for domain adaptive detection instead of proposing a new DA strategy. The proposed meta-learning strategy is compatible with all existing UDA methods and it enhances their overall performance. The performance improvement is possible because meta-learning learns the model learning process over multiple training iterations. As a result, meta-learning optimizes the adaptation process by achieving fine DA updates without overfitting or getting stuck at local optima. In meta-

learning, there are two steps; 1) base/inner learning - an inner learning algorithm that performs task-specific optimization. 2) meta/outer learning - an outer learning algorithm that improves base learning meta-parameters to satisfy the outer/meta objective. Thus, a meta-learning pipeline performs optimization at two levels to improve model learning, such as model initialization or model optimizer, to achieve meta-objectives such as learning speed or few-shot learning performance [25, 3, 38]. As meta-objectives are more generic (i.e., model initialization, model optimizer), this can be extended to any existing algorithm resulting in improved performance. For the UDA detection, the meta-objectives are to minimize the supervised loss and reduce the domain shift source and target. However, performing meta-learning in an UDA detection setting is challenging for two reasons: i) object detection method such as FasterRCNN is a computationally heavy framework and calculating meta-objectives for all training samples is intractable to perform meta updates. ii) Unlabeled target images provide no supervised loss to perform base/inner learning updates [28]. To overcome these challenges, we propose an online meta-learning paradigm where instead of performing meta-update after iterating over all the training samples, we perform intermittent meta-update during training. To demonstrate the effectiveness of the proposed method, we evaluate it on visible and thermal benchmark datasets and adaptation protocols and achieves state-of-the-art performance in all datasets. Moreover, ours is the first work to explore unsupervised domain adaptation for thermal object detection. Our main contributions are summarized as follows:

- We introduce an algorithm agnostic meta-learning framework for thermal object detection in an unsupervised domain adaptation setting.
- We propose an online meta-learning strategy which performs online meta-adaptation resulting in a short and tractable computation graph.
- We empirically demonstrate the algorithm agnostic nature of our meta-learning framework over the existing domain adaptation algorithm and proposed architecture in the UDA setting, producing state-of-the-art performance on the KAIST and DSIAC datasets.

2. Related work

Object detection. Object detection is a fundamental problem being explored by the computer vision community for a long time due to its widespread real-world applications. Classical methods perform object detection based on object proposals obtained from selective search [47], superpixel grouping [27] and HOG detector [9]. The rise of deep CNNs shifted the object detection paradigm and resulted in state-of-the-art detectors. CNN-based detectors can be broadly classified into two categories i) One-stage detector

and ii) Two-stage detector. One-stage detectors are YOLO [37] and SSD [32], whereas two-stage detectors are RCNN [18], Fast-RCNN [17], Faster-RCNN [38]. One-stage detectors perform object classification and bounding box regression in a single pipeline. In contrast, the two-stage detectors perform object detection at two stages. In the first stage, a region proposal network is used to generate object proposals and in the second stage, object proposal undergoes object classification and bounding box regression. However, all of these state-of-the-art detectors’ performance drops under domain shift.

Thermal object detection. Thermal object detection plays a vital role in detecting objects in surveillance and military operation [14]. In [23], the SSD architecture is used to detect objects in the thermal domain. [26] proposed to use the YOLO architecture to detect objects in the thermal domain. Dai *et al.* [8] proposed TIRNet to detect objects in thermal IR images to provide more sensory information for autonomous driving. In order to exploit both visible and thermal domains, Devaguptapu *et al.* [11] proposed a detection framework where they fuse visible and thermal features at a high level to capture more information resulting in better detection. Later in [35], they propose a self-training method to enhance the performance in the thermal domain using both visible and thermal images. Note that all of these works have neglected to address a more practical scenario where we have access to a large-scale labeled visible domain image dataset and adapt the detector to the unlabeled thermal domain images.

Unsupervised domain adaptive object detection. In object detection, Chen *et al.* [6] was the first to explore unsupervised domain adaptation settings. In particular, Chen *et al.* [6] proposed DA Faster-RCNN network, which performs adversarial domain training to mitigate the domain shift at the image and instance levels. Later, Saito *et al.* [41] noted that weak alignment of the global features and strong alignment of the local features plays a significant role in adaptation. Cai *et al.* [4] performed domain adaptive detection using a mean-teacher framework to utilize the unlabeled target data better. Recently, Sindagi *et al.* [45] proposed the use of weather priors for adapting detectors to different weather conditions. Zhu *et al.* [49] performed a region mining strategy in order to perform a region-level alignment and showed its benefits compared to conventional domain adversarial training. In addition, there are many other works that have addressed domain adaptive object detection in 2D [36, 40, 48] and 3D [43, 19] domains. However, no works have explored unsupervised domain adaptation settings for thermal object detection. In this paper, we investigate unsupervised domain adaptation for thermal object detection.

Meta-learning. In conventional deep learning, for a given task, models are optimized according to task-specific loss resulting in minimum prediction error [25, 3, 38]. How-

ever, meta-learning provides an alternative paradigm where the model learns to learn over multiple training episodes [46]. In other words, meta-learning is a process of learning to the learn algorithm over multiple training episodes. The meta-learning landscape can be divided into three parts – meta-optimizer, meta-representation, and meta-objective. Meta-optimizer is the choice of optimizer used to learn how the optimization works in the outer loop of meta-learning [13, 20]. Meta-representation aims what meta-knowledge is to be learned and updated in the process of meta-learning [13]. Finally, the meta-objective is the goal of the meta-learning task to be achieved at the end of training [29, 13, 34]. Therefore in this work, we investigate the meta-learning framework for the UDA detection setting, where the meta-representation is the initial condition of the detector and the meta-objective is the detection and adaptation losses.

3. Proposed method

3.1. Preliminaries

Conventional unsupervised domain adaptation methods assume that both source and target data are available while adapting a model for the target domain. Formally in the UDA setting, we denote the labeled source domain as $D_s = \{X_s^n, y_s^n\}_{n=1}^{N_s}$ where X_s^n denotes the n^{th} source image and y_s^n denotes the corresponding object category and bounding box ground truth. In the target domain, the unlabeled target images are denoted as $D_t = \{X_t^n\}_{n=1}^{N_t}$ where X_t^n denotes the n^{th} unlabelled target image. Following the standard domain adaptive detection works, we employ Faster-RCNN (Θ) [38] with the VGG16 [44] backbone as our detector. Unsupervised domain adaptive detection aims to train a detector on the labeled source images and exploit unlabelled target images to reduce the domain shift. To achieve this, we propose an algorithm agnostic meta-learning framework for the thermal object detector that optimizes the model initial condition for domain adaptive detection.

3.2. Detection framework

The Faster-RCNN pipeline consists of three main components: an encoder, a Region Proposal Network (RPN) and a region classification and regression network (RCNN). For a given image X^n , the encoder extracts the feature map and then on top of the extracted feature map, RPN generates class agnostic object region proposals. Following that, the RCNN network classifies object category and predict bounding box offset for the object proposal. The training objective of the detector pipeline is to minimize the RPN and RCNN loss as follows:

$$\mathcal{L}_{det}(X^n, Y^n) = \mathcal{L}_{cls}^{rpn} + \mathcal{L}_{reg}^{rpn} + \mathcal{L}_{cls}^{rcnn} + \mathcal{L}_{reg}^{rcnn}. \quad (1)$$

where \mathcal{L}_{cls}^{rpn} and \mathcal{L}_{cls}^{rcnn} are the object classification loss

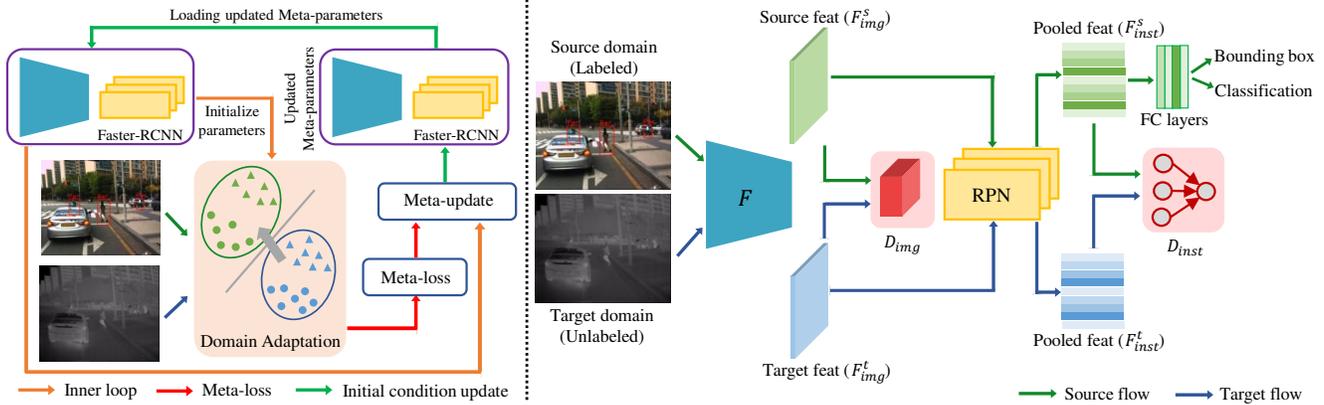


Figure 2. Overview of meta-learning strategy is presented on the left side and adaptation framework is presented on the right side. The meta-learning performs bi-level optimization, where the inner loop performs domain adaptation and the outer loop performs meta-update with respect to meta-loss. The updated meta-parameters are loaded as the initial condition of Faster-RCNN and this process is repeated. The adaptation framework performs adversarial feature alignment between source and target domain at the image and instance level.

[38] and L_{reg}^{rpn} and L_{reg}^{rcnn} are the bounding box regression loss [38] for RPN and RCNN network.

3.3. Image and instance level adaptation

Faster-RCNN [38] is a two-stage detector that performs detection at two levels. When a source trained Faster-RCNN encounters images from the target domain, the performance drops due to domain shift, affecting the detector at two levels. These two levels of the detector are image level and instance level. Image level represents the encoder feature output and instance-level represents the RPN feature output. To mitigate the domain shift, we employ an adversarial domain classifier at both image and instance levels. The adversarial domain classifier helps to align distribution shift, resulting in domain invariant features at the image and instance levels. Briefly, performing adversarial alignment at the image level ensures global feature alignment, such as the shift in image style, illumination. Performing adversarial alignment at the instance level ensures local feature alignment, such as the shift in object size, style, viewpoint, etc. In our work, we have extended the discriminator architecture proposed in DA Faster-RCNN [6] to obtain a stronger and robust classifier which helps in better feature alignment. Architecture details are presented in the supplementary material.

First, let us denote the image-level domain classifier as \mathcal{D}_{img} which classifies the input encoded features as source or target domain. For given source (X_s^n) and target (X_t^n) domain images, the encoder extracted feature map are denoted as $F_{img}^s, F_{img}^t \in \mathbb{R}^{C \times H \times W}$. Feeding F_{img}^s, F_{img}^t to \mathcal{D}_{img} outputs a prediction map of size $H \times W$ with domain labels are set to 1 and 0 for source and target domain respectively. The least squared loss is used to supervise the domain classifier with domain label $y_d \in \{0, 1\}$ and the loss function can be written as:

$$\mathcal{L}_{img}(X_s^n, X_t^n) = - \sum_{h=1}^H \sum_{w=1}^W y_d (1 - \mathcal{D}_{img}(F_{img}^s(h,w)))^2 + (1 - y_d) (\mathcal{D}_{img}(F_{img}^t(h,w)))^2. \quad (2)$$

Second, let us denote instance-level domain classifier as \mathcal{D}_{inst} which classifies the RPN pooled features as source or target domain. For given source (X_s^n) and target (X_t^n) domain images, the RPN pooled features are denoted as $F_{inst}^s, F_{inst}^t \in \mathbb{R}^{C \times D}$. We feed F_{inst}^s, F_{inst}^t to \mathcal{D}_{inst} which outputs a prediction map of size D with domain labels set to 1 and 0 for source and target domain, respectively. The least squared loss is used to supervise the domain classifier and the loss function can be written as:

$$\mathcal{L}_{inst}(X_s^n, X_t^n) = - \sum_{d=1}^D y_d (1 - \mathcal{D}_{inst}(F_{inst}^s(h,w)))^2 + (1 - y_d) (\mathcal{D}_{inst}(F_{inst}^t(h,w)))^2. \quad (3)$$

To achieve the domain alignment, we utilize the Gradient Reversal Layer (GRL) [15], which flips the gradient sign after propagating gradient through the domain classifier. Therefore, when minimizing 2 and 3 for image and instance-level domain classifiers, the GRL helps in achieving equilibrium. In the equilibrium condition, the input features are domain invariant and the domain classifier cannot differentiate the source and target features. Furthermore, we opt for least-squares loss instead of using binary-cross entropy loss, as it is shown to be working better [33]. Hence, the total domain adaptation loss is formulated as follows

$$\mathcal{L}_{da}(D_s, D_t) = \mathcal{L}_{img}(D_s, D_t) + \mathcal{L}_{inst}(D_s, D_t). \quad (4)$$

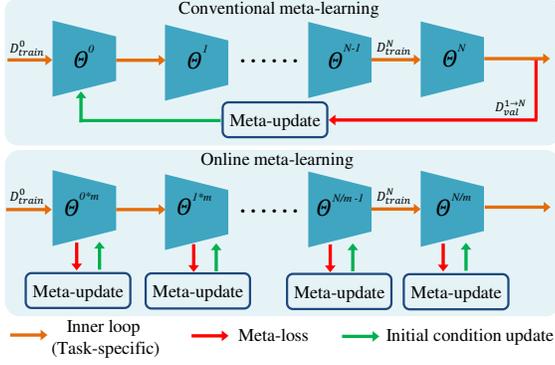


Figure 3. In conventional meta-learning (top-row), the meta-loss is computed for the model obtained from the inner loop after iterating over the complete train set. As a result, this leads to a long and intractable computation graph. In the proposed method, we compute meta-loss at regular interval m during training resulting in a short and tractable computation graph.

3.4. Overall training objective

In the UDA setting, we have labeled samples from the source domain D_s and unlabeled samples from the target domain D_t . The objective of the detector is to learn from labeled source data by minimizing supervised loss (L_{det}). In addition, for domain adaptation, the detector should be domain invariant and this can be realized by reducing the domain shift by minimizing the adaptation loss L_{da} obtained from source data and target data. Hence, the overall loss for unsupervised domain adaptation setting is defined as:

$$\mathcal{L}_{uda}(D_s, D_t) = \mathcal{L}_{det}(D_s) + \mathcal{L}_{da}(D_s, D_t). \quad (5)$$

The degree of domain alignment depends on the model optimization strategy. Thus, meta-learning the detector's initial condition helps in achieving fine DA updates, resulting in an optimally adapted detector.

3.5. Online meta-adaption

Conventional meta-learning [13] the initial condition can be expressed as a bi-level optimization problem, where the inner loop optimizes according to the task-specific loss and the outer algorithm optimizes the meta-parameters with respect to meta-loss as shown in Figure 2. Thus, meta-learning the initial condition is formulated as:

$$\Theta = \underset{\Theta}{\operatorname{argmin}} \underbrace{\mathcal{L}_{\text{outer}}(\underbrace{\mathcal{L}_{\text{inner}}(\Theta, D_{\text{tr}}), D_{\text{val}}}_{\text{Inner-level}})}_{\text{Outer-level}}, \quad (6)$$

where Θ is the meta-parameters which is initialized model parameters; D_{tr} and D_{val} are train and validation dataset; L_{inner} denotes the task-specific loss on its training set and the L_{outer} denotes validation loss obtained after inner optimization. The overall objective of Eqn (6) is to obtain an optimal initial condition such that the validation loss is minimum. Extending this to UDA setting, we create a train and validation dataset for source and target domain denoted as

Algorithm 1 Online meta-adaptation for UDA

Require: $D_s^{tr}, D_t^{tr}, D_s^{val}, D_t^{val}$

Require: α, β : meta learning rate hyperparameters

- 1: randomly initialize Θ
 - 2: **while** not done **do**
 - 3: **for** m **do**
 - 4: Sample batch of $D_s^{tr}, D_t^{tr}, D_s^{val}, D_t^{val}$
 - 5: Evaluate $\mathcal{L}_{uda}(D_s^{tr}(n), D_t^{tr}(n))$ using Equation (5)
 - 6: Compute adapted parameters with gradient descent:
 - 7: $\Theta'_i = \Theta - \alpha \nabla_{\Theta} \mathcal{L}_{uda}$
 - 8: Compute Meta-loss for Θ'_i using Equation (5):
 - 9: $\mathcal{L}_{uda}^{met LOSS} += \mathcal{L}_{uda}(D_s^{val}, D_t^{val})$
 - 10: **end for**
 - 11: Update $\Theta \leftarrow \Theta - \beta \nabla_{\Theta} \mathcal{L}_{uda}^{met LOSS}$
 - 12: **end while**
-

D_s^{tr}, D_t^{tr} and D_s^{val}, D_t^{val} respectively from D_s and D_t . In the inner loop, the task-specific loss corresponds to L_{uda} obtained from the train set D_s^{tr} and D_t^{tr} and is computed as follows:

$$\Theta'_n = \Theta - \alpha \nabla_{\Theta} \mathcal{L}_{uda}(D_s^{tr}(n), D_t^{tr}(n)), \quad (7)$$

where n corresponds to the n^{th} sample from source and target training set and α is the inner loop learning rate. In the outer loop, the meta-loss is computed on the validation set for the inner loop model, which is obtained after fully iterating over the training set. Following that, the initial condition of the detector (i.e. meta-parameters) are updated with respect to meta loss as follows

$$\Theta = \Theta - \beta \nabla_{\Theta} \overbrace{\sum_{n=1}^N \mathcal{L}_{uda}(D_s^{val}, D_t^{val})}^{\text{Meta-loss}}, \quad (8)$$

where β is the meta-learning rate. Thus, we learn to learn the optimization process, resulting in fine DA updates without overfitting or getting stuck at local optima [13, 28]. However, meta-learning is not compatible with the domain adaptive detection framework. Because storing all the inner-loop computation graphs in the detection pipeline is computationally heavy and backpropagating through them is intractable. Thus, we propose an online meta-domain adaptation strategy for the detection pipeline, which performs online meta-adaptation resulting in a short and tractable computation graph. In other words, we extended the meta-learning paradigm to perform on-the-fly meta-updates by optimizing inner and outer loops for intermittent steps as shown in Figure 3. As per Algorithm 1, we perform online meta-adaptation for UDA setting by alternatively optimizing inner and outer loops at short intervals m . This avoids the long computational graphs and provides stable optimization for DA updates. Moreover, the on-the-fly paradigm understands a better association between the initial condition and meta-loss compared to conventional meta-learning. Thus, the online meta-learning ensures gradual optimization and achieves proper fine-tuning for the initial condition resulting in an enhanced adapted detector with

more robustness and generalizing capability.

4. Experiments and results

In this section, we evaluate the proposed method to empirically show its effectiveness on two different adaptation scenarios with visible to thermal domain shift experiments: 1. Automatic Target Recognition [1], and 2. Multi-spectral Pedestrian Detection [22].

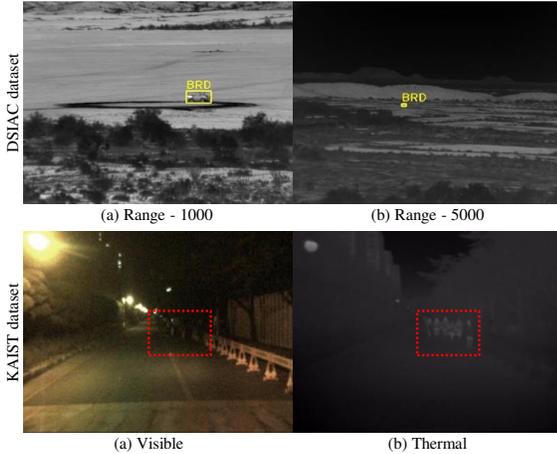


Figure 4. Sample images from the DSIAC thermal dataset at ranges 1000 and 5000 are shown on the top row. We can observe that the object at Range 5000 is very small compared to range 1000 and is not even visible to the naked eyes. However, detecting these objects is crucial for military surveillance. Sample images from the KAIST dataset are shown in the bottom row where the red box denotes the region with pedestrians. We can observe that thermal images can capture more salient features than visible images during nighttime resulting in better image representation.

Automatic target recognition. Automatic Target Recognition (ATR) is an algorithm’s ability to perform real-time target recognition using multiple sensory information. ATR is a well-known problem setup and has been employed in unmanned systems in automated surveillance and military operations [1]. Because these unmanned aerial vehicles (UAVs) and unmanned ground vehicles (UGVs) have multiple sensors to capture different data modes facilitating ATR algorithms. These data are from the visible and thermal domain incorporated to achieve an accurate and robust ATR system. However, most publicly available datasets have labeled visible images and lack labeled samples for thermal images. Hence, it is important to train the detector for the thermal domain to enable ATR in surveillance and military operations.

We implement our method for the ATR problem setting and evaluate our method on the publicly available DSIAC dataset provided by the US Army Night Vision and Electronic Sensors Directorate (NVESD) [1]. The DSIAC dataset contains 106 GB of visible data and 207 GB of Middle Wavelength Infrared (MWIR) data. It contains eight classes, among which two are civilian vehicles and remain-

ing six are military vehicles : ‘Pickup’, ‘Sport vehicle’, ‘BTR70’, ‘BRDM2’, ‘BMP2’, ‘T72’, ‘ZSU23’, ‘2S3’. This dataset was collected during day and night time, with each video sequence containing 1800 video frames. The distance between cameras and targets are varied from 500 to 5000 meters at an interval of 500 meters. Following the conventional UDA setting, we assume we have access to labeled visible data and unlabeled thermal data. For both domains, we sample every ten frames from the dataset video sequence for the ranges 1000-5000 at interval 1000. Further, we assign 75 % of the sampled data as the training set and 25 % of the sampled data as the test set for each range. Thus in this work, we investigate the domain shift between visible and thermal domains at different ranges for the ATR problem.

Multi-spectral pedestrian detection. The KAIST Multi-Spectral dataset [22] contains 95,000 8-bit paired thermal and visible images. This dataset is collected during day and night using a FLIR A35 microbolometer LWIR camera with 320×256 pixels resolution. It contains a standard train-test split of 76000 train images and 19000 test images with only annotation available for the pedestrian class. Following the conventional UDA setting, we assume we have access to labeled visible and unlabeled thermal data and adapt the detector to the thermal domain.

Implementation details. We adopt the unsupervised domain adaptation setting for all our experiments, where we have access to labeled visible data and unlabeled thermal target data. By default, our base detector is Faster-RCNN and is initialized with pre-trained VGG weights. The shorter side of the input images are resized to 600 pixels while maintaining the aspect ratio. We perform random horizontal flip and subtract the image mean for all input images as part of data augmentation. During training, the batch size is set equal to 1. We set the domain adaptation loss weight λ equal to 0.1. The learning rate hyperparameter α and β for the inner and outer loop meta-learning are set equal to 0.001. Momentum is set equal to 0.9 for five epochs and then decreases the learning rate to 0.0001. In meta-learning, alternatively optimizing inner and outer loops for a short interval m is set equal to 3. We train the network for ten epochs and use the mean average precision (mAP) metric as an evaluation metric.

4.1. Quantitative comparison

We compare our method with the existing baselines and show our proposed method’s effectiveness under different domain shift settings. Note that to the best of our knowledge, this is the first work addressing unsupervised domain adaptation for thermal detectors.

Automatic target recognition. In Table 1, we report the performance of our method against DA Faster-RCNN baseline for different ranges. Range - 1000 and 5000 represent the distance at which the target has been captured with re-

Table 1. Quantitative results (mAP) for visible \rightarrow thermal adaptation for the DSIAC dataset. Source only: Trained on visible domain and tested on thermal domain; Oracle: Trained on thermal domain and tested on thermal domain

Method	Range			Pkup Trck	Sprt Vech	BTR70	BRDM2	BMP2	T72	ZSU23	2S3	mAP
	1000	5000	All									
Source Only	✓	×	×	26.9	20.4	81.6	23.9	28.7	46.3	6.0	29.4	32.2
Source Only	×	✓	×	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Source Only	×	×	✓	3.3	4.5	9.5	10.3	0.7	13.3	0.2	5.0	5.2
DA-Faster	✓	×	×	9.4	30.0	79.6	38.9	37.8	44.5	3.6	44.8	36.1
Ours	✓	×	×	16.6	21.2	81.0	67.4	41.1	51.5	1.5	55.5	41.9
DA-Faster	×	✓	×	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Ours	×	✓	×	0.0	0.0	3.0	0.0	0.0	0.0	0.0	6.1	1.1
DA-Faster	×	×	✓	4.2	11.7	12.1	11.4	0.3	9.1	0.2	0.3	6.2
Ours	×	×	✓	9.5	14.0	16.6	12.7	1.4	16.4	0.2	2.0	9.1
Oracle	✓	×	×	100	100	100	100	100	100	100	100	100
Oracle	×	✓	×	21.1	21.8	35.8	21.0	30.0	37.9	1.0	40.1	26.1
Oracle	×	×	✓	63.0	66.3	77.1	68.9	66.8	69.1	80.2	79.9	71.4

Table 2. Quantitative results (mAP) for visible \rightarrow thermal adaptation on the KASIT dataset.

Method	Meta-learn	Person	mAP
Source Only	×	9.1	9.1
DA-Faster	×	11.9	11.9
Ours	×	23.2	23.2
DA-Faster	✓	13.7	13.7
Ours	✓	24.6	24.6
Oracle	×	43.9	43.9

spect to the camera. Range - ‘‘All’’ includes the range from 1000-5000 at intervals of 1000 meters. In addition, the oracle experiment denotes training and testing on the visible domain, whereas the source-only experiments indicate training on the visible domain and testing on the thermal domain. From Figure 4 for Range-1000, we can observe that the targets are big and easy to distinguish. As a result, we obtain oracle performance as 100 mAP. However, visible to thermal domain shift affects the detector by a large margin in the source-only experiment as shown in Table 1. From Figure 4 for Range-1000, we can observe that the targets are very small and difficult to distinguish. Therefore, the oracle performance for Range - 5000 is only 26.8 and the corresponding source-only performance is 0 mAP. For Range - ‘‘All’’, the oracle and source-only performances are 71.4 mAP and 5.2 mAP, respectively. As can be seen from Table 1, domain shift causes catastrophic performance degradation. Hence, we argue that mitigating the domain shift effect plays a crucial role in deploying thermal detectors in real-world applications. Our meta-learning strategy minimizes the domain shift by learning optimal DA updates for adaptation.

As shown in Table 1, for Range-1000, our proposed method performs better than DA-Faster RCNN by 6.2 mAP. For Range-5000, even though after adaptation DA Faster-RCNN performance was 0 mAP, our proposed adaptation strategy ensures the optima adaptation updates, resulting in 1.1 mAP. For Range - ‘‘All’’, our model achieves 40% better mAP than the DA Faster-RCNN method. Thus, we empir-

ically demonstrate that meta-learning the initial condition of the detector is improving the adapted detector’s performance.

Multi-spectral pedestrian detection. In the KAIST dataset, we perform pedestrian detection using Faster-RCNN which produces 43.9 mAP as the oracle performance. In the source-only experiment, the performance drops by 34.8 mAP due to the domain shift as shown in Table 2. DA Faster-RCNN performs adversarial feature alignment to mitigate the domain shift resulting in 11.9 mAP with an improvement of 2.8 mAP from the source-only performance. Following our method, we obtain 21.7 mAP with an improvement of 41.1 mAP from the source-only performance. This difference in improvement shows the effectiveness of a strong discriminator even though our method is an extension of the DA Faster-RCNN approach. However, by applying meta-learning, the initial condition-based adaptation ensures optimal adaption, resulting in improved performance for DA Faster-RCNN and our method by 13.7 mAP and 24.6 mAP, respectively.

Ablation study. We study the adaptation impact of our proposed meta-learning the initial condition strategy for DA Faster-RCNN and our framework. Table 3 presents the ablation analysis done on the DSIAC dataset for all ranges. From Table 3, we can infer that DA Faster-RCNN and our framework without meta-learning produce 6.2 mAP and 8.2 mAP, respectively. The improvement of our framework over DA Faster-RCNN is due to the utilization of a stronger discriminator resulting in a better feature alignment. Furthermore, employing the meta-learning the initial condition strategy for DA Faster-RCNN and our framework, we obtain 14.5 % and 10.9 % improvement for DA Faster-RCNN and our framework, respectively. This performance improvement using meta-learning supports our argument that meta-learning helps in learning to learn the adaptation process by updating the detector’s initial condition.

Qualitative comparison. We visualize the detection performance of source only model, DA Faster-RCNN, our

Table 3. Ablation study for meta-learning strategy in DSIAC dataset Range - "All" for visible \rightarrow thermal adaptation.

Method	Meta-learn	Pkup Trck	Sprt Vech	BTR70	BRDM2	BMP2	T72	ZSU23	2S3	mAP
Source Only	×	3.7	5.8	11.5	10.1	1.6	9.1	0.2	2.8	5.6
DA-Faster	×	4.2	11.7	12.1	11.4	0.3	9.1	0.2	0.3	6.2
Ours	×	10.4	12.9	12.1	12.9	1.2	14.4	0.4	1.5	8.2
DA-Faster	✓	2.7	10.4	21.4	9.4	1.0	12.0	0.4	0.1	7.1
Ours	✓	9.5	14.0	16.6	12.7	1.4	16.4	0.2	2.0	9.1
Oracle	×	63.0	66.3	77.1	68.9	66.8	69.1	80.2	79.9	71.4

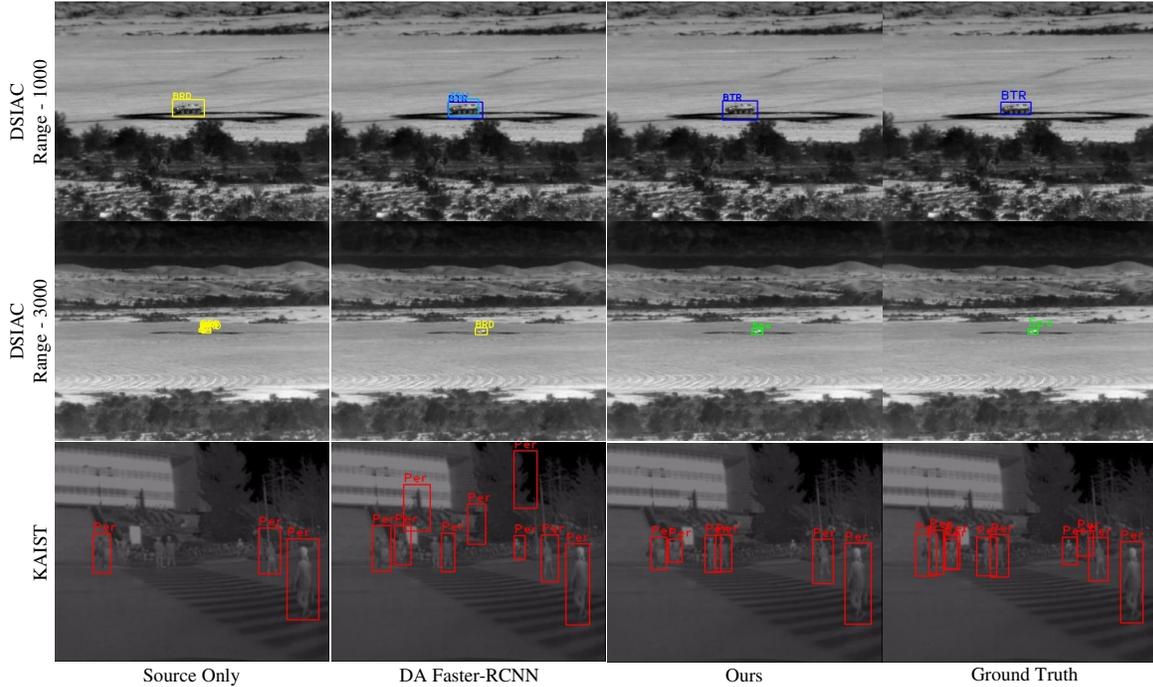


Figure 5. Qualitative detection results. In the DSIAC dataset, source only and DA Faster-RCNN produce false-positive predictions, whereas our method recognizes the object correctly. Similarly, in the KAIST dataset, our method reduces false-positive as well as produces a high-quality prediction. Because meta-learning helps in achieving fine DA updates resulting in a more robust and generalized detector.

method with respect to the ground truth in Figure 5. The visualization is presented for experiments on the DSIAC dataset Range-1000 and All, KAIST dataset in first, second and third row, respectively as shown in Figure 5. In the first and second row, we can observe that the source-only model produces false positive detection due to the domain shift. Further, due to adaptation, DA Faster-RCNN recognizes the object correctly but with a few false-positive predictions. Similarly, for the KAIST dataset (third row), our method produces less miss detection compared to DA Faster-RCNN. Thus, our quantitative and qualitative analysis shows the effectiveness of the proposed method on the DSIAC and KAIST datasets.

5. Conclusion

We presented an unsupervised domain adaptive thermal object detection framework for real-world applications. Specifically, we introduced an algorithm agnostic meta-learning framework applicable for existing domain adaptation techniques. Furthermore, we proposed an online meta-

domain adaptation compatible with the detection framework, which performs online meta-adaptation resulting in a short and tractable computation graph. Finally, we demonstrated the algorithm agnostic nature of our meta-learning framework over the existing domain adaptation algorithm and proposed architecture in the UDA setting. Our method produces state-of-the-art thermal detection performance on the KAIST and DSIAC datasets.

6. Acknowledgment

Research was sponsored by the Army Research Office and was accomplished under Cooperative Agreement Number W911NF-20-2-0224. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Us army night vision and electronic sensors directorate (nvesd). <https://www.dsiac.org/databases/atr-algorithm-development-image-database/>, 2010.
- [2] Ntsb. preliminary report highway hwy18mh010. 2018.
- [3] Vijay Badrinarayanan, Ankur Handa, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *arXiv preprint arXiv:1505.07293*, 2015.
- [4] Qi Cai, Yingwei Pan, Chong-Wah Ngo, Xinmei Tian, Lingyu Duan, and Ting Yao. Exploring object relation in mean teacher for cross-domain detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11457–11466, 2019.
- [5] 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.
- [6] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. *2018 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3339–3348, 2018.
- [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.
- [8] Xuerui Dai, Xue Yuan, and Xueye Wei. Tirnet: Object detection in thermal infrared images for autonomous driving. *Applied Intelligence*, 51(3):1244–1261, 2021.
- [9] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, pages 886–893. Ieee, 2005.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [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.
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [14] Rikke Gade and Thomas B Moeslund. Thermal cameras and applications: a survey. *Machine vision and applications*, 25(1):245–262, 2014.
- [15] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495*, 2014.
- [16] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- [17] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [18] Ross B Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. corr abs/1311.2524 (2013). *arXiv preprint arXiv:1311.2524*, 2013.
- [19] Deepti Hegde, Vishwanath Sindagi, Velat Kilic, A Brinton Cooper, Mark Foster, and Vishal Patel. Uncertainty-aware mean teacher for source-free unsupervised domain adaptive 3d object detection. *arXiv preprint arXiv:2109.14651*, 2021.
- [20] Rein Houthoofd, Richard Y Chen, Phillip Isola, Bradley C Stadie, Filip Wolski, Jonathan Ho, and Pieter Abbeel. Evolved policy gradients. *arXiv preprint arXiv:1802.04821*, 2018.
- [21] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 749–757, 2020.
- [22] 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.
- [23] Rohan Ippalapally, Sri Harsha Mudumba, Meghana Adkay, and Nandi Vardhan H. R. Object detection using thermal imaging. In *2020 IEEE 17th India Council International Conference (INDICON)*, pages 1–6, 2020.
- [24] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 480–490, 2019.
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [26] Mate Krišto, Marina Ivacic-Kos, and Miran Pobar. Thermal object detection in difficult weather conditions using yolo. *IEEE Access*, 8:125459–125476, 2020.
- [27] Alex Levinshtein, Cristian Sminchisescu, and Sven Dickinson. Optimal contour closure by superpixel grouping. In *European Conference on computer vision*, pages 480–493. Springer, 2010.
- [28] Da Li and Timothy Hospedales. Online meta-learning for multi-source and semi-supervised domain adaptation. In *European Conference on Computer Vision*, pages 382–403. Springer, 2020.
- [29] Yiyang Li, Yongxin Yang, Wei Zhou, and Timothy Hospedales. Feature-critic networks for heterogeneous do-

- main generalization. In *International Conference on Machine Learning*, pages 3915–3924. PMLR, 2019.
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [31] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. *International journal of computer vision*, 128(2):261–318, 2020.
- [32] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [33] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017.
- [34] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- [35] Farzeen Munir, Shoaib Azam, and Moongu Jeon. Sstn: Self-supervised domain adaptation thermal object detection for autonomous driving. *arXiv preprint arXiv:2103.03150*, 2021.
- [36] Poojan Oza, Vishwanath A Sindagi, Vibashan VS, and Vishal M Patel. Unsupervised domain adaptation of object detectors: A survey. *arXiv preprint arXiv:2105.13502*, 2021.
- [37] Joseph Redmon, Santosh Kumar Divvala, Ross B Girshick, and Ali Farhadi. You only look once: unified, real-time object detection. corr abs/1506.02640 (2015). *arXiv preprint arXiv:1506.02640*, 2015.
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015.
- [39] R. Retting and S. Schwatz. Governors highway safety association pedestrian traffic fatalities by state (2017 preliminary data). <https://www.ghsa.org/sites/default/files/2018-02/pedestrians18.pdf>, 2017.
- [40] Aruni RoyChowdhury, Prithvijit Chakrabarty, Ashish Singh, SouYoung Jin, Huaizu Jiang, Liangliang Cao, and Erik Learned-Miller. Automatic adaptation of object detectors to new domains using self-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 780–790, 2019.
- [41] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. *CoRR*, abs/1812.04798, 2018.
- [42] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019.
- [43] Cristiano Saltori, Stéphane Lathuilière, Nicu Sebe, Elisa Ricci, and Fabio Galasso. Sf-uda 3d: Source-free unsupervised domain adaptation for lidar-based 3d object detection. In *2020 International Conference on 3D Vision (3DV)*, pages 771–780. IEEE, 2020.
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [45] Vishwanath A Sindagi, Poojan Oza, Rajeev Yasarla, and Vishal M Patel. Prior-based domain adaptive object detection for hazy and rainy conditions. In *European Conference on Computer Vision*, pages 763–780. Springer, 2020.
- [46] Sebastian Thrun and Lorien Pratt. Learning to learn: Introduction and overview. In *Learning to learn*, pages 3–17. Springer, 1998.
- [47] Koen EA Van de Sande, Jasper RR Uijlings, Theo Gevers, and Arnold WM Smeulders. Segmentation as selective search for object recognition. In *2011 international conference on computer vision*, pages 1879–1886. IEEE, 2011.
- [48] Vibashan VS, Vikram Gupta, Poojan Oza, Vishwanath A Sindagi, and Vishal M Patel. Mega-cda: Memory guided attention for category-aware unsupervised domain adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4516–4526, 2021.
- [49] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 687–696, 2019.