

Unsupervised RGB-to-Thermal Domain Adaptation via Multi-Domain Attention Network

Lu Gan, Connor Lee, and Soon-Jo Chung

Abstract—This work presents a new method for unsupervised thermal image classification and semantic segmentation by transferring knowledge from the RGB domain using a multi-domain attention network. Our method does not require any thermal annotations or co-registered RGB-thermal pairs, enabling robots to perform visual tasks at night and in adverse weather conditions without incurring additional costs of data labeling and registration. Current unsupervised domain adaptation methods look to align global images or features across domains. However, when the domain shift is significantly larger for cross-modal data, not all features can be transferred. We solve this problem by using a shared backbone network that promotes generalization, and domain-specific attention that reduces negative transfer by attending to domain-invariant and easily-transferable features. Our approach outperforms the state-of-the-art RGB-to-thermal adaptation method in classification benchmarks, and is successfully applied to thermal river scene segmentation using only synthetic RGB images. Our code is made publicly available at <https://github.com/ganlumomo/thermal-uda-attention>.

I. INTRODUCTION

Cameras are critical for robot perception as they provide dense measurements and rich environmental information. However, most existing vision models are developed for cameras operating in the visible spectrum due to their ubiquity and the accessibility of large-scale RGB datasets [1], [2]. Although these models allow robotic systems such as autonomous vehicles (AV) to work well in ideal conditions with sufficient illumination, their performance is largely degraded at night and in adverse conditions. Thermal cameras, on the other hand, detect electromagnetic waves beyond the visible spectrum that penetrate through dust, smoke, and light fog, enabling around-the-clock robotic operations.

One popular approach towards robust vision is to leverage thermal images in conjunction with RGB via multi-spectral sensor fusion. These methods have largely benefited from recent interests in AV technology, resulting in curated datasets [3], [4] being made publicly available. Notable examples are GAFF [5] and CFT [6], two multi-spectral object detection networks trained on paired RGB-thermal image datasets for feature extraction and fusion. In particular, the fusion network in [6] sees a 25% performance improvement over a single RGB branch on the FLIR-aligned dataset [3]. Urban semantic segmentation has also been improved for nighttime and adverse weather after integrating thermal capabilities [7], [8], [9], [10]. However, these models

*This work is funded by Ford Motor Company and in part by the Office of Naval Research. The authors are with the Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA {ganlu, clee, sjchung}@caltech.edu




Fig. 1: Our RGB-to-thermal unsupervised domain adaptation (UDA) leverages knowledge learned from a synthetic annotated RGB dataset to perform semantic segmentation on thermal river scenes without requiring thermal annotations.

are fully-supervised, using annotated images or co-registered RGB-thermal pairs which are expensive to acquire and small in scale [11]. In non-AV applications, the lack of thermal data and cost of labeling hinder the development of thermal vision models, especially when current vision models, like Transformers, have been trending larger [12].

To overcome this issue, we look to leverage existing large-scale RGB datasets to learn thermal models via unsupervised domain adaptation (UDA) techniques. UDA aims to transfer the knowledge learned in a labeled source domain to an unlabeled target domain [13]. Although most UDA methods focus on domains from different environments but within the same modality (mainly RGB images), such as GTAV-to-Cityscapes [14], [15], the underlying assumption that a domain-invariant feature representation exists also applies to cross-modal data, especially for semantic-related tasks.

In this work, we aim to transfer knowledge learned from labeled RGB images to unlabeled thermal images (Fig. 1). This is challenging for two reasons: First, cross-modal domains have larger domain gaps and more dissimilar features compared to domains within same modalities. UDA methods that match global images or feature distributions of both domains can hurt generalization and lead to *negative transfer* in which untransferable features are forcefully aligned [16], [17], [18]. Second, UDA methods based on generative adversarial networks (GANs) need a large amount of target data to be well-trained [13] which is not a guarantee in the thermal domain.

We surmount these challenges by designing a multi-domain attention network with a shared backbone and domain-specific attention for RGB-to-thermal adaptation.

This shared backbone promotes generalization across domains, prevents feature over-alignment, and relaxes the thermal dataset size requirement. For feature alignment, we train the target-specific attention using adversarial learning to attend to and transfer more domain-invariant and transferable features among all shared features to alleviate negative transfer. The main contributions of our work are as follows:

- We establish an unsupervised RGB-to-thermal domain adaptation method using a multi-domain attention network and adversarial attention learning.
- We evaluate our method on thermal image classification tasks and outperform the state-of-the-art RGB-to-thermal adaptation approach on two benchmarks.
- We demonstrate the versatility of our approach, leveraging it to perform thermal river scene segmentation, and to the best of our knowledge, are the first to utilize synthetic RGB data for thermal semantic segmentation.

II. RELATED WORK

Unsupervised Domain Adaptation: UDA has been successfully applied to a variety of vision tasks including image classification [19], [20], [21], [22], [23], semantic segmentation [15], [14], [24] and 2D/3D object detection [25], [26]. Domain alignment is the fundamental principle of UDA, and can be achieved by two main methodologies: domain mapping and domain-invariant feature learning [13]. Domain mapping can be viewed as pixel-level alignment which maps images from one domain to another via image translation. For instance, PixelDA [27] and CyCADA [15] map source training data into the target domain using conditional GANs and train the downstream model on the fake target data. Pixel-level alignment can remove the domain differences in the input space to some extent but such differences are primarily low-level [13]. Other works achieve domain adaptation by domain-invariant feature learning or feature-level alignment. By mapping source and target input data to the same feature distribution, a downstream predictor trained on such domain-invariant features from source can also work well on the target domain. This is typically done by minimizing a distance defined on distributions [21], or by adversarial training via a domain discriminator that attempts to distinguish between source and target features [19], [20], [22], [14], [23]. Our method is similar to these works and can be viewed as an instance of the general pipeline in [20] by leveraging multi-domain network and attention mechanisms.

RGB-to-Thermal UDA: Despite the success of UDA on visible images, adapting visible models to thermal remains a challenge due to their large domain gap. Existing RGB-to-thermal adaptation works like MS-UDA [9] and HeatNet [10] distill knowledge from a segmentation network pretrained on RGB datasets to their dual-stream network by pseudo-labeling RGB-thermal image pairs. As the pseudo-labels are generated for the RGB image in a pair, the main domain gap here is intra-modal, between the pretraining dataset and RGB images in the paired dataset, rather than inter-modal.

Our work is most related to SGADA [23] and Marnissi *et al.* [26] which aim to transfer knowledge from RGB to ther-

mal without requiring thermal annotations or RGB-thermal pairs. For pedestrian detection, Marnissi *et al.* incorporates alignment at difficult levels into Faster R-CNN [28] using adversarial training. SGADA extends ADDA [20] with a self-training procedure. For pseudo-labeling, not only the model prediction and confidence are considered, but also the prediction and confidence from the domain discriminator. It achieves the best results on MS-COCO [2] to FLIR ADAS [3] adaptation benchmark but depends on the quality of pseudo-labels generated by ADDA.

Attention Networks: Attention mechanisms allow models to dynamically attend to certain parts of the input that are more effective for a task, and become important concepts in neural networks. Attention can be grouped into different types, including sequence attention, channel attention [29], and spatial attention [30], etc. For domain adaptation, Wang *et al.* [17] and Zhang *et al.* [18] propose transferable attention networks using self-attention mechanisms to highlight transferable features. The spatial attention they employed attend to different regions in a feature map. Instead, we use channel-wise attention [29] to attend to different feature maps and use residual adapters [31] to align them, with the intuition that certain types of features are more transferable than others. The transferability difference in feature types (i.e., channels) should be focused on more than in feature regions (i.e., spatial locations) for cross-modal domains.

III. PROPOSED METHOD

A. Multi-Domain Attention Network

Our method draws ideas from multi-domain learning [31] and task attention mechanisms in multi-task learning [32]. These works use shared weights and domain/task-specific parameters to learn both general and domain/task-specific representations. They demonstrate that sharing weights across domains/tasks promotes generalization ability. In contrast, we use domain-specific attention with adversarial learning to facilitate domain-invariant feature extraction and alignment for domain adaptation.

Our multi-domain attention network consists of an encoder-decoder backbone, shared by both source and target domains, with domain-specific attention modules attached at various stages of the encoder. For UDA classification (Fig. 2), the architecture consists of the shared backbone and classifier (blue), source-specific (green), and target-specific (red) attention modules. Hypothesizing that different sensor modality favors different types of features, we use channel-wise attention, i.e., Squeeze-and-Excitation (SE) [29], to highlight more domain-invariant and easily-transferable feature maps among all shared features, and use residual adapters [31] to align them across domains.

Let $F_c \in \mathbb{R}^{h \times w \times C'}$ denote a convolutional layer of C' kernels of size $h \times w$ operating on C' input channels, we have $F_c : x \rightarrow f$, $x \in \mathbb{R}^{H' \times W' \times C'}$, $f \in \mathbb{R}^{H \times W \times C}$, where $f = [f_1, f_2, \dots, f_C]$ represent C output feature maps. A SE module [29] first “squeezes” f into a low-dimensional channel descriptor $d \in \mathbb{R}^{\frac{C}{r}}$ with reduction ratio r . This is done using global average pooling followed by a fully connected

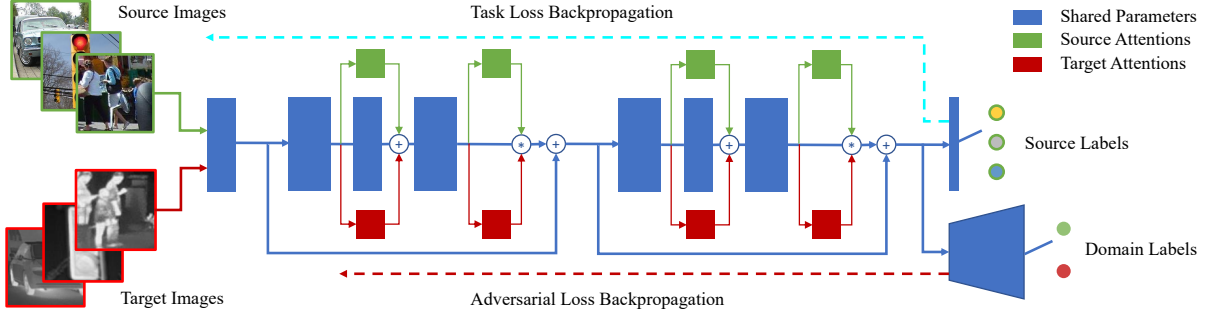


Fig. 2: The network architecture and training procedure of our proposed unsupervised RGB-to-thermal domain adaptation method. The specific architecture is shown for image classification task.

(FC) layer with ReLU activations. The channel descriptor is then transformed into channel-wise weight coefficients $s = [s_1, s_2, \dots, s_C]$, $s_c \in (0, 1)$ through another FC layer and a sigmoid function. Finally, s is used to “excite” different feature maps in f by feature channel reweighting: $\tilde{f}_c = s_c \cdot f_c$. In our network, we use domain-specific SE blocks operating on the shared feature maps right before the residual addition in residual blocks, as shown in Fig. 2.

By attaching domain-specific SE modules to the shared backbone network, they have the ability to accentuate more domain-invariant and transferable features in the shared features while attenuate less-transferable ones. To further align the reweighted features across domains, we leverage residual adapters [31] to directly and dynamically adapting the shared feature extractors F_c to domain-specific feature extractors G_d . Specifically, $d = c$ in our network.

Given a shared convolutional layer of C kernels $F_c \in \mathbb{R}^{h \times w \times C'}$, a domain-specific convolutional layer with D filters $G_d \in \mathbb{R}^{h \times w \times C'}$ can be simply constructed as an affine transformation of F_c using only a small amount of additional parameters $\alpha = \{\alpha_{dc}\}$:

$$G_d = \sum_{c=1}^C \alpha_{dc} F_c. \quad (1)$$

Here, $\alpha \in \mathbb{R}^{D \times C}$ are the trainable residual adapter parameters [31]. This linear parameterization reduces constructing G_d for each domain to the shared F_c with a small amount of domain-specific parameters α . The works [31], [33] further show that α can be reparameterized and implemented as a convolutional layer of 1×1 filters connected in parallel with the shared convolutional layer. In our network, we add residual adapters to the middle 3×3 convolutional layer in residual blocks for feature alignment, as shown in Fig. 2.

We emphasize the differences of our attention modules from those in [31], [32]. The multi-domain learning in [31] and multi-task learning in [32] are essentially supervised. Their objective is to learn a domain/task-invariant feature representation f_{inv} and domain/task-specific attention θ_a, θ_b , so that $\theta_a(f_{inv}) = f_a$, $\theta_b(f_{inv}) = f_b$, where f_a and f_b are features tailored for domain/task A and B respectively. In contrast, for our UDA problem, we learn discriminative

features f_s for a given task using supervised training in the source domain and target-specific attention θ_t using adversarial training for feature alignment, i.e. $\theta_s(f_{sh}) = f_s$, $\theta_t(f_{sh}) = f_{t \rightarrow s}$, where θ_s and θ_t are source and target attention, f_{sh} and $f_{t \rightarrow s}$ are the shared features and the target features aligned with f_s respectively.

B. Adversarial Attention Learning

To perform unsupervised domain adaptation for a network M , we train shared network weights θ_{sh} and source-specific attention weights θ_s using labeled data from the source domain and train the task-specific attention weights θ_t adversarially in an alternating fashion.

Let $\mathcal{D}_s = \{(x_s^i, y_s^i)\}_{i=1}^{n_s}$ and $\mathcal{D}_t = \{(x_t^j)\}_{j=1}^{n_t}$ denote the annotated training data in the source domain and the unlabeled target training data respectively. Then, θ_{sh} and θ_s can be trained with supervision by minimizing the standard cross-entropy loss. For a classification problem, the loss can be written as:

$$\mathcal{L}_{task}(x_s, y_s) = - \sum_{c=1}^C \mathbb{1}_{[c=y_s]} \log M(x_s; \theta_{sh+s}), \quad (2)$$

where (x_s, y_s) are source data-label pairs drawn from \mathcal{D}_s , $\mathbb{1}_{[x]}$ is an indicator function so that $\mathbb{1}_{[x]} = 1$ if $x = 1$, and 0 otherwise, and C is the number of categories.

We train the target-specific attention in our network adversarially by forcing the target attention to attend to domain-invariant features from the shared features and further align them with the source feature distributions. Adversarial attention learning can be achieved by approaching the following minimax game [34], [20] between the target-specific attention θ_t and a domain discriminator D :

$$\min_{\theta_t} \max_D \mathcal{L}(D, \theta_t) = \mathbb{E}_{x_s \sim \mathcal{D}_s} \log D(f_s) + \mathbb{E}_{x_t \sim \mathcal{D}_t} \log(1 - D(f_t)), \quad (3)$$

where f_s and f_t are source and target features from the entire encoder with weights θ_{sh+s} and θ_{sh+t} , respectively.

The minimax loss in (3) is split into two objectives, where the domain discriminator plays the adversarial role and attempts to distinguish between source features f_s and

Algorithm 1 Multi-Domain Attention Network for Unsupervised Domain Adaptation

- 1: **Input:** Training data: $\mathcal{D}_s, \mathcal{D}_t$,
 - 2: Network: $M = \{\theta_{sh}, \theta_s, \theta_t\}$, Discriminator: D
 - 3: Learning rate: α, β, γ
 - 4: Initialize M^0, D^0
 - 5: **for** $n = 1$ to N **do**
 - 6: Sample batch data (x_s, y_s) from \mathcal{D}_s , and x_t from \mathcal{D}_t
 - 7: $l_{task} \leftarrow \mathcal{L}_{task}(x_s, y_s)$ ▷ Evaluate (2)
 - 8: $\theta_{sh+s}^n = \theta_{sh+s}^{n-1} - \alpha \nabla_{\theta_{sh+s}^{n-1}} l_{task}$
 - 9: $l_{adv} \leftarrow \mathcal{L}_{adv}(x_t, D^{n-1})$ ▷ Evaluate (5)
 - 10: $\theta_t^n = \theta_t^{n-1} - \beta \nabla_{\theta_t^{n-1}} l_{adv}$
 - 11: $l_{dis} \leftarrow \mathcal{L}_{dis}(x_s, x_t, M^n)$ ▷ Evaluate (4)
 - 12: $D^n = D^{n-1} - \gamma \nabla_{D^{n-1}} l_{dis}$
 - 13: **end for**
 - 14: **Output:** M^N, D^N
-

target features f_t by minimizing the following loss:

$$\mathcal{L}_{dis}(x_s, x_t, M) = -\log D(f_s) - \log(1 - D(f_t)), \quad (4)$$

and the target-specific attention is trained to fool the domain discriminator and increase domain confusion by minimizing an adversarial loss:

$$\mathcal{L}_{adv}(x_t, D) = -\log D(f_t). \quad (5)$$

The three-step training procedure of our multi-domain attention network is given in Algorithm 1.

Advantages of this alternating training are twofold. First, when training θ_{sh+s} using $\mathcal{L}_{task}(x_s, y_s)$, it reduces to training a supervised source model and the feature extractor learns to extract the most discriminative features for the given task. When training θ_t with fixed θ_{sh} , θ_t learns to select and adapt the most domain-invariant ones among the discriminative features, leading to better adaptation performance. Second, it eliminates a weighting hyperparameter for two loss functions and makes the training procedure more stable.

C. Self-Training

We further fine-tune the model with a single self-training step using pseudo-labels generated for the target training data. Like [23], we save the prediction and the corresponding confidence of the model M trained in Sec. III-B for all unlabeled target training samples. In the meantime, the prediction and confidence of the domain discriminator D are also recorded. For target samples that successfully fool the discriminator (D predicts them as source samples with a high confidence), we assign them pseudo-labels according to the model prediction. For those target samples that the discriminator recognize but with low domain confidence, we also include them in pseudo-labeling. The pseudo-labels are further filtered by the model prediction confidence.

In this stage, we only train the target-specific attention parameters using a cross-entropy loss in a supervised setup:

$$\mathcal{L}_{st}(x_t, \hat{y}_t) = -\sum_{c=1}^C \mathbb{1}_{[c=\hat{y}_t]} \log M(x_t; \theta_t), \quad (6)$$




Fig. 3: Examples of the prepared data from MS-COCO [2] and M³FD Detection datasets [4].

TABLE I: Data statistics of MS-COCO to M³FD dataset.

		Bus	Car	Light	Motor.	People	Truck	Total
MS-COCO	Train	3887	36830	12139	6330	200831	7232	267249
	Val	189	1623	603	229	8331	315	11290
M ³ FD	Train	441	12969	1902	382	8770	696	25160
	Val	55	1621	238	48	1096	87	3141
	Test	55	1620	237	47	1096	86	3141

where \hat{y}_t is the generated pseudo-label for target training data x_t . This way, we can further improve the performance in target while keeping the performance in source, so that we have a single unified model that performs well on both source and target data. This learning-without-forgetting [35] property is another benefit of our multi-domain attention network.

IV. EXPERIMENTAL RESULTS

A. Implementation

For a fair comparison, we employ the same backbone architecture used in other methods, i.e. a ResNet-50 pretrained on ImageNet [1]. We use a FC classifier and a FC discriminator for image classification, and use an Atrous Spatial Pyramid Pooling [36] decoder and a fully-convolutional discriminator for semantic segmentation. Parameters are all updated using the ADAM optimizer with $\beta_1 = 0.5, \beta_2 = 0.999$ and weight decay of 2.5×10^{-5} . The learning rate α, β, γ in Algorithm 1 is set to as $1 \times 10^{-4}, 1 \times 10^{-5}$ and 1×10^{-3} , respectively. All experiments are conducted on a NVIDIA Quadro RTX 8000 GPU with 48GB memory.

B. Ablation Study

To investigate the effects of different attention modules and training strategies on adaptation performance, we conduct a thorough ablation study on 9 training combinations resulting from two types of attention modules, i.e. with (\checkmark) and without the residual adapter/SE module, and 3 different strategies to train them:

- 1) θ_{sh+s+t} : We jointly train all network parameters θ using the sum of \mathcal{L}_{task} in (2) and \mathcal{L}_{adv} in (5), reducing the three training steps in Algorithm 1 to only alternatively training the model M and domain discriminator D .
- 2) $\theta_{sh}, \theta_{s+t}$: We alternatively train the shared parameters θ_{sh} and all domain-specific parameters θ_{s+t} . In this setting, only θ_{sh} is updated using \mathcal{L}_{task} , while θ_{s+t} are adversarially trained using a cross-entropy domain loss in [38] instead of \mathcal{L}_{adv} in Algorithm 1.
- 3) θ_{sh+s}, θ_t : The training procedure given in Algorithm 1.




Fig. 4: Task-specific attention visualization. From left to right in each block: input image, feature map with highest SE weight, feature map with the lowest SE weight, and Grad-CAM [37] of residual adapters for that class.

TABLE II: Ablation study of different attention modules and training strategies on MS-COCO to FLIR ADAS dataset.

	Adapter	SE	Bicycle	Car	Person	Average
θ_{sh+s+t}	✓		89.43	97.14	88.89	91.83
		✓	91.72	93.79	83.96	89.83
$\theta_{sh}, \theta_{s+t}$	✓	✓	87.82	94.34	91.52	91.23
			81.84	96.36	96.66	91.62
$\theta_{sh}, \theta_{s+t}$	✓	✓	82.03	93.18	95.84	90.35
			79.54	96.69	96.99	91.07
θ_{sh+s}, θ_t	✓		90.57	97.22	89.83	92.54
		✓	85.75	97.48	95.85	93.03
	✓	✓	89.20	96.87	95.59	93.88

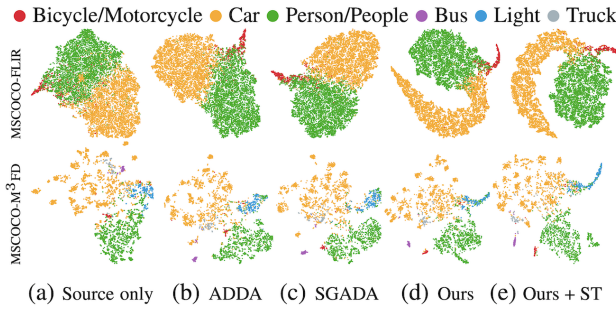


Fig. 5: The t-SNE visualization of the encoded features of all target test samples by different methods (ST: Self-training).

Table II lists the ablation study results using top-1 accuracy. From the table, alternatively training θ_{sh+s} and θ_t has better performance compared with the other two training strategies, and with both residual adapter and SE, it achieves the best result among all configurations. This observation aligns with the discussion in Sec. III-B. In the following experiments, we use setting 3 with both attention modules for our method.

C. Unsupervised Thermal Image Classification

MS-COCO to FLIR ADAS: We first compare our method with SGADA [23] which achieves the best performance on MS-COCO to FLIR ADAS classification benchmark [39] and several other the state-of-the-art general UDA methods. MS-COCO [2] is a large-scale RGB dataset and FLIR ADAS [3] is a popular thermal image dataset for urban environments. We use the dataset prepared by [23] including three categories, i.e., bicycle, car and person, in this experiment. Same as [23], we train our network for 15 epochs with a batch size of 32. Per-class accuracy of all methods are given

in Table III, where the proposed method outperforms other methods by a significant margin even without self-training.

MS-COCO to M³FD: MS-COCO to FLIR ADAS dataset has 633440 unannotated target samples [23]. To further evaluate the adaptation performance when target training samples are scarce, we prepare a new RGB-to-thermal adaptation benchmark using MS-COCO and M³FD [4] including 6 categories for evaluation, following the data preparation process in [23]. Examples and statistics of the prepared dataset are given in Fig. 3 and Table I. Due to less training data, we train all networks for 30 epochs using a batch size of 32. We have similar observations from Table IV as from previous experiment, except that all methods outperform the target only model which shows the effectiveness of UDA when sufficient annotated data is unavailable.

Experiment Analysis: We visualize the feature representations for all test samples on target domain using t-SNE [40] in Fig. 5, where the better feature separation in (d) and (e) shows our method can learn discriminative features for the given task. To examine the effectiveness of attention modules in our method, we further visualize the trained target-specific attentions in Fig. 4 by plotting the features they attend to. For SE modules, we plot the feature map with the highest and the lowest attention weight in the first residual block. As for residual adapters, we plot the Grad-CAM [37] of the last task-specific adapter layer.

We have several interesting observations. First, for bicycle and person categories, the feature maps that the SE highlights the most tend to have high activation on the object contour, which suggests that contour-sensitive features are more domain-invariant and transferable between RGB and thermal domains, aligning with the conclusion in [41]. Second, we notice that cars in thermal images are usually brighter at the bottom due to the high temperature in those regions, as opposed to cars in RGB images which appear darker at the bottom due to shadows. The feature maps that the SE module attends to eliminate this phenomenon and appear visually similar to that of cars from an RGB image. Those observations show the effectiveness of our attention network in extracting domain-invariant and transferable features.

D. Unsupervised Thermal River Scene Segmentation

We present an effective and inexpensive approach for thermal semantic segmentation by adapting from synthetic RGB images using the proposed method, and test on thermal river scene segmentation. We collect 8 sequences of thermal

TABLE III: Top-1 accuracy for MS-COCO to FLIR ADAS.

Method	Bicycle	Car	Person	Average
Source only	69.89	83.89	86.52	80.10
Pixel-DA [27]	62.53	89.99	76.73	76.42
DTA [42]	75.45	97.65	92.45	88.52
MCD-DA [21]	81.71	94.90	91.83	89.48
DANN [19]	78.16	95.07	96.24	89.82
CDAN [22]	78.16	97.10	94.82	90.03
ADDA [20]	86.67	96.95	89.10	90.90
SGADA [23]	87.13	94.44	92.03	91.20
Ours	89.20	96.87	95.59	93.88
Ours + ST	89.63	97.06	96.03	94.24
Target only	87.59	98.78	96.35	94.24

TABLE IV: Top-1 accuracy for MS-COCO to M³FD.

Method	Bus	Car	Light	Motor.	People	Truck	Average
Source only	63.64	76.98	91.14	4.26	94.07	56.98	64.51
MCD-DA [21]	89.09	76.98	95.36	76.59	93.89	30.23	77.00
DANN [19]	89.09	82.72	51.90	68.09	92.15	74.42	76.4
CDAN [22]	89.09	88.58	72.15	46.81	93.98	45.35	72.7
ADDA [20]	96.36	85.86	60.34	51.06	76.73	87.21	76.26
SGADA [23]	94.55	87.22	70.04	51.06	77.01	81.40	76.88
Ours	90.91	85.37	72.57	74.47	93.80	51.16	78.05
Ours + ST	90.91	84.26	85.65	70.21	95.44	56.98	80.57
Target only	94.55	92.53	83.12	21.28	90.24	20.93	67.11

images at 60 Hz using a hand-held FLIR ADK Longwave Infrared (LWIR) thermal camera with an NUC Ruby Mini PC at Big Bear Lake, CA. We sample images every 100 frames from the collected 48676 sequential frames and form an unlabeled training set of 486 thermal images. As our ultimate goal is to enable the nighttime coastline exploration ability of our aerial robots [43], [44] by thermal river segmentation, we manually annotate 282 diverse test images with pixel-level ground truth water labels for evaluation. Examples of collected thermal images are shown in Fig. 1 (4th column).

Due to the lack of annotated RGB dataset for natural scenes similar to our riverine environment, we generate synthetic RGB images with automatically obtained semantic labels using the AirSim simulator [45]. To that end, we use a publicly available simulation environment, i.e. the Landscape Mountains, and simulate a drone platform with a mounted RGB camera to follow a simple survey trajectory around rivers using the built-in simple flight controller. Following our thermal camera, we set the simulated RGB camera to have 75-degree FoV and capture 640×480 images. We acquire RGB images and the corresponding semantic labels at 2Hz, and obtain a synthetic labeled river scene RGB dataset of 1357 samples. We further convert the RGB images to grayscale and invert the intensity values (except for the foliage class), resulting in training samples visually close to our thermal images. Examples of the synthetic RGB images, semantic labels, and inverted grayscale images are shown in the first three columns of Fig 1.

We train the network for 50 epochs using a batch size of 8 without performing self-training for simplicity. From




Fig. 6: Qualitative results of our unsupervised thermal river segmentation model adapted from synthetic RGB data.

TABLE V: Thermal segmentation performance before and after our adaptation using Intersection over Union (IoU).

	Non-water	Water	Average
Source only	78.33	32.90	55.62
Ours adapted model	85.67	54.77	70.22

Table V, our adapted model obtains a performance gain of 14.6% mIoU and 21.87% water-class IoU over the source only model, which is trained using the inverted grayscale images. The effectiveness of our method can be also seen in Fig. 1 and Fig. 6, where the adapted model corrects a large portion of false positive foliage prediction. This experiment demonstrates that thermal vision models can be effectively learned from synthetic RGB data using the proposed method without any manual annotations, even in the source domain.

V. CONCLUSION

This work presented an unsupervised RGB-to-thermal domain adaptation method using multi-domain attention network and adversarial attention learning, and demonstrated its effectiveness on both image classification and semantic segmentation tasks. Vision models adapted by our method achieved a large performance gain over source-only models, and performed on-par with supervised models trained on target. The proposed method can enable robots thermal vision ability without incurring the exorbitant costs of data labeling. In addition, our adaptation method is designed to keep the source performance, i.e. learn without forgetting, providing a unified vision model for both RGB and thermal images.

ACKNOWLEDGMENT

We thank Sushruth Nagesh, Gaurav Pandey, and Sandhya Bhaskar of Ford Motor Company for helpful discussions.

REFERENCES

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.* Ieee, 2009, pp. 248–255.
- [2] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *Proc. European Conf. Comput. Vis.* Springer, 2014, pp. 740–755.
- [3] "Free teledyne flir thermal dataset for algorithm training." [Online]. Available: <https://www.flir.com/oem/adas/adas-dataset-form/>
- [4] J. Liu, X. Fan, Z. Huang, G. Wu, R. Liu, W. Zhong, and Z. Luo, "Target-aware dual adversarial learning and a multi-scenario multi-modality benchmark to fuse infrared and visible for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 5802–5811.
- [5] H. Zhang, E. Fromont, S. Lefèvre, and B. Avignon, "Guided attentive feature fusion for multispectral pedestrian detection," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 72–80.
- [6] F. Qingyun, H. Dapeng, and W. Zhaokui, "Cross-modality fusion transformer for multispectral object detection," *arXiv preprint arXiv:2111.00273*, 2021.
- [7] Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, "MFNet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.* IEEE, 2017, pp. 5108–5115.
- [8] W. Zhou, S. Dong, C. Xu, and Y. Qian, "Edge-aware guidance fusion network for rgb–thermal scene parsing," in *Proc. AAAI Conf. Artif. Intell.*, vol. 36, no. 3, 2022, pp. 3571–3579.
- [9] Y.-H. Kim, U. Shin, J. Park, and I. S. Kweon, "Ms-uda: Multi-spectral unsupervised domain adaptation for thermal image semantic segmentation," *IEEE Robot. Autom. Lett.*, vol. 6, no. 4, pp. 6497–6504, 2021.
- [10] J. Vertens, J. Zürn, and W. Burgard, "HeatNet: Bridging the day-night domain gap in semantic segmentation with thermal images," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.* IEEE, 2020, pp. 8461–8468.
- [11] Z. Küttik and G. Algan, "Semantic segmentation for thermal images: A comparative survey," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 286–295.
- [12] K. Han, Y. Wang, H. Chen, X. Chen, J. Guo, Z. Liu, Y. Tang, A. Xiao, C. Xu, Y. Xu, et al., "A survey on vision transformer," *IEEE Trans. Pattern Anal. Machine Intell.*, 2022.
- [13] G. Wilson and D. J. Cook, "A survey of unsupervised deep domain adaptation," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 11, no. 5, pp. 1–46, 2020.
- [14] Y.-H. Tsai, W.-C. Hung, S. Schuster, K. Sohn, M.-H. Yang, and M. Chandraker, "Learning to adapt structured output space for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 7472–7481.
- [15] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "CyCADA: Cycle-consistent adversarial domain adaptation," in *International conference on machine learning*. Pmlr, 2018, pp. 1989–1998.
- [16] H. Zhao, R. T. Des Combes, K. Zhang, and G. Gordon, "On learning invariant representations for domain adaptation," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7523–7532.
- [17] X. Wang, L. Li, W. Ye, M. Long, and J. Wang, "Transferable attention for domain adaptation," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 01, 2019, pp. 5345–5352.
- [18] C. Zhang, Q. Zhao, and Y. Wang, "Transferable attention networks for adversarial domain adaptation," *Information Sciences*, vol. 539, pp. 422–433, 2020.
- [19] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *International conference on machine learning*. PMLR, 2015, pp. 1180–1189.
- [20] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 7167–7176.
- [21] K. Saito, K. Watanabe, Y. Ushiku, and T. Harada, "Maximum classifier discrepancy for unsupervised domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 3723–3732.
- [22] M. Long, Z. Cao, J. Wang, and M. I. Jordan, "Conditional adversarial domain adaptation," *Proc. Advances Neural Inform. Process. Syst. Conf.*, vol. 31, 2018.
- [23] I. B. Akkaya, F. Altinel, and U. Halici, "Self-training guided adversarial domain adaptation for thermal imagery," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2021, pp. 4322–4331.
- [24] H. Gao, J. Guo, G. Wang, and Q. Zhang, "Cross-domain correlation distillation for unsupervised domain adaptation in nighttime semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 9913–9923.
- [25] Q. Xu, Y. Zhou, W. Wang, C. R. Qi, and D. Anguelov, "SPG: Unsupervised domain adaptation for 3d object detection via semantic point generation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 15446–15456.
- [26] M. A. Marnissi, H. Fradi, A. Sahbani, and N. E. B. Amara, "Unsupervised thermal-to-visible domain adaptation method for pedestrian detection," *Pattern Recognition Letters*, vol. 153, pp. 222–231, 2022.
- [27] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 3722–3731.
- [28] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *Proc. Advances Neural Inform. Process. Syst. Conf.*, vol. 28, 2015.
- [29] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 7132–7141.
- [30] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. European Conf. Comput. Vis.*, 2018, pp. 3–19.
- [31] S.-A. Rebuffi, H. Bilen, and A. Vedaldi, "Learning multiple visual domains with residual adapters," *Proc. Advances Neural Inform. Process. Syst. Conf.*, vol. 30, 2017.
- [32] K.-K. Maninis, I. Radosavovic, and I. Kokkinos, "Attentive single-tasking of multiple tasks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 1851–1860.
- [33] S.-A. Rebuffi, H. Bilen, and A. Vedaldi, "Efficient parametrization of multi-domain deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 8119–8127.
- [34] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *J. Mach. Learning Res.*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [35] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 40, no. 12, pp. 2935–2947, 2017.
- [36] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 40, no. 4, pp. 834–848, 2017.
- [37] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 618–626.
- [38] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko, "Simultaneous deep transfer across domains and tasks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 4068–4076.
- [39] "Mscoco to flir adas benchmark (unsupervised domain adaptation)." [Online]. Available: <https://paperswithcode.com/sota/unsupervised-domain-adaptation-on-mscoco-to/>
- [40] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [41] J. Chen, Z. Liu, D. Jin, Y. Wang, F. Yang, and X. Bai, "Light transport induced domain adaptation for semantic segmentation in thermal infrared urban scenes," *IEEE Trans. Intell. Transport. Syst.*, 2022.
- [42] S. Lee, D. Kim, N. Kim, and S.-G. Jeong, "Drop to adapt: Learning discriminative features for unsupervised domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 91–100.
- [43] J. Yang, A. Dani, S.-J. Chung, and S. Hutchinson, "Vision-based localization and robot-centric mapping in riverine environments," *J. Field Robot.*, vol. 34, no. 3, pp. 429–450, 2017.
- [44] K. Meier, S.-J. Chung, and S. Hutchinson, "River segmentation for autonomous surface vehicle localization and river boundary mapping," *J. Field Robot.*, vol. 38, no. 2, pp. 192–211, 2021.
- [45] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics*, 2017. [Online]. Available: <https://arxiv.org/abs/1705.05065>