

Multi-Granularity Alignment Domain Adaptation for Object Detection

Wenzhang Zhou¹, Dawei Du², Libo Zhang^{1,3*}, Tiejian Luo¹, Yanjun Wu³

¹University of Chinese Academy of Sciences, Beijing, China

²Kitware, Inc., NY, USA

³Institute of Software, Chinese Academy of Sciences, Beijing, China

<https://github.com/tiankongzhang/MGADA>

Abstract

Domain adaptive object detection is challenging due to distinctive data distribution between source domain and target domain. In this paper, we propose a unified multi-granularity alignment based object detection framework towards domain-invariant feature learning. To this end, we encode the dependencies across different granularity perspectives including pixel-, instance-, and category-levels simultaneously to align two domains. Based on pixel-level feature maps from the backbone network, we first develop the omni-scale gated fusion module to aggregate discriminative representations of instances by scale-aware convolutions, leading to robust multi-scale object detection. Meanwhile, the multi-granularity discriminators are proposed to identify which domain different granularities of samples (i.e., pixels, instances, and categories) come from. Notably, we leverage not only the instance discriminability in different categories but also the category consistency between two domains. Extensive experiments are carried out on multiple domain adaptation scenarios, demonstrating the effectiveness of our framework over state-of-the-art algorithms on top of anchor-free FCOS and anchor-based Faster R-CNN detectors with different backbones.

1. Introduction

Owing to the emergence of deep learning [41], modern object detection methods [19, 23, 24, 27, 34] have achieved remarkable progress based on large-scale annotated datasets. However, such domain constrained models often fail in new environments without labeled training data.

To tackle this problem, a feasible solution is to reduce the disparity between label-rich source domain and label-agnostic target domain by *unsupervised domain adaptation* in an adversarial manner [8]. Specifically, the domain discriminator is introduced to identify whether the image is

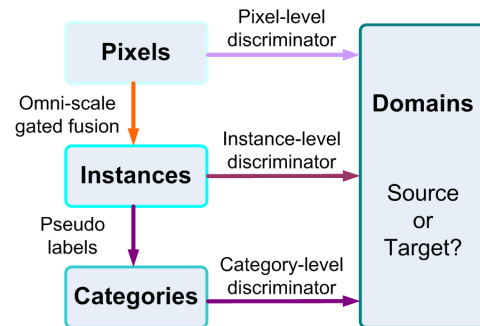


Figure 1. Our framework to encode dependencies across multiple granularities including pixel-, instance-, category-level.

from source domain or target domain; while the object detector learns domain-invariant features to confuse the discriminator [29]. However, classic domain adaptation frameworks suffer from scale variations in cluttered background, resulting in limited performance. Due to convolution layers with fixed kernels in the network, it is difficult to capture accurate features of objects with various scales and aspect ratios. For small objects, the features are convolved from a large region with too much background; for large objects, convolutions cover a small part and lack global structural information.

On the other hand, for better adaptation in target domain, some researchers employ various feature alignment strategies from different granularity perspectives, i.e., instance-, pixel- and category-level. *Instance-level alignment* [3, 20] relies on pooled features of detection proposals to help train the domain discriminator. However, instance-level pooling operation may distort the features of objects with various scales and aspect ratios. In contrast, *pixel-level alignment* [14, 18] focuses on aligning lower-level features that account for each pixel to handle cross-domain variations of objects and the background. However, there exists a large gap between pixel-level features for different scales of objects with the same category. Recently, *category-level alignment* [39, 40] leverages the categorical discriminability in two domains to deal with the hard aligned instances.

*Corresponding author (libo@iscas.ac.cn)

However, these works pay more attention on the consistency between the image-level and instance-level predictions.

To solve the above issues, we propose a unified multi-granularity alignment based object detection framework by using unsupervised domain adaptation. As shown in Figure 1, we encode the dependencies in different granularity perspectives including pixel-, instance-, and category-levels to align source domain and target domain, which is not a rough combination of previous single-granularity alignment techniques. To adapt to various instances, our omni-scale gated fusion selects the most plausible convolutions from low-resolution and high-resolution streams to extract the features. Concretely, we first estimate coarse detections as the guidance based on pixel-level backbone feature maps. Then, parallel convolutions are activated to aggregate discriminative representations of instances with similar scales and aspect ratios. In this way, the following object detection head can predict multi-scale objects more accurately. Meanwhile, we introduce a new category-level discriminator to consider not only the instance discriminability in different categories but also the category consistency between source and target domains. To supervise the category-level discriminator, we assign pseudo labels to important instances with high confidence from object detection. In summary, we construct the multi-granularity discriminators in three granularities of samples (*i.e.*, pixels, instances, and categories). Thus complementary information in different granularity can support each other and achieve better domain adaptation performance.

To verify the effectiveness of our method, we conduct comprehensive experiments on different domain adaptation scenarios (*i.e.*, Cityscapes [4], FoggyCityscapes [30], Sim10k [17], KITTI [10], PASCAL VOC [7], Clipart [16] and Watercolor [16]). The proposed framework is evaluated on top of anchor-free FCOS [34] and anchor-based Faster R-CNN [27] with VGG-16 [32] and ResNet-101 [11] backbones, achieving state-of-the-art performance on different datasets. For example, our method achieves 43.8% mAP score adapting from the source domain Cityscapes [4] to the target domain FoggyCityscapes [30] using FCOS [34], which is 3.6% better than the second best method CFA [14].

Contributions. 1) We propose the multi-granularity alignment framework to encode dependencies across pixel-, instance- and category-level granularities for adaptive object detection, which can be applied in different object detectors. 2) The omni-scale gated fusion module is designed to extract a discriminative representation in terms of objects with different scales and aspect ratios. 3) The category-level discriminator models both instance discriminability in different categories and category consistency between source domain and target domain. 4) Our method achieves the state-of-the-art performance on five domain adaptation applications.

2. Related Work

2.1. Object Detection

CNNs based object detection methods can be generally grouped into anchor-based and anchor-free frameworks. Anchor-based detectors use a series of anchor boxes with different scales and aspect ratios to generate detection proposals, and then apply a network to classify and regress each candidate object. Faster-RCNN [27] develops the region proposal network (RPN) to generate proposals efficiently. FPN [23] introduces a new top-down architecture with lateral connections to capture multi-scale feature maps. In contrast, anchor-free methods rely on keypoints to represent objects. CornerNet [19] is the pioneering work to detect an object bounding box as a pair of the top-left and bottom-right corners. Recently, FCOS [34] leverages the fully convolutional networks to predict labels and bounding box coordinates of each pixel in feature maps. In this work, we build our domain adaptation framework on two representative detectors that are widely used in previous domain adaptation methods, *i.e.*, anchor-based Faster-RCNN [27] and anchor-free FCOS [34].

2.2. Unsupervised Domain Adaptation

Given the labeled source data and unlabeled target data, unsupervised domain adaptation in object detection attracts the interest of researchers. Ganin and Lempitsky [8] perform domain adaptation for classification networks through standard backpropagation training. Inspired by [9], detection networks are optimized by adversarial learning [35,43]. They apply a domain discriminator to distinguish the feature differences between source and target domains, and a gradient reversal layer to reduce the feature distribution differences between domains. Inoue *et al.* [16] propose the cross-domain weakly supervised object detection method by fine-tuning the detector on two types of artificially and automatically generated samples. Saito *et al.* [29] develop a strong-weak distribution alignment method that adjusts the ability of distribution alignment on local and global levels. Recently, Zheng *et al.* [45] generate the attention map with the predicted object categories and sizes to choose regions of objects. In contrast, with the aid of coarse detection, our omni-scale gated fusion module aggregates instance features by scale-aware convolutions to adapt to multi-scale objects in a soft-decision way.

2.3. Alignment Strategies in Domain Adaptation

As discussed in the introduction section, to improve upon domain-level alignment, various feature alignment schemes are applied in other finer levels, *i.e.*, instance-level [3, 20], pixel-level [14, 18] and category-level [6, 15, 26, 37, 39, 40].

Chen *et al.* [3] deal with the domain shift on two lev-

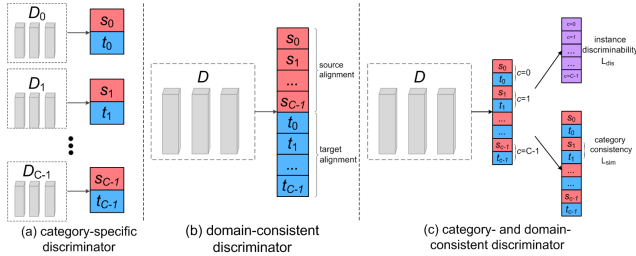


Figure 2. Illustration of different category-level discriminators D , where s_c and t_c denote the c -th category ($c = 0, 1, \dots, C - 1$) in source domain and target domain respectively. (a) Category-specific discriminators for each category [6, 15, 26]. (b) Domain-consistent discriminator to distinguish different categories within one domain [37]. (c) Our category- and domain-consistent discriminator to consider both instance discriminability in different categories and category consistency between two domains.

els including image-level (e.g., image style and illumination) and instance-level (e.g., object appearance and size). Li *et al.* [20] propose the spatial attention pyramid network to capture context information of objects at different scales. Kim *et al.* [18] design the multi-domain-invariant representation learning to encourage the unbiased semantic representation through adversarial learning. Hsu *et al.* [14] propose a center-aware alignment based domain adaptation method to focus on pixel-wise objectness.

Category-level alignment. In terms of category-level alignment, some works [6, 15, 26] design a category-specific discriminator for each category and focus on classification between source and target domains based on pseudo labels (see Figure 2(a)). It is difficult to learn discriminative category-wise representation among multiple discriminators. Wang *et al.* [37] retain one discriminator to distinguish different categories within one domain (see Figure 2(b)). However, it consider little about the consistency of feature subspaces in the same category across two domains. Besides, a categorical regularization method is developed in [39] to locate crucial image regions and important instances to reduce the domain discrepancy. Similarly, Xu *et al.* [40] seek for category-level domain alignment by enhancing intra-class compactness and inter-class separability. It builds the graph based on the Euclidean distances between different category prototypes, where the feature subspaces follow the Gaussian distribution.

In contrast, our category-level discriminator does not rely on the Gaussian distribution assumption but selects important instances to model the subspaces based on an adaptive threshold. Then we model both instance discriminability in different categories and category consistency between two domains (see Figure 2(c)). Moreover, based on the merits of feature alignment at different levels, our method is a unified domain adaptation framework by taking all the granularities into consideration.

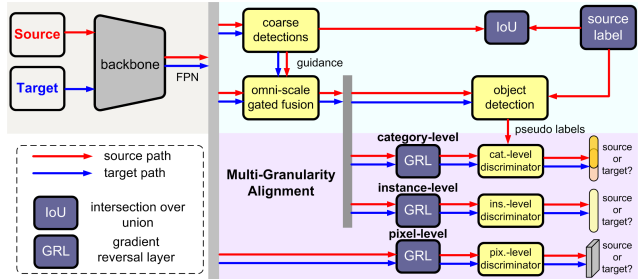


Figure 3. Architecture of our domain adaptive object detection network. Note that object detection heads and discriminators have different sizes of outputs with regard to different detectors.

3. Multi-Granularity Alignment

As shown in Figure 3, given images from the source domain s and target domain t , we first compute the base feature maps using the backbone. Then, the pixel-level features are merged to generate discriminative representations of multi-scale instances by the omni-scale gated fusion module. Based on merged features, the objects can be estimated by the object detection head more accurately. Meanwhile, we introduce the multi-granularity discriminators to distinguish the feature distribution between source and target domains in different perspectives, including pixel-level, instance-level and category-level.

Notably, our method can be applied in different detectors (e.g., anchor-based Faster-RCNN [27] and anchor-free FCOS [34]) and backbones (e.g., VGG-16 [32] and ResNet-101 [11]). Without loss of generality, we first take FCOS [34] as an example, and then explain how our method is applied in Faster-RCNN [27]. For the FCOS detector [34], we extract the last three stages of backbone feature maps and combine them into multi-level feature maps F^k , $k \in \{3, 4, 5, 6, 7\}$ using the FPN representation [23].

3.1. Omni-Scale Gated Object Detection

Most previous domain adaptation methods focus on designing discriminators at specific level and attentive regions. However, the point representation in anchor-free models [14, 25] is difficult to extract robust and discriminative feature in cluttered background, while the AlignROI operation in anchor based models [13, 29] may distort the features of objects with various scales and aspect ratios.

To solve this issue, we employ the omni-scale gated fusion to adapt to various instances with different scales and aspect ratios. Concretely, with the scale guidance from coarse detections, the most plausible convolutions with different kernels are selected to extract the compact features of instances in terms of object scale. Thus it can be applied in different detectors.

Scale guidance. Followed by the multi-level feature maps F^k , we can predict the candidate object boxes \tilde{b}^k by using a series of convolutional layers. According to [28], we use

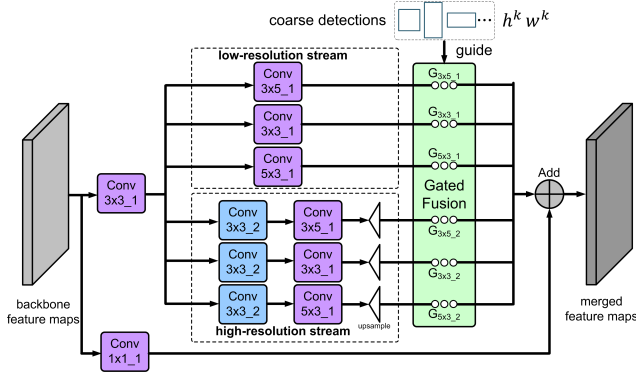


Figure 4. Omni-scale gated fusion module for the FCOS detector [34]. “3x3.2” in the blue rectangles denotes the 3×3 convolutional layer with stride 2.

the cross-entropy Intersection over Union (IoU) loss [42] to regress the bounding boxes of objects in foreground pixels, *i.e.*,

$$\mathcal{L}_{\text{gui}} = - \sum_k \sum_{(i,j)} \ln(\text{IoU}(\tilde{b}_{i,j}^k, b_{i,j}^k)), \quad (1)$$

where $\text{IoU}(\cdot, \cdot)$ is the function to calculate the IoU score between predicted box \tilde{b}^k and ground-truth box b^k . For each pixel (i, j) in the feature map, the corresponding box can be defined as a 4-dimensional vector $b_{i,j}^k = (x_{t_{i,j}}, x_{b_{i,j}}, x_{l_{i,j}}, x_{r_{i,j}})$, denoting the distances between current location and the top, bottom, left and right bounds of ground-truth box respectively. Thus we can calculate the normalized object scale (*i.e.*, width w^k and height h^k) at each level as

$$\begin{cases} w_{i,j}^k = (\tilde{x}_{r_{i,j}} + \tilde{x}_{l_{i,j}}) / \text{stride}^k, \\ h_{i,j}^k = (\tilde{x}_{b_{i,j}} + \tilde{x}_{t_{i,j}}) / \text{stride}^k, \end{cases} \quad (2)$$

where stride^k denotes how many steps we are moving in each step in convolution¹. As defined in the FCOS detector [34], the feature maps at each level are used to individually detect objects with different scales in the range $\{[-1, 64], [64, 128], [128, 256], [256, 512], [512, +\infty]\}$. Therefore, the majority of object scales is less than 8, *i.e.*, $w^k \leq 8, h^k \leq 8$. For simplicity, we omit the superscript k and write F for F^k and \tilde{b} for \tilde{b}^k in the following sections.

Omni-scale gated fusion. To adapt to various scales of objects \tilde{b} with different aspect ratios, we design the omni-scale gated fusion module that is composed of both *low-resolution* and *high-resolution* streams. As shown in Figure 4, the *low-resolution* stream contains three parallel convolutional layers with different kernels $\omega \in \{3 \times 3, 3 \times 5, 5 \times 3\}$, which is used to extract features of small objects ($w^k \leq 5, h^k \leq 5$). In the *high-resolution* stream, we first apply the 3×3 convolutional layers with stride 2 to expand the

¹We have $\{(k, \text{stride}) | (3, 8), (4, 16), (5, 32), (6, 64), (7, 128)\}$.

receptive field and then convolutional layers with kernels ω to deal with large objects ($w^k > 5, h^k > 5$).

After that, we introduce the gate mask G to weight each convolutional layer based on the predicted coarse boxes \tilde{b} ,

$$G_\omega = \frac{\exp(\tau(o_\omega - \hat{o}))}{\sum_\omega \exp(\tau(o_\omega - \hat{o}))}, \quad (3)$$

where τ is the temperature factor. $o_\omega = \text{IoU}(\tilde{b}, \omega)$ denotes the overlap between the predicted box and the convolution kernel ω . \hat{o} is the maximal overlap among them. Finally, we can merge the pixel-level features to exploit the scale-wise representation of instances, *i.e.*,

$$M = \sum_\omega F_{3 \times 3} \odot G_\omega + F_{1 \times 1}, \quad (4)$$

where \odot denotes the element-wise product. F_ω denotes the feature maps after the convolutional layer with kernel ω .

Object detection. After obtaining the merged feature maps M , we predict the categories and bounding boxes of objects. In the FCOS network [34], the object detection heads consist of classification, centerness and regression branches. The classification and centerness branches are optimized by the focal loss [24] \mathcal{L}_{cls} and cross-entropy loss [34] \mathcal{L}_{ctr} respectively. The regression branch is optimized by the IoU loss [42] \mathcal{L}_{reg} . The loss function for object detection is defined as

$$\mathcal{L}_{\text{det}} = \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{ctr}} + \mathcal{L}_{\text{reg}}. \quad (5)$$

More details of the above loss functions refer to [34].

3.2. Multi-Granularity Discriminators

As discussed in the introduction section, we apply the multi-granularity discriminators to identify whether the sample belongs to the source domain or target domain in different perspectives including pixels, instances and categories. The difference between two domains is reduced by the Gradient Reversal Layer (GRL) [8] that transfers reverse gradient when optimizing the object detection network. The discriminator consists of four stacked convolution-groupnorm-relu layers and an additional 3×3 convolutional layer. In the following, we describe our multi-granularity discriminators in detail.

Pixel-level and instance-level discriminators. Pixel- and instance-level discriminators are used to perform pixel-level and instance-level alignment of feature maps respectively. As shown in Figure 3, given the input multi-level features F and merged feature M , L_{pix} and L_{ins} employ the same loss function of the discriminator defined in Eq. (6). Similar to the previous work [14], we use the same loss function, denoted as \mathcal{L}_{pix} and \mathcal{L}_{ins} . For example, the loss of the pixel-level discriminator D^{pix} is defined as

$$\begin{aligned} \mathcal{L}_{\text{pix}} = & - \sum_{(i,j)} y_{i,j}^{\text{pix}} \log D^{\text{pix}}(F^s(i, j)) \\ & + (1 - y_{i,j}^{\text{pix}}) \log(1 - D^{\text{pix}}(F^t(i, j))), \end{aligned} \quad (6)$$

where $F(i, j)$ denotes the feature at pixel (i, j) in the feature map. We have the domain label $y_{i,j}^{\text{pix}} = 1$ if the pixel is from source domain and 0 otherwise.

Category-level discriminator. As shown in Figure 2(c), our category-level discriminator is used to keep semantic consistency between different domain distribution. Specifically, we predict the category and domain labels of pixel (i, j) in each image based on the output feature map $\hat{M} \in \mathbb{R}^{H \times W \times 2C}$, where H and W are the height and width respectively, $2C$ represents the total number of categories for source and target domains.

Since there is no ground-truth to supervise the category-level discriminator, we assign pseudo labels to important samples with high confidence from object detection (see Sec. 3.1). In practice, given a batch of input images, we can output the category probability map P using the object detection heads, and compute the maximum category probability over all levels \bar{P} . Let \mathcal{S} denote the set of selected instances such that its probability is greater than the threshold, *i.e.*, $\mathcal{S} = \{(i, j) | P_{i,j} > \theta_{\text{cat}} \bar{P}\}$. Then the instances in different categories are classified by Eq. (7), while the same category in two domains is aligned by Eq. (9):

- To keep **instance discriminability in different categories**, we separate the category distribution by using the following loss function:

$$\mathcal{L}_{\text{dis}} = -\frac{1}{|\mathcal{S}|} \sum_{(i,j) \in \mathcal{S}} \sum_{c=0}^{C-1} \hat{y}_{i,j,c}^{\text{dis}} \log(p_{i,j,c}^{\text{dis}}). \quad (7)$$

By normalizing the confidence over the domain channel, $p_{i,j,c}^{\text{dis}}$ represents the probability of the c -th category of the pixel, *i.e.*,

$$p_{i,j,c}^{\text{dis}} = \frac{\exp(\hat{M}_{i,j,2c} + \hat{M}_{i,j,2c+1})}{\sum_{c=0}^{C-1} \exp(\hat{M}_{i,j,2c} + \hat{M}_{i,j,2c+1})}, \quad (8)$$

where $\hat{M}_{i,j,2c}$ and $\hat{M}_{i,j,2c+1}$ denote the confidence of the c -th category in source and target domains respectively (*c.f.* Figure 2(c)). $\hat{y}^{\text{dis}} \in \mathbb{R}^{H \times W \times C}$ is the pseudo category label. We have $\hat{y}_{i,j,c}^{\text{dis}} = 1$ if the instance at (i, j) in \hat{M} is an important one of the c -th category and $\hat{y}_{i,j,c}^{\text{dis}} = 0$ otherwise.

- **Category consistency in two domains.** After classifying instances of different categories, we need to further determine which domain the instance comes from. With the GRL [8], the loss function can be written as

$$\mathcal{L}_{\text{sim}} = -\frac{1}{|\mathcal{S}|} \sum_{(i,j) \in \mathcal{S}} \sum_{m=0}^{2C-1} \hat{y}_{i,j,m}^{\text{sim}} \log(p_{i,j,m}^{\text{sim}}), \quad (9)$$

where $y^{\text{sim}} \in \mathbb{R}^{H \times W \times 2C}$ is the pseudo domain label. Similarly, we have $\hat{y}_{i,j,m}^{\text{sim}} = 1$ if the instance at (i, j)

in \hat{M} is an important one of the $\lfloor \frac{m}{2} \rfloor$ -th category in specific domain and $\hat{y}_{i,j,m}^{\text{sim}} = 0$ otherwise. The domain probability p^{sim} is

$$p_{i,j,m}^{\text{sim}} = \frac{\exp(\hat{M}_{i,j,m})}{\sum_{v=m-\%2+1}^{m-\%2+1} \exp(\hat{M}_{i,j,v})}, \quad (10)$$

where $\%_0$ is the remainder function. The loss function of the category-level discriminator D^{cat} is written as

$$\mathcal{L}_{\text{cat}} = \lambda_{\text{dis}} \mathcal{L}_{\text{dis}} + \lambda_{\text{sim}} \mathcal{L}_{\text{sim}}, \quad (11)$$

where λ_{dis} and λ_{sim} are the balancing factors.

3.3. Overall Loss Function

As discussed above, the omni-scale gated object detection network is supervised by \mathcal{L}_{gui} and \mathcal{L}_{det} . Meanwhile, the multi-granularity discriminators are optimized in different granularities, including pixel-level \mathcal{L}_{pix} , instance-level \mathcal{L}_{ins} and category-level \mathcal{L}_{cat} . In summary, the overall loss function is defined as

$$\mathcal{L} = \underbrace{(\mathcal{L}_{\text{gui}} + \mathcal{L}_{\text{det}})}_{\text{object detection}} + \alpha \cdot \underbrace{(\mathcal{L}_{\text{pix}} + \mathcal{L}_{\text{ins}} + \mathcal{L}_{\text{cat}})}_{\text{multi-granularity discriminators}}, \quad (12)$$

where α is the balancing factor between object detection and multi-granularity discriminators.

3.4. Implementation Details

Extension of our framework. To extend our framework to Faster-RCNN [27], we use the backbone features with the stride 16 to collect the base feature maps F . Since Faster-RCNN [27] is a two-stage object detection method, we directly use the Region Proposal Network (RPN) to predict the coarse candidate boxes, supervised by the original RPN loss $\mathcal{L}_{\text{gui}} = \mathcal{L}_{\text{rpn}}$. Similarly, we use the classification and regression branches in the object detection heads to estimate the categories and bounding boxes of objects, defined as $\mathcal{L}_{\text{det}} = \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{reg}}$. Note that the RPN in Faster-RCNN [27] only predicts the top K proposals. To fuse the feature maps with different convolutional layers, we first concatenate the feature map after each convolutional layer and then extract the features for each proposal by the ROIAlign operation. Finally, the merged features are determined by the corresponding object scales according to the RPN outputs. The detailed architecture of our method upon Faster-RCNN [27] can be found in the supplementary materials.

Optimization strategy. We train the proposed network in two stages empirically. First, we disable the category-level discriminator and train the remaining network without multi-scale augmentation. Second, we fine-tune the whole network by adding the category-level discriminator and multi-scale augmentation. The model is trained with learning rate of 0.005, momentum of 0.9, and weight decay of 0.0001. The balancing factors in Eq. (11) are set as $\lambda_{\text{dis}} = 1.0$, $\lambda_{\text{sim}} = 0.1$, and α in Eq. (12) is set as 0.1.

Method	Detector	Backbone	person	rider	car	truck	bus	train	mbike	bicycle	mAP
Baseline	Faster-RCNN	VGG-16	17.8	23.6	27.1	11.9	23.8	9.1	14.4	22.8	18.8
DAF [3]	Faster-RCNN	VGG-16	25.0	31.0	40.5	22.1	35.3	20.2	20.0	27.1	27.6
SC-DA [46]	Faster-RCNN	VGG-16	33.5	38.0	48.5	26.5	39.0	23.3	28.0	33.6	33.8
MAF [12]	Faster-RCNN	VGG-16	28.2	39.5	43.9	23.8	39.9	33.3	29.2	33.9	34.0
SW-DA [29]	Faster-RCNN	VGG-16	29.9	42.3	43.5	24.5	36.2	32.6	30.0	35.3	34.3
DAM [18]	Faster-RCNN	VGG-16	30.8	40.5	44.3	27.2	38.4	34.5	28.4	32.2	34.6
MOTR [1]	Faster-RCNN	ResNet-50	30.6	41.4	44.0	21.9	38.6	40.6	28.3	35.6	35.1
CST [44]	Faster-RCNN	VGG-16	32.7	44.4	50.1	21.7	45.6	25.4	30.1	36.8	35.9
PD [38]	Faster-RCNN	VGG-16	33.1	43.4	49.6	22.0	45.8	32.0	29.6	37.1	36.6
CDN [33]	Faster-RCNN	VGG-16	35.8	45.7	50.9	30.1	42.5	29.8	30.8	36.5	36.6
SFOD-Masoic-Defoggy [22]	Faster-RCNN	VGG-16	34.1	44.4	51.9	30.4	41.8	25.7	30.3	37.2	37.0
ATF [13]	Faster-RCNN	VGG-16	34.6	46.5	49.2	23.5	43.1	29.2	33.2	39.0	37.3
SW-Faster-ICR-CCR [39]	Faster-RCNN	VGG-16	32.9	43.8	49.2	27.2	45.1	36.4	30.3	34.6	37.4
SCL [31]	Faster-RCNN	VGG-16	31.6	44.0	44.8	30.4	41.8	40.7	33.6	36.2	37.9
CFFA [45]	Faster-RCNN	VGG-16	43.2	37.4	52.1	34.7	34.0	46.9	29.9	30.8	38.6
GPA [40]	Faster-RCNN	ResNet-50	32.9	46.7	54.1	24.7	45.7	41.1	32.4	38.7	39.5
SAPNet [20]	Faster-RCNN	VGG-16	40.8	46.7	59.8	24.3	46.8	37.5	30.4	40.7	40.9
UMT [5]	Faster-RCNN	VGG-16	56.5	37.3	48.6	30.4	33.0	46.7	46.8	34.1	41.7
MeGA-CDA [36]	Faster-RCNN	VGG-16	37.7	49.0	52.4	25.4	49.2	46.9	34.5	39.0	41.8
CDG [21]	Faster-RCNN	VGG-16	38.0	47.4	53.1	34.2	47.5	41.1	38.3	38.9	42.3
ours	Faster-RCNN	VGG-16	43.9	49.6	60.6	29.6	50.7	39.0	38.3	42.8	44.3
oracle	Faster-RCNN	VGG-16	46.5	51.3	65.2	32.6	49.9	34.2	39.6	45.8	45.6
SST-AL [25]	FCOS	-	45.1	47.4	59.4	24.5	50.0	25.7	26.0	38.7	39.6
CFA [14]	FCOS	VGG-16	41.9	38.7	56.7	22.6	41.5	26.8	24.6	35.5	36.0
CFA [14]	FCOS	ResNet-101	41.5	43.6	57.1	29.4	44.9	39.7	29.0	36.1	40.2
ours	FCOS	VGG-16	45.7	47.5	60.6	31.0	52.9	44.5	29.0	38.0	43.6
ours	FCOS	ResNet-101	43.1	47.3	61.5	30.2	53.2	50.3	27.9	36.9	43.8
oracle	FCOS	VGG-16	50.1	46.4	68.0	33.7	54.5	38.7	30.7	39.7	45.2
oracle	FCOS	ResNet-101	46.6	45.4	66.1	33.6	54.1	62.9	29.0	37.1	46.9

Table 1. Weather adaptation detection results from Cityscapes to FoggyCityscapes.

4. Experiments

In this section, we compare our method upon different detectors (FCOS [34] and Faster-RCNN [27]) and backbones (VGG-16 [32] and ResNet-101 [11]) with state-of-the-art domain adaptation methods. Moreover, we conduct a detailed ablation study to analyze the influence of important components in our model. Following [3], all the methods are evaluated using the mean average precisions (mAP) at the IoU threshold of 0.5.

4.1. Datasets

Following [3], the experiments are carried out on 7 datasets including Cityscapes [4], FoggyCityscapes [30], Sim10k [17], KITTI [10], PASCAL VOC [7], Clipart [16] and Watercolor [16].

For weather adaptation, Cityscapes [4] is a dataset of outdoor street scenes in normal weather, including 2,975 images for training set and 500 images for validation set with 50 different cities. As a natural target domain, FoggyCityscapes [30] is a fog weather outdoor street scene dataset synthesized on Cityscapes [4]. Sim10k [17] contains 10k images of the synthetic driving scene from the game video Grand Theft Auto V (GTA5). Thus the adaptation from Sim10k [17] to Cityscapes [4] can be used for evaluation in synthetic-to-real adaptation. Similar to Cityscapes [4], KITTI [10] is another popular scene dataset with 7,481 im-

ages in the training set. We verify the cross-camera adaptation ability from KITTI [10] to Cityscapes [4]. Note that only the class *car* is considered in synthetic-to-real and cross-camera adaptations. Besides, we evaluate domain adaptation methods on dissimilar domains, *i.e.*, from PASCAL VOC [7] with real images to Clipart [16] and Watercolor [16] with artistic images. Note that we use 15k images in PASCAL VOC 2007 and 2012 training and validation sets as the source domain.

4.2. Result Analysis

As presented in Table 1, Table 2, and Table 3, we compare our method with other state-of-the-art methods in various domain adaptation scenarios. Meanwhile, we provide the performance of the baseline Faster-RCNN [27] without adaptation. The “oracle” results indicate that we remove the discriminators in our network and then train and evaluate it on the target domain.

Cityscapes→FoggyCityscapes. In Table 1, we evaluate our method on weather adaptation datasets from Cityscapes [4] to FoggyCityscapes [30]. By using FCOS [34], our method achieves 3.6% gain over the second best CFA [14] with ResNet-101 backbone and more gain with VGG-16 backbone. By using Faster-RCNN [27], our method still obtains better performance than the recent CDG [21]. In addition, our method performs slightly worse than the oracle results with different detection backbones, indicating

Method	Detector	Backbone	mAP (<i>car</i>)
Baseline	Faster-RCNN	VGG-16	30.1/30.2
DAF [3]	Faster-RCNN	VGG-16	39.0/38.5
MAF [12]	Faster-RCNN	VGG-16	41.1/41.0
ATF [13]	Faster-RCNN	VGG-16	42.8/42.1
SC-DA [46]	Faster-RCNN	VGG-16	43.0/42.5
UMT [5]	Faster-RCNN	VGG-16	43.1/-
SFOD-Mosaic [22]	Faster-RCNN	VGG-16	43.1/44.6
CST [44]	Faster-RCNN	VGG-16	44.5/43.6
MeGA-CDA [36]	Faster-RCNN	VGG-16	44.8/43.0
SAPNet [20]	Faster-RCNN	VGG-16	44.9/43.4
CDN [33]	Faster-RCNN	VGG-16	49.3/44.9
ours	Faster-RCNN	VGG-16	49.8/45.2
oracle	Faster-RCNN	VGG-16	66.9
SST-AL [25]	FCOS	-	51.8/45.6
CFA [14]	FCOS	VGG-16	49.0/43.2
CFA [14]	FCOS	ResNet-101	51.2/45.0
ours	FCOS	VGG-16	54.6/48.5
ours	FCOS	ResNet-101	54.1/46.5
oracle	FCOS	VGG-16	72.3
oracle	FCOS	ResNet-101	71.3

Table 2. Synthetic-to-Real/Cross-camera adaptation detection results from Sim10k/KITTI to Cityscapes.

the effectiveness of our model.

Sim10k/KITTI→**Cityscapes**. We provide the results on the synthetic-to-real adaptation datasets, where Sim10k [17] is the source domain and Cityscapes [4] is the target domain. As shown in Table 2, our method achieves the best accuracy of 54.6% with VGG-16 backbone and 54.1% with ResNet-101 backbone respectively. Compared with CFA [14] using FCOS [34], our method obtains 5.6% gain with VGG-16 and 2.9% gain with ResNet-101 respectively. We notice there is a huge gap between the results of our method and oracle. This is because of a significant domain shift between synthetic Sim10k [17] and real Cityscapes [4].

We also present the comparison between our method and state-of-the-arts on cross-camera adaptation datasets. KITTI [10] is the source domain while Cityscapes [4] is the target domain. Compared with CFA [14], our method acquires 5.3% and 1.5% gains with VGG-16 [32] and ResNet-101 respectively, showing state-of-the-art performance using different backbones. It is worth mentioning that FCOS [34] with ResNet performs slightly inferior to the VGG counterpart. This is maybe because VGG features are more suitable than ResNet features for adaptation from Sim10k/KITTI to Cityscapes.

PASCAL VOC→**Clipart/Watercolor**. In addition, we evaluate our method using Faster-RCNN [27] with ResNet-101 on real-to-artistic adaptation datasets from PASCAL VOC [7] to Clipart and Watercolor [16]. According to Table 3, our method obtains the best mAP score of 44.8% and 58.1% on Clipart and Watercolor respectively, outperforming the second best UMT [5] slightly on Clipart. Since there

Method	Detector	Backbone	mAP
Baseline	Faster-RCNN	ResNet-101	27.8/44.6
SW-DA [29]	Faster-RCNN	ResNet-101	38.1/53.3
SCL [31]	Faster-RCNN	ResNet-101	41.5/55.2
DBGL [2]	Faster-RCNN	ResNet-101	41.6/53.8
ATF [13]	Faster-RCNN	ResNet-101	42.1/54.9
PD [38]	Faster-RCNN	ResNet-101	42.1/56.9
SAPNet [20]	Faster-RCNN	ResNet-101	42.2/55.2
UMT [5]	Faster-RCNN	ResNet-101	44.1/58.1
ours	Faster-RCNN	ResNet-101	44.8/58.1
oracle	Faster-RCNN	ResNet-101	-/55.4

Table 3. Real-to-Artistic adaptation detection results from PASCAL VOC to Clipart/Watercolor².

exists a significant class imbalance (*i.e.*, the label *car*, *cat* and *dog* have much fewer images than other labels), our method even performs better than the oracle result on Watercolor [16]. By using our multi-granularity discriminators, the training samples in source domain can facilitate training an accurate detection network.

4.3. Ablation Study

To study the effectiveness of important modules in our network, we conduct an ablation study on domain adaptation from Cityscapes [4] to FoggyCityscapes [30]. We use FCOS [34] as the base detector with VGG-16 backbone for all the variants. As shown in Figure 5, the visual results indicate that the proposed omni-scale gated fusion and category-level discriminator reduce false positives and negatives for object detection in adaptive domains.

Effectiveness of omni-scale gated fusion. To verify the ability of our method to deal with scale variations, we compare the mAP scores of our method and CFA [14] in terms of object scale. According to COCO metrics, AP^S , AP^M and AP^L denote the mAP scores such that object area is in the range $[0, 32^2]$, $(32^2, 96^2]$, and $(96^2, +\infty)$ respectively.

In Table 4, the performance at all scales is considerably improved by using the omni-scale gated fusion in the object detection network compared with the baseline method, *i.e.*, 39.3% vs. 36.8%. If we remove the omni-scale gated fusion module before the object detection heads, the performance of “ours (w/o gated fusion)” is reduced by 2.3%. We also notice that our method performs better than CFA [14] at all scales, especially at large scale. We speculate that the omni-scale gated fusion module can deal with various scales of objects to generate a more discriminative representation for object detection.

Moreover, we discuss the influence of coarse detection guidance for omni-scale gated fusion and object detection heads. If we use naive average fusion “ours (w/ average fusion)” or 1×1 convolution fusion “ours (w/ conv fusion)”,

²There are no oracle results for the Clipart [16] dataset because we use all the images in the Clipart [16] dataset as the target domain.

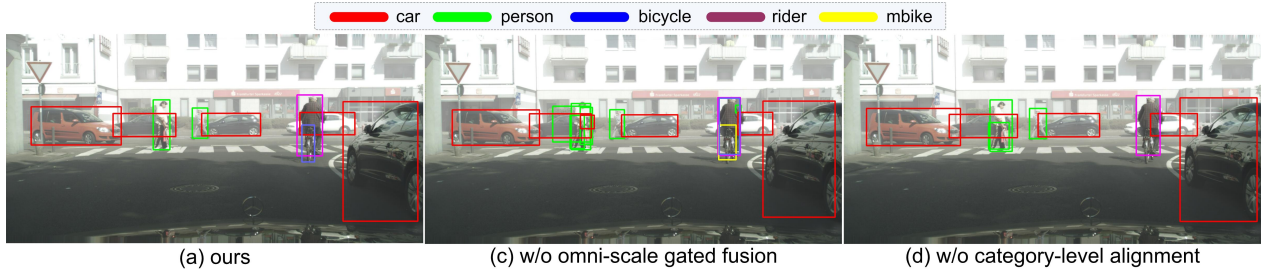


Figure 5. Visual comparison between our method and its variants.

method	mAP	AP ^S	AP ^M	AP ^L
CFA [14]	36.0	8.3	36.7	61.6
ours (w/o all)	36.8	7.2	37.7	64.1
ours (w/o category-level dis.)	39.3	8.7	40.5	64.4
ours (w/o gated fusion)	41.3	8.5	39.1	70.6
ours (w/ all)	43.6	10.1	43.1	72.5
ours (w/ average fusion)	42.1	11.5	40.7	68.9
ours (w/ conv fusion)	41.5	11.2	40.1	71.5
ours (w/ gated fusion)	43.6	10.1	43.1	72.5

Table 4. Effectiveness of multi-scale object detection. Different variants of our method are constructed by removing important modules in the network. Moreover, two baseline fusion strategies are compared with our omni-scale gated fusion.

the performance is lower than our method using the omni-scale gated fusion. It demonstrates that the coarse detection guidance is crucial to select the most plausible convolution for multi-scale feature aggregation.

Effectiveness of category-level discriminators. If we reduce the multi-granularity discriminators to classical discriminators by removing the category-level discriminator D^{cat} in Eq. (11) from our method, we can observe the sharp drop of 4.3% mAP of the baseline (39.3 vs. 43.6), as presented in Table 4. It indicates the importance of our proposed discriminator.

To further demonstrate the superiority of our category-level discriminator D^{cat} , we add three most related discriminators including D^{ins} [14], D^{grp} [15], and D^{cls} [37] in our network with the baseline discriminators in Eq. (6). D^{cen} [14] considers the center-aware distribution alignment of pixel-level instances with its multi-scale extension. As shown in Figure 2(a), D^{grp} [15] utilizes category-level adversarial discriminator to decrease the differences within each category between source and target domain. As shown in Figure 2(b), D^{cls} [37] expands the binary domain labels by using class information, and preserves intra-domain structures of source and target domains. From Table 5, we obtain only 39.3% mAP score by using the traditional pixel-level discriminator D^{pix} in Eq. (6). By using either D^{cen} or D^{grp} in our method, the performance is improved with less than 2% gain. Although D^{cls} can further improve the performance slightly, our method achieves a considerable gain of near 5%. This is attributed to more instance discriminability

discriminator	baseline	D^{cen} [14]	D^{grp} [15]	D^{cls} [37]	D^{cat} (ours)
mAP	39.3	40.5	40.7	41.1	43.6

Table 5. Comparison between different discriminators. Existing discriminators including D^{ins} , D^{grp} , and D^{cls} are added in our network with the baseline discriminators in Eq. (6).

method	CFA [14]	ours	SCL [31]	SAPNet [20]	ours
detector	FCOS	FCOS	Faster-RCNN	Faster-RCNN	Faster-RCNN
# of Params (M)	177	283	580	556	255
FPS	17.5	10.0	11.8	25.2	21.4

Table 6. Comparison of computational complexity.

in different categories over two domains in our method.

Computational complexity. In addition, we provide the comparison of computational complexity between our method and other SOTA works in Table 6. Note that the discriminators are removed in testing phase for most methods except SCL [31]. By using anchor-free FCOS, our multi-granularity alignment framework performs the best with reasonable increased complexity over its primary contender CFA [14]; while our method has less parameters on top of anchor-based Faster-RCNN than two recent methods SCL [31] and SAPNet [20].

5. Conclusions

In this work, we encode the multi-granularity dependencies among pixel-, instance- and category-level information to align the feature distribution of source domain and target domain in a more accurate way. Notably, the proposed omni-scale gated fusion module can exploit instance features among multi-scale feature maps with most plausible convolutions. Meanwhile, the multi-granularity discriminators can distinguish instances in different categories over two domains. The experiment shows the superiority of the above designs in our framework on top of different detectors and backbones for domain adaptive object detection.

Acknowledgement and Declaration of Conflicting Interests. Dr. Wu was supported by the Key Research Program of Frontier Sciences, CAS, Grant No. ZDBS-LY-JSC038. Dr. Zhang was supported by Youth Innovation Promotion Association, CAS (2020111). Dr. Du and his employer received no financial support for the research, authorship, and/or publication of this article.

References

- [1] 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 *CVPR*, pages 11457–11466, 2019. 6
- [2] Chaoqi Chen, Jiongcheng Li, Zebiao Zheng, Yue Huang, Xinghao Ding, and Yizhou Yu. Dual bipartite graph learning: A general approach for domain adaptive object detection. In *ICCV*, pages 2703–2712, 2021. 7
- [3] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster R-CNN for object detection in the wild. In *CVPR*, pages 3339–3348, 2018. 1, 2, 6, 7
- [4] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016. 2, 6, 7
- [5] Jinhong Deng, Wen Li, Yuhua Chen, and Lixin Duan. Unbiased mean teacher for cross-domain object detection. In *CVPR*, pages 4091–4101, 2021. 6, 7
- [6] Liang Du, Jingang Tan, Hongye Yang, Jianfeng Feng, Xiangyang Xue, Qibao Zheng, Xiaoqing Ye, and Xiaolin Zhang. SSF-DAN: separated semantic feature based domain adaptation network for semantic segmentation. In *ICCV*, pages 982–991, 2019. 2, 3
- [7] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *IJCV*, 88(2):303–338, 2010. 2, 6, 7
- [8] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, volume 37, pages 1180–1189, 2015. 1, 2, 4, 5
- [9] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17:59:1–59:35, 2016. 2
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, pages 3354–3361, 2012. 2, 6, 7
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 2, 3, 6
- [12] Zhenwei He and Lei Zhang. Multi-adversarial faster-rcnn for unrestricted object detection. In *ICCV*, pages 6667–6676, 2019. 6, 7
- [13] Zhenwei He and Lei Zhang. Domain adaptive object detection via asymmetric tri-way faster-rcnn. In *ECCV*, volume 12369, pages 309–324, 2020. 3, 6, 7
- [14] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *ECCV*, volume 12354, pages 733–748, 2020. 1, 2, 3, 4, 6, 7, 8
- [15] Lanqing Hu, Meina Kan, Shiguang Shan, and Xilin Chen. Unsupervised domain adaptation with hierarchical gradient synchronization. In *CVPR*, pages 4042–4051, 2020. 2, 3, 8
- [16] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, pages 5001–5009, 2018. 2, 6, 7
- [17] Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, Karl Rosaen, and Ram Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? In *ICRA*, pages 746–753, 2017. 2, 6, 7
- [18] Taekyung Kim, Minki Jeong, Seunghyeon Kim, Seokeon Choi, and Changick Kim. Diversify and match: A domain adaptive representation learning paradigm for object detection. In *CVPR*, pages 12456–12465, 2019. 1, 2, 3, 6
- [19] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, volume 11218, pages 765–781, 2018. 1, 2
- [20] Congcong Li, Dawei Du, Libo Zhang, Longyin Wen, Tiejian Luo, Yanjun Wu, and Pengfei Zhu. Spatial attention pyramid network for unsupervised domain adaptation. In *ECCV*, volume 12358, pages 481–497, 2020. 1, 2, 3, 6, 7, 8
- [21] Shuai Li, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Category dictionary guided unsupervised domain adaptation for object detection. In *AAAI*, 2021. 6
- [22] Xianfeng Li, Weijie Chen, Di Xie, Shicai Yang, Peng Yuan, Shiliang Pu, and Yueting Zhuang. A free lunch for unsupervised domain adaptive object detection without source data. In *AAAI*, 2021. 6, 7
- [23] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, 2017. 1, 2, 3
- [24] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2999–3007, 2017. 1, 4
- [25] Muhammad Akhtar Munir, Muhammad Haris Khan, M. Saqib Sarfraz, and Mohsen Ali. Synergizing between self-training and adversarial learning for domain adaptive object detection. In *NeurIPS*, 2021. 3, 6, 7
- [26] Sujoy Paul, Yi-Hsuan Tsai, Samuel Schulter, Amit K. Roy-Chowdhury, and Manmohan Chandraker. Domain adaptive semantic segmentation using weak labels. In *ECCV*, volume 12354, pages 571–587, 2020. 2, 3
- [27] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *TPAMI*, 39(6):1137–1149, 2017. 1, 2, 3, 5, 6, 7
- [28] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian D. Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, pages 658–666, 2019. 3
- [29] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, pages 6956–6965, 2019. 1, 2, 3, 6, 7
- [30] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *IJCV*, 126(9):973–992, 2018. 2, 6, 7

- [31] Zhiqiang Shen, Harsh Maheshwari, Weichen Yao, and Marios Savvides. SCL: towards accurate domain adaptive object detection via gradient detach based stacked complementary losses. *CoRR*, abs/1911.02559, 2019. [6](#), [7](#), [8](#)
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. [2](#), [3](#), [6](#), [7](#)
- [33] Peng Su, Kun Wang, Xingyu Zeng, Shixiang Tang, Dapeng Chen, Di Qiu, and Xiaogang Wang. Adapting object detectors with conditional domain normalization. In *ECCV*, volume 12356, pages 403–419, 2020. [6](#), [7](#)
- [34] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *ICCV*, pages 9626–9635, 2019. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [35] Yi-Hsuan Tsai, Kihyuk Sohn, Samuel Schulter, and Manmohan Chandraker. Domain adaptation for structured output via discriminative patch representations. In *ICCV*, pages 1456–1465, 2019. [2](#)
- [36] 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 *CVPR*, pages 4516–4526, 2021. [6](#), [7](#)
- [37] Haoran Wang, Tong Shen, Wei Zhang, Ling-Yu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *ECCV*, volume 12359, pages 642–659, 2020. [2](#), [3](#), [8](#)
- [38] Aming Wu, Yahong Han, Linchao Zhu, and Yi Yang. Instance-invariant domain adaptive object detection via progressive disentanglement. *TPAMI*, 2021. [6](#), [7](#)
- [39] Chang-Dong Xu, Xing-Ran Zhao, Xin Jin, and Xiu-Shen Wei. Exploring categorical regularization for domain adaptive object detection. In *CVPR*, pages 11721–11730, 2020. [1](#), [2](#), [3](#), [6](#)
- [40] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *CVPR*, pages 12352–12361, 2020. [1](#), [2](#), [3](#), [6](#)
- [41] Yongjun Xu, Xin Liu, Xin Cao, Changping Huang, Enke Liu, Sen Qian, Xingchen Liu, Yanjun Wu, Fengliang Dong, Cheng-Wei Qiu, et al. Artificial intelligence: A powerful paradigm for scientific research. *The Innovation*, 2(4):100179, 2021. [1](#)
- [42] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas S. Huang. Unitbox: An advanced object detection network. In *MM*, pages 516–520, 2016. [4](#)
- [43] Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, and Fang Wen. Cross-domain correspondence learning for exemplar-based image translation. In *CVPR*, pages 5142–5152, 2020. [2](#)
- [44] Ganlong Zhao, Guanbin Li, Ruijia Xu, and Liang Lin. Collaborative training between region proposal localization and classification for domain adaptive object detection. In *ECCV*, volume 12363, pages 86–102, 2020. [6](#), [7](#)
- [45] Yangtao Zheng, Di Huang, Songtao Liu, and Yunhong Wang. Cross-domain object detection through coarse-to-fine feature adaptation. In *CVPR*, pages 13763–13772, 2020. [2](#), [6](#)
- [46] Xinge Zhu, Jiangmiao Pang, Ceyuan Yang, Jianping Shi, and Dahua Lin. Adapting object detectors via selective cross-domain alignment. In *CVPR*, pages 687–696, 2019. [6](#), [7](#)