

Improving Single Domain-Generalized Object Detection: A Focus on Diversification and Alignment

Muhammad Sohail Danish¹ Muhammad Haris Khan¹ Muhammad Akhtar Munir^{1, 2}
M. Saquib Sarfraz^{3, 4} Mohsen Ali²

¹Mohamed bin Zayed University of Artificial Intelligence, ²Information Technology University of Punjab,

³Mercedes-Benz Tech Innovation, ⁴Karlsruhe Institute of Technology

muhammad.sohail@mbzuai.ac.ae, muhammad.haris@mbzuai.ac.ae

Abstract

In this work, we tackle the problem of domain generalization for object detection, specifically focusing on the scenario where only a single source domain is available. We propose an effective approach that involves two key steps: diversifying the source domain and aligning detections based on class prediction confidence and localization. Firstly, we demonstrate that by carefully selecting a set of augmentations, a base detector can outperform existing methods for single domain generalization by a good margin. This highlights the importance of domain diversification in improving the performance of object detectors. Secondly, we introduce a method to align detections from multiple views, considering both classification and localization outputs. This alignment procedure leads to better generalized and well-calibrated object detector models, which are crucial for accurate decision-making in safety-critical applications. Our approach is detector-agnostic and can be seamlessly applied to both single-stage and two-stage detectors. To validate the effectiveness of our proposed methods, we conduct extensive experiments and ablations on challenging domain-shift scenarios. The results consistently demonstrate the superiority of our approach compared to existing methods. Our code and models are available at: <https://github.com/msohaildanish/DivAlign>.

1. Introduction

In recent years, we have witnessed remarkable performance improvements in supervised object detection [39, 47, 48, 53]. The success of these methods rely on the assumption that the training and testing data are sampled from the same distribution. However, in many real-world applications, such as autonomous driving, this assumption is often violated and these object detectors usually show degraded performance due to a phenomenon called domain-

shift [3, 5, 49]. Shifts in real-world domains are typically caused by environmental alterations, like varying weather and time conditions. These changes manifest in diverse contrasts, brightness levels, and textures among others.

A prominent line of research that attempts to alleviate the impact of domain-shift is known as unsupervised domain adaptation (UDA) [20, 24, 42, 49, 63, 68, 72]. Given the labeled data from the source domain and unlabelled data from the target domain, the aim of UDA methods is to align the source and target data distributions so that the trained model can generalize well to the target domain [67]. An obvious limitation of UDA methods is that they require pre-collecting data and re-training the model for different target domains. Collecting data, even without annotation, for all possible domain shifts and training the model when shift happens is difficult and sometime not possible.

To cope with the domain-shift problem, a more realistic albeit challenging problem is domain generalization [2, 11, 23, 26, 32, 34, 35, 40, 58]. The goal of *domain generalization* (DG) is to learn a generalizable model typically from multiple source domains, available during training, that can perform well on the unseen target domain. A dominant class of DG methods attempts to learn domain-invariant feature space across multiple source domains [1, 11, 33, 35]. The performance of these methods is sensitive to the number of available source domains [7, 70]. In many realistic scenarios, acquiring labeled data from multiple source domains is often costly and time-consuming, eventually restricting the potential utilization of such methods. For these reasons, generalizing from a single-source domain is a more practical setting.

Recent surveys [29, 70] reveal that very little work has been done on DG for object detection [38, 60], despite the fact that object detectors occupy an important position in many safety-critical and security-sensitive applications e.g., autonomous driving. An object detector should be able to provide accurate as well as calibrated detections in different

out-domain scenarios. Towards this end, we study the problem of single-domain generalized object detection (Single-DGOD). In the Single-DGOD setting, a single source domain is available for training, and the goal is to learn an object detector that can generalize well to multiple unseen target domains. Note that, many existing DG methods cannot be applied to solve this task since both multiple source domains and domain-level annotations are unavailable.

We are inspired by the class of DG methods for classification [46, 55, 66] that show that simulating novel domains during model training allows segregating domain-specific and semantics-oriented features. Consequently, the spurious correlations between the input and model predictions are suppressed. We begin by leveraging this idea of augmenting our input examples to realize the simulation of novel domains, thereby increasing the diversity of single-source training domain. In particular, we use the common off-the-shelf visual corruptions coupled with a simple training scheme to build a strong Single-DGOD baseline. The augmentation strategy aims to disrupt surface-level statistical patterns that are specific to a particular domain while preserving high-level semantic concepts that are common across domains. In addition to diversifying the single source domain, we propose a novel approach that aligns detection across different views of the image. The introduced method aligns both the class prediction distributions and predicted bounding box coordinates between augmented views of the image. By ensuring such a consistency of detections, we show that this not only helps in better out-domain detections but importantly also results in better calibrated models.

To summarize, our work makes two main contributions:

- We show that, using a set of carefully selected augmentations, the base model can surpass the performance of state-of-the-art published single-DGOD methods. To our knowledge, achieving such a strong baseline with known augmentation methods has not been shown before in the single domain object detection generalization.
- Our second contribution is an effective classification and localization alignment method that is shown to further improve detections in unseen domain and results in implicitly well calibrated detector model.

Our approach is detector-agnostic and can be seamlessly applied to both single-stage and two-stage object detectors.

2. Related Work

Domain adaptive object detection: Several domain adaptive detection methods aim to reduce the pixel-level or feature-level gap between the source and target domain [36, 44]. For instance, the works of [3, 10, 21, 24, 52, 65] utilized the labeled images generated from CycleGAN [71] to achieve pixel-level consistency. Likewise, [3, 5, 10, 19, 49, 51] proposed various mechanisms to progressively obtain a feature-level consistency. A common component

of these methods is domain adversarial network [15] for achieving feature alignment. The work of [5] performed domain alignment on both global features (backbone features) and local features (Instance-level RoI features). Since then, several methods aimed to further improve these alignments by leveraging multi-scale [19], contextual [3, 49], spatial attention [31], category attention [57], and cross-domain topological relations [4] information. A few methods resorted to pseudo-labelling of instances to include class-level discriminative information. For instance, [24] constructed pseudo-labels by putting a threshold on top-1 accuracy over the predicted class of a detection. Towards refining the noisy pseudo-labels, [27] developed a robust training mechanism coupled with an independent classification module. With a similar goal, [28] proposed a weak self-training method and an adversarial background regularization module. Recently, [42] developed an uncertainty-guided criterion for selecting pseudo-labels in the target domain. We note that almost all of the aforementioned methods require access to both source and (unlabelled) target domain data coupled with an adaptation step to perform well in the target domain. As a consequence, these methods cannot be used to solve single-domain generalized object detection.

Single source domain generalization: A more challenging and recently proposed setting in DG is single source domain generalization. It is a more practical setting compared to multi-source DG since the availability of data from multiple different domains is a difficult requirement to fulfill for many real-world applications. The goal is to learn a generalized model using only a single source domain that can perform well on many unseen target domains. An intuitive approach for this setting is to increase the diversity of single source data by proposing different data augmentation techniques. Qiao *et al.* [46] leveraged adversarial training to construct difficult examples which facilitate model generalization. Earlier attempts involve training a label classifier and a domain classifier in a joint manner by their corresponding perturbations. Similarly, [50] proposed to impose Wasserstein constraint in the feature space to generate adversarial samples from synthetic target distribution [56]. Recently, Wu *et al.* [60] proposed a cyclic-disentangled self-distillation approach for single domain-generalized object detection. Wang *et al.* [59] proposed a style-complement module that synthesizes examples from domains complementary to the source domain. Towards exploring the impact of the normalization layer on DG performance, Wang *et al.* [14] developed a generic normalization approach that learns both the standardization and rescaling statistics by neural networks. The aforementioned methods have shown promising results for the image classification task, however, they are not readily applicable to object detection since it has both classification and localization.

Domain diversification for domain generalization: Sev-




Figure 1. Overall architecture of our proposed method. At the core is a baseline detector, Here a two-stage detector Faster-RCNN[48] is depicted, comprising of backbone, region proposal network (RPN), and ROI alignment (RA). To improve the single domain generalization of the baseline detector, we propose to diversify the single source domain and also align the diversified views by minimizing losses at both classification and regression outputs.

eral DG methods for object recognition have proposed various data augmentation techniques in an effort to improve the diversity of source domain(s). By using a CNN generator with different losses, Zhou *et al.* [69] synthesized novel examples from simulated domains. A few methods attempt to suppress the intrinsic style bias of CNNs. For instance, to focus on the contents and styles of images, [43] proposed content-biased and style-biased networks, respectively. To constantly seek diverse styles, [25] proposed a style synthesis process that is formulated as a monotone submodular optimization. Among the more recent methods, [13] explores diversifying source domain by perturbing channel statistics of low-level features for single domain-generalized object detector. In [54] the authors introduced the semantic augmentation of images by utilizing text-prompts relevant to potential target domain concepts in a vision-language model framework to improve the robustness of object detectors in unseen target domains. Along similar lines, [12] adapted the source model using textual descriptions of the target domain in a zero-shot manner to enhance the cross-domain generalizability in object detectors. Note that, both [54] and [12] incorporate the prior target domain knowledge using textual prompts. However, we follow domain generalization settings and leverage no prior knowledge about the target domain. Further, we diversify the source domain in input space and then align the detections in both original and diversified images.

3. Proposed Methodology

3.1. Preliminaries

Problem Settings: Let \mathcal{D}_s be the single source domain containing N^s labelled training examples $\{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{i=1}^{N^s}$, where $\mathbf{x}_i \in \mathbb{R}^{H \times W \times C}$ is an image and $\mathbf{y}_i = \{k_i, \mathbf{b}_i\}$ is the corresponding label with bounding box coordinates $\mathbf{b}_i \in \mathbb{R}^4$ and the associated class category $k_i \in \{1, \dots, K\}$. K is the number of class categories and H , W and C represent the

height, width, and number of channels, respectively. Let $\{\mathcal{D}_t\}_{t=1}^T$ be the set of T (unseen) target domains. Our goal is to learn an object detector using training examples from \mathcal{D}_s that generalizes well on test examples from \mathcal{D}_t . We assume that both \mathcal{D}_s and \mathcal{D}_t share the same label space.

Object Detection: Let \mathcal{F}_{det} be an object detection model, that takes input image \mathbf{x} and outputs Z detections $\{(P(\hat{k}_n|\mathbf{x}), \hat{\mathbf{b}}_n)\}_{n=1}^Z$, where \hat{k}_n denotes the class and bounding box $\hat{\mathbf{b}}_n$ predictions for n^{th} proposal, respectively. We take the Faster R-CNN [48] model as \mathcal{F}_{det} to detail our method, but note that our method is detector-agnostic and can be easily applied to other object detection paradigms, as shown in our results. We first overview Faster R-CNN to make the method section self-contained. Faster R-CNN is a popular and widely used two-stage object detector. The first stage comprises of a backbone feature extractor layer which generates a feature map and a region proposal network (RPN) which predicts a set of object proposals or regions of interest (RoIs) from the feature map. The second stage begins with ROI-Alignment, which extracts a fixed-size feature representation for each ROI and then feeds it to a classifier and a regressor to predict class confidence scores and bounding box coordinates, respectively.

Single domain-generalized object detection: In this case, we assume that the object detection \mathcal{F}_{det} has been trained on the examples taken from the source domain \mathcal{D}_s and then tested on examples from $\{\mathcal{D}_t\}_{t=1}^T$. Akin to image classification models, object detection models like Faster-RCNN, which exhibits state-of-the-art performance when testing data and training data are from the same domain (i.e. in-domain scenario), but suffers from performance degradation when domains are different (i.e. out-of-domain scenario) (Figure 4). The presence of multiple domains during training penalizes the object detector in learning shortcuts which are typically domain-specific and instead encourages learning domain-invariant features. Nevertheless, the multi-domain data with the ground truth are difficult requirements

to fulfill in several realistic settings and thus we are reduced to a single-domain generalization problem.

Domain invariant object detection: We first re-formulate the definition of the domain invariance for a single source domain-generalized object detector by including a classifier and a regressor. Following [61], let’s consider a visual corruption function $\phi(\cdot)$ as a domain transformation function, in (continuous) domain space \mathcal{D} , which takes a source image from \mathcal{D}_s and transforms it into a different domain \mathcal{D}_ϕ where $\phi \sim \Phi$ (is a set of possible transformations). By making this assumption, we define domain-invariance for single domain-generalized object detector as follows:

Definition 3.1 (Domain Invariance for Object Detection)

Assuming that, for an input image \mathbf{x} , an object detection model \mathcal{F}_{det} predicts class probability distribution $\hat{\mathbf{p}}_n$ and bounding box coordinates $\hat{\mathbf{b}}_n \in \mathbb{R}^4$ for the n^{th} proposal. Let \mathbf{x}^s be an image from \mathcal{D}_s and $\mathbf{x}^\phi = \phi(\mathbf{x}^s)$ be the transformation of \mathbf{x}^s , denoted as \mathbf{x}^ϕ , where $\phi \sim \Phi$. The model \mathcal{F}_{det} is domain invariant if:

$$\hat{\mathbf{p}}_n^s = \hat{\mathbf{p}}_n^\phi \tag{1}$$

$$1 - \text{IoU}(\hat{\mathbf{b}}_n^s, \hat{\mathbf{b}}_n^\phi) = 0 \tag{2}$$

3.2. Proposed Method

Through developing realizations of Eq. (1) and Eq. (2), we intend to improve the domain generalization of object detectors, when the training data from only a single source domain is available. So, our method has two main components: 1) diversification of the source domain (sec. 3.2.1) and 2) a detection alignment mechanism for improving generalization and calibration of detectors (sec. 3.2.2). Fig. 1 displays the overall architecture of our proposed method.

3.2.1 Diversifying Single Source Domain

Our proposed approach is motivated by various techniques used in diversifying source domains for object classification [26, 50, 56, 69]. Deep convolutional networks often learn shortcuts instead of actual semantics, but diversifying training domains can reduce this problem. Since obtaining multiple domains is difficult, we simulate them through data augmentation. This introduces diversity in the training data and improves the network’s ability to learn robust object representations. We achieve this by perturbing the single source domain images using a set of existing visual corruptions Φ . During training, we augment every image in mini-batch by using image augmentation function $\phi(\cdot)$ which randomly samples a corruption method from Φ and apply to the input image. We then pass both original clean images along with their augmented versions to the model. Specifically, we investigate the visual corruptions set used in image

classification domain [9] that simply combines ImageNet-C with Fourier transform-based corruptions, thereby providing a total of 22 corruptions with 5 severity levels.

We classify image corruptions used for augmentation into five groups: Weather, Noise, Blur, Digital, and Fourier. In the case of Blur corruption, we employ various blur functions such as Gaussian, motion, glass, and defocus to smooth the pixel values in the image. Noise introduces different types of noise, including shot, impulse, speckle, and Gaussian, to perturb the image. Digital corruptions modify the pixel intensities by adjusting brightness, contrast, and saturation, or by changing the image resolution using JPEG compression, pixelation, and elastic transformation. Fig. 2 shows the impact of such augmentations. To ensure a fair comparison, we exclude weather-based transformations from our augmentation pool. Fourier-based corruptions such as phase scaling, constant amplitude, and High Pass Filter exploit the well-known property of Fourier transformation in that the phase component represents high-level semantics of an image, while the amplitude component covers low-level statistics or texture information. Perturbing these components enables us to reconstruct a new transformed image. Phase scaling introduces various artifacts by scaling the phase component in the Fourier spectrum. Constant amplitude replaces the amplitude component with a fixed vector, resulting in the loss of textual information and retaining only high-level semantics. We exclude the constant amplitude augmentation as it destroys the instance-level features. We note that this simple technique of diversifying source domain during object detector training outperforms baselines and competing methods by visible margins in challenging out-of-domain scenarios. As a result, though simple, as it makes minimal modifications to the baseline, it provides a solid baseline in single domain-generalized object detection.

3.2.2 Aligning Diversified Domains

While the single-source domain diversification approach provides a strong baseline, it does not enforce the main requirement for the Domain invariance (Def. 3.1); that is the predictions on diversified images x^ϕ do not align with predictions on original images x^s . To quantify this, we measured mAP of the model predictions over original images, and mAP of the same model’s predictions on diversified images. Regardless of whichever diversification is performed, there is considerable performance gap, indicating the lack of generalization. This is depicted in Fig. 3. This gap can be explained when considering the following two constraints. (1.) *Object Classification Constraint*, Eq.(1), states that model’s predicted class probability distribution for the object in the source image, x^s , should be same for the corresponding object in diversified image x^ϕ . (2.) *Object Local-*

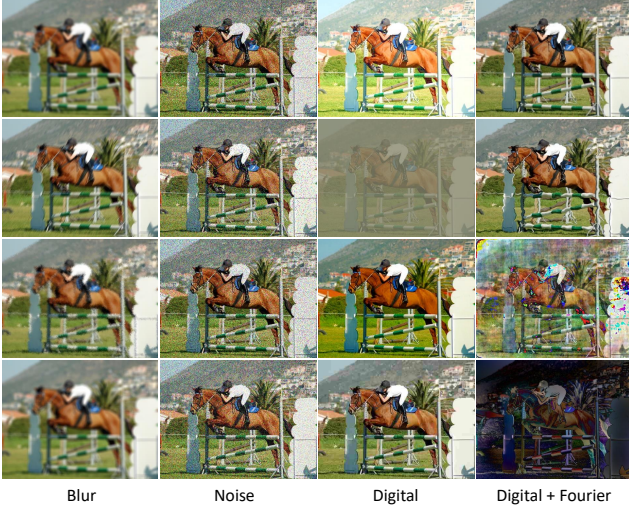


Figure 2. Examples of augmentations for domain diversification.

ization Constraint, Eq.(2), states that is the objects should be predicted at same location (and of same size) on both diversified and original version of the images. Note that here these constraints are not related to the accuracy of the prediction, but that any prediction made should be same.

Unlike the image classification, where there is only one class prediction vector, enforcing such constraints in object detection is more challenging. We propose a strategy to overcome this challenge. Based on definition 3.1, on domain-invariance for object detection, we align the detections from two views i.e. clean and diversified (augmented) images both from the classification and the regression outputs. We take the widely used two-stage object detector as an example i.e Faster R-CNN. Given \mathbf{x}_i^s or \mathbf{x}_i^ϕ as input, we can obtain a feature map $F \in \mathbb{R}^{m \times w \times h}$ from the backbone, where m , h , and w denote the number of channels, height and width of the feature map, respectively. This feature map is then passed through region proposal network (RPN) to obtain object proposals $O \in \mathbb{R}^{Z \times 4}$, where Z is the number of proposals. Next, RoI-Alignment (RA) is performed using O i.e $A = \text{RA}(O, F) \in \mathbb{R}^{Z \times m}$ to obtain a fixed length feature representation for each proposal. Detection loss is given by:

$$\mathcal{L}_{\text{det}} = \sum_{n=1}^Z L_{\text{det}}(\Psi(A_n), \mathbf{y}_n). \quad (3)$$

Here Ψ indicates the classifier and a bounding box regressor and \mathbf{y}_n is the ground truth including class label and bounding box coordinates for the n^{th} proposal. The L_{det} consists of smooth-L1 loss and the negative log-likelihood loss.

Aligning classification: To enforce *Object Classification constraint for object detector generalization*, we propose to measure the Kullback-Leibler (KL) divergence between the classifier output probability vectors $\hat{\mathbf{p}}_n^s$ and $\hat{\mathbf{p}}_n^\phi$ correspond-

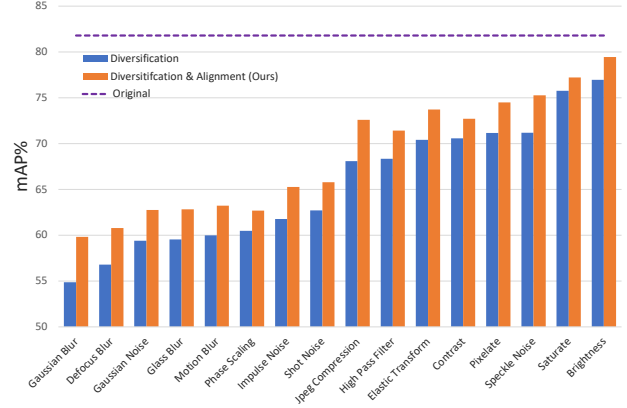


Figure 3. We augment the validation set from the source domain by one augmentation at a time and report the performance of the strong baseline model trained on all these augmentations. There is noticeable gap between the performance on original and diversified images. Our alignment losses allows reducing the gap between the performance on original and diversified images.

ing to \mathbf{x}_i^s and \mathbf{x}_i^ϕ . It is likely that RPN will produce a different set of proposals from \mathbf{x}_i^s and \mathbf{x}_i^ϕ i.e. $O^s \neq O^\phi$. These proposals might not be spatially consistent. To achieve 1-1 correspondence between the proposals from \mathbf{x}_i^s and \mathbf{x}_i^ϕ , we perform predictions using the region proposals from the \mathbf{x}_i^s and features pooled from the feature map for \mathbf{x}_i^ϕ . That is we pass the feature map F^ϕ corresponding to \mathbf{x}_i^ϕ with proposals from O^s : $A^{\phi'} = \text{RA}(O^s, F^\phi)$. O^s : $A^{\phi'} = \text{RA}(O^s, F^\phi)$, obtaining $\hat{\mathbf{p}}_n^{\phi'}$. The difference between the predictions is minimized using the KL-divergence:

$$\mathcal{L}_{\text{cal}} = \sum_{n=1}^Z \text{KL}(\hat{\mathbf{p}}_n^s \parallel \hat{\mathbf{p}}_n^{\phi'}). \quad (4)$$

Aligning localization: To make the *Object Localization constraint for generalization* part of the training process, we maximize the IoU between $\hat{\mathbf{b}}_n^s$ and $\hat{\mathbf{b}}_n^{\phi'}$. This is realized by minimizing the L2-squared norm between $\hat{\mathbf{b}}_n^s$ and $\hat{\mathbf{b}}_n^{\phi'}$:

$$\mathcal{L}_{\text{ral}} = \|\hat{\mathbf{b}}_n^s - \hat{\mathbf{b}}_n^{\phi'}\|_2^2 \quad (5)$$

Our overall training objective is given by:

$$\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{det}} + \alpha \mathcal{L}_{\text{cal}} + \beta \mathcal{L}_{\text{ral}} \quad (6)$$

where α and β are the hyperparameters for balancing the contributions of \mathcal{L}_{cal} and \mathcal{L}_{ral} .

4. Experiments

Datasets: Real to Artistic consists of four datasets, including the Pascal VOC which contains real images from 20 different object classes and acts as a source domain, and remaining three are Clipart1k, Watercolor2k and Comic2k

which offer artistic and comic images and are used as target domains. Clipart1k shares the same 20 classes as Pascal VOC, while Watercolor2k and Comic2k consist of 6 classes each, which are subsets of the Pascal VOC classes. Following [24], we combine Pascal VOC2007 and VOC2012 trainval set resulting in 16551 images for training and 5K images from VOC2007 test for in-domain evaluation. The Clipart1k consists of 1K and Watercolor2k and Comic2k each contain 2K images which are equally split into training and test sets. Following [24], we select all images from Clipart1k and test sets from Watercolor2k and Comic2k for out-domain evaluations. **Urban Scene Detection** is a self-driving dataset introduced by [60]. It offers images of 5 different weather conditions including Daytime Sunny, Night Clear, Night Rainy, Dusk Rainy and Daytime Foggy. The images are collected from three different datasets: Berkeley Deep Drive 100k (BDD-100k)[64], FoggyCityscapes[8] and Adverse-Weather[17]. Rainy scenes are synthetic images generated by simulating rain on BDD-100k images. Daytime Sunny serves as a source domain consisting 26,518 images out of which 8,313 are used for in-domain evaluation whereas Night Clear, Night Rainy, Dusk Rainy and Daytime Foggy are used as target domains containing 26,158, 2,494, 3,501, and 3,775 images, respectively.

Implementation and Training Details: For two-stage object detector as a baseline, we use Faster R-CNN[48] with ResNet101[18] as feature backbone, while for single-stage detector, we use FCOS [53] with ResNet50-FPN[37] as feature backbone. In all experiments our models are trained using SGD for 18k iterations with initial learning rate of 0.2 and dropped by a factor of 10 after 12K and 16k iterations. The batch size is set to 16. All results are reported using mean average precision (mAP) at IoU@0.5. For reporting calibration performance, we use the detection expected calibration error (D-ECE) [30] and reliability diagrams [16].

Real to artistic: We evaluate our method on challenging real to artistic domain-shift where the distribution shift is large i.e. from PASCAL VOC to Clipart1k, Watercolor2k, and Comic2k. Table 1 shows that diversifying the training domain helps the model to generalize quite well by obtaining a significant gain of 8.5%, 8.5%, and 5.3% from the Faster R-CNN baseline in Clipart1k, Watercolor2k, and Comic2k, respectively. With our proposed alignment losses, we are able to boost the overall performance by 13.2%, 12.9%, and 14.3%. We also evaluated Normalization Perturbation (NP) [13] method on this benchmark and found it inferior to our method. Tables 2 and 3 reveal the class-wise results in Watercolor2k and Comic2k shifts. See supplementary for class-wise results on Clipart1k.

Urban scene detection: Table 4 shows the results on different weather conditions manifesting various shifts scenarios. Due to large shifts like Daytime-Foggy and Night Raining the objects become unclear. We note that,

Method	VOC	Clipart	Watercolor	Comic
Faster R-CNN	81.8	25.7	44.5	18.9
NP [13]	79.2	35.4	53.3	28.9
Diversification (div.)	82.1	34.2	53.0	24.2
div. + \mathcal{L}_{cal}	82.1	36.2	53.9	28.7
div. + \mathcal{L}_{ral}	80.7	35.0	53.8	28.7
div. + \mathcal{L}_{cal} + \mathcal{L}_{ral} (Ours)	80.1	38.9	57.4	33.2

Table 1. Performance comparison with baseline and possible ablations, mAP@0.5(%) reported. The model is trained on Pascal VOC and tested on Clipart1k, Watercolor2k and Comic2k.

Method	bike	bird	car	cat	dog	person	mAP
Faster R-CNN	85.7	42.5	36.4	29.0	18.7	54.5	44.5
Diversification (div.)	87.1	51.7	53.6	35.1	23.6	63.6	52.5
div. + \mathcal{L}_{cal}	83.9	50.4	48.5	40.3	36.5	63.3	53.8
div. + \mathcal{L}_{ral}	83.3	52.3	50.7	39.2	32.7	65.2	53.9
div. + \mathcal{L}_{cal} + \mathcal{L}_{ral} (Ours)	90.4	51.8	51.9	43.9	35.9	70.2	57.4

Table 2. Class-wise AP(%) comparison of baseline and proposed method on Pascal VOC to Watercolor2k scenario.

Method	bike	bird	car	cat	dog	person	mAP
Faster R-CNN	39.7	9.1	23.9	9.1	9.1	22.2	18.9
Diversification (div.)	41.7	12.3	29.0	13.2	20.6	36.5	25.5
div. + \mathcal{L}_{cal}	49.0	15.8	27.0	20.2	21.8	38.4	28.7
div. + \mathcal{L}_{ral}	53.3	13.5	29.0	19.1	20.5	36.7	28.7
div. + \mathcal{L}_{cal} + \mathcal{L}_{ral} (Ours)	54.1	16.9	30.1	25.0	27.4	45.9	33.2

Table 3. Class-wise AP(%) comparison of baseline and proposed method on Pascal VOC to Comic2k scenario.

by only diversifying single domain, we obtain 0.5%, 7%, 6.3%, and 2.5% on Night Clear (NC), Dusk Rainy (DR), Night Rainy (NR), and Daytime Foggy (DF), respectively. Our method beats all baselines and state-of-the-art method from [54]. Over the baseline, our method delivers the highest gain of 8.4% on Night Rainy while 8.1%, 4.1% and 3.6% gains on Dusk Rainy, Daytime Foggy and Night Clear respectively.

Method	DS	NC	DR	NR	DF
Faster R-CNN	51.8	38.9	30.0	15.7	33.1
SW[45]	50.6	33.4	26.3	13.7	30.8
IBN-Net[62]	49.7	32.1	26.1	14.3	29.6
IterNorm[22]	43.9	29.6	22.8	12.6	28.4
ISW[6]	51.3	33.2	25.9	14.1	31.8
Wu <i>et al.</i> [60]	56.1	36.6	28.2	16.6	33.5
Vidit <i>et al.</i> [54]	51.3	36.9	32.3	18.7	38.5
Diversification	50.6	39.4	37.0	22.0	35.6
Ours	52.8	42.5	38.1	24.1	37.2

Table 4. Results, mAP@0.5(%) reported on a multi-weather scenario where the model is trained on Daytime Sunny (DS) and tested on Night-Clear (NC), Night-Rainy (NR), Dusk-Rainy (DR) and Daytime-Foggy (DF). The results for SW, IBN-Net, IterNorm, ISW, S-DGOD are directly taken from [60].

Model Calibration: Table 5, reports the calibration performance (D-ECE) of the base model and the impact of us-




Figure 4. Our method (diversification and alignment) results in both considerable improvement in domain generalization and out-of-domain calibration. Diversification with the Label Smoothing (LS) or Temperature Scaling (TS) improves calibration but overall lower mAP indicates lacking in generalization. Note that our method does not have an explicit model calibration mechanism. (left) mAP: higher the better, (right) D-ECE: lower the better.

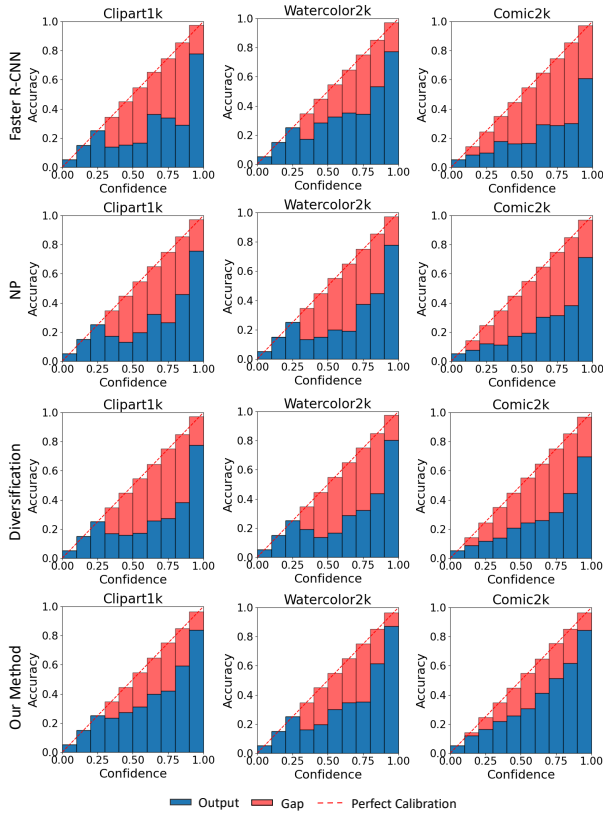


Figure 5. Reliability Diagram for different target domains.

ing our diversification and alignment approach under various shift scenarios. As seen, simple diversification of the domain, while improving generalization, negatively affects the calibration for both in-domain and out-domain scenarios. This is to be expected since model is trained using Negative log-likelihood loss in \mathcal{L}_{det} (Eq. 3) which is in general attributed to making models highly confident, resulting in

poor calibration. On the diversified images, where information might have been reduced or content has become noisy (w.r.t original image), high confidence predictions affects the calibration. In contrast, as detailed in (Sec. 3.2.2), our alignment step implicitly reduces the overconfidence in predictions (Fig. 5). Overall, as our results show, we achieve not only better generalization (higher mAP) but also better calibration (lower D-ECE). To further make apparent the the

Method	Artistic Shifts			Urban Scene			
	Clipart	Watercolor	Comic	NR	DR	NC	DF
Faster R-CNN	11.9	18.5	15.4	31.5	29.3	27.9	25.8
Diversification (div.)	14.5	21.4	17.4	33.0	30.2	28.9	25.7
Ours	10.7	14.4	14.3	29.3	24.9	15.8	20.6

Table 5. Comparison of calibration performance using D-ECE metric (%) on Real to artistic shifts and in urban scene detection.

utility of our alignment approach in achieving better calibration, We also include the impact of using existing simple calibration and generalization techniques such as temperature scaling (TS) [16] or label smoothing (LS) [41] with our diversification approach. Fig. 4 depicts the improvements our presented methods achieve in multiple shift scenarios.

Results with single-stage detector: To show the applicability of our proposed method, we instantiate it in one-stage detector, namely FCOS [53] (Table 6). In comparison to FCOS, our method delivers a significant gain of 13.0%, 10.7% and 15.8% on Clipart1k, Watercolor2k, and Comic2k shifts, respectively.

Comparison with domain-adaptive detection methods: Table 7 compares our method with some unsupervised domain-adaptive object detectors using the same backbone. Note that these methods assume access to target domain and directly train using unlabelled target samples. In contrast even though our method does not have access to the target domain dataset at training time, it still shows better generalization results than many domain-adaptive detection methods. For instance, in watercolor and comic

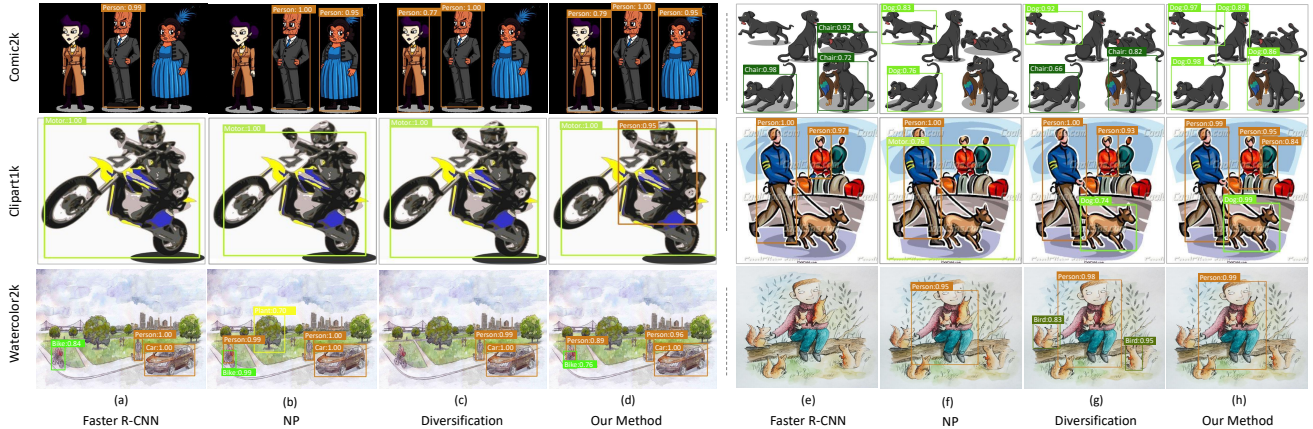


Figure 6. Qualitative results of baseline (Faster-RCNN), only diversifying domain, and our method.

Method	VOC	Clipart	Watercolor	Comic
FCOS	78.1	24.4	44.3	15.4
Diversification (div.)	79.6	31.7	48.8	25.2
div. + \mathcal{L}_{cal}	80.1	35.4	52.6	29.4
div. + \mathcal{L}_{ral}	77.5	29.8	50.3	24.0
div. + \mathcal{L}_{cal} + \mathcal{L}_{ral} (Ours)	77.5	37.4	55.0	31.2

Table 6. Performance comparison with single-stage baseline, mAP@0.5(%) reported. The model is trained on Pascal VOC and tested on Clipart1k, Watercolor2k and Comic2k.

shifts, our model outperforms all methods by delivering a gain of 3.2% and 2.5%, respectively. Results for the existing domain adaptive methods are taken from [4].

Method	Clipart	Watercolor	Comic
DA-Faster[5]	19.8	46.0	-
SWDA[49]	38.1	53.3	27.4
HTCN[3]	40.3	-	-
DBGL[4]	41.6	53.8	29.7
Our Method	38.5	57.0	32.2

Table 7. Comparison with domain adaptive detectors in mAP@0.5(%)

Qualitative Results: We sample some results for the qualitative comparison in Fig. 6. Aligning the predicted confidence and predicted localization between the original and diversified image, results in improved generalizability as visible from columns (c & f) where objects missed by baseline, diversification and NP[13] method, were properly localized and classified. Note that diversification actually results in more false positive detections with high confidence than baseline. This is visible in the column g where objects have been falsely classified as chairs and birds.

Effect of different types of augmentations: We analyze the impact of different types of visual corruptions obtained from Φ in our method. First, we only diversify VOC-dataset using ImageNet-C corruption and then with Fourier-based

Method	Clipart	Watercolor	Comic
Faster R-CNN	25.7	44.5	18.9
Ours - w/Fourier-based	36.1	54.8	29.1
Ours - w/ImageNet-C	36.8	55.5	30.6
Ours	38.9	57.4	33.2

Table 8. Ablation study with different augmentations used for diversifying source domain in our method, mAP@0.5(%) reported.

corruption. As reported in Table 8, ImageNet-C corruption has a slight edge over Fourier corruption, however, using both provides the best results. This suggests that obtaining diversity through both types of corruption is complementary. **Aligning Classification vs aligning localization:** Not aligning one of the prediction tasks across the diversified domains results in considerable degradation of the results (Table 1). In general, the results indicate that, as expected, the alignment of both these tasks is equally vital for achieving cross-domain generalizable object detectors.

5. Conclusion

We proposed a detector-agnostic, simple, and effective approach to improve the generalization and calibration of the object detectors given only a single source domain for training. To this end, we first explored effective image augmentation methods for domain diversification and developed a strong baseline for single-DGOD. Our results show that by carefully selecting a set of augmentations, a base detector can outperform SOTA methods for single domain generalization by a good margin. Secondly, we proposed a novel method to align the detections across the original and diversified images. This alignment procedure leads to better generalized and well-calibrated object detectors, which are crucial for accurate decision-making in safety-critical applications. Extensive results on various challenging shift scenarios corroborate the effectiveness and applicability of our approach.

References

- [1] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chelappa. Metareg: Towards domain generalization using meta-regularization. In *Advances in Neural Information Processing Systems*, pages 998–1008, 2018. 1
- [2] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 1
- [3] 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. 1, 2, 8
- [4] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2703–2712, 2021. 2, 8
- [5] 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. 1, 2, 8
- [6] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. pages 11575–11585, 2021. 6
- [7] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11580–11590, 2021. 1
- [8] 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 *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [9] Ilke Cugu, Massimiliano Mancini, Yanbei Chen, and Zeynep Akata. Attention consistency on visual corruptions for single-source domain generalization. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4164–4173, 2022. 4
- [10] Jinhong Deng, Wen Li, Yuhua Chen, and Lixin Duan. Unbiased mean teacher for cross-domain object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4091–4101, 2021. 2
- [11] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *Advances in Neural Information Processing Systems*, pages 6450–6461, 2019. 1
- [12] Mohammad Fahes, Tuan-Hung Vu, Andrei Bursuc, Patrick Pérez, and Raoul de Charette. Pøda: Prompt-driven zero-shot domain adaptation. In *ICCV*, 2023. 3
- [13] Qi Fan, Mattia Segu, Yu-Wing Tai, Fisher Yu, Chi-Keung Tang, Bernt Schiele, and Dengxin Dai. Towards robust object detection invariant to real-world domain shifts. In *The Eleventh International Conference on Learning Representations*, 2023. 3, 6, 8
- [14] Xinjie Fan, Qifei Wang, Junjie Ke, Feng Yang, Boqing Gong, and Mingyuan Zhou. Adversarially adaptive normalization for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8208–8217, 2021. 2
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016. 2
- [16] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017. 6, 7
- [17] M. Hassaballah, Mourad A. Kenk, Khan Muhammad, and Shervin Minaee. Vehicle detection and tracking in adverse weather using a deep learning framework. *IEEE Transactions on Intelligent Transportation Systems*, 22(7):4230–4242, 2021. 6
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [19] Zhenwei He and Lei Zhang. Multi-adversarial faster-rcnn for unrestricted object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6668–6677, 2019. 2
- [20] 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 *European Conference on Computer Vision*, pages 733–748. Springer, 2020. 1
- [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 *The IEEE Winter Conference on Applications of Computer Vision*, pages 749–757, 2020. 2
- [22] Lei Huang, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. Iterative normalization: Beyond standardization towards efficient whitening. pages 4869–4878, 2019. 6
- [23] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. 2020. 1
- [24] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5001–5009, 2018. 1, 2, 6
- [25] Juwon Kang, Sohyun Lee, Namyup Kim, and Suha Kwak. Style neophile: Constantly seeking novel styles for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7130–7140, 2022. 3

- [26] Muhammad Haris Khan, Talha Zaidi, Salman Khan, and Fahad Shehbaz Khan. Mode-guided feature augmentation for domain generalization. 2021. 1, 4
- [27] Mehran Khodabandeh, Arash Vahdat, Mani Ranjbar, and William G Macready. A robust learning approach to domain adaptive object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 480–490, 2019. 2
- [28] Seunghyeon Kim, Jaehoon Choi, Taekyung Kim, and Chang-ick Kim. Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6092–6101, 2019. 2
- [29] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. 1
- [30] Fabian Küppers, Jan Kronenberger, Amirhossein Shantia, and Anselm Haselhoff. Multivariate confidence calibration for object detection. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020. 6
- [31] Congcong Li, Dawei Du, Libo Zhang, Longyin Wen, Tiejian Luo, Yanjun Wu, and Pengfei Zhu. Spatial attention pyramid network for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 481–497. Springer, 2020. 2
- [32] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 1
- [33] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 1
- [34] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *ICCV*, 2019. 1
- [35] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 1
- [36] Wanyi Li, Fuyu Li, Yongkang Luo, Peng Wang, et al. Deep domain adaptive object detection: a survey. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1808–1813. IEEE, 2020. 2
- [37] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 936–944, 2017. 6
- [38] Hong Liu, Pinhao Song, and Runwei Ding. Towards domain generalization in underwater object detection. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 1971–1975. IEEE, 2020. 1
- [39] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*. Springer, 2016. 1
- [40] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18, 2013. 1
- [41] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems*, 2019. 7
- [42] Muhammad Akhtar Munir, Muhammad Haris Khan, M Sarfraz, and Mohsen Ali. Ssal: Synergizing between self-training and adversarial learning for domain adaptive object detection. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 2
- [43] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8690–8699, 2021. 3
- [44] 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. 2
- [45] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 6
- [46] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 2
- [47] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 1
- [48] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1, 3, 6
- [49] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6956–6965, 2019. 1, 2, 8
- [50] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 2, 4
- [51] Zhiqiang Shen, Harsh Maheshwari, Weichen Yao, and Marios Savvides. Scl: Towards accurate domain adaptive object detection via gradient detach based stacked complementary losses. *arXiv preprint arXiv:1911.02559*, 2019. 2
- [52] Waqas Sultani, Wajahat Nawaz, Syed Javed, Muhammad Sohail Danish, Asma Saadia, and Mohsen Ali. Towards low-cost and efficient malaria detection. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20687–20696, 2022. 2
- [53] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 1, 6, 7
- [54] Vidit Vidit, Martin Engilberge, and Mathieu Salzmann. Clip the gap: A single domain generalization approach for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3219–3229, 2023. 3, 6
- [55] Riccardo Volpi and Vittorio Murino. Addressing model vulnerability to distributional shifts over image transformation sets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7980–7989, 2019. 2
- [56] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems*, pages 5334–5344, 2018. 2, 4
- [57] 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. 2
- [58] Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. 2020. 1
- [59] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to diversify for single domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 834–843, 2021. 2
- [60] Aming Wu and Cheng Deng. Single-domain generalized object detection in urban scene via cyclic-disentangled self-distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 847–856, 2022. 1, 2, 6
- [61] Zehao Xiao, Jiayi Shen, Xiantong Zhen, Ling Shao, and Cees Snoek. A bit more bayesian: Domain-invariant learning with uncertainty. In *Proceedings of the 38th International Conference on Machine Learning*, pages 11351–11361. PMLR, 2021. 4
- [62] Pan Xingang, Luo Ping, Shi Jianping, and Tang Xiaou. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, 2018. 6
- [63] Minghao Xu, Hang Wang, Bingbing Ni, Qi Tian, and Wenjun Zhang. Cross-domain detection via graph-induced prototype alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12355–12364, 2020. 1
- [64] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2(5):6, 2018. 6
- [65] Yixin Zhang, Zilei Wang, and Yushi Mao. Rpn prototype alignment for domain adaptive object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12425–12434, 2021. 2
- [66] Long Zhao, Ting Liu, Xi Peng, and Dimitris Metaxas. Maximum-entropy adversarial data augmentation for improved generalization and robustness. *Advances in Neural Information Processing Systems*, 33:14435–14447, 2020. 2
- [67] Sicheng Zhao, Xiangyu Yue, Shanghang Zhang, Bo Li, Han Zhao, Bichen Wu, Ravi Krishna, Joseph E Gonzalez, Alberto L Sangiovanni-Vincentelli, Sanjit A Seshia, et al. A review of single-source deep unsupervised visual domain adaptation. *IEEE Transactions on Neural Networks and Learning Systems*, 2020. 1
- [68] Yangtao Zheng, Di Huang, Songtao Liu, and Yunhong Wang. Cross-domain object detection through coarse-to-fine feature adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13766–13775, 2020. 1
- [69] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. 2020. 3, 4
- [70] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1
- [71] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2
- [72] 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. 1