

# CD-FKD: Cross-Domain Feature Knowledge Distillation for Robust Single-Domain Generalization in Object Detection

Junseok Lee<sup>1,\*</sup>, Sungho Shin<sup>2,\*</sup>, Seongju Lee<sup>3</sup>, and Kyoobin Lee<sup>3,†</sup>

**Abstract**—Single-domain generalization is essential for object detection, particularly when training models on a single source domain and evaluating them on unseen target domains. Domain shifts, such as changes in weather, lighting, or scene conditions, pose significant challenges to the generalization ability of existing models. To address this, we propose Cross-Domain Feature Knowledge Distillation (CD-FKD), which enhances the generalization capability of the student network by leveraging both global and instance-wise feature distillation. The proposed method uses diversified data through downscaling and corruption to train the student network, whereas the teacher network receives the original source domain data. The student network mimics the features of the teacher through both global and instance-wise distillation, enabling it to extract object-centric features effectively, even for objects that are difficult to detect owing to corruption. Extensive experiments on challenging scenes demonstrate that CD-FKD outperforms state-of-the-art methods in both target domain generalization and source domain performance, validating its effectiveness in improving object detection robustness to domain shifts. This approach is valuable in real-world applications, like autonomous driving and surveillance, where robust object detection in diverse environments is crucial.

## I. INTRODUCTION

In recent years, deep learning-based object detection technologies have made remarkably advanced, playing a critical role in various visual perception tasks [1]–[4]. These technologies achieve strong performance when the training and testing data share the same domain distribution, but this is rarely the case in real-world scenarios. In fields such as autonomous driving, industrial automation, and video surveillance, environmental changes caused by lighting, weather, or time of day result in domain shifts, which significantly degrade the performance of trained models [5], [6].

To address this, methodologies such as unsupervised domain adaptation (UDA) [7], [8] and domain generalization (DG) [9], [10] have been proposed. UDA aligns data distributions between the source and target domains to ensure model performance. However, it requires access to target domain data, limiting its ability to generalize across environments. DG focuses on training models to maintain performance on unseen

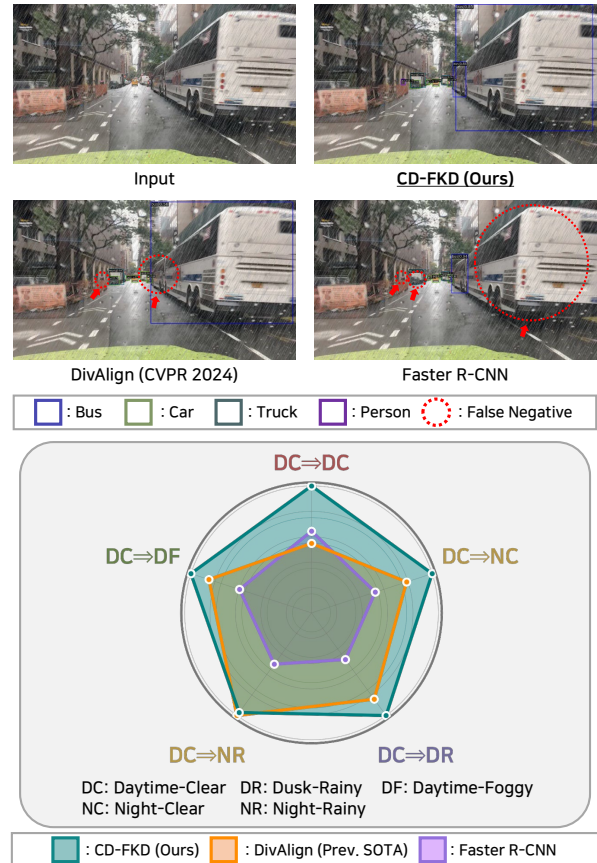


Fig. 1. Overview of results of our proposed CD-FKD. The top panel qualitatively compares our method with DivAlign [11] and Faster R-CNN [12], on an Dusk-Rainy. The bottom panel shows a radar chart comparing relative performance.

target domains by leveraging diverse source domain data, making it a more challenging problem compared with UDA. Traditional DG methods aim to learn shared representations across multiple source domains or enhance data diversity via augmentation. However, these approaches often need multiple source domains, which can be impractical owing to time and cost constraints.

Single-domain generalization (SDG) focuses on improving performance across diverse target domains using only a single source domain, offering a practical solution to the challenges of domain generalization [13], [14]. Unlike domain generalization (DG) methods, which require multiple source domains or domain-level annotations and thus increase complexity, SDG reduces both data requirements and costs.

Traditional approaches typically rely on data augmentation

\* These authors contributed equally to this work.

<sup>1</sup> Junseok Lee is with the Advanced Robotics Lab, LG Electronics, Seoul, Republic of Korea.

<sup>2</sup> Sungho Shin is with the Place AI Platform, Naver, Gyeonggi-do, Republic of Korea

<sup>3</sup> Seongju Lee and Kyoobin Lee are with the Department of AI Convergence, Gwangju Institute of Science and Technology (GIST), Gwangju, Republic of Korea.

† Corresponding author: Kyoobin Lee (E-mail: kyoobinlee@gist.ac.kr)

and feature disentanglement. Feature disentanglement separates domain-invariant (object-centric) features from domain-specific (background-centric) features, using only the invariant features for object detection. However, this neglects the surrounding background, limiting the detector’s ability to understand the full context of the image. Although data augmentation can improve generalization across environments, it has been shown to reduce performance on the source domain [15]. To address these issues, the proposed method enables the detector to extract object-centric features while maintaining a comprehensive understanding of the image’s context. This not only enhances generalization performance on target domains but also improves detection on the source domain.

We propose a novel methodology that effectively operates across both the source and target domains by understanding causal features and learning the global context of images. To achieve this, we introduce **Cross-Domain Feature Knowledge Distillation (CD-FKD)** for single domain generalized object detection. The proposed method comprises two core components: (1) cross-domain knowledge distillation (KD) framework and (2) cross-domain feature distillation loss. The cross-domain KD framework bridges domain gaps between teacher and student networks. The teacher network receives original source domain data, whereas the student network learns from diversified source data with various scales and corruptions.

The primary goal of CD-FKD is to train the network to effectively extract object-centric features even in data with various distortions. The proposed feature distillation process comprises two key steps: (1) global feature distillation, which enables the network to learn the global context of an image and (2) instance-wise feature distillation, which helps the network focus on object-specific features. The instance-wise feature distillation component ensures that the network effectively extracts object-centric features, whereas the global feature distillation component focuses on learning the overall context of the image. These distillation techniques prevent the network from overfitting to a specific domain and enable robust learning under various distortion conditions. Consequently, the student network is capable of performing reliably under diverse distortions while maintaining strong performance on the source domain data. This approach enhances the generalization between the source and target domains, enabling robust object detection across a wide range of domain environments. As illustrated in Figure I, the proposed method achieves superior SDG performance compared with state-of-the-art (SOTA) methods. In summary, the main contributions of this study are as follows:

- We propose CD-FKD, a novel cross-domain FKD method for single-domain generalized object detection.
- We evaluate the proposed method against other SDG approaches on the SDG benchmark dataset and demonstrate that the proposed method outperforms previous methods. Also, we present comprehensive analyses that contribute to an improved understanding of the proposed framework.

## II. RELATED WORK

### A. Single Domain Generalization

SDG focuses on training models with a single source domain while ensuring they generalize well to unseen target domains, a key task when collecting data from multiple domains is difficult. SDG methods are divided into data augmentation [13], [16] and domain-invariant feature learning [17]. Data augmentation generates synthetic samples to improve robustness against distribution shifts, whereas domain-invariant feature learning aims to extract stable features across domains.

Recent studies on single-domain generalized object detection have progressed with methods such as S-DGOD [17], which uses contrastive learning and self-distillation to improve generalization. CLIP-Gap [18] leverages CLIP for text-based augmentations to enhance object detection without requiring target domain data. G-NAS [19] reduces overfitting using differentiable NAS and OoD-aware G-loss. UFR [20] improves generalization with scene-level causal attention and object-level prototypes. Danish et al. [11] introduced carefully selected augmentations to diversify the source domain and a classification and localization alignment method to enhance out-of-domain detections. Li et al. [21] proposed a dynamic object-centric perception network with prompt learning to adapt to image complexity and improve cross-domain generalization. Despite these advancements, we propose a novel cross-domain feature distillation approach within a knowledge distillation (KD) framework that outperforms existing methods.

### B. Knowledge Distillation for Object Detection

KD was originally proposed for model compression in image classification, transferring knowledge from large teacher models to smaller student models [22]. Initially designed for image classification [23], [24], KD has been expanded to object detection for both compression and performance enhancement [25], [26]. By replicating the capabilities of the teacher model with fewer resources, KD has proven to be a versatile strategy.

In object detection, KD improves compact model performance by enabling them to emulate key features from larger models. FineGrained [27] introduces feature imitation for performance enhancement, whereas DeFeat [28] separates foreground and background features for independent distillation. FGD [29] and ScaleKD [30] improve detection accuracy by leveraging focal-global features and scale-aware knowledge transfer, respectively. CrossKD [31] uses cross-head distillation to improve performance by mimicking teacher predictions.

We propose a KD approach for SDG in object detection. The proposed method allows the student model to detect objects in images with reduced visibility and noise by mimicking the ability of the teacher model to detect objects in clear images. This improves performance under challenging conditions, making the model robust for unseen target domains.

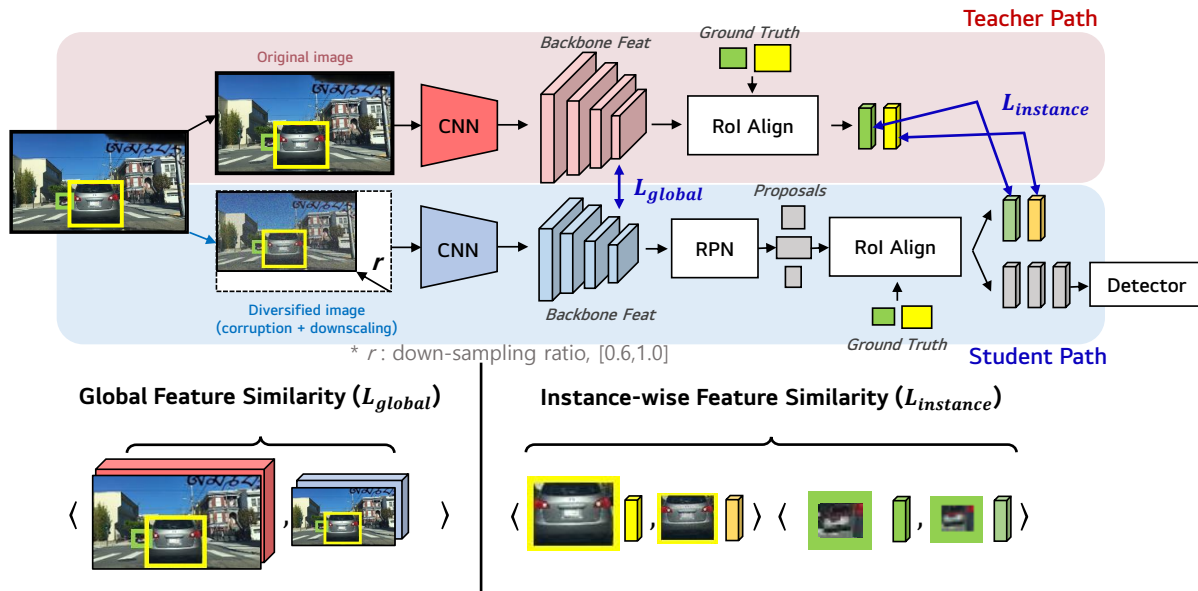


Fig. II. Illustration of the proposed single-domain generalized object detection using cross-domain FKD. As a KD framework, the frozen teacher network (represented in pink color) receives source domain data, whereas the student network (represented in blue sky color) is provided with downscaled and corrupted source domain data.

### III. METHOD

We introduce CD-FKD, a simple yet effective single-domain generalized object detection framework. First, we provide an overview of the proposed framework in Section III-A. Second, we introduce the method for generating diversified source domain image in Section III-B. Third, we describe the global feature distillation method in Section III-C. Finally, we present the instance-specific feature distillation approach in Section III-D.

#### A. Cross Domain Distillation Framework

The proposed framework uses a self-distillation structure with two identical detectors. We employ the widely recognized Faster R-CNN, a popular two-stage detector for single-domain object detection. The key idea is to leverage the differences in input images between the teacher and student networks, maximizing the benefits of this approach. In this framework, the teacher network receives clear, high-resolution images, enabling it to extract detailed features. In contrast, the student network is trained with corrupted and downscaled images, which makes detecting small objects more challenging. Despite these difficulties, the student network learns to handle such scenarios, becoming more robust. Training on these challenging inputs exposes the student network to difficult detection situations, promoting more rigorous learning. Through distillation, the student learns feature representations from the teacher, improving its ability to detect fine-grained details in tough environments. This process helps the student network become more resilient to corruption and better at small object detection.

As shown in Figure II, the teacher network receives original source domain data, while the student network is trained on diversified source domain data with various scales and



Fig. III. Examples of corrupted source domain data

corruptions. The training begins by pre-training the teacher network on the source domain data. During distillation, the teacher network's parameters are frozen, and it performs inference on the clear source domain data. Meanwhile, the student network learns object classification and localization from the diversified data, benefiting from the teacher's knowledge.

### B. Diversified Source Domain Data

The proposed approach leverages diversified source domain images ( $D_\varphi$ ) by applying various downscaling and corruption techniques to the source domain images ( $D_s$ ). Object detectors face challenges in target domains due to significant differences from the source domain, such as changes in image quality and object sizes. To address this, the proposed method applies downscaling and corruption to create discrepancies between the source and target domains.

We downscale the source image resolution to generate low-resolution images, while preserving the original high-resolution image for the teacher network. The teacher network receives the high-resolution image, and the student network is trained with the downsampled images. This encourages the student network to learn to detect small objects effectively, even from low-resolution data. Additionally, various corruptions are applied to the source domain images during training, enhancing generalization and preventing overfitting, as demonstrated in previous works like SDG [11]. These corruptions, such as object occlusion and image blurring, can challenge the network's ability to detect objects. To mitigate this, the teacher network receives clear source domain data, while the student network is trained on corrupted data. This setup allows the student network to learn from noisy data, extracting features that reflect those learned by the teacher network, ultimately improving generalization to unseen target domains.

We apply generic corruptions that are independent of the target domains. As shown in Figure III, a total of 15 different corruptions are applied using ImageNet-C [32]. These corruptions are applied with equal probability, and the intensity levels (ranging from 1 to 5) are distributed evenly across the augmentations.

### C. Global Feature Distillation

Global feature distillation is a technique that guides the student network to focus on important areas related to object detection, rather than focusing on the noise in diversified images. Diversified images distort object features, leading to semantic mismatches in the feature space. To address this, the student network is trained to learn semantically consistent features from diversified images.

During the distillation process, the features extracted by the teacher network from the original image ( $x^s$ ) serve as a reference for training the student network. The student network is trained to align its features, extracted from the diversified image ( $x^\varphi$ ), with this reference. Specifically, we obtain the backbone features of the teacher, denoted as  $F_T^s = f_T(x^s)$ , and the backbone features of the student, denoted as  $F_S^\varphi = f_S(x^\varphi)$ , where  $f_T(\cdot)$  and  $f_S(\cdot)$  represent the backbones of the teacher and student networks, respectively. To align  $F_S^\varphi$  with  $F_T^s$ , a global feature distillation loss function ( $\mathcal{L}_{global}$ ) is employed.

In the global feature distillation process, feature embeddings ( $F_T^s, F_S^\varphi$ ) are extracted from the final backbone layers of both the teacher and student networks, which use ResNet-101.  $F_T^s$  has a fixed size, whereas  $F_S^\varphi$  is resized according to the

scale-down ratio (0.6 to 1.0) of the input resolution. To align the feature map sizes of  $F_S^\varphi$  and  $F_T^s$ , bilinear interpolation is applied to  $F_S^\varphi$  to match its size with  $F_T^s$ .

The two feature embeddings are flattened, and cosine similarity loss is applied to maximize the cosine similarity between the two features. The cosine similarity between  $F_T^s$  and  $F_S^\varphi$  is calculated using the formula below:

$$\mathcal{L}_{global}(F_T^s, F_S^\varphi) = \sum_{i=1}^N \left( 1 - \frac{F_{T,i}^s \cdot F_{S,i}^\varphi}{\|F_{T,i}^s\| \cdot \|F_{S,i}^\varphi\|} \right) \quad (1)$$

where  $N$  denotes the number of train images,  $i$ -th refers to the image index.

### D. Instance-Wise Feature Distillation

Instance-wise feature distillation is a technique that focuses on the relationship between objects in diversified images and their corresponding objects in original images, excluding the background. The student network is provided with an image that has been diversified compared with the original image. These diversified images often suffer from occlusions or blurring, and owing to the smaller object sizes, they lack the visual information needed for the detector to accurately recognize the objects. Therefore, the goal of instance-wise feature distillation is to address the problem of reduced object visibility in corrupted images. The main objective is to use the region of interest (RoI) to make the RoI features of the diversified image as similar as possible to those of the original image. To achieve this, instance-wise feature distillation guides the student network to learn to extract features from the diversified image that are similar to those in the original image.

To implement this approach, the backbone features of the teacher ( $F_T^s$ ) and student ( $F_S^\varphi$ ) are processed using RoI Align (RA) with the ground truth bounding box values (GT) as the RoIs. This process extracts the instance features ( $I_T^s$ ) of the teacher and instance features ( $I_S^\varphi$ ) of the student, which are defined as follows:  $I_T^s = \text{RA}(F_T^s, \text{GT})$ ,  $I_S^\varphi = \text{RA}(F_S^\varphi, \text{GT})$ .  $I_T^s$  and  $I_S^\varphi$  are then aligned to ensure that corresponding objects are properly matched. All aligned features are flattened, and a cosine similarity loss is applied to calculate the similarity between the corresponding objects. The cosine similarity between the object features in the teacher and student networks is calculated as follows:

$$\mathcal{L}_{instance}(I_T^s, I_S^\varphi) = \sum_{i=1}^N \sum_{j=1}^O \left( 1 - \frac{I_{T,i,j}^s \cdot I_{S,i,j}^\varphi}{\|I_{T,i,j}^s\| \cdot \|I_{S,i,j}^\varphi\|} \right) \quad (2)$$

where  $N$  denotes the number of images,  $i$ -th refers to the image index,  $O$  is the number of instances in the image, and  $j$ -th denotes to the instance index.

This approach enables the student network to extract features from diversified images guided by clean-image features, while global and instance-wise distillation ensures it captures both overall context and individual instance knowledge.

Finally, the proposed model comprises the sum of the

TABLE I

SINGLE DOMAIN GENERALIZATION OBJECT DETECTION RESULTS. AVERAGE RESULTS ARE CALCULATED USING THE RESULTS FROM FOUR TARGET DOMAINS TO COMPARE THE GENERALIZATION ABILITY. THE RESULTS OF THE COMPARED MODELS ARE OBTAINED FROM THEIR RESPECTIVE PAPERS. **BOLD/UNDERLINE** INDICATES THE BEST/SECOND-BEST RESULT.

Method	Daytime-Clear	Night-Clear	Dusk-Rainy	Night-Rainy	Daytime-Foggy	Average
Faster R-CNN [12]	54.9	36.6	27.9	12.1	32.1	27.2
IBN-Net [33]	49.7	32.1	26.1	14.3	29.6	25.5
SW [34]	50.6	33.4	26.3	13.7	30.8	26.1
IterNorm [35]	43.9	29.6	22.8	12.6	28.4	23.4
ISW [36]	51.3	33.2	25.9	14.1	31.8	26.3
S-DGOD [17]	56.1	36.6	28.2	16.6	33.5	28.7
CLIP-Gap [18]	51.3	36.9	32.3	18.7	38.5	31.6
G-NAS [19]	58.4	<u>45.0</u>	35.1	17.4	36.4	33.5
PDDOC [21]	53.6	38.5	33.7	19.2	39.1	32.6
DivAlign [11]	52.8	42.5	<u>38.1</u>	<b>24.1</b>	37.2	<u>35.5</u>
UFR [20]	<u>58.6</u>	40.8	33.2	19.2	<u>39.6</u>	33.2
CD-FKD (Ours)	<b>62.7</b>	<b>47.3</b>	<b>42.3</b>	<u>23.4</u>	<b>40.2</b>	<b>38.3</b>

loss function for object localization and classification ( $\mathcal{L}_{det}$ ), global feature distillation loss ( $\mathcal{L}_{global}$ ), and instance-wise feature distillation ( $\mathcal{L}_{instance}$ ). Our overall training objective is as follows:

$$\mathcal{L}_{total} = \mathcal{L}_{det} + \alpha\mathcal{L}_{global} + \beta\mathcal{L}_{instance} \quad (3)$$

where  $\alpha$  and  $\beta$  denote hyperparameters that contribute toward balancing  $\mathcal{L}_{global}$  and  $\mathcal{L}_{instance}$ .

#### IV. EXPERIMENTS

##### A. Experimental Setup

**Dataset.** We evaluated the proposed method on the diverse weather dataset, a single domain generalization benchmark dataset in urban scenes built by [17]. The dataset comprises five different weather domains: Daytime-Clear, Night-Clear, Dusk-Rainy, Night-Rainy, and Daytime-Foggy. Daytime-Clear is used as the source domain for training, while the other four weather domains are used as target domains for testing. Daytime-Foggy comprises 19,395 images used for training and 8,313 images for testing. The four target domains used solely for evaluation are as follows: Night-Clear with 26,158 images, Dusk-Rainy with 3,501 images, Night-Rainy with 2,494 images, and Daytime-Foggy with 3,775 images. This dataset is an urban scene dataset containing seven object classes: bus, bike, car, motor, person, rider, and truck.

**Implementation Details.** We use Faster R-CNN [12] with a ResNet101-FPN [37] as the feature backbone, implemented in mmdetection [38], as the detector. The backbone is initialized with pre-trained weights from ImageNet. The model is trained using the stochastic gradient descent (SGD) optimizer with a learning rate of 0.01, momentum of 0.9, and weight decay of 0.0001. The batch size is set to 4. The values of  $\alpha$  and  $\beta$  in the loss function are both set to 1.0.

##### B. Comparison with State-of-the-arts

Our method is compared with normalization-based SDG works [33]–[36] and recent SOTA single-domain generalized detectors [11], [17]–[21]. For the baseline, we use Faster R-CNN [12] initialized with ImageNet pre-trained weights. We evaluate the performance of the proposed method on the

source domain, Daytime-Clear, and assess its generalization capability on unseen target domains, including Night-Clear, Dusk-Rainy, Night-Rainy, and Daytime-Foggy. The evaluation metric used is mean average precision (mAP), and we report results for mAP@0.5.

As summarized in Table I, the proposed method, CD-FKD, achieves the good performance not only on the source domain but also across all target domains. The average performance across the four target domains (Night-Clear, Dusk-Rainy, Night-Rainy, and Daytime-Foggy) achieves a mAP@0.5 of 38.3%, representing an improvement of 11.1% mAP@0.5 over the Faster R-CNN baseline. Furthermore, the proposed method outperforms the previous best, DivAlign [11], by 2.8% mAP@0.5. Remarkably, the proposed method enhances generalization to target domains without compromising performance on the source domain, surpassing previous methods on both source and target domains. These experimental results validate that the proposed approach significantly improves performance on the source domain while also enhancing generalization to unseen domains.

Figure IV visually compares the detection results of the Faster R-CNN [12], DivAlign [11], and the proposed method. The figure shows the detection results of each model for the Night-Clear, Dusk-Rainy, Night-Rainy, and Daytime-Foggy scenes from left to right. In contrast to the bottom-row results of the proposed method, the top-row Faster R-CNN and middle-row DivAlign results display both false negatives (indicated by red) and false positives (indicated by yellow).

**Results on Night-Clear Scene.** Table II summarizes the results for the Night-Clear scene. Due to reduced visibility at night, object detection becomes significantly more challenging compared with Daytime-Clear. In this scenario, the proposed method achieves the best performance with 47.3% mAP, outperforming previous SOTA methods across six object categories, except for *bike*. Low light conditions in night scenes lead to confusion between visually similar categories, such as *bike* and *motor*.

**Results on Dusk-Rainy Scene.** Table III summarizes the results for the Dusk-Rainy scene, where low light and rain affect detection performance. The proposed method



Fig. IV. Qualitative evaluation results of the model’s generalization ability on the Night-Clear, Dusk-Rainy, Night-Rainy, and Daytime-Foggy scenes. The top-row images show the results of Faster R-CNN [12]. The middle-row images show the results of DivAlign [11]. The bottom-row images show the results of our method. Red circles and arrows indicate false negatives, while yellow circles and arrows indicate false positives.

TABLE II  
QUANTITATIVE RESULTS (%) ON THE NIGHT-CLEAR SCENE.

Method	bus	bike	car	motor	person	rider	truck	mAP
Faster R-CNN [12]	37.5	36.6	62.0	14.0	43.6	28.8	41.6	37.7
IBN-Net [33]	37.8	27.3	49.6	15.1	39.2	27.1	38.9	32.1
SW [34]	38.7	29.2	49.8	16.6	31.5	28.0	40.2	33.4
IterNorm [35]	38.5	23.5	38.9	15.8	26.6	25.9	38.1	29.6
ISW [36]	38.5	28.5	49.6	15.4	31.9	27.5	41.3	33.2
S-DGOD [17]	40.6	35.1	50.7	19.7	34.7	32.1	43.4	36.6
CLIP-Gap [18]	37.7	34.3	58.0	19.2	37.6	28.5	42.9	36.9
G-NAS [19]	46.9	40.5	67.5	26.5	50.7	35.4	47.8	45.0
PDDOC [21]	40.9	35.0	59.0	21.3	40.4	29.9	42.9	38.5
UFR [20]	43.6	38.1	66.1	14.7	49.1	26.4	42.9	36.9
Ours	47.8	40.1	71.8	27.2	55.2	38.8	50.1	47.3

TABLE III  
QUANTITATIVE RESULTS (%) ON THE DUSK-RAINY SCENE.

Method	bus	bike	car	motor	person	rider	truck	mAP
Faster R-CNN [12]	34.6	22.3	63.6	7.9	24.2	13.2	39.4	29.3
IBN-Net [33]	37.0	14.8	50.3	11.4	17.9	13.3	38.4	26.1
SW [34]	35.2	16.7	50.1	10.4	20.1	13.0	38.8	26.3
IterNorm [35]	32.9	14.1	38.9	11.0	15.5	11.6	35.7	22.8
ISW [36]	34.7	16.0	50.0	11.1	17.8	12.6	38.8	25.9
S-DGOD [17]	37.1	19.6	50.9	13.4	19.7	16.3	40.7	28.2
CLIP-Gap [18]	37.8	22.8	60.7	16.8	26.8	18.7	42.4	32.3
G-NAS [19]	44.6	22.3	66.4	14.7	32.1	19.6	45.8	35.1
PDDOC [21]	39.4	25.2	60.9	20.4	29.9	16.5	43.9	33.7
UFR [20]	37.1	21.8	67.9	16.4	27.4	17.9	43.9	33.2
Ours	50.3	30.9	75.7	19.6	40.4	26.4	52.9	42.3

outperforms others, showing significant improvements across most categories, especially *bike*, *person*, and *rider*, which are prone to occlusion from small objects like those influenced by rain or fog. However, confusion between *motor* and *bike* slightly lowers *motor* performance.

**Results on Night-Rainy Scene.** Table IV summarizes the results for the Night-Rainy scene. The Night-Rainy scene is the most challenging condition owing to low light and the effects of rain. Despite these challenges, our method outper-

TABLE IV  
QUANTITATIVE RESULTS (%) ON THE NIGHT-RAINY SCENE.

Method	bus	bike	car	motor	person	rider	truck	mAP
Faster R-CNN [12]	19.8	8.6	30.2	0.3	0.9	4.8	16.6	12.8
IBN-Net [33]	24.6	10.0	28.4	0.9	8.3	9.8	18.1	14.3
SW [34]	22.3	7.8	27.6	0.2	10.3	10.0	17.7	13.7
IterNorm [35]	21.4	6.7	22.0	0.9	9.1	10.6	17.6	12.6
ISW [36]	22.5	11.4	26.9	0.4	9.9	9.8	17.5	14.1
S-DGOD [17]	24.4	11.6	29.5	9.8	10.5	11.4	19.2	16.6
CLIP-Gap [18]	28.6	12.1	36.1	9.2	12.3	9.6	22.9	18.7
G-NAS [19]	28.6	9.8	38.4	0.1	13.8	9.8	21.4	17.4
PDDOC [21]	25.6	12.1	35.8	10.1	14.2	12.9	22.9	19.2
UFR [20]	29.9	11.8	36.1	9.4	13.1	10.5	23.3	19.2
Ours	34.3	14.2	49.6	2.1	17.2	15.6	30.6	23.4

TABLE V  
QUANTITATIVE RESULTS (%) ON THE DAYTIME-FOGGY SCENE.

Method	bus	bike	car	motor	person	rider	truck	mAP
Faster R-CNN [12]	30.2	26.6	54.9	27.2	35.1	37.6	19.4	33.0
IBN-Net [33]	29.9	26.1	44.5	24.4	26.2	33.5	22.4	29.6
SW [34]	30.6	36.2	44.6	25.1	30.7	34.6	23.6	30.8
IterNorm [35]	29.7	21.8	42.4	24.4	26.0	33.3	21.6	28.4
ISW [36]	29.5	26.4	49.2	27.9	30.7	34.8	24.0	31.8
S-DGOD [17]	32.9	28.0	48.8	29.8	32.5	38.2	24.1	33.5
CLIP-Gap [18]	36.2	34.2	57.9	34.0	38.7	43.8	25.1	38.5
G-NAS [19]	32.4	31.2	57.7	31.9	38.6	38.5	24.5	36.4
PDDOC [21]	36.1	34.5	58.4	33.3	40.5	44.2	26.2	39.1
UFR [20]	36.9	35.8	61.7	33.7	39.5	42.2	27.5	39.6
Ours	36.6	33.2	62.6	34.6	42.9	44.2	27.5	40.2

forms others, achieving the highest performance. However, *motor* exhibits lower performance due to confusion with *bike*.

**Results on Daytime-Foggy Scene.** Table V summarizes the results for the Daytime-Foggy scene, where fog causes occlusions and blurring, complicating detection. The proposed method outperforms others across most categories. However, foggy conditions increase confusion between visually similar categories like *bus* and *truck*, and *bike* and *motor*.

TABLE VI

ABLATION STUDY RESULTS (%) OF THE PROPOSED CD-FKD. DC, NC, DR, NR, AND DF REPRESENT DAYTIME-CLEAR, NIGHT-CLEAR, DUSK-RAINY, NIGHT-RAINY, AND DAYTIME-FOGGY, RESPECTIVELY.

Methods	Corrupt&Down	FKD	DC	NC	DR	NR	DF	Avg.
Faster R-CNN	✗	✗	54.9	36.6	27.9	12.1	32.1	27.2
Ours	✓	✗	58.8	41.1	33.9	14.9	34.0	30.9
Ours	✓	$\mathcal{L}_{\text{glo}}$	62.0	46.3	41.7	21.4	39.5	37.2
Ours	✓	$\mathcal{L}_{\text{ins}}$	62.5	46.7	42.1	21.9	39.7	37.6
Ours	✓	$\mathcal{L}_{\text{glo}}+\mathcal{L}_{\text{ins}}$	<b>62.7</b>	<b>47.3</b>	<b>42.3</b>	<b>23.4</b>	<b>40.2</b>	<b>38.3</b>

TABLE VII

ABLATION STUDY ON THE SOURCE DOMAIN TEST SET COMPARES FASTER R-CNN AND CD-FKD, EVALUATING CD-FKD WITH ONLY CORRUPTION (CORRUPT) AND WITH BOTH CORRUPTION AND DOWNSCALING (CORRUPT&DOWN).

Method	mAP	mAP <sub>50</sub>	mAP <sub>s</sub>	mAP <sub>m</sub>	mAP <sub>l</sub>
Faster R-CNN	29.5	54.9	8.2	32.4	54.6
+ Corrupt	34.5	62.1	13.2	37.7	<b>57.6</b>
+ Corrupt&Down	<b>34.8</b>	<b>62.7</b>	<b>13.6</b>	<b>38.0</b>	<b>57.6</b>

### C. Ablation Study

We conducted an ablation study to analyze the impact of different components of the proposed method. First, we evaluated the effects of corruption and downscaling (Corrupt&Down) on the input data. Second, we performed an ablation analysis to evaluate the impact of the FKD, specifically  $\mathcal{L}_{\text{global}}$  and  $\mathcal{L}_{\text{instance}}$ .

Table VI summarizes the results. "Avg." represents the average performance across the four target domains. The application of corruption and downscaling to the source images significantly improved performance compared to the baseline, demonstrating its effectiveness in enhancing generalization. Moreover, using both  $\mathcal{L}_{\text{global}}$  and  $\mathcal{L}_{\text{instance}}$  distillation methods resulted in a 38.3% mAP, a notable improvement over the baseline. Each distillation method individually contributed to the performance gains, not only for the target domains but also for the source domain. Through distillation, our model learned to replicate the feature representations from the teacher network, improving adaptability to the source domain while avoiding overfitting to the diversified data.

Table VII presents the results of another ablation study, examining the effect of corruption and downscaling on the input images. The proposed method, when applied to the student network, significantly outperformed the baseline. Notably, the downscaling technique achieved a 0.4% improvement in mAP<sub>s</sub> and a 0.3% improvement in mAP<sub>m</sub> compared to corruption-only experiments, underscoring the method's effectiveness in improving small object detection.

Additionally, Figure V compares the heatmaps generated by the baseline and the proposed CD-FKD. The baseline-generated heatmap mainly highlights irrelevant background areas. In contrast, the proposed method reduces background focus and places more attention on the objects. This shows that the proposed approach focuses on objects instead of the background in unseen domains, thereby improving




Fig. V. Heatmap visualization for target domain scenes. The left column displays the original images, the middle column presents the results from Faster R-CNN, and the right column shows the results from our method. The heatmaps highlight the areas where the model focuses, with regions of higher attention marked by a redder hue.

generalization performance.

TABLE VIII

THE IMPACT OF THE HYPERPARAMETERS  $\alpha$  AND  $\beta$  OF  $\mathcal{L}_{\text{global}}$  AND  $\mathcal{L}_{\text{instance}}$  IN THE DISTILLATION TRAINING PHASE.

$\alpha/\beta$	DC	NC	DR	NR	DF	Avg.
0.5/1.5	62.5	46.8	42.1	23.1	40.1	38.0
<b>1.0/1.0</b>	<b>62.7</b>	<b>47.3</b>	<b>42.3</b>	<b>23.4</b>	<b>40.2</b>	<b>38.3</b>
1.5/0.5	62.1	46.7	41.8	22.9	40.1	37.9

In Table VIII, we applied the hyperparameters  $\alpha$  and  $\beta$  to the  $\mathcal{L}_{\text{global}}$  and  $\mathcal{L}_{\text{instance}}$  loss functions during the feature distillation training phase from the teacher network to the student network. According to the results from the ablation study in the main manuscript, when only one of the two loss functions ( $\mathcal{L}_{\text{global}}$  or  $\mathcal{L}_{\text{instance}}$ ) was used,  $\mathcal{L}_{\text{instance}}$  contributed slightly more to performance improvement. However, the final results show that applying both  $\mathcal{L}_{\text{global}}$  and  $\mathcal{L}_{\text{instance}}$  resulted in the largest performance gain.

Therefore, we further analyze the effects of both loss functions. Table VIII reports the performance achieved by applying  $\alpha$  and  $\beta$  values of 0.5/1.5, 1.0/1.0, and 1.5/0.5. Our evaluation shows that the optimal  $\alpha$  and  $\beta$  values of 1.0 and 1.0 strike an effective balance between transferring global features through  $\mathcal{L}_{\text{global}}$  and instance-wise features through  $\mathcal{L}_{\text{instance}}$ . This balance maximized the learning efficiency of the student network.

## V. CONCLUSION

In this study, we introduced a novel approach for single-domain generalized object detection using CD-FKD. Our proposed method effectively addressed the domain shift problem and significantly improved generalization to unseen target domains using only a single source domain. We applied global and instance-wise feature distillation, enabling the model to extract robust and meaningful object-centric features while maintaining detection performance even under severe corruptions. Experimental results demonstrated that CD-FKD consistently outperformed existing methods under adverse weather conditions. Furthermore, an ablation study confirmed the effectiveness of each component in the proposed method, leading to performance improvements in both unseen target and source domains.

## ACKNOWLEDGEMENT

This research was supported by the National Research Council of Science & Technology (NST) grant by the Korea government (MSIT) (No. GTL25041-000)

## REFERENCES

- [1] Y. Cai, T. Luan, H. Gao, H. Wang, L. Chen, Y. Li, M. A. Sotelo, and Z. Li, "Yolov4-5d: An effective and efficient object detector for autonomous driving," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–13, 2021.
- [2] J. Park, J. Lee, S. Moon, and K. Lee, "Deep learning based detection of missing tooth regions for dental implant planning in panoramic radiographic images," *Applied Sciences*, vol. 12, no. 3, p. 1595, 2022.
- [3] J. Lee, J. Park, S. Lee, S.-Y. Moon, and K. Lee, "Automated diagnosis for extraction difficulty of maxillary and mandibular third molars and post-extraction complications using deep learning," *Scientific Reports*, vol. 15, no. 1, p. 19036, 2025.
- [4] J. Lee, S. Lee, J. Kim, J. Park, and K. Lee, "Robust maritime object detection under adverse conditions via joint semantic learning without extra computational overhead," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 8166–8173.
- [5] R. Geirhos, C. R. Temme, J. Rauber, H. H. Schütt, M. Bethge, and F. A. Wichmann, "Generalisation in humans and deep neural networks," *Advances in neural information processing systems*, vol. 31, 2018.
- [6] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, "Do imagenet classifiers generalize to imagenet?" in *International conference on machine learning*. PMLR, 2019, pp. 5389–5400.
- [7] Y.-J. Li, X. Dai, C.-Y. Ma, Y.-C. Liu, K. Chen, B. Wu, Z. He, K. Kitani, and P. Vajda, "Cross-domain adaptive teacher for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 7581–7590.
- [8] C. Chen, Z. Zheng, X. Ding, Y. Huang, and Q. Dou, "Harmonizing transferability and discriminability for adapting object detectors," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8869–8878.
- [9] F. M. Carlucci, A. D'Innocente, S. Bucci, B. Caputo, and T. Tommasi, "Domain generalization by solving jigsaw puzzles," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 2229–2238.
- [10] Z. Huang, H. Wang, E. P. Xing, and D. Huang, "Self-challenging improves cross-domain generalization," in *Computer vision—ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part II 16*. Springer, 2020, pp. 124–140.
- [11] M. S. Danish, M. H. Khan, M. A. Munir, M. S. Sarfraz, and M. Ali, "Improving single domain-generalized object detection: A focus on diversification and alignment," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17732–17742.
- [12] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1137–1149, 2016.
- [13] Y. Wang, Zijian fand Luo, R. Qiu, Z. Huang, and M. Baktashmotlagh, "Learning to diversify for single domain generalization," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 834–843.
- [14] C. Wan, X. Shen, Y. Zhang, Z. Yin, X. Tian, F. Gao, J. Huang, and X.-S. Hua, "Meta convolutional neural networks for single domain generalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4682–4691.
- [15] P. Kirichenko, M. Ibrahim, R. Balestrieri, D. Bouchacourt, S. R. Vedantam, H. Firooz, and A. G. Wilson, "Understanding the detrimental class-level effects of data augmentation," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [16] R. Volpi, H. Namkoong, O. Sener, J. C. Duchi, V. Murino, and S. Savarese, "Generalizing to unseen domains via adversarial data augmentation," *Advances in neural information processing systems*, vol. 31, 2018.
- [17] A. Wu and C. 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*, 2022, pp. 847–856.
- [18] V. Vedit, M. Engilberge, and M. 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*, 2023, pp. 3219–3229.
- [19] F. Wu, J. Gao, L. Hong, X. Wang, C. Zhou, and N. Ye, "G-nas: Generalizable neural architecture search for single domain generalization object detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 6, 2024, pp. 5958–5966.
- [20] Y. Liu, S. Zhou, X. Liu, C. Hao, B. Fan, and J. Tian, "Unbiased faster r-cnn for single-source domain generalized object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 28 838–28 847.
- [21] D. Li, A. Wu, Y. Wang, and Y. Han, "Prompt-driven dynamic object-centric learning for single domain generalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17 606–17 615.
- [22] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [23] S. Shin, J. Lee, J. Lee, Y. Yu, and K. Lee, "Teaching where to look: Attention similarity knowledge distillation for low resolution face recognition," in *European Conference on Computer Vision*. Springer, 2022, pp. 631–647.
- [24] O. S. EL-Assiouti, G. Hamed, D. Khattab, and H. M. Ebied, "Hdkd: Hybrid data-efficient knowledge distillation network for medical image classification," *Engineering Applications of Artificial Intelligence*, vol. 138, p. 109430, 2024.
- [25] G. Bang, K. Choi, J. Kim, D. Kum, and J. W. Choi, "Radardistill: Boosting radar-based object detection performance via knowledge distillation from lidar features," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 491–15 500.
- [26] L. Zhao, J. Song, and K. A. Skinner, "Crkd: Enhanced camera-radar object detection with cross-modality knowledge distillation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 470–15 480.
- [27] T. Wang, L. Yuan, X. Zhang, and J. Feng, "Distilling object detectors with fine-grained feature imitation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4933–4942.
- [28] J. Guo, K. Han, Y. Wang, H. Wu, X. Chen, C. Xu, and C. Xu, "Distilling object detectors via decoupled features," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 2154–2164.
- [29] Z. Yang, Z. Li, X. Jiang, Y. Gong, Z. Yuan, D. Zhao, and C. Yuan, "Focal and global knowledge distillation for detectors," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4643–4652.
- [30] Y. Zhu, Q. Zhou, N. Liu, Z. Xu, Z. Ou, X. Mou, and J. Tang, "Scalekd: Distilling scale-aware knowledge in small object detector," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19 723–19 733.
- [31] J. Wang, Y. Chen, Z. Zheng, X. Li, M.-M. Cheng, and Q. Hou, "Crosskd: Cross-head knowledge distillation for object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 520–16 530.
- [32] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," *Proceedings of the International Conference on Learning Representations*, 2019.
- [33] X. Pan, P. Luo, J. Shi, and X. Tang, "Two at once: Enhancing learning and generalization capacities via ibn-net," in *Proceedings of the european conference on computer vision (ECCV)*, 2018, pp. 464–479.
- [34] X. Pan, X. Zhan, J. Shi, X. Tang, and P. Luo, "Switchable whitening for deep representation learning," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 1863–1871.
- [35] L. Huang, Y. Zhou, F. Zhu, L. Liu, and L. Shao, "Iterative normalization: Beyond standardization towards efficient whitening," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4874–4883.
- [36] S. Choi, S. Jung, H. Yun, J. T. Kim, S. Kim, and J. 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*, 2021, pp. 11 580–11 590.
- [37] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [38] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "MMDetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.