

Object-Aware Domain Generalization for Object Detection

Wooju Lee*, Dasol Hong*, Hyungtae Lim†, Hyun Myung†

Urban Robotics Lab, School of Electrical Engineering,
Korea Advanced Institute of Science and Technology, Republic of Korea
 {dnwn24, ds.hong, shapelim, hmyung}@kaist.ac.kr

Abstract

Single-domain generalization (S-DG) aims to generalize a model to unseen environments with a single-source domain. However, most S-DG approaches have been conducted in the field of classification. When these approaches are applied to object detection, the semantic features of some objects can be damaged, which can lead to imprecise object localization and misclassification. To address these problems, we propose an object-aware domain generalization (OA-DG) method for single-domain generalization in object detection. Our method consists of data augmentation and training strategy, which are called *OA-Mix* and *OA-Loss*, respectively. OA-Mix generates multi-domain data with multi-level transformation and object-aware mixing strategy. OA-Loss enables models to learn domain-invariant representations for objects and backgrounds from the original and OA-Mixed images. Our proposed method outperforms state-of-the-art works on standard benchmarks. Our code is available at <https://github.com/WoojuLee24/OA-DG>.

Introduction

Modern deep neural networks (DNNs) have achieved human-level performances in various applications such as image classification and object detection (He et al. 2016; Dosovitskiy et al. 2021; Carion et al. 2020; Ren et al. 2015). However, DNNs are vulnerable to various types of domain shifts, which have not been seen in the source domain (Dan and Thomas 2019; Lee and Myung 2022; Michaelis et al. 2019; Wu and Deng 2022). Even a small change in the domain can have disastrous results in real-world scenarios such as autonomous driving (Michaelis et al. 2019; Wu and Deng 2022). Thus, DNNs should be robust against the domain shifts to be applied in real-world applications.

Domain generalization (DG) aims to generalize a model to unseen target domains by using only the source domain (Hendrycks et al. 2019; Kim et al. 2021; Yao et al. 2022; Zhou et al. 2022). However, most DG methods rely on multiple source domains and domain annotations, which are generally unavailable (Kim et al. 2021; Lin et al. 2021; Yao

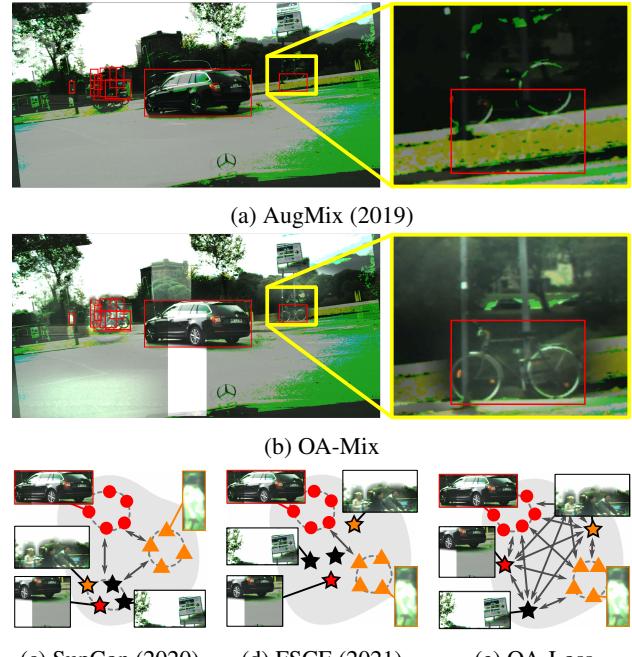


Figure 1. (a) Existing data augmentation methods can damage the semantic features of objects. (b) OA-Mix generates multiple domains while preserving semantic features. (c)-(e) The dotted line, arrow, circle, triangle, and star mean pull, push, car, person, and background, respectively. (c) The background contains different semantics (red and orange stars), but these are treated identically and pulled. (d) The method ignores semantic relationships between background instances. (e) OA-Loss considers semantic relations of background instances from multi-domain.

et al. 2022). Single-domain generalization (S-DG) achieves DG without any additional knowledge about multiple domains (Wan et al. 2022; Wang et al. 2021b).

Data augmentation has been successfully applied for S-DG in image classification (Hendrycks et al. 2019; Modas et al. 2022). The methods can be used to generate multi-source domains from a single-source domain. However, object detection addresses multiple objects in an image. When data augmentation methods for S-DG are applied

*These authors contributed equally.

†Corresponding authors: Dr. Hyungtae Lim and Prof. Hyun Myung
 Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

to object detection, object annotations may be damaged. As shown in Figure 1(a), spatial and color transformations can damage the positional or semantic features of objects. Michaelis et al. (2019) avoided this problem with style-transfer augmentations that do not change object locations. However, this approach does not leverage a rich set of transformations in image classification, thus limiting the domain coverage. Therefore, data augmentation for single-source domain generalization for object detection (S-DGOD) should include various transformations without damaging object annotation.

Recently, contrastive learning methods have been proposed to reduce the gap between the original and augmented domains (Kim et al. 2021; Yao et al. 2022). The methods address inter-class and intra-class relationships from multi-domain. However, they only train the relationships in object classes and do not consider background class for object detection (Sun et al. 2021) as shown in Figure 1(d). Because the object detector misclassifies foreground as background in OOD, considering the background class is required to classify the presence of objects in OOD. Therefore, the contrastive learning method should consider both the foreground and background classes for object detectors to classify objectness in out-of-distribution.

In this study, we propose object-aware domain generalization method (OA-DG) for S-DGOD to overcome the aforementioned limitations. Our method consists of OA-Mix for data augmentation and OA-Loss for reducing the domain gap. OA-Mix consists of multi-level transformations and object-aware mixing. Multi-level transformations introduce local domain changes within an image and object-aware mixing prevents the transformations from damaging object annotation. OA-Mix is the first object-aware approach that allows a rich set of image transformations for S-DGOD.

OA-Loss reduces the gap between the original and augmented domains in an object-aware manner. For foreground instances, the method trains inter-class and intra-class relations to improve object classification in out-of-distribution. Meanwhile, the background instances that belong to the same class can partially include the foreground objects of different classes. OA-Loss allows the model to learn the semantic relations among background instances in multi-domain as illustrated in Figure 1(e). To the best of our knowledge, the proposed method first trains the semantic relations for both the foreground and background instances to achieve DG with contrastive methods.

Our proposed method shows the best performance on common corruption (Michaelis et al. 2019) and various weather benchmarks (Wu and Deng 2022) in an urban scene. The contributions can be summarized as follows.

- We propose OA-Mix, a general and effective data augmentation method for S-DGOD. It increases image diversity while preserving important semantic features with multi-level transformations and object-aware mixing.
- We propose OA-Loss that reduces the domain gap between the original and augmented images. OA-Loss enables the model to learn semantic relations of foreground and background instances from multi-domain.

- Extensive experiments on standard benchmarks show that the proposed method outperforms state-of-the-art methods on unseen target domains.

Related Work

Data Augmentation for Domain Generalization

Recently, many successful data augmentation methods for S-DG have been proposed for image classification. AugMix (Hendrycks et al. 2019) mixes the results of the augmentation chains to generate more diverse domains. Modas et al. (2022) generated diverse domains using three primitive transformations with smoothness and strength parameters. However, direct application of these methods to object detection can damage the locations or semantic annotations of objects, as shown in Figure 1(a). Geirhos et al. (2018) avoided this issue by using style transfer (Gatys, Ecker, and Bethge 2016) that does not damage object annotations and improved robustness to synthetic corruptions. However, data augmentation methods for S-DGOD should generate diverse domains with a broad range of transformations rather than a limited set of transformations. The proposed OA-Mix uses an object-aware approach to create various domains without damaging object annotations.

Contrastive Learning for Domain Generalization

The contrastive learning methods for domain generalization build domain-invariant representations from multi-domain with sample-to-sample pairs (Yao et al. 2022; Kim et al. 2021; Li et al. 2021). The methods pull positive sample pairs of multi-domain and push negative samples away. PCL (Yao et al. 2022) arranges representative features and sample-to-sample pairs from multi-domain. SelfReg (Kim et al. 2021) aligns positive sample pairs and regularizes features with a self-supervised contrastive manner. PDEN (Li et al. 2021) progressively generates diverse domains and arranges features with a contrastive loss. However, these contrastive learning methods address semantic relations of sample-to-sample pairs only in object classes, not the background class.

Single-Domain Generalization for Object Detection

Recently, several S-DGOD methods have been proposed. CycConf (Wang et al. 2021a) enforces the object detector to learn invariant representations across the same instances under various conditions. However, CycConf requires annotated video datasets, which are not generally given. CSDS (Wu and Deng 2022) propagates domain-invariant features to the detector with self-knowledge distillation. However, the model does not diversify the single-domain with data augmentation, which potentially can be improved. Vudit et al. (2023) utilized a pre-trained vision-language model to generalize a detector, but textual hints about the target domains should be given. Our proposed method does not require prior knowledge about target domains and achieves S-DGOD in an object-aware approach.

Method

We propose OA-DG method for S-DGOD. Our approach consists of OA-Mix, which augments images, and OA-Loss,

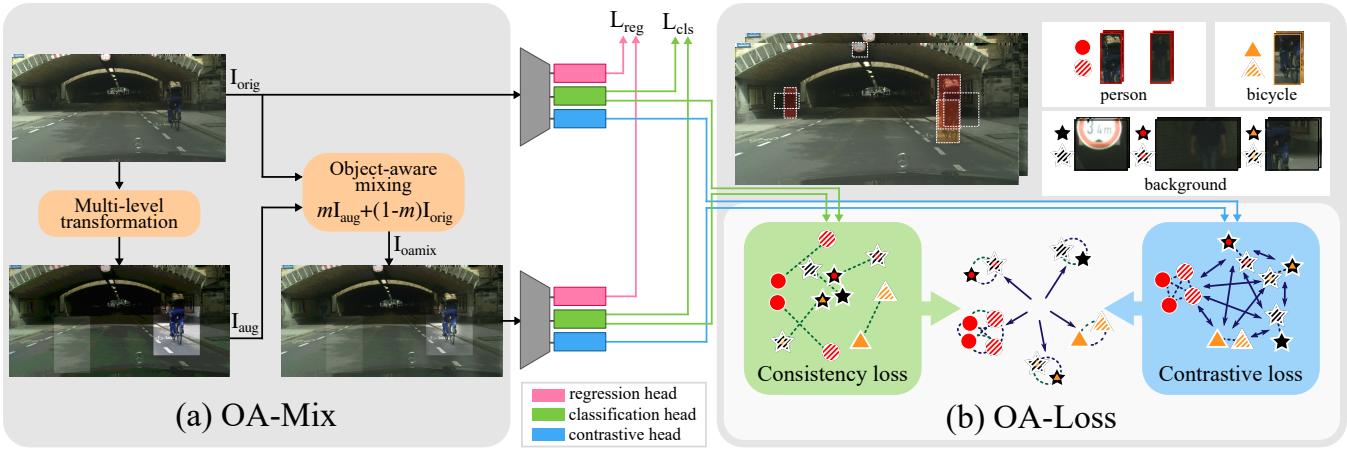


Figure 2. Overview of OA-DG method. (a) OA-Mix transforms an image at multiple levels and mixes the transformed image with the original image in an object-aware manner. The original and augmented images are fed into the object detectors that share weights. (b) OA-Loss trains the object detectors to learn the domain-invariant representations from the original and augmented images. Circles, triangles, and stars represent car, person, and background classes, respectively. An augmented instance has white slanted lines. Another shape within a star indicates the partially included object in the background. The consistency loss aligns the original and augmented instances (dotted line). The contrastive loss aligns the positive pairs (dotted line) and repulses the negative pairs (arrow).

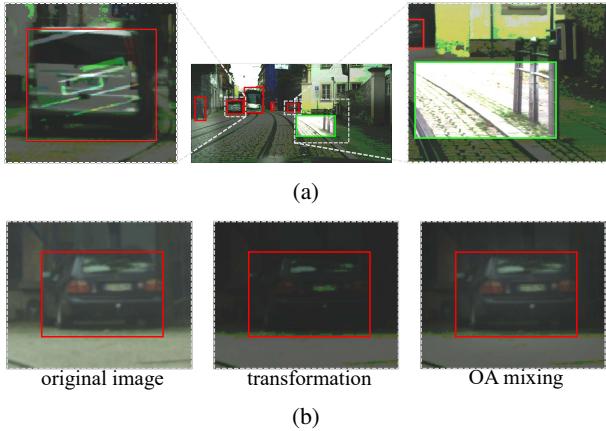


Figure 3. Overview of OA-Mix method. (a) Multi-level transformations generate locally diverse changes. Red and green boxes denote foreground and random regions, respectively. (b) Object-aware mixing strategy preserves semantic features against image transformations.

which achieves DG based on these augmented images. OA-DG can be applied to both one-stage and two-stage object detectors. The overview of OA-DG is illustrated in Figure 2.

OA-Mix

Diversity and affinity are two important factors to be considered in data augmentation for S-DGOD. Diversity refers to the variety of features in the augmented image. Affinity refers to the distributional similarity between the original and augmented images. In other words, effective data augmentation methods should generate diverse data that does not deviate significantly from the distribution of the origi-

nal data. This section decomposes OA-Mix into two components: multi-level transformation for diversity and object-aware mixing strategy for affinity.

Multi-Level Transformations OA-Mix enhances domain diversity by applying locally diverse changes within an image. Locally diverse changes can be achieved with a multi-level transformation strategy. As shown in Figure 3(a), an image is randomly divided into several regions such as foreground, background, and random box regions. Then, different transformation operations, such as color and spatial transformations, are randomly applied to each region. Spatial transformation operations are applied at the foreground level to preserve the location of the object. As a result, multi-level transformation enhances the domain diversity of the augmented image without damaging the object locations

Object-Aware Mixing In object detection, each object in an image has different characteristics, such as size, location, and color distribution. Depending on these object-specific characteristics, transformations can damage the semantic features of an object. For example, Figure 3(b) shows that image transformation can damage the semantic features of objects with monotonous color distribution. Previous works mix the original and augmented images at the image level to mitigate the degradation of semantic features. However, the method of mixing the entire image with the same weight does not sufficiently utilize object information.

Therefore, we calculate the saliency score s for each object based on the saliency map to consider the properties of each object. The saliency score is calculated as $\mathbf{S} \in \mathbb{R}^{h \times w}$ from object with size $h \times w$:

$$s = \frac{1}{hw} \sum_{x=1}^w \sum_{y=1}^h \mathbf{S}_{x,y}. \quad (1)$$

Specifically, the saliency map is a spatial representation of the spectral residual, which is the unpredictable frequency in an image. The object with a high saliency score has clear semantic signals. In contrast, the object with a low score has weak semantic signals as shown in Figure 3(b). Object-aware mixing strategy increases the mixing weight of the original image for objects with low saliency scores, thereby preventing semantic feature damage. Specifically, for each area P of the image, it linearly combines the original and augmented images with the mixing weight m :

$$\mathbf{I}_{\text{oamix}} = \sum_P \left(m \mathbf{I}_{\text{aug}}^P + (1 - m) \mathbf{I}_{\text{orig}}^P \right),$$

where m is sampled from different distributions depending on saliency score s . As a result, the strategy enhances the affinity of the augmented image, mitigating the negative effects of transformations.

OA-Loss

OA-Loss is designed to train the domain-invariant representations between the original and OA-Mixed domains in an object-aware approach. The object-aware method arranges instances in multi-domain according to the semantics of the instances. OA-Loss does not depend on any object detection framework and can be applied to both one-stage and two-stage detection frameworks.

Review of Supervised Contrastive Learning The supervised contrastive learning methods construct positive pairs for the same class and negative pairs for different classes (Khosla et al. 2020; Sun et al. 2021; Yao et al. 2022). The methods align intra-class and repulse inter-class instances in the embedding space. Previous methods for object detection (Sun et al. 2021) utilized only the foreground instances and ignored the semantic relations among background instances. In contrast, we explore the meaning of background instances for domain-invariant representations.

Meaning of Background Instances In object detection, each instance feature is labeled as the background if all intersection of unions (IoUs) with the ground-truth set are less than an IoU threshold, and foreground otherwise. Instance features correspond to region proposals and grid cell features in two-stage and one-stage detection frameworks, respectively.

The background instances can partially contain the foreground object of different classes. However, existing supervised contrastive learning methods (Khosla et al. 2020; Kang et al. 2021; Li et al. 2022; Yao et al. 2022) regard different background instances identically and form positive pairs, leading to false semantic relations of foreground objects. Therefore, the relationship between background instances is required to be defined in an object-aware manner.

OA Contrastive Learning OA-Loss trains the semantic relations among instances in an object-aware approach to reduce the multi-domain gap. For foreground instances, the method pulls the same class and pushes different classes away to improve object classification in multi-domain (Sun et al. 2021). For background instances, the method pushes

background instances away from each other except for the augmented instances as shown in Figure 2. OA-Loss trains the model to discriminate between background instances containing foreground objects of various classes. As a result, our approach reflects the meaning of the background class, which is difficult to define as a single class.

In addition, the negative pairs of background instances repulse each other, helping the model to output different features. OA-Loss allows the model to generate various background features and trains the multi-domain gaps for various features, further improving generalization capability of the object detector.

To incorporate positive and negative sample pairs into contrastive learning, we encode the instance features into contrastive features \mathbf{Z} . Following supervised contrastive learning (Kim et al. 2021; Sun et al. 2021), a contrastive branch is introduced parallel to the classification and regression branches as shown in Figure 2. The contrastive branch consists of two layer multilayer perceptrons (MLPs). The contrastive branch generates feature set \mathbf{Z} from the multi-level regions of OA-Mix and proposed regions of the model. The multi-level regions enable detectors to learn semantic relations in a wider variety of domains.

We set each contrastive feature $\mathbf{z}_i \in \mathbf{Z}$ as an anchor. The positive set of each feature \mathbf{z}_i is defined as:

$$\mathbf{Z}_i^{\text{pos}} = \begin{cases} \{\mathbf{z} | y(\mathbf{z}) = y(\mathbf{z}_i), \forall \mathbf{z} \in \mathbf{Z} \setminus \{\mathbf{z}_i\}\} & \text{if foreground} \\ \{\mathbf{z}'_i\} & \text{otherwise,} \end{cases} \quad (2)$$

where y is the label mapping function for feature \mathbf{z} . \mathbf{z}' is the augmented feature from \mathbf{z} . The positive set $\mathbf{Z}_i^{\text{pos}}$ is defined as the same class for foreground anchors and the same instance for background anchors. The proposed contrastive loss L_{ct} is defined as:

$$L_{ct} = \frac{1}{|\mathbf{Z}|} \sum_{\mathbf{z}_i \in \mathbf{Z}} L_{\mathbf{z}_i}, \quad (3)$$

$$L_{\mathbf{z}_i} = \frac{-1}{|\mathbf{Z}_i^{\text{pos}}|} \sum_{\mathbf{z}_j \in \mathbf{Z}_i^{\text{pos}}} \log \frac{\exp(\tilde{\mathbf{z}}_i \cdot \tilde{\mathbf{z}}_j / \tau)}{\sum_{\mathbf{z}_k \in \mathbf{Z} \setminus \{\mathbf{z}_i\}} \exp(\tilde{\mathbf{z}}_i \cdot \tilde{\mathbf{z}}_k / \tau)}, \quad (4)$$

where $|\mathbf{Z}|$ and $|\mathbf{Z}_i^{\text{pos}}|$ are the cardinalities of \mathbf{Z} and $\mathbf{Z}_i^{\text{pos}}$, respectively. τ is the temperature scaling parameter and $\tilde{\mathbf{z}}_i = \frac{\mathbf{z}_i}{\|\mathbf{z}_i\|}$ denotes normalized features. The contrastive loss optimizes the feature similarity to align instances for the positive set $\mathbf{Z}_i^{\text{pos}}$, and repulse otherwise.

We also designed the consistency loss to reduce the gap between the original and augmented domains at logit-level. To make consistent output in multi-domain (Hendrycks et al. 2019; Modas et al. 2022), we adopt Jensen-Shannon (JS) divergence as the consistency loss to reduce the gap between the original and OA-Mixed domains. JS divergence is a symmetric and smoothed version of the Kullback–Leibler (KL) divergence. The consistency loss is defined as:

$$L_{cs} = \frac{1}{2} (\mathbf{KL}[\mathbf{p} || \mathbf{M}] + \mathbf{KL}[\mathbf{p}_+ || \mathbf{M}]), \quad (5)$$

where \mathbf{p} and \mathbf{p}_+ are model predictions from original and OA-Mixed images, respectively. $\mathbf{M} = \frac{1}{2}(\mathbf{p} + \mathbf{p}_+)$ is the

Method	Clean	Noise			Blur			Weather				Digital				mPC	
		Gauss.	Shot	Impulse	Defocus	Glass	Motion	Zoom	Snow	Frost	Fog	Bright	Contrast	Elastic	Pixel	JPEG'	
Standard	42.2	0.5	1.1	1.1	17.2	16.5	18.3	2.1	2.2	12.3	29.8	32.0	24.1	40.1	18.7	15.1	15.4
+ Data augmentation																	
Cutout	42.5	0.6	1.2	1.2	17.8	15.9	18.9	2.0	2.5	13.6	29.8	32.3	24.6	40.1	18.9	15.6	15.7
Photo [†]	42.7	1.6	2.7	1.9	17.9	14.1	18.7	2.0	2.4	16.5	36.0	39.1	27.1	39.7	18.0	16.4	16.9
AutoAug-det [†]	42.4	0.9	1.6	0.9	16.8	14.4	18.9	2.0	1.9	16.0	32.9	35.2	26.3	39.4	17.9	11.6	15.8
AugMix	39.5	5.0	6.8	5.1	18.3	18.1	19.3	6.2	5.0	20.5	31.2	33.7	25.6	37.4	20.3	19.6	18.1
Stylized	36.3	4.8	6.8	4.3	19.5	18.7	18.5	2.7	3.5	17.0	30.5	31.9	22.7	33.9	22.6	20.8	17.2
OA-Mix (ours)	42.7	7.2	9.6	7.7	22.8	18.8	21.9	5.4	5.2	23.6	37.3	38.7	31.9	40.2	22.2	20.2	20.8
+ Loss function																	
SupCon	43.2	7.0	9.5	7.4	22.6	20.2	22.3	4.3	5.3	23.0	37.3	38.9	31.6	40.1	24.0	20.1	20.9
FSCE	43.1	7.4	10.2	8.2	23.3	20.3	21.5	4.8	5.6	23.6	37.1	38.0	31.9	40.0	23.2	20.4	21.0
OA-Loss (ours) = OA-DG	43.4	8.2	10.6	8.4	24.6	20.5	22.3	4.8	6.1	25.0	38.4	39.7	32.8	40.2	23.8	22.0	21.8

[†]We followed the searched policies (Zoph et al. 2020).

Table 1. Comparison with state-of-the-art methods on Cityscapes-C. For each corruption type, the average performance was calculated. mPC is an average performance of 15 corruption types.

mixture probability of predictions. The loss improves the generalization ability to classify the objects in OOD.

Training Objectives Our method trains the base detector in an end-to-end manner. Our OA-Loss L_{OA} consists of consistency loss L_{cs} and contrastive loss L_{ct} .

$$L_{OA} = L_{cs} + \gamma L_{ct}, \quad (6)$$

where γ is hyperparameter. OA-Loss can be added to the original loss L_{det} for the general detector. The joint training loss is defined as

$$L = L_{det} + \lambda L_{OA}, \quad (7)$$

where λ balances the scales of the losses. The joint training loss allows the model to learn object semantics and domain-invariant representations from OA-Mixed domains.

Experiments

In this section, we evaluate the robustness of our method against out-of-distribution. We also conduct ablation studies to verify the effectiveness of proposed modules.

Datasets

We evaluated the DG performance of our method in an urban scene for common corruptions and various weather conditions. Cityscapes-C (Michaelis et al. 2019) is a test benchmark to evaluate object detection robustness to corrupted domains. Cityscapes-C provides 15 corruptions on five severity levels. The corruptions are divided into four categories: noise, blur, weather, and digital. These corruption types are used to measure and understand the robustness against OOD (Dan and Thomas 2019). The common corruptions should be used only to evaluate the robustness of the model and are strictly prohibited to be used during training.

Diverse Weather Dataset (DWD) is an urban-scene detection benchmark to assess object detection robustness to various weather conditions. DWD collected data from

BDD-100k (2020), FoggyCityscapes (2018), and Adverse-Weather (2020) datasets. It consists of five different weather conditions: daytime-sunny, night-sunny, dusk-rainy, night-rainy, and daytime-foggy. Training should be conducted only using the daytime-sunny dataset and robustness is evaluated against other adverse weather datasets.

Evaluation Metrics

Following (Michaelis et al. 2019; Wu and Deng 2022), we evaluate the domain generalization performance of our method in various domains using mean average precision (mAP). The robustness against OOD is evaluated using mean performance under corruption (mPC), which is the average of mAPs for all corrupted domains and severity levels.

$$mPC = \frac{1}{N_C} \sum_{C=1}^{N_C} \frac{1}{N_S} \sum_{S=1}^{N_S} P_{C,S}, \quad (8)$$

where $P_{C,S}$ is the performance on the test data corrupted by corruption C at severity level S , N_C and N_S are the number of corruption and severity, respectively. N_C and N_S are set to 15 and 5 in Cityscapes-C; and 4 and 1 in DWD, respectively.

Robustness on Common Corruptions

Table 1 shows the performance of the state-of-the-art models on clean and corrupted domains. The baseline model is Faster R-CNN with ResNet-50 and feature pyramid networks (FPN). Each model is trained only on the clean domain and evaluated on both the clean and corrupted domains. Temperature scaling parameter τ for contrastive loss is set to 0.06. We set λ and γ to 10 and 0.001. More details can be seen in the supplementary material.

Cutout (Zhong et al. 2020) and photometric distortion (Redmon and Farhadi 2018) are data augmentation methods that can improve the generalization performance

	mAP					mPC
	Daytime-sunny	Night-sunny	Dusk-rainy	Night-rainy	Daytime-foggy	
Baseline	48.1	34.4	26.0	12.4	32.0	26.2
SW	50.6	33.4	26.3	13.7	30.8	26.1
IBN-Net	49.7	32.1	26.1	14.3	29.6	25.5
IterNorm	43.9	29.6	22.8	12.6	28.4	23.4
ISW	51.3	33.2	25.9	14.1	31.8	26.3
SHADE	-	33.9	29.5	16.8	33.4	28.4
CDSD	56.1	36.6	28.2	16.6	33.5	28.7
SRCD	-	36.7	28.8	17.0	35.9	29.6
Ours	55.8	38.0	33.9	16.8	38.3	31.8

Table 2. Comparison with state-of-the-art methods on Di-
verse Weather Dataset. mPC is an average performance of
OOD weathers: night-sunny, dusk-rainy, night-sunny, and
daytime-foggy.

of object detectors. However, they showed limited improvements in performance on corrupted domains. AutoAugment (Zoph et al. 2020) explores augmentation policies for the generalization of object detection. The method improved performance in the clean domain, but there was no significant performance improvement in the corrupted domain.

AugMix (Hendrycks et al. 2019) and stylized augmentation (Geirhos et al. 2018) are effective for S-DG in image classification. AugMix generates augmented images with unclear object locations, leading to a performance drop on the clean domain. Stylized augmentation improved performance on the corrupted domains with style transfer. However, the method did not consider affinity with the original domain, which decreased performance in the clean domain.

In contrast, OA-Mix maintains the clean domain performance with an object-aware mixing and improves the DG performance with multi-level transformations. Furthermore, we evaluated contrastive methods combined with OA-Mix. SupCon (Khosla et al. 2020) and FSCE (Sun et al. 2021) are contrastive learning methods that can reduce the multi-domain gap. The methods improved the clean performance, but the performance for the corrupted domain was not improved significantly. OA-Loss trains the semantic relations in an object-aware approach and achieved higher performance in both clean and corrupted domains. Consequently, OA-DG showed the best performance with 43.4 mAP and 21.8 mPC in the clean and corrupted domains, respectively.

Robustness on Various Weather Conditions

Table 2 shows the DG performances in real-world weather conditions through DWD. We follow the settings in the CDSD (Wu and Deng 2022) for DWD. We used Faster R-CNN with ResNet-101 backbone as a base object detector. Temperature scaling hyperparameter τ is set to 0.07. We set λ and γ to 10 and 0.001, respectively, for Faster R-CNN. More details can be seen in the supplementary material.

OA-DG is compared with state-of-the-art S-DGOD methods. SW (Pan et al. 2019), IBN-Net (Pan et al. 2018), IterNorm (Huang et al. 2019), ISW (Choi et al. 2021), and

Transformations	Mixing strategy	mAP	mPC
✗	✗	42.2	15.4
Single-level	✗	39.9	16.8
Multi-level	✗	40.5	17.5
Multi-level	Standard	41.6	19.7
Multi-level	Object-aware	42.7	20.8

Table 3. Ablation analysis of our proposed OA-Mix on Cityscapes-C dataset.

Consistency	Contrastive		mAP	mPC
	Target	Rule		
✗	✗	✗	42.7	20.8
✓	✗	✗	43.1	21.2
✓	fg	class-aware	43.4	21.3
✓	fg+bg	class-aware	43.3	21.1
✓	fg+bg	object-aware	43.4	21.8

Table 4. Ablation analysis of our proposed contrastive method on Cityscapes-C. Our method considers all the semantic relations of the same instance, foreground class, background class, and background instances. fg and bg denote foreground and background instances, respectively.

SHADE (Zhao et al. 2022) are feature normalization methods to improve DG. Compared with the baseline without any S-DGOD approaches, the feature normalization methods improve performance in certain weather conditions, but do not consistently enhance the performance for OOD. CDSD (Wu and Deng 2022) and SRCD (Rao et al. 2023) improved the performances in all weather conditions with domain-invariant features. Our proposed OA-DG achieved the top performance with 31.8 mPC in OOD weather conditions. Compared with the baseline, our method improves domain generalization in all weather conditions, especially in corrupted environments such as dusk-rainy and daytime-foggy.

Ablation Studies

All ablation studies evaluate DG performance in Cityscapes-C with the same settings as Table 1. Details and more ablation studies can be found in the supplementary material.

OA-Mix This experiment validates the impact of individual components within OA-Mix. Table 3 shows the clean performance and robustness according to multi-level transformations and object-aware mixing strategy. The single-level transformation improved performance on corrupted domains compared with the baseline. The multi-level transformation transforms the image at the local level and improved mPC compared with the single-level transformation. However, both transformations showed lower clean performance than the baseline due to the damage to the semantic features of objects. Standard mixing mitigates the negative effects of transformations. However, it does not utilize object information, which leads to lower clean performance than the baseline. Only the object-aware mixing strategy showed a clean performance above the baseline and



Figure 4. From the left to the right, the detection results of the clean, shot noise, defocus blur, and frost domains are visualized. Compared with the baseline, OA-DG detects the object more accurately in diverse corrupted domains.

Faster R-CNN		YOLOv3	
Baseline	OA-DG	Baseline	OA-DG
mAP	42.2	43.4 (+1.2)	34.6
mPC	15.4	21.8 (+6.4)	12.4
FLOPs (G)	408.7	409.0	194.0
			199.5

Table 5. Performance and computational complexity of Faster R-CNN and YOLOv3 on Cityscapes.

achieved the best mPC performance as well.

Contrastive Methods In Table 4, we verify the effectiveness of OA-Loss on Cityscapes-C. The consistency loss improved mPC by reducing the domain gap between the original and augmented instances. Then, we conducted an ablation study of contrastive learning according to the target and rule. The class-aware methods simply pull intra-class and push inter-class away. Both class-aware methods improved clean performance, but did not improve performance on the corrupted domains. In contrast, OA-Loss allows the model to train multi-domain gaps in an object-aware manner for all the instances. Our OA-Loss was improved by 0.7 mAP and 1.0 mPC in both clean and corrupted domains, respectively, validating the superiority of the object-aware approach.

Object Detector Architectures Table 5 shows the generalization capability of OA-DG to object detection framework. We conducted additional experiments with YOLOv3, a popular one-stage object detector. Although YOLOv3 used various bag of tricks such as photometric distortion to improve its generalization capability, OA-DG improved its performance in both clean and corrupted domains. This implies that our method is not limited to a specific detector architecture and can be applied to general object detectors.

Qualitative Analysis

Visual Analysis Figure 4 shows the results for the baseline and our model at clean and corrupted domains on Cityscapes. In the clean domain, both the baseline and our model accurately detect objects, but the baseline fails to detect most objects in the corrupted domain. Compared with the baseline, our model accurately detects small objects and improves object localization in OOD.

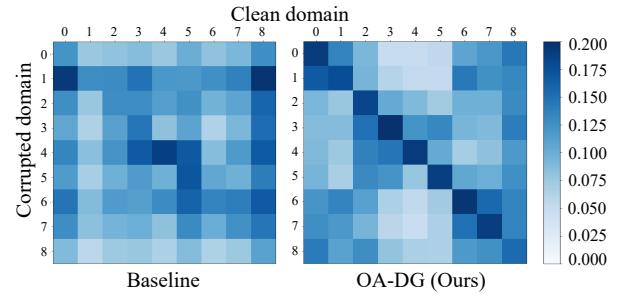


Figure 5. Feature correlations of the baseline and OA-DG between the clean and corrupted domains. Zero to seven on each axis denote the object classes and eight denotes the background class. Compared with the baseline, our method has higher correlations between the same class in OOD.

Feature Analysis Figure 5 shows the feature correlations between the clean and shot noise domains on Cityscapes. The feature correlations are measured as the average of cosine similarity for each class and normalized in the x -axis. Following (Ranasinghe et al. 2021), we use the features of the penultimate layer in the classification head. In Figure 5, the x -axis and y -axis represent the clean and corrupted domains, respectively. The comparison between the corrupted and clean domains shows that the baseline has little correlation between the same classes. This hinders the baseline from detecting objects in the corrupted domain, as shown in Figure 4. Compared with the baseline, the OA-DG method has higher correlations between the same class, but lower correlations with other classes.

Conclusion

In this study, we propose object-aware domain generalization (OA-DG) method for single-domain generalization in object detection. OA-Mix generates multi-domain with mixing strategies and preserved semantics. OA-Loss trains the domain-invariant representations for foreground and background instances from multi-domain. Our experimental results demonstrate that the proposed method can improve the robustness of object detectors to unseen domains.

Acknowledgements

This work was supported in part by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by Korea government (MSIT) (No.2020-0-00440, Development of Artificial Intelligence Technology that Continuously Improves Itself as the Situation Changes in the Real World). This work was supported in part by Korea Evaluation Institute Of Industrial Technology (KEIT) grant funded by the Korea government (MOTIE) (No.20023455, Development of Cooperate Mapping, Environment Recognition and Autonomous Driving Technology for Multi Mobile Robots Operating in Large-scale Indoor Workspace). The students are supported by the BK21 FOUR from the Ministry of Education (Republic of Korea).

References

Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, 213–229. Springer.

Choi, S.; Jung, S.; Yun, H.; Kim, J. T.; Kim, S.; and Choo, J. 2021. 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*, 11580–11590.

Dan, H.; and Thomas, D. 2019. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*.

Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.

Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2414–2423.

Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2018. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*.

Hassaballah, M.; Kenk, M. A.; Muhammad, K.; and Minaee, S. 2020. Vehicle detection and tracking in adverse weather using a deep learning framework. *IEEE Transactions on Intelligent Transportation Systems*, 22(7): 4230–4242.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 770–778.

Hendrycks, D.; Mu, N.; Cubuk, E. D.; Zoph, B.; Gilmer, J.; and Lakshminarayanan, B. 2019. AugMix: A simple data processing method to improve robustness and uncertainty. In *International Conference on Learning Representations*.

Huang, L.; Zhou, Y.; Zhu, F.; Liu, L.; and Shao, L. 2019. Iterative normalization: Beyond standardization towards efficient whitening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4874–4883.

Kang, B.; Li, Y.; Xie, S.; Yuan, Z.; and Feng, J. 2021. Exploring balanced feature spaces for representation learning. In *International Conference on Learning Representations*.

Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.

Kim, D.; Yoo, Y.; Park, S.; Kim, J.; and Lee, J. 2021. SelFReg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9619–9628.

Lee, W.; and Myung, H. 2022. Adversarial attack for asynchronous event-based data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1237–1244.

Li, L.; Gao, K.; Cao, J.; Huang, Z.; Weng, Y.; Mi, X.; Yu, Z.; Li, X.; and Xia, B. 2021. Progressive domain expansion network for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 224–233.

Li, T.; Cao, P.; Yuan, Y.; Fan, L.; Yang, Y.; Feris, R. S.; Indyk, P.; and Katabi, D. 2022. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6918–6928.

Lin, C.; Yuan, Z.; Zhao, S.; Sun, P.; Wang, C.; and Cai, J. 2021. Domain-invariant disentangled network for generalizable object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8771–8780.

Michaelis, C.; Mitzkus, B.; Geirhos, R.; Rusak, E.; Bringmann, O.; Ecker, A. S.; Bethge, M.; and Brendel, W. 2019. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv preprint arXiv:1907.07484*.

Modas, A.; Rade, R.; Ortiz-Jiménez, G.; Moosavi-Dezfooli, S.-M.; and Frossard, P. 2022. PRIME: A few primitives can boost robustness to common corruptions. In *Proceedings of the European Conference on Computer Vision*, 623–640. Springer.

Pan, X.; Luo, P.; Shi, J.; and Tang, X. 2018. Two at once: Enhancing learning and generalization capacities via IBN-net. In *Proceedings of the European Conference on Computer Vision*, 464–479.

Pan, X.; Zhan, X.; Shi, J.; Tang, X.; and Luo, P. 2019. Switchable whitening for deep representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1863–1871.

Ranasinghe, K.; Naseer, M.; Hayat, M.; Khan, S.; and Khan, F. S. 2021. Orthogonal projection loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12333–12343.

Rao, Z.; Guo, J.; Tang, L.; Huang, Y.; Ding, X.; and Guo, S. 2023. SRCD: Semantic reasoning with compound domains for single-domain generalized object detection. *arXiv preprint arXiv:2307.01750*.

Redmon, J.; and Farhadi, A. 2018. YOLOv3: An Incremental Improvement. *arXiv:1804.02767*.

Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28.

Sakaridis, C.; Dai, D.; and Van Gool, L. 2018. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126: 973–992.

Sun, B.; Li, B.; Cai, S.; Yuan, Y.; and Zhang, C. 2021. FSCE: Few-shot object detection via contrastive proposal encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7352–7362.

Vudit, V.; Engilberge, M.; and Salzmann, M. 2023. CLIP the gap: A single domain generalization approach for object detection. *arXiv preprint arXiv:2301.05499*.

Wan, C.; Shen, X.; Zhang, Y.; Yin, Z.; Tian, X.; Gao, F.; Huang, J.; and Hua, X.-S. 2022. Meta convolutional neural networks for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4682–4691.

Wang, X.; Huang, T. E.; Liu, B.; Yu, F.; Wang, X.; Gonzalez, J. E.; and Darrell, T. 2021a. Robust object detection via instance-level temporal cycle confusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9143–9152.

Wang, Z.; Luo, Y.; Qiu, R.; Huang, Z.; and Baktashmotlagh, M. 2021b. Learning to diversify for single domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 834–843.

Wu, A.; and Deng, C. 2022. 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*, 847–856.

Yao, X.; Bai, Y.; Zhang, X.; Zhang, Y.; Sun, Q.; Chen, R.; Li, R.; and Yu, B. 2022. PCL: Proxy-based contrastive learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7097–7107.

Yu, F.; Chen, H.; Wang, X.; Xian, W.; Chen, Y.; Liu, F.; Madhavan, V.; and Darrell, T. 2020. BDD100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2636–2645.

Zhao, Y.; Zhong, Z.; Zhao, N.; Sebe, N.; and Lee, G. H. 2022. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *Proceedings of the European Conference on Computer Vision*, 535–552. Springer.

Zhong, Z.; Zheng, L.; Kang, G.; Li, S.; and Yang, Y. 2020. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 13001–13008.

Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zoph, B.; Cubuk, E. D.; Ghiasi, G.; Lin, T.-Y.; Shlens, J.; and Le, Q. V. 2020. Learning data augmentation strategies for object detection. In *Proceedings of the European Conference on Computer Vision*, 566–583. Springer.