

Boosting Object Detection with Zero-Shot Day-Night Domain Adaptation

Zhipeng Du*

Department of Informatics
 King's College London
 zhipeng.du@kcl.ac.uk

Miaojing Shi†

College of Electronic and Information Engineering
 Tongji University
 mshi@tongji.edu.cn

Jiankang Deng

Department of Computing
 Imperial College London
 j.deng16@imperial.ac.uk

Abstract

Detecting objects in low-light scenarios presents a persistent challenge, as detectors trained on well-lit data exhibit significant performance degradation on low-light data due to low visibility. Previous methods mitigate this issue by exploring image enhancement or object detection techniques with real low-light image datasets. However, the progress is impeded by the inherent difficulties about collecting and annotating low-light images. To address this challenge, we propose to boost low-light object detection with zero-shot day-night domain adaptation, which aims to generalize a detector from well-lit scenarios to low-light ones without requiring real low-light data. Revisiting Retinex theory in the low-level vision, we first design a reflectance representation learning module to learn Retinex-based illumination invariance in images with a carefully designed illumination invariance reinforcement strategy. Next, an interchange-redecomposition-coherence procedure is introduced to improve over the vanilla Retinex image decomposition process by performing two sequential image decompositions and introducing a redecomposition cohering loss. Extensive experiments on ExDark, DARK FACE, and CODaN datasets show strong low-light generalizability of our method. Our code is available at <https://github.com/ZPDu/DAI-Net>.

1. Introduction

Object detection, aiming at identifying and localizing objects in an image, is a fundamental and well-investigated task in computer vision. Advanced detectors have achieved remarkable progress thanks to large-scale datasets such as COCO [30] and Open Images [25] for general object detection as well as WIDER FACE [75] for human face detection [11, 17, 37, 38]. However, these methods encounter severe performance degradation on low-light images resulting from dark environments, inadequate lighting, and exposure

*Work partially done during internship at Huawei London Research.

†Corresponding author

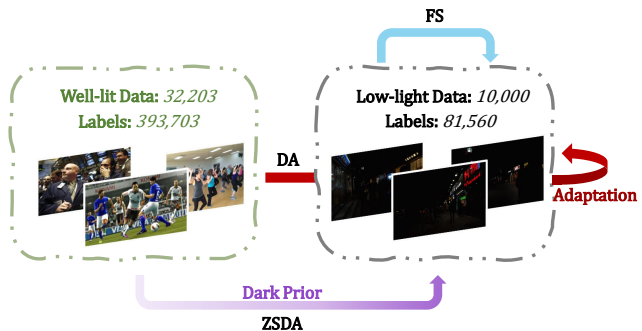


Figure 1. Left and right are respective well-lit and low-light datasets. FS: fully-supervised learning directly using low-light images and labels; DA: conventional domain adaptation with the access to low-light target domain data (and labels); ZSDA: zero-shot day-night domain adaptation without the access of domain data but only knowing the dark scenario as a prior.

time [72]. Issues such as low visibility, color distortion, and noises arise in low-light images and impede the accuracy of object detectors. To tackle this challenge, image enhancement methods are normally investigated to enhance the visibility of the scene under low light [4, 6, 16, 23, 44, 72–74]. Object detectors can benefit either through retraining or testing with light-enhanced images. Another direction is that detectors trained from well-lit images are fine-tuned on low-light images [58, 65–67].

Whilst existing image enhancement methods appear to be effective, they rely on a significant amount of low-light images collected from the real world. Many of them have to be trained with paired low-light and well-lit images [69, 70]. On the other hand, for dark object detection methods, the utilization of low-light images is also necessary [8, 65]. In contrast to the well-lit data, low-light images in established benchmarks [39, 77] are considerably sparser in amount and limited in terms of collecting scenarios. At the same time, annotating bounding boxes in low-light images is also undesirable due to low visibility. For instance, there are 32,203 images and 393,703 labelled faces in WIDER FACE, but only 10,000 images and 81,560 faces in DARK FACE. These difficulties in collecting and annotating low-

light data hinder the development of low-light image enhancement and object detection.

To circumvent the requirement for object detection in low-light scenarios, we propose to work in a zero-shot day-night domain adaptation setting [27, 40] as shown in Fig. 1, where object detectors are trained in the source domain with only well-lit images available and are evaluated in the low-light target domain with no images provided. The divergence between the well-lit source domain and the low-light target domain stems from the illumination variation and the corruptions it brings [4]. It is essential to emphasize that we are aware of the low-light scenario as our generalization target in this problem.

In this paper, we introduce a novel DArk-Illuminated Network, dubbed as DAI-Net, for low-light object detection. Given well-lit images from the source domain, we first utilize a physics-inspired low-illumination synthesis pipeline [8] to generate synthesized low-light images and form pairs with original well-lit images. Next, we revisit the Retinex theory [26] in low-level vision for the high-level detection task by decomposing low-light images into domain-invariant (*i.e.* image reflectance) and domain-specific (*i.e.* image illumination) information, where only the former should be retained to learn a generalizable detector. To this end, we build our framework upon an established object detection pipeline (*e.g.* DSFD [29], YOLOv3 [54]) and propose a reflectance representation learning module as an additional decoder. This module decodes reflectance-related illumination-invariant information from well-lit images and synthetic low-light images. It is optimized with pseudo ground truth produced by a pre-trained Retinex decomposition network (*e.g.* RetinexNet [69]) and is reinforced with a specifically designed illumination invariance reinforcement strategy. Next, we design an interchange-redemption-coherence procedure to improve the Retinex-based image decomposition process. It performs two times image decomposition sequentially by interchanging the decomposed well-lit/low-light reflectance for reconstructing the images and then redempting them. A redemption cohering loss is introduced to promote consistency between the produced reflectances in the two decompositions, so that the learned reflectance representation is stable and accurate.

In summary, to reduce the need for real-world low-light data, we propose to boost low-light object detection with zero-shot day-night domain adaptation. The contribution of this paper is threefold:

- We introduce a reflectance representation learning module additional to an established object detector to enhance its illumination-invariance. Specifically, an illumination invariance reinforcement strategy is designed to strengthen the learning process.
- We propose an interchange-redemption-coherence procedure to improve the vanilla image decomposition

process. A redemption cohering loss is introduced to maintain the consistency between the decomposition reflectances in sequence.

- We conduct extensive experiments on ExDark for general object detection, DARK FACE for face detection, and CODaN for image classification. Experimental results show that our method outperforms the state-of-the-art in multiple settings by large margins.

2. Related Works

2.1. Object Detection

Mainstream object detectors can be roughly classified into two categories: single-stage detectors, *e.g.* SSD [34], YOLO [54] and FCOS [61]; and two-stage detectors, *e.g.* Faster R-CNN [55] and R-FCN [10]. Single-stage detectors aim to directly predict object bounding boxes and class labels in a single step. Two-stage detectors first generate a set of region proposals and then perform classification and bounding box regression to refine the proposals. Recently, based on the transformer structure, DETR [5] and many DETR-based approaches [28, 33, 46, 84] are designed to further enhance general object detection.

In contrast, face detectors [11, 17, 21, 29, 36, 50, 60] are specialized detectors designed to detect human faces. Derived from general object detection, most face detection methods apply single-stage detectors to generate bounding boxes and corresponding classification scores simultaneously [11, 17, 29, 50, 60]. Many works focus on developing the model architectures [11, 17, 29, 35], anchor sampling/matching schemes [47, 60, 81], as well as feature enhancement techniques [21, 31, 50].

The recent prosperity of object detection should also be accredited to the availability of large-scale datasets. These datasets, such as COCO [30], Open Images [25] for general object detection, and WIDER FACE [75], FDDB [22] for face detection, allowing researchers with ample annotated examples to train and assess their detection models.

2.2. Low-Light Images

Low-Light Image Enhancement. Low-light image enhancement has been largely explored owing to the advent of deep learning. For instance, Guo *et al.* [16] propose zero-DCE to estimate image-specific light-enhancement curves without reference images. Many works [2, 4, 23, 32, 64, 69, 70, 78] base themselves on the Retinex theory [26], which assumes that an image can be decomposed into reflectance and illumination and can be reconstructed via element-wise multiplication of the latter two. Images can thus be enhanced by two major ways: using the reflectance as enhanced image [4, 18, 44, 64], or reconstructing an enhanced image with adjusted illumination [2, 23, 32, 69, 70, 80]. For instance, Wei *et al.* [69] conduct light enhancement and de-

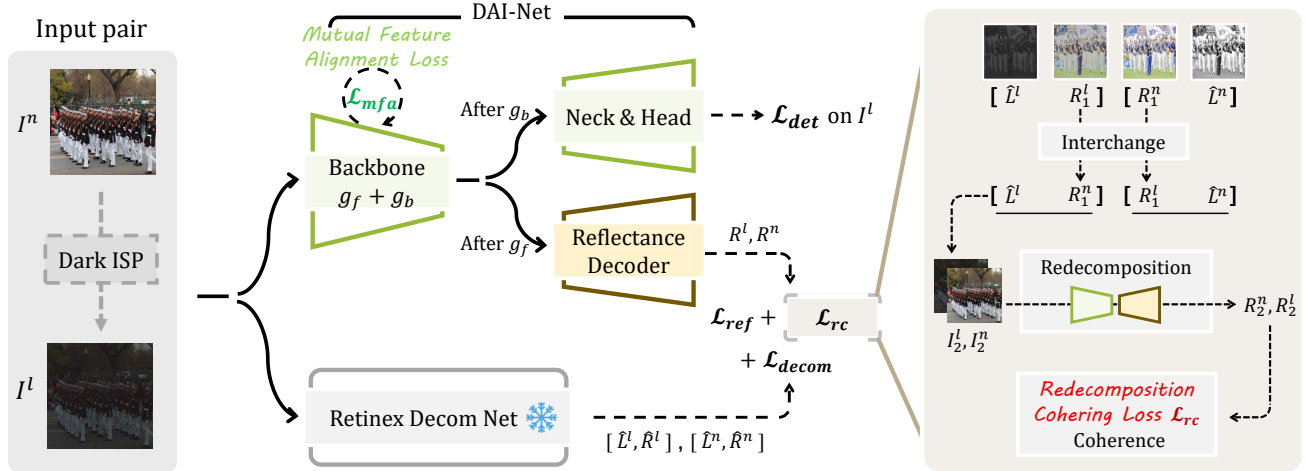


Figure 2. The structure overview of our method: we input a well-lit image and its corresponding synthesized low-light image as a pair into our framework. The pretrained Retinex decomposition net, shown as the gray block on the bottom, is frozen and only used during training to infer reflectance and illumination pseudo ground truth $\hat{R}^l, \hat{R}^n, \hat{L}^l, \hat{L}^n$ to supervise the reflectance decoder (Sec. 3.2). Specifically, illumination pseudo ground truth and the first-round reflectance predictions R_1^l, R_1^n are forwarded into the proposed interchange-redecomposition-coherence procedure on the right block for reconstructing and redecomposing second-round reflectance predictions R_2^l, R_2^n , as well as calculating redecomposition cohering loss \mathcal{L}_{rc} (Sec. 3.3). The inference is the same as the original detector that our method is built on, represented as green blocks.

noising based on reflectance and illumination maps decomposed from a Retinex decomposition model. Liu *et al.* [32] unroll the optimizing process of Retinex-based model and find the desired network structure with a cooperative bilevel search strategy. Cai *et al.* [4] simultaneously model corruptions in reflectance and illumination through a transformer structure.

Low-light Object Detection. Low-light object detection methods can be mainly categorized into three groups: detection-by-enhancement [16, 44, 52], enhancement-for-detection [45, 59, 67] and low-light detector learning strategies [8, 19, 58, 65]. In detection-by-enhancement, low-light enhancement methods are employed to generate illuminated images before detection [16, 44, 52]. In enhancement-for-detection, representative works are [19, 59] which propose image restoration training pipelines for improving object detection performance. For low-light detector learning strategies, multi-model merging [58], multi-task auto encoding transformation [8] and unsupervised domain adaptation framework [65] are introduced. In addition, there also have been efforts to construct low-light datasets, such as Nightowls [51], ExDark [39], DARK FACE [77], NOD [48]. However, these datasets are not as adequate as well-lit ones. Our method falls into the third category and addresses the concern for low-light data shortage in a zero-shot day-night domain adaptation setting.

Dark Domain Learning. Many perception tasks such as object segmentation and detection in dark scenarios have been investigated under domain transfer learning, which can be broadly categorized into domain adaptation and generalization. Domain adaptation (DA) methods train a model on

well-lit source domain and adapt it to the accessible low-light target domain data. Representative attempts include synthesizing low-light images [1, 8, 15, 49, 56, 57, 71], self-supervised learning for aligning well-lit and low-light domain distributions [65–67], merging components learned in two domains [58, 62], and many other multi-stage strategies [9, 12, 24, 57, 65, 79]. In contrast, domain generalization (DG) differs from DA by generalizing to unseen domains without knowing the knowledge of the target domain [13, 14, 53, 63]. The DG methods aim to provide general solutions to a broad range of potential target domains instead of focusing on one single domain. For low-light scenarios, a setting coined as zero-shot day-night domain adaptation (ZSDA) is introduced in [27, 40], which can actually be regarded as a special case of DA: real low-light data are inaccessible but the target domain is known to be the low-light scenario. Our work falls into the realm of ZSDA and we are the first to focus on ZSDA detection. A prior work, [65], also achieves low-light detection but belongs to unsupervised DA, which leverages unlabeled real low-light data, different from ZSDA.

3. Method

3.1. Overview

We train an object detector on well-lit images and generalize it to low-light images. Following [27, 40], this is a zero-shot day-night domain adaptation (ZSDA) problem from well-lit source domain to low-light target domain. The primary factors leading to the disparity between the two domains include illumination and the corruptions brought by

illumination degradation [4].

To resolve this gap, we focus on acquiring domain-invariant representations from the source domain that can transcend the target domain. This is specifically interpreted as learning illumination-invariant information, *i.e.* the reflectance representation, in the ZSDA setting. We illustrate the proposed framework in Fig. 2. During training, the framework consists of the proposed DArk-Illuminated Network (green and yellow blocks in Fig. 2) and a pretrained Retinex decomposition net (gray block in Fig. 2). The input to both networks is the same, which is a pair of a well-lit source domain image and its low-light counterpart synthesized by a physics-inspired low-light degradation synthesis pipeline [8], termed as Dark ISP in Fig. 2. Dark ISP is a non-deep method and is only used to create paired input for training our framework. It is noteworthy that training the detector directly on synthetic low-light images leads to much worse results than training on well-lit images (see Sec. 4.4). The proposed DArk-Illuminated Network (DAI-Net) is built on an established object detector with an additional decoder for reflectance representation learning (Sec. 3.2). A pretrained Retinex decomposition net is leveraged to provide pseudo ground truth for this decoder. An interchange-redecomposition-coherence procedure (Sec. 3.3) is introduced to further strengthen the reflectance representation learning.

3.2. Learning Retinex-based Reflectance Representation

We first revisit the Retinex theory [26], an image I can be decomposed into the reflectance R and illumination L , $I = R \cdot L$. The theory posits that the visibility of an image is influenced by illumination, while the reflectance remains unchanged. Sometimes the corruptions caused by the illumination variation are also jointly modeled, whilst we simply term them as part of illumination for convenience (Sec. 4.3). For object detection with ZSDA, we consider the reflectance as the illumination-invariant counterpart, the acquisition of reflectance knowledge becomes instrumental in achieving an illumination-invariant detector. To this end, we design a reflectance representation learning module to enhance the detector’s resilience to low-light effect.

Considering image decomposition is rather a low-level vision task, we split the detector backbone at the second Conv layer, and denote this frontal part of the backbone as g_f and the whole backbone as g_b . The output of g_f , *i.e.* feature F , encodes low-level information extracted by shallow layers, which is suitable for decoding the reflectance. We branch off the reflectance decoder after g_f , shown as the yellow block in Fig. 2. The decoder is a light module consisting of two Conv+ReLU layers. Since the detection head shares g_f with the reflectance decoder, the extracted illumination-invariant features also benefit the object detec-

tion task. Notice that we find learning image decomposition on the fly with object detection is unstable and could fail in some scenarios. We leverage a pretrained Retinex decomposition net (see Sec. 3.4 for details) to generate reflectance and illumination pseudo ground truth $\{\hat{R}, \hat{L}\}$ to supervise reflectance decoder in a more stable manner.

Illumination invariance reinforcement. We further introduce an illumination invariance reinforcement scheme to reinforce the illumination invariance of the detector from feature level. Illumination discrepancy between the paired well-lit and synthesized low-light input images can lead to different feature distributions. But the two images are in essence with the same semantic information. Since our target is to learn illumination invariant representation, we require the output features F from g_f forwarded into the reflectance decoder to be closely aligned between well-lit and low-light images. We explicitly match the well-lit and low-light features F^n and F^l extracted from g_f by designing a mutual feature alignment loss as follows:

$$\mathcal{L}_{mfa} = \mathcal{KL}(F^n || F^l) + \mathcal{KL}(F^l || F^n), \quad (1)$$

where $\mathcal{KL}(\cdot || \cdot)$ refers to KL-Divergence, F^l and F^n are flattened and spatially averaged features from g_f w.r.t. well-lit and synthesized low-light images, respectively.

3.3. Interchange-Redecomposition-Coherence

To further enhance the reflectance learning, we aim at designing a stronger image decomposition process. Given a pair of low-light image I^l and well-lit image I^n , a typical Retinex-based image decomposition process [4, 23, 32, 64, 70, 78] decomposes them into corresponding reflectance and illumination, *i.e.* low-light reflectance R_1^l and illumination L^l for I^l and well-lit reflectance R_1^n and illumination L^n for I^n . Both reflectance R_1^l and R_1^n (they should be ideally the same) should be interchangeable with each other to reconstruct I^n, I^l when combined with the corresponding illumination map L^n, L^l . Thanks to this interchangeability, we can add a constraint to strengthen the image decomposition and reflectance representation learning. One intuitive approach is to impose a penalty on the images that are reconstructed using the interchanged reflectance once they deviate from the original inputs, as proposed in [69]. This is considered as a vanilla penalty loss constraint. To fully harness the information produced in image decomposition, we propose an interchange-redecomposition-coherence procedure, which is depicted in the right of Fig. 2. First, we interchange the reflectance between well-lit and low-light images and reconstruct the images as $I_2^l = R_1^n \cdot L^l, I_2^n = R_1^l \cdot L^n$. The reconstructed images can be decomposed into a second round. As the emphasis of our DAI-Net is to learn the illumination-invariant part of the image (the reflectance), we decompose reflectances R_2^n, R_2^l from I_2^n, I_2^l using the same reflectance decoding branch in DAI-Net

as in the first round and request them to be coherent to corresponding reflectances R_1^l, R_1^n decomposed in the first round. We, therefore, introduce the redecomposition cohering loss:

$$\mathcal{L}_{rc} = \|R_1^n - R_2^l\|_1 + \|R_1^l - R_2^n\|_1 \quad (2)$$

Compared with the vanilla penalty loss, \mathcal{L}_{rc} involves a redecomposition process and takes full advantage of the interchangeability of reflectance.

3.4. Network Training

Retinex decomposition net. The Retinex decomposition net is an off-the-shelf network that can be based on any image decomposition networks (*e.g.* [4, 69]). We can train this network from scratch using our paired input (default setting), or directly load pretrained weights that are publicly available. Afterwards, we freeze it to only infer the reflectance and illumination of the input as pseudo ground truth $\{\hat{R}, \hat{L}\}$ during the learning of the DAI-Net. We choose a basic structure [69] by default, as perfect pseudo labels are not a necessity for enabling good detection performance.

DAI-Net. The DAI-Net consists of a detection branch and a reflectance decoding branch. For the former, we utilize the detection losses of the selected detector, denoted by \mathcal{L}_{det} . While the objective functions of reflectance decoding branch are summarized into three parts. The first part consists of the two proposed losses $\mathcal{L}_{mfa}, \mathcal{L}_{rc}$ (Eq. 1 and 2). The other two are:

Reflectance learning loss. We supervise the reflectance decoder output R by the pseudo ground truth \hat{R} through a reflectance learning loss $\mathcal{L}_{ref} = \text{MAE}(R, \hat{R}) + (1 - \text{SSIM}(R, \hat{R}))$ where MAE is the Mean Absolute Error and SSIM is the Structural Similarity Index Measure. (R, \hat{R}) is realized as (R^l, \hat{R}^l) or (R^n, \hat{R}^n) in practice.

Image decomposition loss. We further strengthen the reflectance learning through image decomposition loss. We apply a representative decomposition loss \mathcal{L}_{decom} , as in [69], upon \hat{L} and R . Specifically, the loss is an integration of image reconstruction loss \mathcal{L}_{recon} , invariant reflectance loss \mathcal{L}_{ir} , and illumination smoothness loss \mathcal{L}_{smooth} . \mathcal{L}_{recon} is to let $R \cdot \hat{L}$ reconstruct the input image I , specifically, $R^l \cdot \hat{L}^l$ for I^l and $R^n \cdot \hat{L}^n$ for I^n ; \mathcal{L}_{smooth} and \mathcal{L}_{ir} are computed between the paired input. $\mathcal{L}_{ir} = \text{MSE}(R^l, R^n) + (1 - \text{SSIM}(R^l, R^n))$ enforces the predicted well-lit reflectance and low-light reflectance to be the same, in the form of a combination of Mean Squared Error and Structure Similarity Index Measure [68]. Therefore, we have

$$\mathcal{L}_{decom} = \mathcal{L}_{recon} + \lambda_{smooth}\mathcal{L}_{smooth} + \lambda_{ir}\mathcal{L}_{ir}, \quad (3)$$

where $\lambda_{smooth}, \lambda_{ir}$ represents loss weight for corresponding loss, respectively.

The total objective function of DAI-Net can be formulated as:

Category	Method	mAP(%)
WIDER FACE→DARK FACE test set using DSFD		
Generalization	Faster-RCNN [55]	1.7
	SSH [50]	6.9
	RetinaFace [11]	8.6
	SRN [7]	9.0
	SFA [41]	9.3
	ParamidBox [60]	12.5
	SmallHardFace [83]	16.1
	DSFD [29]	16.1
Zero-shot Adaptation	CICov [27]	18.4
	Sim-MinMax [40]	25.7
	DAI-Net	28.0
Fully Supervised	Fine-tuned DSFD [65]	46.0
	Fine-tuned DAI-Net	52.9
COCO→DARK FACE validation set [8] using YOLOv3		
Pretrained and Tuning	YOLO _N	48.3
	YOLO _N +MBLLEN [43]	51.6
	YOLO _N +KIND [82]	51.6
	YOLO _N +Zero-DCE [16]	54.2
	YOLO _L	54.0
	MAET [8]	55.8
	DAI-Net	57.0

Table 1. Comparison with state of the art on DARK FACE test and validation sets.

$$\mathcal{L} = \mathcal{L}_{det} + \lambda_{mfa}\mathcal{L}_{mfa} + \lambda_{rc}\mathcal{L}_{rc} + \mathcal{L}_{ref} + \mathcal{L}_{decom} \quad (4)$$

where λ_{mfa} refers to loss weight for mutual feature alignment loss in Eq. 1 and λ_{rc} refers to loss weight for redecomposition cohering loss in Eq. 2.

4. Experiments

In this section, we first provide settings and results on three tasks: face detection (Sec. 4.1), object detection (Sec. 4.2) and image classification (Sec. 4.3) in darkness. The first two tasks validate the proposed detection method, while the last one shows the generalizability of our approach. Furthermore, we ablate the proposed components and conduct corresponding analyses.

4.1. Face Detection in Darkness

Settings. We conduct experiments by selecting HLA-Face [65] and MAET [8] as comparable methods. They both reproduce and report the results of a number of methods in low-light settings. Based on their configurations, we select WIDER FACE [75] and COCO 2017 [30] as two separate well-lit source domains, which contain over 100K and 30K images, respectively. We test the trained models on the target domain DARK FACE containing around 10K images. These are denoted as WIDER FACE/COCO→DARK FACE (from source to target). We consistently adopt the same base face detector DSFD [29]/YOLOv3 [54] and plug

Method	Bicycle	Boat	Bottle	Bus	Car	Cat	Chair	Cup	Dog	Motorbike	People	Table	Total
YOLO _N	71.8	64.5	63.9	81.6	76.8	55.4	49.7	56.8	63.8	61.8	65.7	40.5	62.7
+KinD [82]	73.4	68.1	65.5	86.2	78.3	63.0	56.9	62.7	68.2	67.1	69.6	48.2	67.3
+Zero-DCE [16]	79.5	71.3	70.4	89.0	80.7	68.4	65.7	68.6	75.4	67.2	76.2	51.1	72.0
YOLO _L	78.2	70.8	72.3	88.1	80.7	67.9	62.4	70.5	74.8	69.4	75.8	50.9	71.6
MAET [8]	81.3	71.6	74.5	89.7	82.1	69.5	65.5	72.6	75.4	72.7	77.4	53.3	74.0
Ours	83.8	75.8	75.1	94.2	84.1	74.9	73.1	79.2	82.2	76.4	80.7	59.8	78.3

Table 2. Comparison with state of the art on ExDark.

Method	Top-1 (%)
MAET [8]	56.48
CICov [27]	60.32
Sim-MinMax [40]	65.87
Ours	68.44

Table 3. Task generalization on CODaN.

DAI-Net on top of it. We train the network following the same training process outlined in [8, 29]. We build Retinex decomposition net using RetinexNet [69] and pretrain it as described in Sec. 3.4 for 10 epochs with a learning rate of $1e-3$. Loss weights λ_{smooth} and λ_{ir} are set to 0.5 and 0.01, respectively. λ_{rc} and λ_{mfa} in Eq. 4 are 0.001 and 0.1. The optimal weights for λ_{smooth} and λ_{ir} can be easily found around their default settings in [69]. We mainly tune λ_{rc} and λ_{mfa} on the WIDER FACE validation set.

We evaluate our method using the commonly used metric in object detection, *i.e.* mean Average Precision(mAP). We compare our results with results presented in HLA-Face [65] and MAET [8] in Table 1. We use the same multi-scale scheme as in [29, 65] to DAI-Net for testing.

WIDER FACE→DARK FACE. We compare with the results in HLA-Face [65] using DSFD as they do. Generalization. In the first group of Table 1, comparable methods are trained on the source domain and directly evaluated on the target domain. DSFD achieves the best generalizability, which we base our framework on. Our method further improves upon DSFD by a large margin.

Zero-shot Adaptation. Under this setting, the methods are aware that the target domain is the dark scenario. We directly compare with ZSDA methods [27, 40] by applying them to DSFD for detection. Ours achieves the best result, showing a strong dark domain generalizability.

Fully Supervised. In this group, the trained models from the first group (Generalization) are fine-tuned in their respective object detection learning manners on the training set of DARK FACE with ground truth labels. Our fine-tuned DAI-Net still outperforms other comparable methods by a large margin. Our method hence contributes to a new baseline for this dataset for future development.

Method		Method									mAP(%)
DISP		RD			DP		\mathcal{L}_{mfa}				
		R	L	R+L	\mathcal{L}_p	\mathcal{L}_{rc}	KL	L1	L2		
-	-	-	-	-	-	-	-	-	-	-	15.2
✓	-	-	-	-	-	-	-	-	-	-	11.6
✓	✓	-	-	-	-	-	-	-	-	-	20.5
✓	-	✓	-	-	-	-	-	-	-	-	16.5
✓	-	-	✓	-	-	-	-	-	-	-	19.0
✓	✓	-	-	-	✓	-	-	-	-	-	21.8
✓	✓	-	-	-	-	✓	-	-	-	-	22.3
✓	✓	-	-	-	✓	✓	-	-	-	-	22.0
✓	✓	-	-	-	-	✓	✓	-	-	-	23.5
✓	✓	-	-	-	-	-	-	✓	-	-	20.7
✓	✓	-	-	-	-	✓	-	-	-	✓	21.1

Table 4. Ablation Study of proposed components on DARK FACE.

COCO→DARK FACE. We compare with MAET [8] using YOLOv3 trained on COCO and finetuned the same way on DARK FACE with labels. The results in MAET [8] are reported on a specifically sampled validation set. For a direct comparison with it, we use its published train/val division to have 5500 training images for finetuning DAI-Net pretrained on COCO and 500 validation images to evaluate it. The result is shown in Table 1, under the category of COCO→DARK FACE. Our model performs clearly better than MAET and other comparable methods reported in [8].

4.2. Object Detection in Darkness

Settings. Following MAET [8], we use YOLOv3 as the object detector and take COCO as the source domain and ExDark [39] as the target domain. ExDark has around 7K dark images. The settings of hyperparameters are consistent as in Sec. 4.1.

Results. In Table 2, we compare our results to the state of the art. YOLO_N and YOLO_L refer to training YOLOv3 based DAI-Net on original well-lit images of COCO and synthesized low-light images generated by the same Dark ISP [8], respectively. Our method shows superior performance on all object categories and the best overall result. For instance, our method enhances mAP on the difficult class *Chair* by 7.6%. This implies our proposed modules




Figure 3. Feature visualization. The original image, backbone features from DSFD [29], and our method are given in the three rows.

in DAI-Net can be easily plugged into general object detectors to boost their performance in dark object detection.

4.3. Generalization to Dark Image Classification

We show the generalizability of our approach beyond the dark object detection task by applying our method to the dark image classification task. This allows us to directly compare with the reported results in ZSDA methods [4, 27].

Settings. Following Sim-MinMax [4], we conduct experiments on a nighttime image classification dataset CODaN [27]. CODaN covers 10000 well-lit training samples, 2500 well-lit test samples for validation, and 2500 low-light test samples. We use the same ResNet-18 [20] backbone. The hyperparameters of our method remain as in Sec. 4.1.

Results. The results on CODaN are shown in Table 3. Our results surpass the previous best by a large margin. This demonstrates ours can be effectively applied to other tasks.

4.4. Ablation Study

We conduct ablation studies on WIDER FACE \rightarrow DARK FACE for face detection in darkness to validate the effectiveness of proposed components. The multi-scale testing strategy is not used for evaluation efficiency. More experiments are provided in supplementary materials.

Baseline. Our baseline model is the vanilla DSFD model [29] where we do not add any component proposed in DAI-Net. We report its result in the first row of Table 4. We apply the Dark ISP [8] to synthesize low-light images, which is denoted as the DISP column in Table 4. Directly training on synthesized images degrades the precision of DSFD from 15.2% to 11.6%, implying that synthesized low-light images are not of the same attributes as real low-light data.

Reflectance Decoding. We add the reflectance decoding branch to the baseline. This is denoted as R under column RD of Table 4. The result is shown in the third row in Table 4. The performance is largely boosted from 15.2% to 20.5%. Adding the detector with an auxiliary head to generate reflectance helps to enhance the illumination-invariant

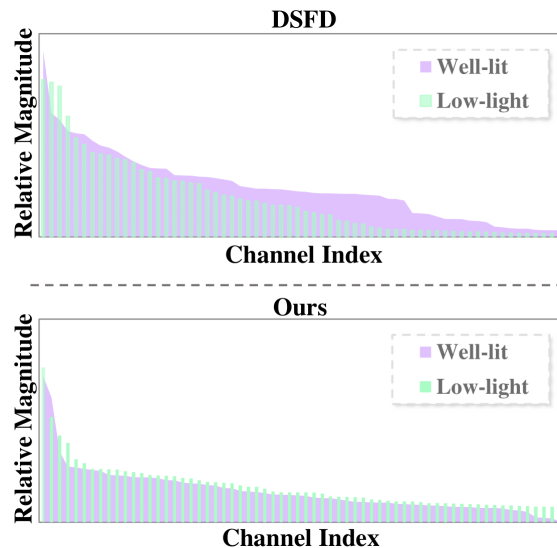


Figure 4. Mean magnitude of feature channels before and after applying our method. We use discrete and continuous curves for the two distributions respectively only for the visualization purpose.

information in the detector, and thus leads to more robust detection under various illuminations.

Learning Target. Based on Retinex theory, we identify the reflectance as an illumination-invariant target. To validate this, we try other variants such as adding auxiliary branches to DAI-Net for decoding illumination or decoding both reflectance and illumination, denoted by respectively adding L and R+L under column RD in Table 4. Both variants are inferior to our original version, which decodes only the reflectance.

Decomposition Processes. We investigate the effectiveness of the interchange-redecomposition-coherence procedure by comparing redecomposition cohering loss with the vanilla penalty loss [69] discussed in Sec. 3.3. In Table 4, we present the results among the sixth to eighth rows.

Image Decomposition Loss. We first apply the representative penalty loss (\mathcal{L}_p , see Sec. 3.3) to replace \mathcal{L}_{rc} in DAI-Net. Table 4 shows that \mathcal{L}_p improves the mAP by 1.3% to

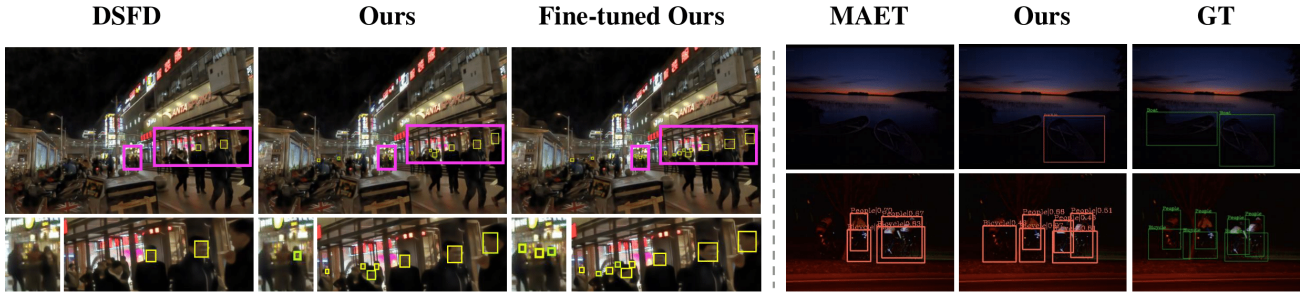


Figure 5. Left: Qualitative examples of DSFD [29], ours and fine-tuned ours (see Sec. 4.1) on DARK FACE. DARK FACE images are enhanced only for visualization. Right: Qualitative examples of MAET [8], ours and the ground truth on ExDark (zoom in for details).

21.8% since it enforces reflectance interchangeability.

Decomposition Cohering Loss. We then utilize the proposed redecomposition cohering loss during training. This is denoted by \mathcal{L}_{rc} under column DP in Table 4. The proposed loss boosts the performance of not using decomposition process from 20.5% to 22.3%, which is superior to the image decomposition loss. This is because we fully exploit the reflectance interchangeability by redecomposing and cohering the sequentially decomposed results.

Simultaneous Usage. We can also use both losses. The result in Table 4 displays that it hardly brings improvement than solely using our redecomposition cohering loss. This indicates that the effect of decomposition loss is essentially covered by the redecomposition cohering loss.

More Redecomposition. We have also tried to implement the decomposition three times. The performance is hardly improved. Considering the computation efficiency, we stick to the two-times decomposition.

Mutual Feature Alignment Loss. The mutual feature alignment loss in Eq. 1 exerts a constraint on illumination invariance on feature level. Under the column \mathcal{L}_{mfa} in Table 4, we show that \mathcal{L}_{mfa} using KL-Divergence as our original design increases mAP from 22.3% to 23.5%, which validates its effectiveness.

Choice of loss functions. We further replace KL-Divergence in Eq. 1 with L1- and L2-Distance, denoted by L1 and L2 in Table 4. Both degrade performance because of their harsh constraints.

Feature Visualization. *Feature map visualization.* We extract backbone features from g_f in face detection model [29] w.r.t. a given image; then transform the features into a single-channel map by applying channel-wise averaging. We visualize these maps in Fig. 3 to show the improvement of our method visually on DARK FACE; furthermore, we visualize them on real paired well-lit and low-light images from the LOL-v2 dataset [76] for the observation of their reflectance consistency. Compared with DSFD [29], our feature map captures clearer and more detailed object information in extremely low-light environments. Though sharing the same architecture and workflow with DSFD during in-

ference, reflectance-aware information is implicitly injected into ours, contributing to higher-quality features in the dark.

Feature channel visualization. We calculate the mean magnitude of every feature channel from all g_f features as [3, 42]. Mean magnitudes in Fig. 4 are calculated on well-lit samples (purple mesh) in WIDER FACE and low-light samples (green histograms) in DARK FACE. In a vanilla DSFD trained on WIDER FACE, the mean magnitudes of channels w.r.t. well-lit and low-light images are misaligned with each other. Some channels of low-light features are suppressed compared with well-lit features. This illustrates how illumination degrades a detector from a feature channel perspective. In DAI-Net, the relative mean magnitudes are consistent across two sets, validating our network learns illumination-invariant information that does not change along illumination.

Qualitative Examples. Some qualitative results on DARK FACE and ExDark are given in Fig. 5. More visualized results are given in the supplementary material.

5. Conclusion

This paper studies dark object detection in a new setting named zero-shot dark domain generalization. We propose a novel DArk-Illuminated Network (DAI-Net) to learn illumination invariant representations from a well-lit source domain. We first introduce a reflectance representation learning module to extract illumination-invariant feature representations from well-lit images. In the module, a reflectance decoding branch and a mutual feature alignment loss are designed to enforce illumination invariance on both image- and feature-level. Next, we improve the Retinex image decomposition process by designing an interchange-redecomposition-coherence procedure. We perform two times sequential decomposition, where the input of the latter is obtained by interchanging reflectance outputs from the former. At last, a redecomposition cohering loss is designed to enforce the consistency between decomposition results over the two sequential decompositions. Experiments on DARK FACE and ExDark validate the strong dark domain generalizability of DAI-Net.

References

- [1] Vinicius F Arruda, Thiago M Paixao, Rodrigo F Berriel, Alberto F De Souza, Claudine Badue, Nicu Sebe, and Thiago Oliveira-Santos. Cross-domain car detection using unsupervised image-to-image translation: From day to night. In *IJCNN*, 2019. [3](#)
- [2] Bolun Cai, Xianming Xu, Kailing Guo, Kui Jia, Bin Hu, and Dacheng Tao. A joint intrinsic-extrinsic prior model for retinex. In *ICCV*, 2017. [2](#)
- [3] Ruisi Cai, Zhenyu Zhang, Tianlong Chen, Xiaohan Chen, and Zhangyang Wang. Randomized channel shuffling: Minimal-overhead backdoor attack detection without clean datasets. *NeurIPS*, 2022. [8](#)
- [4] Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. 2023. [1, 2, 3, 4, 5, 7](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. [2](#)
- [6] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *CVPR*, 2018. [1](#)
- [7] Cheng Chi, Shifeng Zhang, Junliang Xing, Zhen Lei, Stan Z Li, and Xudong Zou. Selective refinement network for high performance face detection. In *AAAI*, 2019. [5](#)
- [8] Ziteng Cui, Guo-Jun Qi, Lin Gu, Shaodi You, Zenghui Zhang, and Tatsuya Harada. Multitask aet with orthogonal tangent regularity for dark object detection. In *ICCV*, 2021. [1, 2, 3, 4, 5, 6, 7, 8](#)
- [9] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3819–3824. IEEE, 2018. [3](#)
- [10] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. *NeurIPS*, 2016. [2](#)
- [11] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020. [1, 2, 5](#)
- [12] Xueqing Deng, Peng Wang, Xiaochen Lian, and Shawn Newsam. Nightlab: A dual-level architecture with hardness detection for segmentation at night. In *CVPR*, 2022. [3](#)
- [13] Zhipeng Du, Jiankang Deng, and Miaojing Shi. Domain-general crowd counting in unseen scenarios. In *AAAI*, 2023. [3](#)
- [14] 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 *ICLR*, 2023. [3](#)
- [15] Huan Gao, Jichang Guo, Guoli Wang, and Qian Zhang. Cross-domain correlation distillation for unsupervised domain adaptation in nighttime semantic segmentation. In *CVPR*, 2022. [3](#)
- [16] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *CVPR*, 2020. [1, 2, 3, 5, 6](#)
- [17] Jia Guo, Jiankang Deng, Alexandros Lattas, and Stefanos Zafeiriou. Sample and computation redistribution for efficient face detection. In *ICLR*, 2022. [1, 2](#)
- [18] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on image processing*, 26(2):982–993, 2016. [2](#)
- [19] Khurram Azeem Hashmi, Goutham Kallempudi, Didier Stricker, and Muhammad Zeshan Afzal. Featenhancer: Enhancing hierarchical features for object detection and beyond under low-light vision. In *ICCV*, 2023. [3](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [7](#)
- [21] Peiyun Hu and Deva Ramanan. Finding tiny faces. In *CVPR*, 2017. [2](#)
- [22] Vidit Jain and Erik Learned-Miller. Fddb: A benchmark for face detection in unconstrained settings. Technical report, UMass Amherst technical report, 2010. [2](#)
- [23] Yeying Jin, Wenhan Yang, and Robby T Tan. Unsupervised night image enhancement: When layer decomposition meets light-effects suppression. In *ECCV*, 2022. [1, 2, 4](#)
- [24] Mikhail Kennerley, Jian-Gang Wang, Bharadwaj Veeravalli, and Robby T Tan. 2pcnet: Two-phase consistency training for day-to-night unsupervised domain adaptive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11484–11493, 2023. [3](#)
- [25] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. [1, 2](#)
- [26] Edwin H Land. The retinex theory of color vision. *Scientific american*, 237(6):108–129, 1977. [2, 4](#)
- [27] Attila Lengyel, Sourav Garg, Michael Milford, and Jan C van Gemert. Zero-shot day-night domain adaptation with a physics prior. In *ICCV*, 2021. [2, 3, 5, 6, 7](#)
- [28] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *CVPR*, 2022. [2](#)
- [29] Jian Li, Yabiao Wang, Changan Wang, Ying Tai, Jianjun Qian, Jian Yang, Chengjie Wang, Jilin Li, and Feiyue Huang. Dsfed: dual shot face detector. In *CVPR*, 2019. [2, 5, 6, 7, 8, 1](#)
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [1, 2, 5](#)
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. [2](#)
- [32] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *CVPR*, 2021. [2, 3, 4](#)

- [33] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. In *ICLR*, 2022. 2
- [34] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 2
- [35] Yang Liu and Xu Tang. Bfbox: Searching face-appropriate backbone and feature pyramid network for face detector. In *CVPR*, 2020. 2
- [36] Yuting Liu, Miaojing Shi, Qijun Zhao, and Xiaofang Wang. Point in, box out: Beyond counting persons in crowds. In *CVPR*, 2019. 2
- [37] Yang Liu, Fei Wang, Jiankang Deng, Zhipeng Zhou, Baigui Sun, and Hao Li. Mogface: Towards a deeper appreciation on face detection. In *CVPR*, 2022. 1
- [38] Yang Liu, Jiankang Deng, Fei Wang, Lei Shang, Xuansong Xie, and Baigui Sun. Damofd: Digging into backbone design on face detection. In *ICLR*, 2023. 1
- [39] Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178:30–42, 2019. 1, 3, 6
- [40] Rundong Luo, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Similarity min-max: Zero-shot day-night domain adaptation. In *ICCV*, 2023. 2, 3, 5, 6
- [41] Shi Luo, Xiongfei Li, Rui Zhu, and Xiaoli Zhang. Sfa: small faces attention face detector. *IEEE Access*, 7:171609–171620, 2019. 5
- [42] Xu Luo, Jing Xu, and Zenglin Xu. Channel importance matters in few-shot image classification. In *ICCV*, 2022. 8
- [43] Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. Mbllen: Low-light image/video enhancement using cnns. In *BMVC*, page 4, 2018. 5
- [44] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *CVPR*, 2022. 1, 2, 3
- [45] Tengyu Ma, Long Ma, Xin Fan, Zhongxuan Luo, and Risheng Liu. Pia: Parallel architecture with illumination allocator for joint enhancement and detection in low-light. In *ACM MM*, 2022. 3
- [46] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional detr for fast training convergence. In *ICCV*, 2021. 2
- [47] Xiang Ming, Fangyun Wei, Ting Zhang, Dong Chen, and Fang Wen. Group sampling for scale invariant face detection. In *CVPR*, 2019. 2
- [48] Igor Morawski, Yu-An Chen, Yu-Sheng Lin, and Winston H Hsu. Nod: Taking a closer look at detection under extreme low-light conditions with night object detection dataset. 2021. 3
- [49] Amitangshu Mukherjee, Ameya Joshi, Anuj Sharma, Chinmay Hegde, and Soumik Sarkar. Generative semantic domain adaptation for perception in autonomous driving. *Journal of big data analytics in transportation*, 4(2-3):103–117, 2022. 3
- [50] Mahyar Najibi, Pouya Samangouei, Rama Chellappa, and Larry S Davis. Ssh: Single stage headless face detector. In *ICCV*, 2017. 2, 5
- [51] Lukáš Neumann, Michelle Karg, Shanshan Zhang, Christian Scharfenberger, Eric Piegert, Sarah Mistr, Olga Prokofyeva, Robert Thiel, Andrea Vedaldi, Andrew Zisserman, et al. Nightowls: A pedestrians at night dataset. In *ACCV*, 2019. 3
- [52] Qingpao Qin, Kan Chang, Mengyuan Huang, and Guiqing Li. Denet: Detection-driven enhancement network for object detection under adverse weather conditions. In *ACCV*, 2022. 3
- [53] Nikhil Reddy, Abhinav Singhal, Abhishek Kumar, Mahsa Baktashmotlagh, and Chetan Arora. Master of all: Simultaneous generalization of urban-scene segmentation to all adverse weather conditions. In *ECCV*, 2022. 3
- [54] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2, 5
- [55] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015. 2, 5
- [56] Eduardo Romera, Luis M Bergasa, Kailun Yang, Jose M Alvarez, and Rafael Barea. Bridging the day and night domain gap for semantic segmentation. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1312–1318. IEEE, 2019. 3
- [57] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *ICCV*, 2019. 3
- [58] Yukihiro Sasagawa and Hajime Nagahara. Yolo in the dark-domain adaptation method for merging multiple models. In *ECCV*, 2020. 1, 3
- [59] Shangquan Sun, Wenqi Ren, Tao Wang, and Xiaochun Cao. Rethinking image restoration for object detection. *NeurIPS*, 2022. 3
- [60] Xu Tang, Daniel K Du, Zeqiang He, and Jingtuo Liu. Pyramidbox: A context-assisted single shot face detector. In *ECCV*, 2018. 2, 5
- [61] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. 2
- [62] Madhu Vankadari, Sourav Garg, Anima Majumder, Swagat Kumar, and Ardhendu Behera. Unsupervised monocular depth estimation for night-time images using adversarial domain feature adaptation. In *ECCV*, 2020. 3
- [63] Vidit Vidit, Martin Engilberge, and Mathieu Salzmann. Clip the gap: A single domain generalization approach for object detection. In *CVPR*, 2023. 3
- [64] Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. Underexposed photo enhancement using deep illumination estimation. In *CVPR*, 2019. 2, 4
- [65] Wenjing Wang, Wenhan Yang, and Jiaying Liu. Hla-face: Joint high-low adaptation for low light face detection. In *CVPR*, 2021. 1, 3, 5, 6
- [66] Wenjing Wang, Xinhao Wang, Wenhan Yang, and Jiaying Liu. Unsupervised face detection in the dark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1250–1266, 2022.

- [67] Wenjing Wang, Zhengbo Xu, Haofeng Huang, and Jiaying Liu. Self-aligned concave curve: Illumination enhancement for unsupervised adaptation. In *ACM MM*, 2022. 1, 3
- [68] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [69] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. 2018. 1, 2, 4, 5, 6, 7
- [70] Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhan Yang, and Jianmin Jiang. Uretinex-net: Retinex-based deep unfolding network for low-light image enhancement. In *CVPR*, 2022. 1, 2, 4
- [71] Xinyi Wu, Zhenyao Wu, Hao Guo, Lili Ju, and Song Wang. Dannet: A one-stage domain adaptation network for unsupervised nighttime semantic segmentation. In *CVPR*, 2021. 3
- [72] Yuhui Wu, Chen Pan, Guoqing Wang, Yang Yang, Jiwei Wei, Chongyi Li, and Heng Tao Shen. Learning semantic-aware knowledge guidance for low-light image enhancement. In *CVPR*, 2023. 1
- [73] Ke Xu, Xin Yang, Baocai Yin, and Rynson WH Lau. Learning to restore low-light images via decomposition-and-enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [74] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *CVPR*, 2022. 1
- [75] Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. Wider face: A face detection benchmark. In *CVPR*, 2016. 1, 2, 5
- [76] Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. From fidelity to perceptual quality: A semi-supervised approach for low-light image enhancement. In *CVPR*, 2020. 8
- [77] Wenhan Yang, Ye Yuan, Wenqi Ren, Jiaying Liu, Walter J Scheirer, Zhangyang Wang, Taiheng Zhang, Qiaoyong Zhong, Di Xie, Shiliang Pu, et al. Advancing image understanding in poor visibility environments: A collective benchmark study. *IEEE Transactions on Image Processing*, 29:5737–5752, 2020. 1, 3
- [78] Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE Transactions on Image Processing*, 30:2072–2086, 2021. 2, 4
- [79] Junjie Ye, Changhong Fu, Guangze Zheng, Danda Pani Paudel, and Guang Chen. Unsupervised domain adaptation for nighttime aerial tracking. In *CVPR*, 2022. 3
- [80] Xunpeng Yi, Han Xu, Hao Zhang, Linfeng Tang, and Jiayi Ma. Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model. In *ICCV*, 2023. 2
- [81] Shifeng Zhang, Xiangyu Zhu, Zhen Lei, Hailin Shi, Xiaobo Wang, and Stan Z Li. S3fd: Single shot scale-invariant face detector. In *ICCV*, 2017. 2
- [82] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *ACM MM*, 2019. 5, 6
- [83] Zhishuai Zhang, Wei Shen, Siyuan Qiao, Yan Wang, Bo Wang, and Alan Yuille. Robust face detection via learning small faces on hard images. In *WACV*, 2020. 5
- [84] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2021. 2