

# Image-Adaptive YOLO for Object Detection in Adverse Weather Conditions

Wenyu Liu<sup>1,2\*</sup>, Gaofeng Ren<sup>3</sup>, Runsheng Yu<sup>4</sup>, Shi Guo<sup>5</sup>, Jianke Zhu<sup>1,2†</sup>, Lei Zhang<sup>3,5</sup>

<sup>1</sup> College of Computer Science and Technology, Zhejiang University

<sup>2</sup> Alibaba-Zhejiang University Joint Institute of Frontier Technologies

<sup>3</sup> DAMO Academy, Alibaba Group

<sup>4</sup> The Hong Kong University of Science and Technology

<sup>5</sup> The HongKong Polytechnic University

{liuwenyu.lwy, jkzhu}@zju.edu.cn, {gaof.ren, runshengyu}@gmail.com, {csshiguo, cslzhang}@comp.polyu.edu.hk

## Abstract

Though deep learning-based object detection methods have achieved promising results on the conventional datasets, it is still challenging to locate objects from the low-quality images captured in adverse weather conditions. The existing methods either have difficulties in balancing the tasks of image enhancement and object detection, or often ignore the latent information beneficial for detection. To alleviate this problem, we propose a novel Image-Adaptive YOLO (IA-YOLO) framework, where each image can be adaptively enhanced for better detection performance. Specifically, a differentiable image processing (DIP) module is presented to take into account the adverse weather conditions for YOLO detector, whose parameters are predicted by a small convolutional neural network (CNN-PP). We learn CNN-PP and YOLOv3 jointly in an end-to-end fashion, which ensures that CNN-PP can learn an appropriate DIP to enhance the image for detection in a weakly supervised manner. Our proposed IA-YOLO approach can adaptively process images in both normal and adverse weather conditions. The experimental results are very encouraging, demonstrating the effectiveness of our proposed IA-YOLO method in both foggy and low-light scenarios. The source code can be found at <https://github.com/wenyu/Image-Adaptive-YOLO>.

## Introduction

CNN-based methods have become prevailed in object detection (Ren et al. 2015; Redmon and Farhadi 2018). They not only have achieved promising performance on benchmark datasets (Deng et al. 2009; Everingham et al. 2010; Lin et al. 2014) but also have been deployed in the real-world applications such as autonomous driving (Wang et al. 2019). Due to the domain shift in input images (Sindagi et al. 2020), general object detection models trained by high quality images often fail to achieve satisfactory results under adverse weather conditions (e.g., foggy and dark light). Narasimhan and Nayar (2002) and You et al. (2015) suggested that an image captured under adverse weather can be decomposed into a clean image and its corresponding weather-specific information, and the image quality degradation in adverse weather

\*This work was done when the author was visiting Alibaba as a research intern.

†corresponding author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: In the real-world foggy condition, our method can adaptively output clearer images with sharper edges around objects’ boundary, and consequently produce higher confidence detection results with fewer missing instances.

is mainly caused by the interaction between weather-specific information and objects, which leads to poor detection performance. Fig. 1 shows an example of object detection in foggy conditions. One can see that if the image can be properly enhanced with respect to the weather condition, more latent information about the original blurred objects and the misidentified objects can be recovered.

To tackle this challenging problem, Huang, Le, and Jaw (2020) employed two subnetworks to jointly learn visibility enhancement and object detection, where the impact of image degradation is reduced by sharing the feature extraction layers. However, it is hard to tune the parameters to balance the weights between detection and restoration during training. Another approach is to dilute the influence of weather-specific information by preprocessing the image with existing methods such as image defogging (Hang et al. 2020; Liu et al. 2019) and image enhancement (Guo et al. 2020). However, complicated image restoration networks have to be included in these methods, which need to be trained separately with pixel-level supervision. This requires to manually label the images for restoration. This problem can also be treated as an unsupervised domain adaptation task (Chen et al. 2018; Hnewa and Radha 2021). Compared with training detectors with clear images (source image), it is assumed that images captured under adverse weathers (target image) have a distribution shift. These methods mostly adopt the domain adaptation principles and focus on aligning the features of two distributions, and the latent information which can be obtained in the process of weather-based image restoration is

usually ignored.

To address the above limitations, we propose an ingenious image-adaptive object detection method, called IA-YOLO. Specifically, we suggest a fully differentiable image processing module (DIP), whose hyperparameters are adaptively learned by a small CNN-based parameter predictor (CNN-PP). The CNN-PP adaptively predicts the DIP’s hyperparameters according to the brightness, color, tone and weather-specific information of the input image. After processed by the DIP module, the interference of weather-specific information in the image can be suppressed while the latent information can be restored. We present a joint optimization scheme to learn the DIP, CNN-PP and YOLOv3 backbone detection network (Redmon and Farhadi 2018) in an end-to-end manner. To enhance the image for detection, CNN-PP is weakly supervised to learn an appropriate DIP through the bounding box annotations. Additionally, we make use of the images in both normal and adverse weather conditions to train the proposed network. By taking advantages of the CNN-PP network, our proposed IA-YOLO approach is able to adaptively deal with images affected by different degrees of weather conditions. Fig. 1 shows an example of the detection result by our proposed method.

The highlights of this work are: 1) an image-adaptive detection framework is proposed, which achieves promising performance in both normal and adverse weather conditions; 2) a white-box differentiable image processing module is proposed, whose hyperparameters are predicted by a weakly supervised parameter prediction network; 3) encouraging experimental results are achieved on both synthetic testbeds (VOC\_Foggy and VOC\_Dark) and real-world datasets (RTTS and ExDark), comparing against the previous methods.

## Related Work

### Object Detection

As a fundamental task in computer vision, object detection has received intensive attentions. The object detection methods can be roughly divided into two categories (Zhao et al. 2019). One category is the region proposal-based methods (Girshick et al. 2014; Girshick 2015; Ren et al. 2015), which first generate regions of interest (RoIs) from the image, and then classify them by training neural networks. Another category is one-stage regression-based approaches such as the YOLO series (Redmon et al. 2016; Redmon and Farhadi 2017, 2018; Bochkovskiy, Wang, and Liao 2020) and SSD (Liu et al. 2016), where the object labels and bounding box coordinates are predicted by a single CNN. In this paper, we employ the classic one-stage detector YOLOv3 (Redmon and Farhadi 2018) as the baseline detector, and improve its performance in adverse conditions.

### Image Adaptation

Image adaptation is widely used in image enhancement. To appropriately enhance an image, some traditional methods (Polesel, Ramponi, and Mathews 2000; Yu and Bajaj 2004; Wang et al. 2021) adaptively calculate the parameters of image transformation according to the corresponding image features. For example, Wang et al. (2021) proposed an

brightness adjustment function that adaptively adjusts the enhancement parameters based on the illumination distribution characteristics of an input image.

To achieve adaptive image enhancement, (Hu et al. 2018; Yu et al. 2018; Zeng et al. 2020) employed a small CNN to flexibly learn the hyperparameters of image transformation. Hu et al. (2018) proposed a post-processing framework with a set of differentiable filters, where deep reinforcement learning (DRL) is used to generate the image operation and filter parameters according to the quality of the current retouched image. Zeng et al. (2020) utilized a small CNN to learn image-adaptive 3D LUTs according to the global context such as brightness, color and tones.

### Object Detection in Adverse Conditions

Compared against general object detection, few research efforts have been made on object detection under adverse weather conditions. A straightforward approach is to preprocess the image by using the classical dehazing or image enhancement methods (Guo et al. 2020; He, Sun, and Tang 2009; Liu et al. 2019; Hang et al. 2020; Qin et al. 2020), which are originally designed to remove fog and enhance image quality. However, the improvement on image quality may not definitely benefit the detection performance. Some prior-based methods (Li et al. 2017; Huang, Le, and Jaw 2020) jointly perform image enhancement and detection to attenuate the interference of adverse weather-specific information. Sindagi et al. (2020) proposed an unsupervised prior-based domain adversarial object detection framework for detection in hazy and rainy conditions. A few methods (Chen et al. 2018; Zhang et al. 2021; Hnewa and Radha 2021) make use of domain adaptation to address this problem. Hnewa and Radha (2021) assumed that there is domain shift between images captured under normal and adverse weather conditions. They designed a multi-scale domain adaptive YOLO that supports domain adaptation in different layers at the feature extraction stage.

### Proposed Method

The images captured in adverse weather conditions have poor visibility due to the interference of weather-specific information, causing difficulties in object detection. To tackle this challenge, we suggest an image-adaptive detection framework by removing weather-specific information and revealing more latent information. As illustrated in Fig. 2, the whole pipeline consists of a CNN-based parameter predictor (CNN-PP), a differentiable image processing module (DIP) and a detection network. We first resize the input image into the size of  $256 \times 256$ , and feed it into CNN-PP to predict DIP’s parameters. Then, the image filtered by DIP module is treated as the input for YOLOv3 detector. We present an end-to-end hybrid data training scheme with detection loss so that the CNN-PP can learn an appropriate DIP to enhance the image for object detection in a weakly supervised manner.

### DIP Module

As in (Hu et al. 2018), the design of image filters should conform to the principle of differentiability and resolution-




Figure 2: The end-to-end training pipeline of the proposed IA-YOLO framework. Our method learns a YOLO with a small CNN-based parameter predictor (CNN-PP), which employs the downsampled input image to predict the hyperparameters of filters in the DIP module. The high-resolution input images are processed by DIP’s filters to aid YOLOv3 achieve high detection performance. The *Defog filter* is only used in foggy conditions.

independent. For the gradient-based optimization of CNN-PP, the filters should be differentiable to allow training the networks by backpropagation. Since CNN will consume a lot of computing resources to process high resolution images (e.g.,  $4000 \times 3000$ ), in this paper we learn the filter parameters from the downsampled low-resolution image of size  $256 \times 256$ , and then apply the same filter to the image of original resolution. Therefore, these filters need to be independent of image resolution.

Our proposed DIP module consists of six differentiable filters with adjustable hyperparameters, including *Defog*, *White Balance*(WB), *Gamma*, *Contrast*, *Tone* and *Sharpen*. As in (Hu et al. 2018), the standard color and tone operators, such as *WB*, *Gamma*, *Contrast* and *Tone*, can be expressed as pixel-wise filters. Therefore, our designed filters can be categorized into *Defog*, *Pixel-wise filters* and *Sharpen*. Among these filters, *Defog filter* is specially designed for foggy scenes. The details are as follows.

**Pixel-wise Filters.** The pixel-wise filters map an input pixel value  $P_i = (r_i, g_i, b_i)$  into an output pixel value  $P_o = (r_o, g_o, b_o)$ , in which  $(r, g, b)$  represent the values of the three color channels red, green and blue, respectively. The mapping functions of the four pixel-wise filters are listed in Table 1, where the second column lists the parameters to be optimized in our approach. *WB* and *Gamma* are simple multiplication and power transformations. Obviously, their mapping functions are differentiable with respect to both the input image and the parameters.

The differentiable contrast filters are designed with an input parameter to set the linear interpolation between the original image and the fully enhanced image. As shown in Table 1, the definition of  $En(P_i)$  in the mapping function is as follows:

$$Lum(P_i) = 0.27r_i + 0.67g_i + 0.06b_i \quad (1)$$

Filter	Parameters	Mapping Function
WB	$W_r, W_g, W_b$ : factors	$P_o = (W_r r_i, W_g g_i, W_b b_i)$
Gamma	$G$ : gamma value	$P_o = P_i^G$
Contrast	$\alpha$ : contrast value	$P_o = \alpha \cdot En(P_i) + (1 - \alpha) \cdot P_i$
Tone	$t_i$ : tone params	$P_o = (L_{t_r}(r_i), L_{t_g}(g_i), L_{t_b}(b_i))$

Table 1: The mapping functions of pixel-wise filters.

$$EnLum(P_i) = \frac{1}{2}(1 - \cos(\pi \times (Lum(P_i))) \quad (2)$$

$$En(P_i) = P_i \times \frac{EnLum(P_i)}{Lum(P_i)} \quad (3)$$

As in (Hu et al. 2018), we design the tone filter as a monotonic and piecewise-linear function. We learn tone filter with  $L$  parameters, represented as  $\{t_0, t_1, \dots, t_{L-1}\}$ . The points of tone curve are denoted as  $(k/L, T_k/T_L)$ , in which  $T_k = \sum_{i=0}^{k-1} t_i$ . Furthermore, the mapping function is expressed by the differentiable parameters, which enable the function to be differentiable with respect to both the input image and the parameters  $\{t_0, t_1, \dots, t_{L-1}\}$  as below

$$P_o = \frac{1}{T_L} \sum_{j=0}^{L-1} clip(L \cdot P_i - j, 0, 1)t_k \quad (4)$$

**Sharpen Filter.** Image sharpening can highlight the image details. Like the unsharpen mask technique (Polesel, Ramponi, and Mathews 2000), the sharpening process can be described as follows:

$$F(x, \lambda) = I(x) + \lambda(I(x) - Gau(I(x))) \quad (5)$$

where  $I(x)$  is the input image,  $Gau(I(x))$  denotes Gaussian filter, and  $\lambda$  is a positive scaling factor. This sharpening operation is differentiable with respect to both  $x$  and  $\lambda$ . Note that the sharpening degree can be tuned for better object detection performance by optimizing  $\lambda$ .

**Defog Filter.** Motivated by the dark channel prior method (He, Sun, and Tang 2009), we design a defog filter with a learnable parameter. Based on the atmospheric scattering model (McCartney 1976; Narasimhan and Nayar 2002), the formation of a hazy image can be formulated as follows:

$$I(x) = J(x)t(x) + A(1 - t(x)) \quad (6)$$

where  $I(x)$  is the foggy image, and  $J(x)$  represents the scene radiance (clean image).  $A$  is the global atmospheric light, and  $t(x)$  is the medium transmission map, which is defined as:

$$t(x) = e^{-\beta d(x)} \quad (7)$$

where  $\beta$  represents the scattering coefficient of the atmosphere, and  $d(x)$  is the scene depth.

In order to recover the clean image  $J(x)$ , the key is to obtain the atmospheric light  $A$  and the transmission map  $t(x)$ . To this end, we first compute the dark channel map and pick the top 1000 brightest pixels. Then,  $A$  is estimated by averaging those 1000 pixels of the haze image  $I(x)$ .

According to Eq. (6), one can derive an approximate solution of  $t(x)$  as follows (He, Sun, and Tang 2009)

$$t(x) = 1 - \min_C \left( \min_{y \in \Omega(x)} \frac{I^C(y)}{A^C} \right) \quad (8)$$

We further introduce a parameter  $\omega$  to control the degree of defogging as follows:

$$t(x, \omega) = 1 - \omega \min_C \left( \min_{y \in \Omega(x)} \frac{I^C(y)}{A^C} \right) \quad (9)$$

Since the above operation is differentiable, we can optimize  $\omega$  through back propagation to make *defog filter* more conducive to foggy image detection.

### CNN-PP Module

In camera image signal processing (ISP) pipeline, some adjustable filters are usually employed for image enhancement, whose hyperparameters are manually tuned by experienced engineers through visual inspection (Mosleh et al. 2020). Generally, such a tuning process is very awkward and expensive to find suitable parameters for a broad range of scenes. To address this limitation, we propose to employ a small CNN as a parameter predictor to estimate the hyperparameters, which is very efficient.

Taking foggy scene as an example, the purpose of CNN-PP is to predict the DIP’s parameters by understanding the global content of the image, such as brightness, color and tone, as well as the degree of fog. Therefore, the down-sampled image is sufficient to estimate these information, which can greatly save the computational cost. Given an input image of any resolution, we simply use bilinear interpolation to downsample it to  $256 \times 256$  resolution. As shown in Fig. 2, the CNN-PP network is composed of five convolutional blocks and two fully-connected layers. Each convolutional block includes a  $3 \times 3$  convolutional layer with stride 2 and a leaky Relu. The final fully-connected layer outputs the hyperparameters for the DIP module. The output channel of these five convolutional layers are 16, 32, 32, 32 and 32, respectively. The CNN-PP model contains only 165K parameters when the total number of parameter is 15.

### Detection Network Module

In this paper, we choose the one-stage detector YOLOv3 as the detection network, which is widely used in practical applications, including image editing, security monitoring, crowd detection and autonomous driving (Zhang et al. 2021). Compared with the previous version, YOLOv3 designs darknet-53 that is composed of successive  $3 \times 3$  and  $1 \times 1$  convolutional layers based on the idea of Resnet (He et al. 2016). It implements multi-scale training by making predictions on multi-scale feature maps, so as to further improve the detection accuracy, especially for small objects. We adopt the same network architecture and loss functions as the original YOLOv3 (Redmon and Farhadi 2018).

### Hybrid Data Training

In order to achieve ideal detection performance in both normal and adverse weather conditions, we adopt a hybrid data training scheme for IA-YOLO. Algorithm 1 summarizes the training process of our proposed method. Each image has a probability of 2/3 to be randomly added with some kind of fog or be transformed to a low-light image before being input to the network for training. With both normal and synthetic low-quality training data, the whole pipeline is trained end-to-end with YOLOv3 detection loss, which ensures all modules in IA-YOLO can adapt to each other. Therefore, the CNN-PP module is weakly supervised by the detection loss without manually labeling ground truth images. The hybrid data training mode ensures that IA-YOLO can adaptively process the image according to the content of each image, so as to achieve high detection performance.

## Experiments

We evaluate the effectiveness of our method in the scenarios of fog and low-light. The filter combination is [*Defog*, *White Balance*(WB), *Gamma*, *Contrast*, *Tone*, *Shapen*], while the *Defog filter* is only used in foggy conditions.

### Implementation Details

We adopt the training protocol of (Redmon and Farhadi 2018) in our proposed IA-YOLO approach. The backbone network for all experiments is Darknet-53. During training, We randomly resize the image to  $(32N \times 32N)$ , where  $N \in [9, 19]$ . Moreover, the data augmentation methods like image flipping, cropping and transformation are applied to expand the training dataset. Our IA-YOLO model is trained by the Adam optimizer (Kingma and Ba 2014) with 80 epochs. The starting learning rate is  $10^{-4}$  and the batch size is 6. IA-YOLO predicts the bounding boxes at three different scales, and three anchors at each scale. We use Tensorflow for our experiments and run it on the Tesla V100 GPU.

### Experiments on Foggy Images

**Datasets** There are few publicly available datasets for object detection in adverse weather conditions, and the data amounts are usually small for training a stable CNN-based detector. To facilitate fair comparison, we build upon the classic VOC dataset (Everingham et al. 2010) a VOC\_Foggy dataset according to the atmospheric scattering model (Narasimhan

---

**Algorithm 1: Image-Adaptive YOLO training procedure**


---

Initialize the CNN-PP network  $P^\theta$  and the YOLOv3 network  $D^\beta$  with random weights  $\theta$  and  $\beta$ .  
Set the training stage:  $num\_epochs = 80$ ,  $batch\_size = 6$ .  
Prepare the normal dataset VOC\_norm\_trainval.  
**for**  $i$  in  $num\_epochs$  **do**  
  **repeat**  
    Take a batch images  $M$  from VOC\_norm\_trainval;  
    **for**  $j$  in  $batch\_size$  **do**  
      **if**  $random.randint(0, 2) > 0$  **then**  
        Generate the foggy image  $M(j)$  (Eq. (10, 11)), where  $A = 0.5$ ,  $k = random.randint(0, 9)$ ,  $\beta = 0.01 * k + 0.05$  //for foggy conditions  
        Generate the low-light image  $M(j)$  by  $f(x) = x^\gamma$ , where  $\gamma = random.uniform(1.5, 5)$  //for low-light conditions  
      **end if**  
    **end for**  
    Compute DIP params by  $P_N = P^\theta(image\_batch)$ ;  
    Perform DIP filter processing by  $image\_batch = DIP(image\_batch, P_N)$ ;  
    Send  $image\_batch$  to YOLOv3 network  $D^\beta$ ;  
    Update the CNN-PP network  $P^\theta$  and the YOLOv3 network  $D^\beta$  according to the YOLOv3 detection loss.  
  **until** all images have been fed into training models  
**end for**

---

Dataset	image	ps	bc	car	bus	mc	Total
V_n_tv	8111	13256	1064	3267	822	1052	19561
V_n_ts	2734	4528	337	1201	213	325	6604
RTTS	4322	7950	534	18413	1838	862	29577

Table 2: Statistics of the used datasets, including V\_n\_tv (VOC\_norm\_trainval), V\_n\_ts (VOC\_norm\_test) and RTTS. The classes are ps (person), bc (bicycle), car, bus, and mc (motorcycle).

and Nayar 2002). Moreover, RTTS (Li et al. 2018) is a relatively comprehensive real-world dataset available in foggy conditions, which has 4,322 natural hazy images with five annotated object classes, namely, person, bicycle, car, bus and motorcycle. To form our training dataset, we select the data containing these five categories to add haze.

For VOC2007\_trainval and VOC2012\_trainval, we filtrate the images which include the above five classes of objects to build VOC\_norm\_trainval. VOC\_norm\_test is selected from the VOC2007\_test in a similar way. We also evaluate our method on RTTS. The statistics of the datasets are summarized into Table 2.

To avoid the computational cost on generating foggy images during the training process, we construct the VOC\_Foggy dataset offline. According to Eqs. (6, 7), the foggy image  $I(x)$  is obtained by

$$I(x) = J(x)e^{-\beta d(x)} + A(1 - e^{-\beta d(x)}) \quad (10)$$

Method		Train data	V_n_ts	V_F_t	RTTS
Baseline	YOLOv3 I	VOC_norm	70.10	31.05	28.82
	YOLOv3 II	Hybrid data	64.13	63.40	30.80
Defog	MSBDN	VOC_norm	/	57.38	30.20
	GridDehaze	VOC_norm	/	58.23	31.42
DA	DAYOLO	Hybrid data	56.51	55.11	29.93
Mul-task	DSNet	Hybrid data	53.29	67.40	28.91
Ours	IA-YOLO	Hybrid data	<b>73.23</b>	<b>72.03</b>	<b>37.08</b>

Table 3: Performance comparison on foggy images. "DA" means Domain Adaption. The right three columns list the mAP on three test datasets, including V\_n\_ts (VOC\_norm\_test), V\_F\_t (VOC\_Foggy\_test) and RTTS.

$d(x)$  is defined as below:

$$d(x) = -0.04 * \rho + \sqrt{\max(row, col)} \quad (11)$$

where  $\rho$  is the Euclidean distance from the current pixel to the central pixel,  $row$  and  $col$  are the number of rows and columns of the image, respectively. By setting  $A = 0.5$  and  $\beta = 0.01 * i + 0.05$ , where  $i$  is a integer number from 0 to 9, ten different levels of fog can be added to each image. Based on the VOC\_norm\_trainval dataset, we generate the VOC\_Foggy\_trainval dataset ten times larger offline. To obtain the VOC\_Foggy\_test dataset, every image in VOC\_norm\_test is randomly processed with fog.

**Experimental Results** To demonstrate the effectiveness of IA-YOLO, we compare our method with baseline YOLOv3, *Defog+Detect*, domain adaption (Hnewa and Radha 2021), and multi-task learning (Huang, Le, and Jaw 2020) on the three testing datasets. For *Defog + Detect*, we employ the defogging method as the preprocessing step and use YOLOv3 trained on VOC\_norm for detection. We select MSBDN (Hang et al. 2020) and GridDehaze (Liu et al. 2019) as the preprocessing methods, which are popular CNN-based dehazing methods. For domain adaption approach, we employ the one-stage multi-scale domain-adaptive detector DAYOLO (Hnewa and Radha 2021) with multiple domain adaptation paths and the corresponding domain classifiers at different scales of YOLOv3. For multi-task learning algorithms, we choose DSNet (Huang, Le, and Jaw 2020) that jointly learns dehazing and detection in adverse weather conditions. We reproduce its detection subnetwork and restoration module by sharing the first five convolution layers of Yolov3, and jointly train two networks with hybrid data.

Table 3 compares the mean average precision (mAP) results (Everingham et al. 2012) between IA-YOLO and other competing algorithms in both the normal and hazy conditions. The second column lists the training data for each method, where "Hybrid data" means the hybrid data training scheme used in our proposed IA-YOLO. Comparing to the baseline (YOLO I), all methods have improvements on both synthetic and real-word foggy weather testing datasets, while only our IA-YOLO has no deterioration in the case of normal scenarios. This is because the previous methods are mainly designed to deal with object detection in foggy weather condition while sacrificing their performance on the normal weather images. For our proposed IA-YOLO method, the CNN-PP and DIP modules are able to adaptively process



Figure 3: Detection results of YOLOv3 II (columns 1, 3 and 5) and our IA-YOLO (columns 2, 4 and 6) on synthetic VOC\_Foggy\_test images (top row) and real-world RTTS foggy images (bottom row). The proposed method learns to reduce the haze and sharpen image edge, which has better detection performance with fewer missed and wrong detections.



Figure 4: Detection results of YOLOv3 II (left column) and our IA-YOLO (right column) on synthetic VOC\_Dark\_test images (top row) and real-world ExDark low-light images (bottom row). The proposed method learns to enhance the image contrast with more details.

the image with different degrees of fog for object detection. As a result, our proposed IA-YOLO approach outperforms all the competing methods on the three testing datasets by a large margin, demonstrating its effectiveness on object detection in adverse weather conditions.

Fig. 3 shows several visual examples of our IA-YOLO method and the baseline YOLOv3 II. Though for some instances, our adaptive DIP module produces some unwilling noises for visual perception, it greatly boosts the local image gradients based on image semantics, and leads to better detection performance.

## Experiments on Low-light Images

**Datasets** PSCAL VOC (Everingham et al. 2010) and the relatively comprehensive low-light detection dataset ExDark (Loh and Chan 2019) both contain ten categories of objects: *Bicycle, Boat, Bottle, Bus, Car, Cat, Chair, Dog, Motorbike, People(Person)*. From the VOC2007\_trainval and VOC2012\_trainval, we filtrated the images which include any of the above ten classes of objects to build VOC\_norm\_trainval. VOC\_norm\_test is selected from

Method		Train data	V_n_ts	V_D_t	E_t
Baseline	YOLOv3 I	VOC_norm	69.13	45.92	36.42
	YOLOv3 II	Hybrid data	65.33	52.28	37.03
Enhance	ZeroDCE	VOC_norm	/	33.57	34.41
DA	DAYOLO	Hybrid data	41.68	21.53	18.15
Mul-task	DSNet	Hybrid data	64.08	43.75	36.97
Ours	IA-YOLO	Hybrid data	<b>70.02</b>	<b>59.40</b>	<b>40.37</b>

Table 4: Performance comparison on low-light images. The right three columns list the mAP on the three test datasets, including V\_n\_ts (VOC\_norm\_test), V\_D\_t (VOC\_Dark\_test) and E\_t (Exdark\_test).

VOC2007\_test in the same way. The total number of images in VOC\_norm\_trainval, VOC2007\_test and ExDark\_test are 12334, 3760 and 2563, respectively.

We synthesize the low-light VOC\_dark dataset based on VOC\_norm through the transformation  $f(x) = x^\gamma$ , where the value of  $\gamma$  is randomly sampled from a uniform distribution within the range  $[1.5, 5]$ , and  $x$  denotes the input pixel intensity.

**Experimental Results** We compare our presented IA-YOLO method with baseline YOLOv3, *Enhance + Detect*, DAYOLO, and DSNet on the three testing datasets. For *Enhance + Detect*, we employ the recent image enhancement method Zero-DCE (Guo et al. 2020) to preprocess the low-light images and use YOLOv3 trained on VOC\_norm for detection. The remaining experimental settings are kept the same as those on foggy images. Table 4 shows the mAP results. It can be seen that our method yields the best results. IA-YOLO improves the baseline YOLO I by 0.89, 13.48 and 3.95 percent on VOC\_norm\_test, VOC\_Dark\_test and ExDark\_test, respectively, and it improves the YOLO II baseline by 4.69, 7.12 and 3.34 percent on those test sets. This demonstrates that our proposed IA-YOLO approach is also effective in low-light conditions.

Fig. 4 shows the qualitative comparisons between IA-YOLO and the baseline YOLOv3 II. It can be observed that our proposed DIP module is able to adaptively increase the contrast of input image and reveal the image details, which are essential to object detection.




Figure 5: Performance comparison of different settings in foggy conditions.

### Ablation Study

To examine the effectiveness of each module in our proposed framework, we conduct ablation experiments on different settings, including hybrid data training scheme, DIP and image adaption. We also evaluate the selection of the proposed differentiable filters on three testing datasets.

The results of the conducted experiments are depicted in Fig. 5. Except for YOLO I trained with VOC\_norm, the rest of experiments employ the same hybrid data training and the experimental settings. It can be seen that hybrid data training, DIP filter preprocessing, and image adaptive methods all can improve the detection performance on both VOC\_Foggy\_test and RTTS compared with YOLO I. IA-YOLO achieves the best results by using all the three modules. YOLOv3 with fixed DIP means that the filter’s hyperparameters are a given set of fixed values, all of which are within a reasonable range. YOLOv3\_deep II is a deeper version of YOLO II by adding eight convolutional layers with learning parameters over 411K. As shown in Fig. 5, our proposed IA-YOLO approach performs better than YOLOv3\_deep II with only 165K additional parameters in CNN-PP. It is worthy mentioning that only the adaptive learning module improves the performance on VOC\_norm\_test compared against YOLO I in the normal weather conditions, while both YOLOv3 II and YOLOv3 with fixed DIP obtain inferior results. This demonstrates that IA-YOLO can adaptively process both the normal images and foggy ones, which is beneficial for the down-streamed detection task.

As shown in Table 5, we conduct quantitative mAP evaluation on filter selection using the three testing datasets. By combining the three sets of filters, *Model D* obtains the best result, demonstrating the effectiveness of these filters. Fig. 6 shows the visual comparisons of several models in Table 5. Compared with *Model C* which brightens and sharpens the image, and *Model B* which defogs the image, the image processed by *Model D* is not only brighter and sharper but also clearer, which makes the foggy objects much more visible. Moreover, we provide some examples on how CNN-PP predicts the DIP module’s parameters in the Appendix.




Figure 6: Detection results on RTTS. From left to right, top to bottom: YOLO II, Model B, Model C and Model D.

Model	Defog Filter	Pixel_wise Filters	Sharpen Filter	V_n_ts	V_F_t	RTTS
A	✓		✓	71.47	70.25	34.88
B	✓	✓		71.43	70.14	34.83
C		✓	✓	70.51	70.09	34.95
D	✓	✓	✓	73.23	72.03	37.08

Table 5: Ablation analysis on the filters in DIP module.

### Efficiency Analysis

In our IA-YOLO framework, we introduce a small CNN-PP learning module with 165K trainable parameters into YOLOv3. IA-YOLO takes 44 ms to detect a  $544 \times 544 \times 3$  resolution image on a single Tesla V100 GPU. It costs only additional 13 ms over the YOLOv3 baseline, while it is 7 ms and 50 ms faster than GridDehaze-YOLOv3 and MSBDN-YOLOv3, respectively. In summary, IA-YOLO only adds 165K trainable parameters while achieving much better performance on all test datasets with comparable running time.

### Conclusion

We proposed a novel IA-YOLO approach to improve object detection in adverse weather conditions, where each input image was adaptively enhanced to obtain better detection performance. A fully differentiable image processing module was developed to restore the latent content by removing the weather-specific information for YOLO detector, whose hyperparameters are predicted by a small convolutional neural network. Moreover, the whole framework was trained in an end-to-end fashion, where the parameter prediction network was weakly supervised to learn an appropriate DIP module through the detection loss. By taking advantages of hybrid training and the parameter prediction network, our proposed approach was able to adaptively handle normal and adverse weather conditions. The experimental results showed that our method performed much better than previous approaches in both the foggy and low-light scenarios.

## Acknowledgments

This work is supported by National Natural Science Foundation of China under Grants (61831015) and Hong Kong RGC RIF Grant (R5001-18). This work is also supported by the funding of "Leading Innovation Team of the Zhejiang Province" (Grant NO. 2018R01017).

## References

- Bochkovskiy, A.; Wang, C.-Y.; and Liao, H.-Y. M. 2020. Yolov4: Optimal speed and accuracy of object detection. arXiv:2004.10934.
- Chen, Y.; Li, W.; Sakaridis, C.; Dai, D.; and Van Gool, L. 2018. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 3339–3348.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 248–255. IEEE.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2): 303–338.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2012. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- Girshick, R. 2015. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 1440–1448.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 580–587.
- Guo, C. G.; Li, C.; Guo, J.; Loy, C. C.; Hou, J.; Kwong, S.; and Cong, R. 2020. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 1780–1789.
- Hang, D.; Jinshan, P.; Zhe, H.; Xiang, L.; Xinyi, Z.; Fei, W.; and Ming-Hsuan, Y. 2020. Multi-Scale Boosted Dehazing Network with Dense Feature Fusion. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*.
- He, K.; Sun, J.; and Tang, X. 2009. Single image haze removal using dark channel prior. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 770–778.
- Hnewa, M.; and Radha, H. 2021. Multiscale Domain Adaptive YOLO for Cross-Domain Object Detection. arXiv:2106.01483.
- Hu, Y.; He, H.; Xu, C.; Wang, B.; and Lin, S. 2018. Exposure: A White-Box Photo Post-Processing Framework. *ACM Transactions on Graphics (TOG)*, 37(2): 26.
- Huang, S.-C.; Le, T.-H.; and Jaw, D.-W. 2020. DSNet: Joint semantic learning for object detection in inclement weather conditions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. arXiv:1412.6980.
- Li, B.; Peng, X.; Wang, Z.; Xu, J.; and Feng, D. 2017. Aodnet: All-in-one dehazing network. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 4770–4778.
- Li, B.; Ren, W.; Fu, D.; Tao, D.; Feng, D.; Zeng, W.; and Wang, Z. 2018. Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28(1): 492–505.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 740–755. Springer.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*, 21–37. Springer.
- Liu, X.; Ma, Y.; Shi, Z.; and Chen, J. 2019. GridDehazeNet: Attention-Based Multi-Scale Network for Image Dehazing. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Loh, Y. P.; and Chan, C. S. 2019. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178: 30–42.
- McCartney, E. J. 1976. Optics of the atmosphere: scattering by molecules and particles. *New York*.
- Mosleh, A.; Sharma, A.; Onzon, E.; Mannan, F.; Robidoux, N.; and Heide, F. 2020. Hardware-in-the-loop end-to-end optimization of camera image processing pipelines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7529–7538.
- Narasimhan, S. G.; and Nayar, S. K. 2002. Vision and the atmosphere. *International Journal of Computer Vision*, 48(3): 233–254.
- Polesel, A.; Ramponi, G.; and Mathews, V. J. 2000. Image enhancement via adaptive unsharp masking. *IEEE Transactions on Image Processing*, 9(3): 505–510.
- Qin, X.; Wang, Z.; Bai, Y.; Xie, X.; and Jia, H. 2020. FFA-Net: Feature fusion attention network for single image dehazing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 11908–11915.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 779–788.
- Redmon, J.; and Farhadi, A. 2017. YOLO9000: better, faster, stronger. In *Proceedings of IEEE/CVF Conference Computer Vision Pattern Recognition (CVPR)*, 7263–7271.



- 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: 91–99.
- Sindagi, V. A.; Oza, P.; Yasarla, R.; and Patel, V. M. 2020. Prior-based domain adaptive object detection for hazy and rainy conditions. In *European Conference on Computer Vision (ECCV)*, 763–780. Springer.
- Wang, G.; Guo, J.; Chen, Y.; Li, Y.; and Xu, Q. 2019. A PSO and BFO-based learning strategy applied to faster R-CNN for object detection in autonomous driving. *IEEE Access*, 7: 18840–18859.
- Wang, W.; Chen, Z.; Yuan, X.; and Guan, F. 2021. An adaptive weak light image enhancement method. In *Twelfth International Conference on Signal Processing Systems*, volume 11719, 1171902. International Society for Optics and Photonics.
- You, S.; Tan, R. T.; Kawakami, R.; Mukaigawa, Y.; and Ikeuchi, K. 2015. Adherent raindrop modeling, detection and removal in video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9): 1721–1733.
- Yu, R.; Liu, W.; Zhang, Y.; Qu, Z.; Zhao, D.; and Zhang, B. 2018. Deepexposure: Learning to expose photos with asynchronously reinforced adversarial learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NeurIPS)*, 2153–2163.
- Yu, Z.; and Bajaj, C. 2004. A fast and adaptive method for image contrast enhancement. In *2004 International Conference on Image Processing, 2004. ICIP'04.*, volume 2, 1001–1004. IEEE.
- Zeng, H.; Cai, J.; Li, L.; Cao, Z.; and Zhang, L. 2020. Learning image-adaptive 3D lookup tables for high performance photo enhancement in real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, S.; Tuo, H.; Hu, J.; and Jing, Z. 2021. Domain Adaptive YOLO for One-Stage Cross-Domain Detection. arXiv:2106.13939.
- Zhao, Z.-Q.; Zheng, P.; Xu, S.-t.; and Wu, X. 2019. Object detection with deep learning: A review. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11): 3212–3232.

## Appendix

### Defog Filter Design

Motivated by the conventional dark channel prior method (He, Sun, and Tang 2009), we design a defog filter with a learnable parameter. In the atmospheric scattering model (McCartney 1976; Narasimhan and Nayar 2002), the formation of a hazy image can be formulated as follows:

$$I(x) = J(x)t(x) + A(1 - t(x)) \quad (12)$$

where  $I(x)$  is the foggy image, and  $J(x)$  represents the scene radiance (clear image).  $A$  is the global atmospheric light, and  $t(x)$  is the medium transmission map.

In order to recover the clear image  $J(x)$ , the key is to obtain the global atmospheric light  $A$  and the medium transmission map  $t(x)$ . To this end, we first compute the dark channel map and pick the top 1000 brightest pixels. Then,  $A$  is estimated by averaging these 1000 pixels in the haze image  $I(x)$ . From Eq. (12), we can derive that

$$\frac{I^C(x)}{A^C} = t(x) \frac{J^C(x)}{A^C} + (1 - t(x)) \quad (13)$$

where  $C$  donates the RGB color channel. By taking two min operations, one on the channels and one on a local patch, in the above equation, we can obtain:

$$\min_C \left( \min_{y \in \Omega(x)} \frac{I^C(y)}{A^C} \right) = t(x) \min_C \left( \min_{y \in \Omega(x)} \frac{J^C(y)}{A^C} \right) + (1 - t(x)) \quad (14)$$

Based on the dark channel prior, we can get that

$$J^{dark}(x) = \min_C \left( \min_{y \in \Omega(x)} J^C(y) \right) = 0 \quad (15)$$

Since  $A^C$  is always positive, Eq. (15) can be written as:

$$\min_C \left( \min_{y \in \Omega(x)} \frac{J^C(y)}{A^C} \right) = 0 \quad (16)$$

By substituting Eq. (16) into Eq. (13), we can obtain:

$$t(x) = 1 - \min_C \left( \min_{y \in \Omega(x)} \frac{I^C(y)}{A^C} \right) \quad (17)$$

We further introduce a parameter  $\omega$  to control the degree of defogging. There is:

$$t(x, \omega) = 1 - \omega \min_C \left( \min_{y \in \Omega(x)} \frac{I^C(y)}{A^C} \right) \quad (18)$$

Since the above operation is differentiable, we can optimize  $\omega$  through back propagation to make *defog filter* more conducive to foggy image detection.

### Experiments

**Experiments on Foggy Images** We compare our method with the baseline YOLOv3 (Redmon and Farhadi 2018), *Defog + Detect* (Hang et al. 2020; Liu et al. 2019), domain adaptation (Hnewa and Radha 2021), and multi-task learning (Huang, Le, and Jaw 2020). For domain adaptation approach, we employ the one-stage multi-scale domain-adaptive detector DAYOLO (Hnewa and Radha 2021) with

multiple domain adaptation paths and the corresponding domain classifiers at different scales of YOLOv3. We set the loss weight  $\lambda = 0.1$  for training, and each batch has 2 images, one from the source domain and the other from the target domain. Other hyperparameters are set the same as in the original paper.

Fig. 7 shows several visual examples of our IA-YOLO method, the baseline YOLOv3 II and the *Defog + Detect* methods. Both GridDehaze (Liu et al. 2019) and MSBDN (Hang et al. 2020) can reduce the haze effect, which is generally beneficial to detection. Our IA-YOLO method not only reduces the haze, but also enhances the local image gradients, which lead to better detection performance.

Fig. 8 shows two examples on how the CNN-PP module predicts DIP’s parameters, including detailed parameter values and the images processed by each sub-filter. The CNN-PP is able to learn a set of DIP parameters for each image according to its brightness, color, tone and weather-specific information. After the input image is processed by the learned DIP module, more image details are revealed, which are conducive to the subsequent detection task.

**Experiments on Low-light Images** The total number of images in VOC\_norm\_trainval, VOC\_norm\_test and Ex-Dark\_test are 12334, 3760 and 2563, respectively. The numbers of instances are listed in Table 6.

We compare our presented method with the baseline YOLOv3, *Enhance + Detect* (Guo et al. 2020), DAYOLO, and DSNet on the three testing datasets. Fig. 9 shows several visual examples of our IA-YOLO method, the baseline YOLOv3 II and the *Enhance + Detect* methods. It can be observed that both Zero-DCE (Guo et al. 2020) and IA-YOLO can brighten the image and reveal the image details. The proposed IA-YOLO can further increase the contrast of the input image, which is essential to object detection.

### Efficiency Analysis

In our proposed IA-YOLO framework, we introduce a learning module of CNN-PP into YOLOv3, which is a small network containing five convolutional layers and two fully connected layers. Table 7 shows the efficiency analysis of some methods used in our experiments. The methods not listed are validated using the YOLOv3 architecture. The second column lists the number of additional parameters over the YOLOv3 model. The third column lists the running time on a  $544 \times 544 \times 3$  resolution image with a single Tesla V100 GPU. It can be seen that IA-YOLO only adds 165K trainable parameters over YOLOv3 while achieving the best performance on all testing with comparable running time. Note that IA-YOLO has fewer trainable parameters than YOLO\_deep II but its running time is longer. This is because that the filtering process in the DIP module incurs additional computation.



Figure 7: Detection results by different methods on real-world RTTS foggy images. From left to right: YOLOv3 II, GirdDehaze + YOLOv3 I, MSBDN + YOLOv3 I and our IA-YOLO. The proposed method learns to reduce the haze and enhance the image contrast, which leads to better detection performance with fewer missed and wrong detections.




Figure 8: Examples of learned DIP module and their filtering outputs. The image-adaptive processing module can output the corresponding filter parameters according to the brightness, color, tone and weather-information of each input image, so as to get better detection performance.

Dataset	person	bicycle	car	bus	motorbike	boat	bottle	cat	chair	dog	Total
Voc_norm_trainval	13256	1064	3267	822	1052	1140	1764	1593	3152	2025	29135
Voc_norm_test	4528	337	1201	213	325	263	469	358	756	489	8939
ExDark_test	2235	418	919	164	242	515	433	425	609	490	6450

Table 6: Statistics of the used datasets.




Figure 9: Detection results of different methods on synthetic VOC\_Dark\_test images (top row), real-world ExDark\_test low-light images (bottom two rows). From left to right: YOLOv3 II, ZeroDCE + YOLOv3 I and our IA-YOLO. The proposed method learns to make the image brighter with more details, which results in better detection performance with fewer missed and wrong detections.

Method	Additional Params	Speed(ms)
YOLOv3	/	31
YOLOv3_deep II	412K	35
ZeroDCE	79K	34
MSBDN	31M	94
GridDehaze	958K	51
IA-YOLO(Ours)	165K	44

Table 7: Efficiency analysis on the compared methods.