

# Contrastive Learning-based Robust Object Detection under Smoky Conditions

Wei Wu<sup>1\*</sup>, Hao Chang<sup>1</sup>, Yonghua Zheng<sup>1</sup>, Zhu Li<sup>2</sup>, Zhiwen Chen<sup>1</sup>, Ziheng Zhang<sup>1</sup>

<sup>1</sup>State Key Laboratory of Integrated Services Networks, Xidian University

<sup>2</sup>Department of Computer Science & Electrical Engineering, University of Missouri, Kansas City

wwu@xidian.edu.cn, chang\_hao98@163.com, 295569449@qq.com,

zhu.li@ieee.org, {1551431202,1102161646}@qq.com

## Abstract

*Object detection is to effectively find out interested targets in images and then accurately determine their categories and positions. Recently many excellent methods have been developed to provide powerful detection capability. However, their performance may degrade significantly under severe weather such as smoky conditions. In this paper, we propose a contrastive learning-based robust object detection algorithm for smoke images. The proposed object detector consists of two modules: contrastive learning module and object bounding box prediction module. The first module learns representation vectors by maximizing agreement between different augmented views of the same smoke image. These representations are then sent to the second module to yield the bounding box for each object. In addition, we also propose a novel affine data augmentation method. Extensive experiments have been conducted on A2I2-Haze dataset which is the first real haze dataset with in-situ smoke measurement aligned to aerial and ground imagery. This dataset is also the only dataset used in the 5<sup>th</sup> UG<sup>2+</sup> challenges of CVPR 2022 for both training and testing. Compared with state-of-the-art methods, evaluation results show the superiority of our proposed object detector.*

## 1. Introduction

Object detection, one of core issues in computer vision fields, aims to effectively find out interested targets in images and then accurately determine their categories and positions [1]. It has some down-stream applications, in particular, in autonomous driving, unmanned aerial vehicle (UAV) camera, surveillance, and visual question answering. So far, some excellent techniques have been proposed especially for object detection, and widely used in many areas. However, in severe weather scenarios such as smoky, hazy, and rainy environments [2-5], the performance of object detection algorithms may degrade

significantly because they usually assume that the inputs are clear images or videos [6]. Therefore, it is crucial to develop robust object detection on outdoor platforms like UAVs under real-world adverse conditions.

In the past decades, object detection has been extensively studied, and a large number of methods have been presented for natural objects (e.g. pedestrians, animals, and plants) and artificial objects (e.g. traffic signs and lights, vehicles, buildings, and bridges). Traditional approaches rely heavily on hand-crafted features extracted from captured images or videos. However, these approaches usually are sensitive to tiny changes in lighting, pose, etc.

Recently, with the rapid development of deep learning, convolutional neural network (CNN) has proved to be an extremely powerful tool in extracting features, spurring researcher's passion in creating object detection methods via deep learning. Currently, deep learning approaches can be roughly divided into two categories, i.e. anchor-based methods [7-20] and anchor-free methods [21-25]. The former represents each object through an axis-aligned bounding box along with its label, while the latter express an object only by a single point at its bounding box center. The representatives of anchor-based approaches are R-CNN series, YOLO series, and SSD, while those of anchor-free methods are CornerNet and CenterNet. However, these excellent algorithms achieve pretty good detection results but dramatically inferior performance under adverse weather conditions.

Smoke, consisting of dust and particles, is a normal atmospheric phenomenon. As we know, similar to haze, the brightness and contrast of images captured in smoky environments are always reduced, producing seriously degraded images. This may lead to a significant decline in the ability of the detection of critical objects. Thus, in this work, we focus on object detection under smoky conditions which is also the goal of Track 1 in the 5<sup>th</sup> UG<sup>2+</sup> challenges of CVPR 2022.

Contrastive learning, belonging to the class of discriminative representation learning, learns a representation by comparing among different samples. This comparison is implemented between positive pairs of

---

\*Corresponding author.

“similar” inputs and negative pairs of “dissimilar” pairs. By comparing them, the objective of contrastive learning that the representation of “similar” samples should be mapped close together while that of “dissimilar” samples should be further away in the embedding space, can be effectively accomplished. It follows that, contrastive learning has a clear advantage that it can neglect the superficial phenomenon of inputs while is able to learn and extract their internal consistency. Thus, we leverage it to seize the essential information of objects among different inputs, to help forge a robust object detection.

In this paper, we propose a contrastive learning-based robust object detection algorithm for smoke images. The proposed algorithm has two modules: contrastive learning module and object bounding box prediction module. We first use the first module, to train a base encoder network and a projection head, so as to learn representations by maximizing agreement between different augmented views of the same smoke image through a contrastive loss in the latent space. The learned representations along with the original smoke image are then fed to the second module, generating three prediction features with different scales. These prediction features are employed to estimate the bounding box of each object. In addition, we also propose a novel affine data augmentation method to simulate UAV view angle changes. In this method, testing data are transformed by perspective transformation, into new images with different camera angles, which are then added to training data.

Our main contributions can be highlighted as follows:

1. In this work, we introduce contrastive learning to effectively seize the internal consistency of objects, and then propose a contrastive learning-based robust object detection algorithm for smoke images.
2. Considering UAV view angle changes usually exist among photos shot by UAVs, we also propose a novel affine data augmentation method to simulate these changes.
3. Experimental results demonstrate our proposed method is the current state-of-the-art (SOTA) in the 5<sup>th</sup> UG<sup>2+</sup> challenges of CVPR 2022.

## 2. Related work

### 2.1. Anchor-based detectors

Anchor-based detectors can be further classified into the following two types: two-stage methods and one-stage methods. Well-known R-CNN series [7-9] belong to the former, while the latter mainly includes YOLO series [10-15], SSD [16], RetinaNet [17], and RefineDet [18]. For the two-stage approaches, in the first stage a huge number of candidate bounding boxes are established through a region proposal network based on a sliding-window mechanism, and in the second stage region-of-interest (RoI) pooling is

leveraged to extract feature maps from each bounding box. In contrast, the one-stage detectors directly regress the bounding boxes, leading to high efficiency while sacrificing accuracy.

R-CNN [7] is one of the earliest and successful object detection methods using CNN. In R-CNN, the traditional handcrafted feature tool is replaced by a CNN-based feature learning process, resulting in a significant performance boost. Its successor important variant, Fast R-CNN [8], improves detection efficiency by a large margin. In this variant, an input image and multiple RoIs are input into a fully convolutional (FC) network. Each RoI is pooled into a fixed-size feature map and then transformed into a fixed-length vector, to output softmax probabilities and per-class bounding box regression offsets. Following Fast R-CNN, Faster R-CNN [9] introduces a region proposal network to construct region proposals, which are then sent into a RoI pooling to obtain proposal feature maps and bounding box positions.

The first version of YOLO series, YOLO [10], was proposed by Redmon *et al.* in 2015. This technique uses a CNN to predict the class and bounding box of each object in a single run. Compared with YOLO, YOLOv2 [11] can detect 9000 different objects with more accuracy and lower complexity. It introduces anchor boxes to predict bounding boxes and develop a new network, namely DarkNet-19. Their subsequent variant is YOLOv3 [12] that is based on DarkNet-53. As its backbone, DarkNet-53 has 53 convolutional layers but without FC layers. Inspired by the pyramid idea, YOLOv3 generates three scales of feature maps to detect different sizes of objects. In 2020, YOLOv4 [13] was proposed to make everyone be able to use a 1080 Ti or 2080 Ti GPU to train a super-fast and accurate object detector. It consists of CSPDarkNet-53 [26] as its backbone, SPP and PAN as its neck, and YOLOv3 as its head. YOLOv5 [14] is also proposed soon after the emergence of YOLOv4. In contrast to its four predecessors, YOLOv5 not only has the network with the smaller size of weight parameters, but also is faster. The latest version of the YOLO series is YOLOX [15]. It chooses YOLOv3 with DarkNet-53 as baseline and exceeds YOLOv5 by 1.8 % AP.

Liu *et al.* proposed SSD [16] in 2016 to detect objects in images using a single deep neural network. In SSD, the output space of bounding boxes is discretized into a set of default boxes over different aspect ratios and scales per feature map location. A simple one-stage object detector called RetinaNet [17], is developed to demonstrate the effectiveness of a new loss function. This loss function is acted as a more effective alternative to previous approaches for dealing with class imbalance. To achieve better accuracy than two-stage methods and maintain comparable efficiency of one-stage methods, RefineDet [18] is specially designed by Zhang *et al.* It consists of an anchor refinement module and an object detection module.

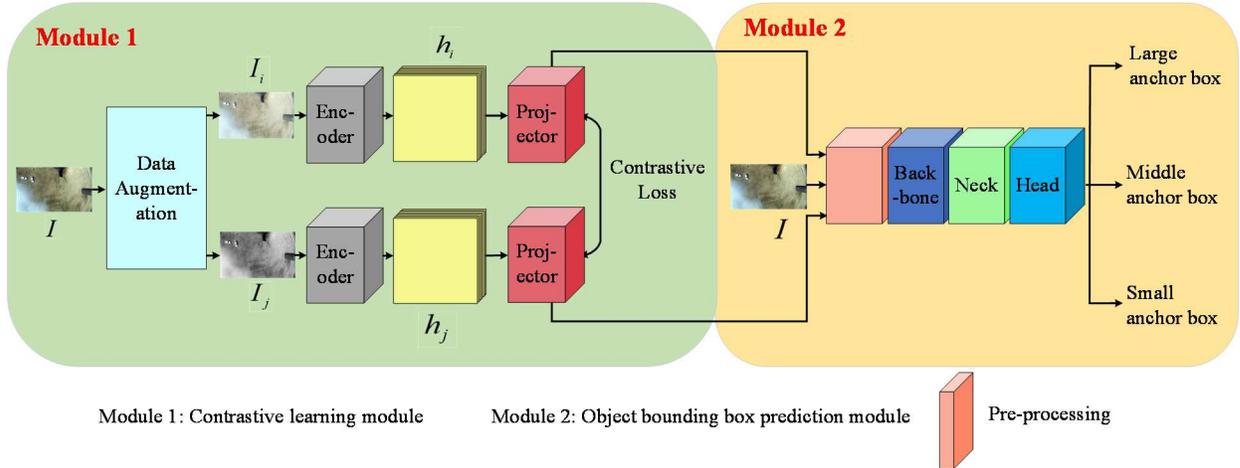


Figure 1: Overall framework of our proposed object detector.

## 2.2. Anchor-free detectors

Substituting key point prediction for bounding box estimation, CornerNet [21] and CenterNet [22] have outperformed their anchor-based counterparts. Without the concept of anchor, Law and Deng proposed CornerNet to detect an object bounding box as a pair of points, the top-left corner, and the bottom-right corner. In CornerNet, a single convolutional network is used to predict heatmaps, and moreover a corner pooling is designed, which is a new type of pooling layer that helps a convolutional network better localize corners of bounding boxes. Besides CornerNet, CenterNet also does not rely on anchors. This anchor-free detector models an object as a single point, i.e. the center of its bounding box. Object size is then regressed directly from image features at the center location. Rashwan *et al.* proposed a new scale and aspect ratio aware architecture, i.e. MatrixNets. This method maps objects with similar sizes and aspect ratios into specialized layers.

## 3. Proposed object detector

Our proposed object detection algorithm mainly contains two modules, one being contrastive learning module, and the other object bounding box prediction module, as illustrated in Fig. 1. The former module is devoted to capturing the intrinsic agreement features among the different augmented versions of each same smoke image, and thus be able to learn their consistent representations. Then, the latter module does not only take original smoke image data, instead takes these representation vectors along with the original smoke image as input, to predict the bounding box for each object.

## 3.1. Contrastive learning module

Fig. 2 shows the structure of the contrastive learning module. Three components in total that are data augmentation, encoder network, and projection head, are included into this module, which is similar to SimCLR [27]. First, in the data augmentation, given a smoke image  $I$  randomly, we employ two image augmentation methods, namely image graying and random image brightness enhancement, to generate its relatedly augmented views. These two augmentations, denoted as  $I_i$  and  $I_j$ , respectively, can be regarded as a positive pair of  $I$ .

Second, the encoder extracts features  $h_i$  and  $h_j$  from the augmented smoke data examples  $I_i$  and  $I_j$ , respectively. In this encoder, we leverage a well-known deep learning network, i.e. ResNet50 [28]. These extracted features are subsequently output as the consistent representation vectors of the data examples.

Third, the last component, i.e. the projector that is also a nonlinear transformation, further maps  $h_i$  and  $h_j$  to their corresponding more abstractive characteristics  $z_i$  and  $z_j$ , respectively. In the projection head, we use not a complex network but two simple MLPs each of which has only one hidden layer. In addition, between those two MLPs we also adopt a ReLU nonlinearity.

Finally, we optimize the contrastive learning module using the following loss function:

$$L_{cont} = \frac{1}{2N} \sum_{k=1}^N [l(2k-1, 2k) + l(2k, 2k-1)] \quad (1)$$

$$l(i, j) = -\log \frac{\text{sim}(z_i, z_j)/\tau}{\sum_{k=1}^{2N} \mathbb{I}_{[k \neq i]} \exp(\text{sim}(z_i, z_k)/\tau)} \quad (2)$$

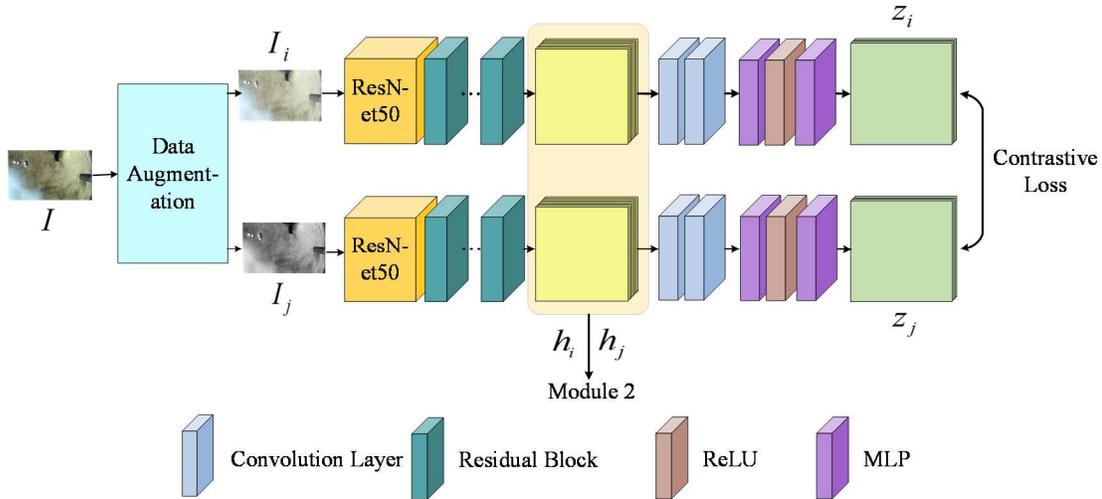


Figure 2: Diagram of the contrastive learning module.

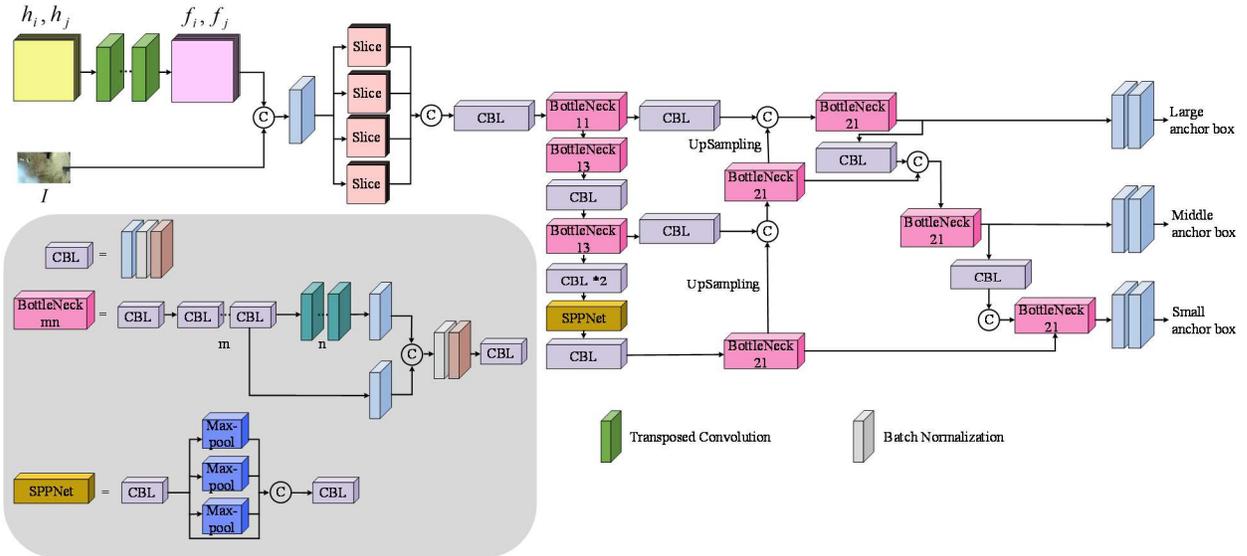


Figure 3: Illustration of the object bounding box prediction module.

where a batch of  $N$  examples are randomly sampled to construct  $2N$  augmented images, in which for  $z_i$  except  $z_j$  the other  $2N-2$  augmented examples are acted as its negative examples.  $\mathbb{I}_{[k \neq i]} \in \{0,1\}$  is an indicator function equaling to 1 if  $k \neq i$ ,  $sim(u, v) = u^T v / \|u\| \|v\|$  represents the dot product between L2 normalized  $u$  and  $v$ , and  $\tau$  denotes a temperature parameter [29-31].

### 3.2. Object bounding box prediction module

As the same suggests, the second module aims to

determine the bounding boxes of objects. This module consists of four components: pre-processing, backbone, neck, and head, as shown in Fig. 1. To be specific, Fig. 3 gives the detailed diagram of the object bounding box prediction module. In the pre-processing part, we perform Transposed convolution to change the input representation vectors  $h_i$  and  $h_j$  into the features  $f_i$  and  $f_j$ . These features  $f_i$  and  $f_j$ , together with the feature  $g$  derived from  $I$ , are then pass through the backbone subnetwork of YOLOv5.

YOLOv5 has four different versions: YOLOv5s,

YOLOv5m, YOLOv5l, YOLOv5x, where they have gradually increasing network depths and widths from YOLOv5s to YOLOv5x. In this work we employ the simplest version YOLOv5s to decrease the complexity. This YOLOv5s is composed of backbone, neck, and head, which are the same with those in [14], respectively. The backbone part mainly includes Focus and bottleneck, where Focus part is composed of convolutional layer, slice, and concatenation, and bottleneck mainly contains CBL in [14] and residual blocks. Next, the neck is used to collect feature maps from different stages, and is also mainly composed of CBL. Finally, the head is adopted to predict the bounding boxes of objects. As the loss function, we employ the loss functions of YOLOv5 to optimize the object bounding box prediction module.

### 3.3. Data augmentation

In this paper, we also propose a novel affine data augmentation method considering UAV view angle changes exist among shooting images. First, we take some photos with similar contents but with different camera angles and different shooting distances. Then, the homography matrix between each pair of similar images is computed by using perspective transformation. After that, we perform these homography matrices on every image of testing dataset, creating new transformed images. Finally, these transformed images are put to use in the training of network, together with original training data.

## 4. Experiments

### 4.1. Baselines and datasets

In order to evaluate the performance of our proposed object detector, in experiments it is compared with widely used SOTA object detection methods including CenterNet [22], YOLOv5 [14], and twelve excellent methods proposed by CVPR 2022 challenge participants. The dataset used for training and testing is A2I2-Haze [32] that is the only dataset adopted in the track of object detection in Haze in the 5<sup>th</sup> UG<sup>2+</sup> challenge of CVPR 2022. A2I2-Haze is the first real haze dataset with in-situ smoke measurement aligned to aerial and ground imagery. This dataset consists of not only a total of 177 paired hazy/clean frame images clipped from 12 videos but also 240 annotated clean images collected from the same sources for training, and 60 other smoke images for testing. In addition, we also employ the proposed data augmentation method. Affine transformation is performed on those test images, to produce 384 transformed images with different angles and different distances, which are then employed as the supplement to the training data.

Table 1: Detection results of the SOTA methods and our proposed algorithm.

Method	AP (%)	AP <sub>50</sub> (%)	AP <sub>75</sub> (%)
YOLOv5 [14]	36.63	55.75	42.29
CenterNet [22]	27.74	44.72	29.59
sl	-	98.43	-
asdfghjkl	-	95.98	-
longpham3105	-	95.69	-
haoxl	-	95.38	-
thlbsj	-	94.99	-
Feiyu_Yao	-	94.75	-
Frank_Yao	-	93.60	-
da2986	-	92.94	-
willer	-	92.35	-
tangweiyi	-	92.35	-
tanghulu	-	92.09	-
tangzixia	-	86.23	-
Proposed	<b>80.77</b>	<b>98.49</b>	<b>95.42</b>

### 4.2. Implementation details

We first use the training data to individually update weights only for the contrastive learning module. Once accomplishing this training process, the built representation vectors along with original images are then used as the input of the object bounding box prediction module to train its network. After all these tasks, the training of our proposed object detector is really finished.

The proposed object detector is trained and tested on NVIDIA GeForce RTX3080 GPU. In the training process of the contrastive learning module, the batch size and the patch size are set to 8 and  $512 \times 512$ , respectively, and the stochastic gradient descent (SGD) optimizer is adopted with the learning rate initialized to 0.05 and decreased by 90 percent at 150, 200, and 250 epochs in a total of 300 epochs. In addition, for the training of the object bounding box prediction module, we set the batch size and the patch size to 32 and  $512 \times 512$ , respectively. Warmup is performed before the training. The learning rate is initialized to 0.01, and then the Cosine learning rate drop strategy is adopted in the training in a total of 500 epochs.

### 4.3. Quantitative evaluation

Table 1 gives the numerical results of object detection produced by the SOTA methods and our proposed detector. For those challenge participants, only their AP<sub>50</sub> results are announced by CVPR 2022, as described in Table 1. From the results in this table, one can see that because neither YOLOv5 nor CenterNet takes measures to adapt to smoky conditions, they both obtain inferior performance. Although each of the twelve challenge participants gets high score in the AP<sub>50</sub>, our proposed algorithm still achieves the best average prediction result referring to object detection among all the evaluated methods.

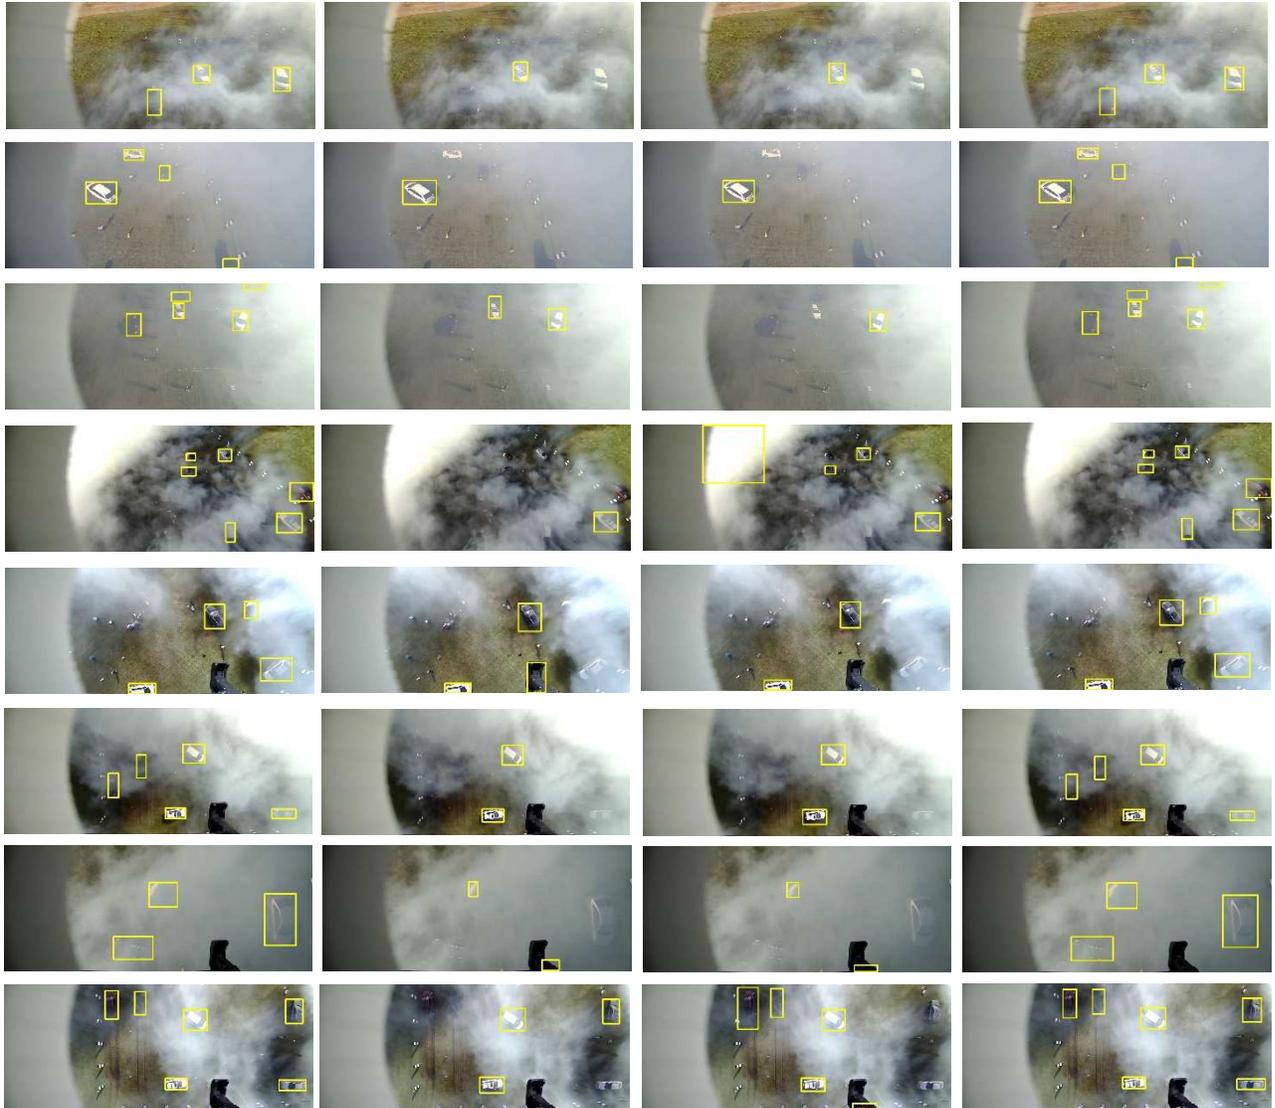


Figure 4: Qualitative comparisons on six real smoke images in A2I2 dataset for GT and three methods, i.e. GT, CenterNet, YOLOv5, and our proposed algorithm, from left to right, respectively.

Table 2: Quantitative results of ablation study.

Method	AP (%)	AP <sub>50</sub> (%)	AP <sub>75</sub> (%)
<i>Baseline1</i>	79.29	98.48	95.25
<i>Baseline2</i>	38.85	55.86	42.57
Proposed	<b>80.77</b>	<b>98.49</b>	<b>95.42</b>

#### 4.4. Quantitative evaluation

Fig. 4 shows the qualitative comparison results

conducted on six chosen test images from A2I2 dataset. From this figure, we can observe that YOLOv5 or CenterNet does not find the target well or produce false detection, so that they get low AP results. Moreover, our proposed algorithm detects the most objects among all the methods.

#### 4.5. Ablation study

In this subsection, we will verify the validity of our

contrastive learning mechanism used in object detection by directly removing Module 1 from the framework in Fig. 1. In addition, the validity of the proposed data augmentation is also evaluated by avoiding employing those transformed data in the training. To clarify the effect of them, we compare our proposed method with the following two baselines: 1) *Baseline1*: We only train Module 2 with the augmented training data. 2) *Baseline2*: We train our proposed method only with original training data. Table 2 compares their object detection results in the ablation study. From these results in Table 2, it can be seen that contrastive learning can learn the internal consistency of objects to further accurately determine target positions. Also, our proposed data augmentation method largely improves the AP results of object detection.

## 5. Conclusions

In this paper, we focused on object detection under smoky conditions and proposed a corresponding robust object detector. Contrastive learning is a strong tool that can extract the internal consistency of “similar” samples by comparing positive pairs of “similar” inputs and negative pairs of “dissimilar” pairs. We thus introduce contrastive learning to help forge a robust object detection. In the proposed algorithm, we first employ a contrastive learning module to effectively learn consistent representations for objects. The second module, namely object bounding box prediction module, takes these representation vectors as input, to determine the bounding boxes of objects. Considering that UAV view angle changes usually exist among photos taken by UAVs, we also propose a novel affine data augmentation method to simulate these angle changes. This augmentation method transforms the testing data via perspective transformation into new transformed images with different camera angles, which are further used for the training of network. In comparison with the current SOTA methods, experimental results indicate that our proposed object detector achieves more accurate detection performance for smoke images, and is the current SOTA in the Track 1 of 5<sup>th</sup> UG<sup>2+</sup> challenges of CVPR 2022.

## 6. Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 61471277 and in part by the 111 Project under Grant B08038.

## References

- [1] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE TPAMI*, 32(9): 1627–1645, 2009.
- [2] W. Yang, R. T. Tan, J. Feng, Z. Guo, S. Yan and J. Liu. Joint rain detection and removal from a single image with contextualized deep networks. *TPAMI*, 42(6): 1377-1393, 2020.
- [3] P. Wang, W. Wu, Z. Li, and Y. Liu. See SIFT in a rain: Divide-and-conquer SIFT key point recovery from a single rainy image. In *VCIP*, pages 1-5, 2021.
- [4] K. Gu, Z. Xia, J. Qiao, W. Lin. Deep dual-channel neural network for image-based smoke detection. *IEEE TMM*, 22(2): 311-323, 2020.
- [5] W. Ren, J. Pan, H. Zhang, X. Cao, and M. -H. Yang. Single image dehazing via multi-scale convolutional neural networks with holistic edges. *IJCV*, 128: 240-259, 2020.
- [6] M. Hnawa and H. Radha. Object detection under rainy conditions for autonomous vehicles. *IEEE Signal Processing Magazine*, 38(1): 53-67, 2021.
- [7] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014.
- [8] R. Girshick. Fast R-CNN. In *ICCV*, pages 1440–1448, 2015.
- [9] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *TPAMI*, 39(6): 1137–1149, 2017.
- [10] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [11] J. Redmon and A. Farhadi. YOLO9000: Better, faster, stronger. In *CVPR*, pages 7263–7271, 2017.
- [12] J. Redmon and A. Farhadi. YOLOv3: An incremental improvement. 2018, arXiv:1804.02767.
- [13] A. Bochkovskiy, C. Wang, and H. M. Liao. Yolov4: Optimal speed and accuracy of object detection. 2004, arXiv preprint arXiv:2004.10934.
- [14] G. Jocher. YOLOv5. 2021, <https://github.com/ultralytics/yolov5>.
- [15] Z. Ge, S. Liu, F. Wang, Z. Li, J. Sun. YOLOX: Exceeding YOLO series in 2021. 2021, arXiv:2107.08430v2.
- [16] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Fu, and A. C. Berg. SSD: Single shot multibox detector. In *ECCV*, pages 21–37, 2016.
- [17] T. Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. *TPAMI*, 42(2): 318–327, 2020.
- [18] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li. Single-shot refinement neural network for object detection. In *CVPR*, pages 4203–4212, 2018.
- [19] J. Dai, Y. Li, K. He, and J. Sun. R-FCN: Object detection via region-based fully convolutional networks. In *NIPS*, pages 379–387, 2016.
- [20] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, pages 2961–2969, 2017.
- [21] H. Law and J. Deng. CornerNet: Detecting objects as paired keypoints. 2018, arXiv:1808.01244.
- [22] X. Zhou, D. Wang, and P. Krahenbuhl. Objects as points. 2019, arXiv:1904.07850v2.
- [23] A. Rashwan, A. Kalra, and P. Poupart. Matrix nets: A new deep architecture for object detection. In *ICCV*, pages 2025–2028, 2019.
- [24] Z. Tian, C. Shen, H. Chen, and T. He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019.

- [25] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin. RepPoints: Point set representation for object detection. In *ICCV*, pages 9657–9666, 2019.
- [26] C. Wang, H. M. Liao, Y. Wu, P. Chen, J. Hsieh, and I. Yeh. CSPNet: a new backbone that can enhance learning capability of CNN. In *CVPRW*, pages 1571–1580, 2020.
- [27] T. Chen, S. Kornblith, M. Norouzi, G. Hinton. A simple framework for contrastive learning of visual representations. 2020, arXiv:2002.05709v3.
- [28] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [29] K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NIPS*, pages 1857–1865, 2016.
- [30] Z. Wu, Y. Xiong, S. X. Yu, and D. Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018.
- [31] A. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. 2018, arXiv preprint arXiv:1807.03748.
- [32] A2I2 haze dataset.  
[http://cvpr2022.ug2challenge.org/dataset22\\_t1.html](http://cvpr2022.ug2challenge.org/dataset22_t1.html)