

# Visible and Clear: Finding Tiny Objects in Difference Map

Bing Cao, Haiyu Yao, Pengfei Zhu\*, and Qinghua Hu

College of Intelligence and Computing, Tianjin University, Tianjin, China  
Tianjin Key Lab of Machine Learning, Tianjin, China  
{caobing,yaohaiyu,zhupengfei,huqinghua}@tju.edu.cn

**Abstract.** Tiny object detection is one of the key challenges for most generic detectors. The main difficulty lies in extracting effective features of tiny objects. Existing methods usually perform generation-based feature enhancement, which is seriously affected by spurious textures and artifacts, making it difficult to make the tiny-object-specific features visible and clear for detection. To address this issue, we propose a self-reconstructed tiny object detection (SR-TOD) framework. We for the first time introduce a self-reconstruction mechanism in the detection model, and discover the strong correlation between it and the tiny objects. Specifically, we impose a reconstruction head in-between the neck of a detector, constructing a difference map of the reconstructed image and the input, which shows high sensitivity to tiny objects. This inspires us to enhance the weak representations of tiny objects under the guidance of the difference maps. Thus, improving the *visibility* of tiny objects for the detectors. Building on this, we further develop a Difference Map Guided Feature Enhancement (DGFE) module to make the tiny feature representation more *clear*. In addition, we further propose a new multi-instance anti-UAV dataset. Extensive experiments demonstrate our effectiveness. The code is available: <https://github.com/Hiyuur/SR-TOD>.

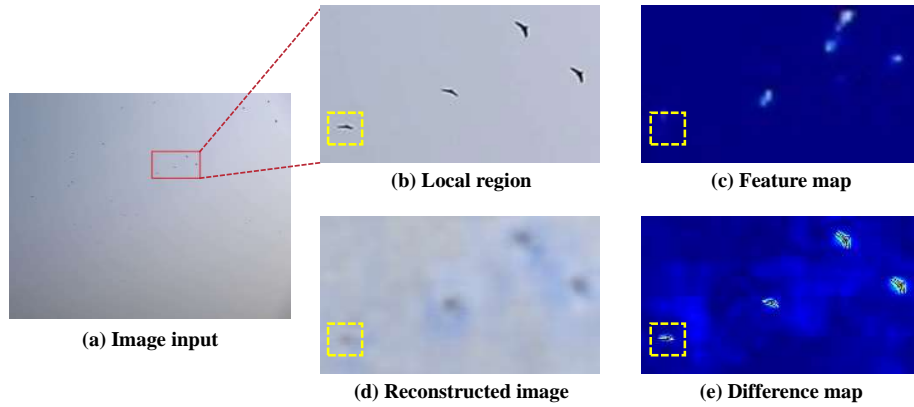
**Keywords:** tiny object detection · self-reconstruction · difference map

## 1 Introduction

The subfield of object detection that identifies and categorizes objects with diminutive dimensions is known as tiny object detection. According to setting in MS COCO [32], an object qualifies as *small* if it occupies an area equal to or less than  $32 \times 32$  pixels. The AI-TOD benchmark [51] refines this definition and categorizes objects as “very tiny” if they span 2 to 8 pixels, “tiny” for 8 to 16 pixels, and “small” for 16 to 32 pixels. We term these objects uniformly as “*tiny*” objects. Tiny objects commonly appear in a variety of real-world applications, such as autonomous driving, industrial inspection, and pedestrian detection, often restricted by considerable imaging distances or the inherently minute size of the objects. Despite its relevance, tiny object detection (TOD) remains an

---

\* Corresponding author



**Fig. 1:** Visualization of the image self-reconstruction mechanism. The results are from Cascade R-CNN [6]. **a** shows the entire image. To demonstrate the effect more clearly, **b** zooms in on the local region in the red box. **c** shows the visual heatmap of the corresponding region in the feature map of FPN P2. **d** is the reconstructed image. **e** is the difference map. The yellow dotted box highlights the tiny drone whose signal is almost wiped out in the feature map.

arduous endeavor, with even state-of-the-art detectors struggling to bridge the performance disparity between tiny and normal-sized object detection [11]. The pursuit of improved methods in this area is of considerable theoretical and practical importance.

Relative to prevalent complications in general object detection, such as object occlusion, tiny object detection presents some unique challenges. The foremost challenge arises from the issue of information loss on tiny objects [11]. Prevailing detection architectures employ backbone networks, like ResNet and others [16, 22, 52], for feature extraction. Nonetheless, these frameworks often implement downsampling operations aimed at eliminating noisy activations and diminishing the spatial resolution of feature maps, a process that unavoidably results in the information loss of tiny objects. In addition, TOD is also impeded by the inherently limited size and scant information content of tiny objects, resulting in substantial information loss throughout the feature extraction phase. Such degradation in object representation significantly hampers the detection head’s capacity to localize and distinguish between tiny objects. Consequently, tiny objects become indiscernible to the detector. In particular, the weak signals of “very tiny” objects are almost wiped out under these conditions, making it difficult for detectors to localize and identify them. As shown in Fig. 1, the feature heatmaps display the activated signals of the detection model towards tiny drones, which is often weak for tiny objects. For instance, the tiny object in the lower left almost disappears, affecting the detection performance. This indicates that many tiny objects are insufficiently visible to detectors. Therefore, the performance of generic detectors dramatically decreases in tiny object detection tasks [51, 56].

Many existing methods [1–3, 28, 37, 39] often use a generation approach with super-resolution architecture to alleviate the low-quality representation issue of tiny objects caused by information loss. These methods typically incorporate generative adversarial networks [20] into the object detection framework, constructing pairs of high-resolution and low-resolution samples. This approach enables the generator to learn to restore the distorted structure of tiny objects, aiming to enhance the features of low-quality tiny objects. Nonetheless, these approaches typically necessitate a substantial volume of medium and large-sized samples, posing a significant challenge in executing super-resolution on tiny objects with weak signals. Furthermore, these methods are inclined to fabricate spurious textures and artifacts, which decreases detection performance [14]. It is worth noting that super-resolution architecture brings a significant amount of computational overhead, complicating the end-to-end optimization [11].

Compared to the inefficient feature enhancement with intricate super-resolution architectures, recapturing lost information in the backbone network is a more intuitive and reasonable strategy. For the first time, we introduce a simple but effective image self-reconstruction mechanism into the object detection framework. The feature maps extracted by the detection model are restored by a reconstruction head, which is constrained by mean square error at the pixel level. Note that, image reconstruction, a task situated within the low-level vision domain, is highly sensitive to pixel changes [7]. Since we reconstruct the input image from the detection model, the image regions that are difficult to restore may correspond to those where structural and textural information is severely lost during the feature extraction of the backbone network, notably tiny objects. With the disparities between the reconstructed and the original images, we can pinpoint the regions that have undergone significant information loss, which in turn provides potential prior knowledge for detecting tiny objects. Therefore, we subtract the original image from the self-reconstructed image to construct a difference map, as shown in Fig. 1. We first discover a strong correlation between the self-reconstructed difference map and tiny objects. The “*very tiny*” objects that are almost eliminated in Fig. 1 can also be clearly shown in the difference map. Most tiny objects in the image have a significant activation in the difference map. Moreover, the difference map also preserves the main structures of the tiny objects. We believe that the difference map manifests the detector’s pixel-level discernment for the areas of interest, as well as the potential location and structure of the tiny objects. Overall, the difference map enables tiny objects that have weak signals to be more *visible*.

Therefore, building upon this discovery, we further incorporate the prior information from the difference map into the object detection model. We develop a simple and effective Difference Map Guided Feature Enhancement module, which calculates an element-wise attention matrix by reweighting the difference map along the channel dimension to perform feature enhancement for tiny objects. Consequently, by converting the reconstruction loss into a constraint specifically targeting tiny objects, we enhance the model’s capacity for detecting such objects, making the tiny object more *clear* to detectors.

Furthermore, we collected a new anti-UAV dataset, named DroneSwarms, which is a typical tiny object detection scenario under various complex backgrounds and lighting conditions. Our DroneSwarms has the smallest average size (about 7.9 pixels) of drones for anti-UAV. Experiments on our dataset and two other datasets demonstrate our superiority against the competing methods.

The contributions of this paper can be summarized as follows:

- (1) We propose a self-reconstructed tiny object detection (SR-TOD) framework, revealing for the first time the robust association between the difference map and tiny objects, thus offering prior information on the position and structure of tiny objects. We effectively convert the typically lost information of tiny objects into actionable prior guidance.
- (2) We design a Difference Map Guided Feature Enhancement (DGFE) module, which improves the feature representation of tiny objects to make them clearer. The DGFE module can be easily and flexibly integrated into generic detectors, effectively improving the performance of tiny object detection.
- (3) We propose a new tiny object detection dataset for anti-UAV, named DroneSwarms, which has the smallest average object size currently. Extensive experiments on our dataset and two other datasets with a large number of tiny objects validate our effectiveness against the competing methods.

## 2 Related Work

**Object detection** algorithms come in various types. Two-stage detectors feed the extracted feature maps into a Region Proposal Network to extract proposals in the first stage. Then, in the second stage, they perform classification and regression tasks based on these proposals, providing high recognition and localization accuracy. Classic two-stage detectors include Fast R-CNN [18], Faster R-CNN [42], Cascade R-CNN [6] *etc.* One-stage detectors directly perform object localization and classification on input images, offering faster processing speeds. Representative examples of one-stage detectors include the YOLO series [4, 40, 41] and RetinaNet [31]. Besides, anchor-free algorithms such as FCOS [49] and FoveaBox [26] predict objects based on central points, while methods like CornerNet [27], Grid R-CNN [34], and RepPoints [55] make predictions based on keypoints. Recent detectors such as DETR [8], Deformable DETR [64], and Sparse R-CNN [47] explore new paradigms for end-to-end object detection. These algorithms commonly use FPN [30] as the neck module, which allows our method to be easily integrated into most generic detectors.

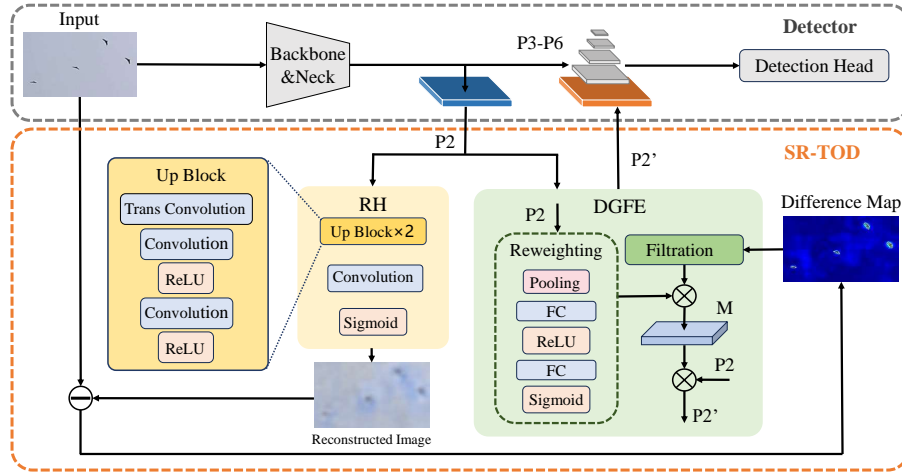
**Tiny Object Detection** is a challenging task for most generic detectors. In recent years, the research in tiny object detection has mainly focused on data augmentation, scale awareness, context modeling, feature imitation, and label assignment [11]. *Data augmentation*: Krisantal *et al.* [25] copy and paste tiny objects to increase the number of samples. DS-GAN [5] designs a new data augmentation pipeline to generate high-quality synthetic data for tiny objects. *Scale awareness*: Lin [30] *et al.* propose the most popular multi-scale network feature pyramid network using a pyramid of feature levels and feature fusion. Singh

*et al.* [46] design the Scale Normalization for Image Pyramids (SNIP) to select some instances for training. PANet [33] enriches the feature hierarchy through bidirectional paths and enhanced deeper features with accurate positioning signals. NAS-FPN [17], Bi-FPN [48] and Recursive-FPN [38] further develop upon the basis of FPN. Gong *et al.* [19] adjust the coupling between adjacent layers of FPN by setting fusion factors to optimize feature fusion. Yang *et al.* [54] design a cascaded sparse query mechanism to effectively utilize high-resolution features to enhance the detection performance of tiny objects while maintaining fast inference speed. *Context modeling*: Chen *et al.* [9] utilizes context region representations containing proposed patches for subsequent recognition. SINet [23] introduces a context-aware RoI pooling layer to maintain contextual information. *Feature imitation*: Many methods [1,2,28] utilize generative adversarial networks to perform super-resolution on tiny objects. Noh *et al.* [37] alleviated the mismatch between the receptive fields of high-resolution features and low-resolution features through dilated convolutions. Deng *et al.* [14] proposed a feature texture transfer module to expand the feature pyramid, enabling the new feature layers to contain more detailed information of tiny objects. *Label assignment*: ATSS [60] adaptively adjusts positive and negative samples based on their statistical features. Xu *et al.* [53] propose a simple yet effective strategy called Receptive Field-based Label Assignment (RFLA) to alleviate the scale-sample imbalance problem in anchor-based and anchor-free detectors.

Many tiny object detection methods have not focused on the key issue of information loss. Feature imitation methods attempt to alleviate this issue through generation but tend to fabricate spurious textures and artifacts [11, 14]. In contrast, we introduce an image self-reconstruction mechanism to identify areas of significant information loss and leverage this prior knowledge to improve tiny object detection performance.

**Anti-UAV Dataset.** Drones, due to their tiny size and low cost among other features, are widely used in inspection and surveillance. However, illegal drone flights have introduced many potential risks and threaten public safety, making anti-UAV measures an important task. Since drones are tiny in size and often fly at medium to high altitudes, they tend to appear very small in images. Therefore, anti-UAV technology is very well suited for the application of tiny object detection.

To our knowledge, there are three publicly available visible light anti-UAV datasets in the current field. MAV-VID [43] consists of 64 video sequences containing a single drone captured from multiple viewpoints, with the targets predominantly located in the central region of the images. The drones are relatively large in scale, with an average size of approximately 166 pixels. Drone-vs-Bird [12] consists of 77 video sequences and is used to distinguish between drone and bird targets. The average size of the drones in the dataset is approximately 28 pixels. DUT Anti-UAV [61] is divided into two subsets: detection and tracking. The majority of the images in this dataset contain single objects and almost no very tiny objects.



**Fig. 2:** Overall model architecture. RH refers to the reconstruction head, DGFE is the difference map guided feature enhancement.  $P2 - P6$  denote the feature maps extracted from FPN.  $P2'$  is the enhanced feature map.  $M$  is the element-wise attention matrix.

### 3 Method

#### 3.1 Overall Architecture

In this paper, we propose a tiny object detection framework based on the image self-reconstruction mechanism, as shown in Fig. 2, to address the major challenge of severe information loss in the feature extraction process of the backbone network. Initially, the image is fed into the backbone network, which extracts features from this input and relays the feature maps to the neck module—typically FPN [30]—to create a multi-scale feature pyramid ranging from  $P2$  to  $P5$ . In line with the architecture of prevalent detectors, the task of tiny object detection is designated to  $P2$ . Consequently, the interfacing of our self-reconstruction mechanism with the detector occurs exclusively via  $P2$ . We feed  $P2$  into the reconstruction head, which will be detailed in Sec. 3.2. The reconstruction head generates images that are dimensionally consistent with the original input. By subtracting the original images from the reconstructed images, taking the absolute values, and averaging across the three color channels, we obtain the difference map. We feed both the difference map and  $P2$  into the Difference Map Guided Feature Enhancement (DGFE) module, which will be detailed in Sec. 3.3. The DGFE module enhances specific tiny object features in  $P2$  based on the prior knowledge of the difference map, resulting in  $P2'$ . This enhanced feature map,  $P2'$ , supplants the original  $P2$  as the bottom level of the feature pyramid, subsequently fed into the detection head. Note that some one-stage detectors [31] rely exclusively on lower-resolution feature maps,  $P3$ , hence our framework can also use  $P3$  for image reconstruction. Owing to the widespread adoption of Feature Pyramid

Networks (FPN) [30] and its variants as the neck module, our framework is readily integrable with the majority of contemporary detection models.

### 3.2 Difference Map

The downsampling process inherent to the feature extraction in the backbone network inevitably results in information loss of objects. This loss is particularly acute for tiny objects due to their restricted sizes. In this case, the weak signals of tiny objects are almost wiped out, making it difficult for the detection head to predict from such low-quality representations [11]. In response to this problem, we reconsidered the attributes of feature maps at different levels within the FPN framework. Given that high-level, low-resolution features are imbued with rich semantic content, while low-level, high-resolution features possess greater local detail and positional information, we favor image reconstruction using the low-level feature maps.

The U-Net [44] architecture, generally used for image reconstruction tasks, bears a similarity to the popular FPN module of object detection models. Therefore, we designed a simple reconstruction head inserted into the top-down path of the FPN, as shown in Fig. 2. Given a low-level feature map  $X$  with  $C$  channels and size of  $H \times W$ , the upsampling result  $Up(X) \in \mathbb{R}^{\frac{C}{2} \times 2H \times 2W}$  can be computed as

$$Up(X) = \delta(Conv2(\delta(Conv1(TranConv(X))))), \quad (1)$$

where  $\delta$  denotes Rectified Linear Unit (ReLU) [36].  $Conv1$  and  $Conv2$  denote the convolutions with kernel size of  $C \times C \times 3 \times 3$ .  $TranConv$  denotes Transpose Convolution [58]. The kernel size of  $TranConv$  is  $\frac{C}{2} \times C \times 4 \times 4$ , and the step length is 2. Given the original image  $I_o \in \mathbb{R}^{3 \times H \times W}$  and the bottom-level feature map  $P2 \in \mathbb{R}^{C \times \frac{H}{4} \times \frac{W}{4}}$ , the reconstruction head shown in Fig. 2 can be formulated as

$$I_r = \sigma(Conv(Up(Up(P2)))), \quad (2)$$

where  $I_r$  is the reconstructed image.  $\sigma$  denotes Sigmoid function and the kernel size of  $Conv$  is  $3 \times C \times 3 \times 3$ . The reconstructed image  $I_r$  obtained after two upsampling operations is of the same shape as the original image  $I_o$ . Therefore, the difference map  $D$  can be computed as

$$D = Mean_{channel}(Abs(I_r - I_o)), \quad (3)$$

where  $Mean_{channel}$  denotes computing the mean value along the channel dimension, and  $Abs$  denotes computing the absolute value of each element.

The optimization of the reconstruction head’s parameters is achieved by computing the Mean Squared Error (MSE) loss between the original and reconstructed images.

### 3.3 Difference Map Guided Feature Enhancement

Following the creation of the difference map via the self-reconstruction mechanism, a critical challenge is the effective utilization of the difference map’s

prior information to bolster tiny object detection capabilities. Considering that the difference map represents the potential location and structural information of tiny objects, we designed a simple plug-and-play module called Difference Map Guided Feature Enhancement (DGFE). The DGFE module calculates an element-wise attention matrix, denoted as  $M \in \mathbb{R}^{C \times H \times W}$ , from the difference map with the aim of executing targeted feature enhancement of tiny objects within  $P2 \in \mathbb{R}^{C \times H \times W}$ , as shown in Fig. 2.

**Filtration.** The inherent discrepancies between the reconstructed and original images result in activation across virtually the entire difference map to varying extents. To filter out most of the noise signals and make the difference map more sharp, we construct a binary difference map  $D_b \in \mathbb{R}^{1 \times 4H \times 4W}$ . For this purpose, we set a learnable threshold  $t$ . Given a difference map  $D \in \mathbb{R}^{1 \times 4H \times 4W}$  with the same size as the original image and considering the back propagation of gradients, the  $Filtration(D) \in \mathbb{R}^{1 \times H \times W}$  in Fig. 2 can be formulated as

$$Filtration(D) = Resize(D_b) + 1 = Resize((Sign(D - t) + 1) \times 0.5) + 1, \quad (4)$$

where  $Sign$  denotes the Sign function, and  $Resize$  denotes resizing  $D_b$  to the same size as  $P2$ . Specifically,  $Resize(D_b) + 1$  can preserve the valuable information originally present in feature map  $P2$ , ensuring that it is not affected by regions in  $D_b$  where the values are 0.

**Reweighting.** Since the difference map only contains spatial information, we need to utilize broadcasting to reweight it along the channel dimension, which helps to sustain the feature diversity. Given the feature map  $P2 \in \mathbb{R}^{C \times H \times W}$ , weights  $Reweighting(P2) \in \mathbb{R}^{C \times 1 \times 1}$  can be computed as

$$Reweighting(P2) = \sigma(MLP(AvgPool(P2)) + MLP(MaxPool(P2))), \quad (5)$$

where,  $AvgPool$  denotes the average pooling along the spatial dimension, and  $MaxPool$  denotes the max pooling. The  $MLP$  includes two fully connected layers and a ReLU function. Therefore, the DGFE module can be formulated as

$$P2' = M \otimes P2 = (Reweighting(P2) \otimes Filtration(D)) \otimes P2, \quad (6)$$

where  $M \in \mathbb{R}^{C \times H \times W}$  denotes the element-wise attention matrix, and  $P2' \in \mathbb{R}^{C \times H \times W}$  denotes the feature map that has been specifically enhanced for tiny objects.

## 4 Experiment

### 4.1 Experimental Setting

**Datasets.** The experiments are conducted on three datasets. The VisDrone2019 [15] encompasses objects spanning 10 categories, and all images are captured from the perspective of drones. The AI-TOD dataset [51] comprises objects across 8 categories, and the images are collected from various datasets containing tiny objects. Our anti-UAV dataset DroneSwarms has the smallest average absolute



**Table 1:** Results on DroneSwarms. RFLA is based on Cascade R-CNN. The best results are highlighted.  $\Delta$  represents the improvement of the best towards its baseline.

Method	AP	AP <sub>0.5</sub>	AP <sub>0.75</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
TridentNet [29]	7.1	14.8	6.0	0.1	12.2	56.6
ATSS [60]	26.2	67.0	14.4	16.5	39.1	57.7
CZ Det [35]	26.5	72.6	11.7	-	-	-
DyHead [13]	26.6	68.2	14.7	16.8	39.8	57.7
RetinaNet [31]	27.3	75.8	11.3	21.0	40.5	56.2
CFINet [57]	33.7	79.3	21.1	26.0	44.3	55.6
Faster R-CNN [42]	35.0	83.9	20.9	27.3	44.3	56.5
Cascade RPN [50]	35.3	83.1	22.5	27.1	45.3	58.2
Cascade R-CNN [6]	36.4	85.0	23.5	28.8	45.7	58.3
RFLA [53]	36.9	86.3	23.4	29.5	45.3	58.0
DetectoRS [38]	37.9	87.4	24.8	30.5	46.9	59.3
Deformable-DETR [64]	32.6	81.9	17.3	25.3	42.3	55.8
DINO [59]	35.4	85.9	20.3	28.8	44.7	57.3
DINO w/ SR-TOD	35.6	86.0	20.7	28.4	44.8	58.2
RetinaNet w/ SR-TOD	28.7	77.2	12.9	22.3	42.1	56.9
Faster R-CNN w/ SR-TOD	36.3	86.4	21.7	29.5	45.1	56.0
Cascade R-CNN w/ SR-TOD	38.3	87.4	25.4	30.8	47.4	<b>59.4</b>
DetectoRS w/ SR-TOD	38.8	87.9	<b>26.3</b>	31.6	<b>47.7</b>	59.0
RFLA w/ SR-TOD	<b>39.0</b>	<b>88.9</b>	25.8	<b>31.8</b>	47.6	59.2
$\Delta$ Improvement	<b>2.1</b>	<b>2.6</b>	<b>1.5</b>	<b>2.3</b>	<b>0.8</b>	<b>1.1</b>

object size of about 7.9 pixels. All selected datasets contain a large number of tiny objects, especially tiny objects with sizes below 16 pixels.

**Implementation Details.** All core codes are built upon MMDetection [10]. Since the DroneSwarms dataset mainly consists of a large number of tiny objects, in order to ensure model convergence, we set the initial learning rate to 0.0025 for all experiments on the DroneSwarms dataset, and use the Stochastic Gradient Descent (SGD) optimizer training models for 20 epochs with 0.9 momenta, 0.0001 weight decay, 2 batch size, and 2 anchor scale. For experiments on VisDrone2019 [15] and AI-TOD [51], we strictly followed all the experiment settings of RFLA [53] on these two datasets, such as using the SGD optimizer for 12 epochs training. And the initial learning rate is set to 0.005 and decays at the 8th and 11th epochs. The ImageNet [45] pre-trained model, ResNet-50, is used as the backbone. All the other parameters of baselines are set the same as default in MMDetection. Our experiments are conducted on Huawei Atlas 800 Training Server with CANN and NVIDIA RTX 3090 GPU.

**Evaluation Metrics.** While the Average Precision (AP) evaluation metric from MS COCO [32] is prevalent among object detection algorithms, it is developed only with generic detectors in mind. This metric broadly categorizes objects with size under 32 pixels as small, employing only the AP<sub>s</sub> for their assessment.

**Table 2:** Results on VisDrone2019. The best results are highlighted.

Method	AP	AP <sub>0.5</sub>	AP <sub>0.75</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
Faster R-CNN [42]	23.9	42.2	23.8	0.1	6.5	21.1
Cascade R-CNN [6]	25.2	42.6	25.9	0.1	7.0	22.5
DetectoRS [38]	26.3	43.9	26.9	0.1	7.5	23.3
RFLA [53]	27.2	48.0	26.6	4.5	13.0	23.6
Faster R-CNN w/ SR-TOD	26.3	46.8	26.0	2.9	11.0	23.7
DetectoRS w/ SR-TOD	27.2	47.1	27.2	2.4	11.7	24.2
Cascade R-CNN w/ SR-TOD	27.3	46.9	<b>27.5</b>	2.3	11.5	<b>24.7</b>
RFLA w/ SR-TOD	<b>27.8</b>	<b>48.8</b>	<b>27.5</b>	<b>4.8</b>	<b>13.2</b>	24.5
$\Delta$ Improvement	<b>0.6</b>	<b>0.8</b>	<b>1.6</b>	<b>0.3</b>	<b>0.2</b>	<b>2.2</b>

Thus, to more effectively demonstrate our method’s performance in detecting tiny objects of diverse sizes, the evaluation metric follows AI-TOD benchmark [51]. Note that AP<sub>vt</sub>, AP<sub>t</sub>, AP<sub>s</sub> are APs for very tiny, tiny, small scales, respectively. The definitions of these scales are introduced in Sec. 1.

## 4.2 Results on DroneSwarms

As shown in Tab. 1, RFLA w/ SR-TOD achieves an AP of 39.0, surpassing the competing methods by 1.1 AP and resulting in an overall performance boost of 2.1 AP. Notably, RFLA w/ SR-TOD registers a substantial performance gain of 2.3 points in both AP<sub>vt</sub> and AP<sub>t</sub>, underscoring the pronounced efficacy of our method in detecting tiny objects. Furthermore, within the Cascade R-CNN framework, Cascade R-CNN with SR-TOD exceeds RFLA by 1.4 AP, indicating a notable improvement of 1.9 AP. Moreover, our method has also manifested clear performance gains when applied to other detectors. One-stage detectors typically perform worse than multi-stage detectors due to the lack of multi-stage regression [15, 53, 56, 63]. Even though the one-stage detector RetinaNet [31] only utilizes the feature map  $P3$  with lower resolution from FPN, our method achieves a 1.4 AP improvement. It is particularly noteworthy that, due to focusing on different challenges of tiny object detection, our framework is qualified to work with other methods like RFLA. Our framework does not conflict with methods like DetectoRS [38], which also enhance feature representation by improving FPN. Compared to the baseline methods, our method achieves consistent improvements. We have also conducted experiments with transformer-based detectors on the DroneSwarms dataset. Due to the slow convergence speed of transformer-based detectors, the schedule for Deformable-DETR [64] is set to 50 epochs as recommended. Obviously, our method shows significant superiority, even compared with the advanced transformer-based detector DINO [59]. Besides, we have also integrated our SR-TOD with DINO, which introduces improvements of 0.2 AP and 0.9 AP<sub>s</sub>. This further indicates our effectiveness in finding tiny objects when integrating with different architectures, including transformer and

**Table 3:** Results on AI-TOD. The best and second results are highlighted.

Method	AP	AP <sub>0.5</sub>	AP <sub>0.75</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
Faster R-CNN [42]	11.7	27.4	8.2	0.0	8.6	23.7
Cascade R-CNN [6]	14.0	31.2	10.7	0.1	10.3	26.2
DetectoRS [38]	14.6	31.8	11.5	0.0	11.0	27.4
RFLA [53]	21.7	50.5	15.3	8.3	21.8	26.3
HANet [21]	<b>22.1</b>	<b>53.7</b>	14.4	<b>10.9</b>	22.2	27.3
Faster R-CNN w/ SR-TOD	20.6	49.8	13.2	7.1	21.3	26.3
RFLA w/ SR-TOD	21.8	50.8	15.4	9.7	21.8	<b>27.4</b>
Cascade R-CNN w/ SR-TOD	21.9	50.6	<b>15.6</b>	9.6	<b>22.4</b>	26.7
DetectoRS w/ SR-TOD	<b>24.0</b>	<b>54.6</b>	<b>17.1</b>	<b>10.1</b>	<b>24.8</b>	<b>29.3</b>
$\Delta$ Improvement	<b>9.4</b>	<b>22.8</b>	<b>5.6</b>	<b>10.1</b>	<b>13.8</b>	<b>1.9</b>

CNN. In addition, these experiments that transformer-based detectors shows limited performance in detecting tiny objects. Recent studies also indicated that DETR lacks effective multi-scale features and that the sparse query paradigm does not adequately cover tiny objects [8, 11]. Moreover, the cost of using high-resolution feature maps beneficial for tiny objects in Transformer-based detectors is prohibitively high due to attention mechanisms [64].

### 4.3 Results on VisDrone2019 and AI-TOD

We further compare our model with the competing methods on VisDrone2019 [15] and AI-TOD [51], as shown in Tab. 2 and Tab. 3.

**VisDrone2019.** RFLA w/ SR-TOD has a improvement of 0.6 AP and achieves 27.8 AP on VisDrone2019 [15], showing a clear advantage. Furthermore, Cascade R-CNN w/ SR-TOD has a performance improvement of 2.1 AP, with the highest performance in AP<sub>s</sub>. This suggests that Cascade R-CNN w/ SR-TOD demonstrates better performance on small objects exceeding 16 pixels in size. Importantly, Cascade R-CNN w/ SR-TOD and RFLA w/ SR-TOD show significant improvements of 1.6 points and 0.9 points in AP<sub>0.75</sub>, respectively. The enhanced cascade structure with SR-TOD effectively leverages multi-stage regression to increase the localization accuracy for tiny objects.

**AI-TOD.** DetectoRS w/ SR-TOD attains an AP of 24 on the AI-TOD dataset [51], exceeding RFLA by a margin of 2.3 AP and significantly surpassing all competitors. All detectors have significant performance improvements. Significantly, Cascade R-CNN w/ SR-TOD surpasses RFLA by 1.3 points in AP<sub>vt</sub> and by 0.6 points in AP<sub>t</sub>, underscoring the substantial impact of our method on tiny object detection within the AI-TOD dataset [51]. Our method outperforms HANet [21] on most metrics (*e.g.*, AP: 22.1  $\rightarrow$  24.0, AP<sub>t</sub>: 22.2  $\rightarrow$  24.8), indicating our superiority. HANet [21] aims to obtain scale-specific feature subspaces with a predicted activation map. Although the shallow features further compensate for the detailed information of the specific scale feature space, it is still hard to

capture tiny objects as their weak signals are not enhanced in shallow layers. Compared to the activation map of HANet, our self-reconstructed difference map is more sensitive to information loss, especially for tiny objects. The consistent improvement on various datasets indicates the generality of SR-TOD. Experiments on more datasets are provided in the supplementary materials.

**Table 4:** Influence of different designs. RH means Reconstruction Head. The results are on DroneSwarms.

RH DGFE	AP	AP <sub>0.5</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
	36.4	85.0	28.8	45.7	58.3
✓	36.5	84.9	28.7	45.9	58.5
✓ ✓	<b>38.3</b>	<b>87.4</b>	<b>30.8</b>	<b>47.4</b>	<b>59.4</b>

**Table 5:** Effectiveness of threshold filtration. T means threshold. The results are on VisDrone2019.

Method	AP	AP <sub>0.5</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
w/o T	27.0	46.6	2.2	11.3	24.2
Fixed T	27.1	46.8	<b>2.3</b>	11.2	24.0
Learnable T	<b>27.3</b>	<b>46.9</b>	<b>2.3</b>	<b>11.5</b>	<b>24.7</b>

**Table 6:** Effectiveness of different feature enhancement methods. EA means element-wise attention, Concat means concatenate and EM means element-wise multiplication.

Method	AP	AP <sub>0.5</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
EM	36.2	84.9	28.8	45.7	58.4
Concat	36.6	85.3	29.0	46.0	58.7
EA	<b>38.3</b>	<b>87.4</b>	<b>30.8</b>	<b>47.4</b>	<b>59.4</b>

**Table 7:** Performance of different designs of difference map. PD means the pixel difference map and HFD means the high-frequency difference map.

Method	AP	AP <sub>0.5</sub>	AP <sub>vt</sub>	AP <sub>t</sub>	AP <sub>s</sub>
baseline	36.4	85.0	28.8	45.7	58.3
PD	38.3	87.4	30.8	47.4	<b>59.4</b>
HFD	<b>38.4</b>	<b>87.6</b>	<b>31.1</b>	<b>47.6</b>	59.2

#### 4.4 Ablation Study and Discussion

**Effectiveness of individual component.** As shown in Sec.3.2 and Sec.3.3, the self-reconstruction mechanism consists of two key components: the Reconstruction Head and the Difference Map Guided Feature Enhancement module. Note that the validation of DGFE relies on the Reconstruction Head, as it requires the construction of a difference map as input. When the reconstruction head is applied alone, we obtain the reconstructed image without constructing a difference map. We add Reconstruction Head and DGFE into the Cascade R-CNN one-by-one. The results are listed in Tab. 4. The DGFE increased 1.8 points in AP over the use of the reconstruction head alone, accompanied by a substantial gain of 2.1 points in AP<sub>vt</sub>. These results conclusively prove that DGFE is both an effective and indispensable module for leveraging the difference map. It should be highlighted that the sole application of the reconstruction head yielded a modest improvement of 0.1 AP, along with an increment of 0.2 points in both AP<sub>t</sub> and AP<sub>s</sub>. This enhancement can be attributed to the fact that integrating image reconstruction constraints within the object detection model not only preserves detection performance but also, to a certain degree, improves the backbone network’s pixel comprehension. This approach could potentially offer fresh inspiration for the field of object detection.

**Effectiveness of different feature enhancement methods.**


In this section, we explore various feature enhancement techniques using the Cascade R-CNN framework to incorporate prior knowledge from the difference map. Beyond employing element-wise attention, we experiment with fusion method that concatenate the binary difference map with the feature map and explore approach that execute direct element-wise multiplication between the two. As shown in Tab. 6, it is evident that element-wise attention markedly surpasses the performance of alternative methods. The fusion method using the concatenate operation has a weak effect, as direct fusion brings more noise and diminishes the spatial positional information of the difference map. The approach that employs element-wise multiplication is deemed too simplistic and basic, as numerous inactive areas in the difference map can cause the obliteration of original features in the feature map, ultimately impairing detection capabilities. Consequently, we opt for element-wise attention in DGFE, which provides a more sophisticated solution.

**Effectiveness of threshold filtration.**


To investigate the effect of thresholds, we also explore methods employing fixed thresholds and those without threshold-based filtering. We conduct experiments on the VisDrone2019 dataset [15], which features intricate backgrounds and diverse categories. The results are shown in Tab. 5. Undoubtedly, the approaches that involve setting a threshold generally exceed the performance of the method without a threshold, thus affirming the utility of thresholding. In addition, the learnable threshold method performs similarly to the fixed threshold method

in  $AP_{vt}$  but shows a 0.3 points improvement in  $AP_t$  and a 0.7 points improvement in  $AP_s$ . This indicates that tiny objects with a size above 8 pixels are more sensitive to thresholds. Therefore, the method with a learnable threshold achieves better performance by finely adjusting the threshold.

**Exploration of the high-frequency difference map.** Recently, some image reconstruction methods [24, 62] model images in the frequency domain. The high-frequency information of an image includes details such as edges and textures, which are precisely the parts that are difficult to recover in image



**Fig. 3:** Visualizations of the pixel difference maps and high-frequency difference maps.



**Fig. 4:** Visualizations on DroneSwarms. The first row is the local regions of input images. The second and third rows are the corresponding feature maps from Cascade R-CNN *w/o* and *w/* SR-TOD.

self-reconstruction. In particular, noise is also considered part of high-frequency information, and very tiny objects can be approximated as noise-like points with a size of only a few pixels. Inspired by this, we further utilize the fast Fourier transform algorithm to separate the high-frequency components of the original input image and the reconstructed image, and construct high-frequency difference maps, as shown in Fig. 3. Although the high-frequency difference map and pixel difference map are quite similar, the high-frequency difference map has more refined contours on some drone objects. Additionally, the high-frequency difference map significantly reduces noise signals. However, this also results in certain smaller drone objects appearing even more faint in the high-frequency difference map. Then we replace the pixel difference map (PD) with the high-frequency difference map (HFD) and conduct experiments on DroneSwarms using Cascade R-CNN w/ SR-TOD. The results are shown in Tab. 7. The HFD has a 0.1 AP improvement compared to PD. Moreover, there is an improvement of 0.3 points in  $AP_{vt}$  and 0.2 points in  $AP_t$ . This demonstrates that the high-frequency difference map can provide more accurate priors for tiny objects compared to the pixel difference map. This indicates that the information severely lost in tiny object feature extraction mainly lies in the high-frequency components. Note that compared to the baseline, the detection performance of using pixel difference map and high-frequency difference map is very similar. Considering computational efficiency, we use the pixel difference map as the basic setting for SR-TOD. However, the high-frequency difference map shows the potential for further exploration of our method. More results are provided in the supplementary materials.

#### 4.5 Visualization Analysis

To demonstrate the effect of DGFE more clearly and intuitively, we perform visualizations on DroneSwarms. The results are shown in Fig. 4. From the results in the second row, it can be seen that before feature enhancement, the detection model lacks sufficient attention to tiny drones. The signals of some tiny objects are very weak and almost wiped out, resulting in invisibility. The results in the third row show that the SR-TOD significantly enhances the feature representation of tiny objects, which makes the tiny objects visible and clear to the detectors.

## 5 Conclusion

In this paper, we analyze the challenge of information loss in tiny object detection and the limitations faced by generative methods attempting to alleviate this issue. To this end, we introduce an image self-reconstruction mechanism, constructing difference maps as the prior information of tiny objects, making the feature more visible to detectors. Then, we further design a Difference Map Guided Feature Enhancement (DGFE) module to improve the feature representation of tiny objects, providing more clear representations. Experiments on our proposed DroneSwarms and two other datasets show the superiority and robustness of the SR-TOD. In the future, we will explore more effective ways to construct more accurate difference maps for tiny objects.

## Acknowledgements

This work was sponsored in part by the National Key R&D Program of China 2022ZD0116500, in part by the National Natural Science Foundation of China (62106171, 62222608, U23B2049, 61925602), in part by the Haihe Lab of ITAI under Grant 22HHXCJC00002, in part by Tianjin Natural Science Funds for Distinguished Young Scholar under Grant 23JCJQJC00270, and in part by the Key Laboratory of Big Data Intelligent Computing, Chongqing University of Posts and Telecommunications under Grant BDIC-2023-A-008. This work was also sponsored by CAAI-CANN Open Fund, developed on OpenI Community.

## References

1. Bai, Y., Zhang, Y., Ding, M., Ghanem, B.: Finding tiny faces in the wild with generative adversarial network. In: CVPR. pp. 21–30 (2018)
2. Bai, Y., Zhang, Y., Ding, M., Ghanem, B.: Sod-mtgan: Small object detection via multi-task generative adversarial network. In: ECCV. pp. 206–221 (2018)
3. Bashir, S.M.A., Wang, Y.: Small object detection in remote sensing images with residual feature aggregation-based super-resolution and object detector network. *Remote Sensing* **13**(9), 1854 (2021)
4. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934 (2020)
5. Bosquet, B., Cores, D., Seidenari, L., Brea, V.M., Mucientes, M., Del Bimbo, A.: A full data augmentation pipeline for small object detection based on generative adversarial networks. *PR* **133**, 108998 (2023)
6. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: CVPR. pp. 6154–6162 (2018)
7. Cao, B., Bi, Z., Hu, Q., Zhang, H., Wang, N., Gao, X., Shen, D.: Autoencoder-driven multimodal collaborative learning for medical image synthesis. *IJCV* **131**, 1–20 (2023)
8. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. pp. 213–229 (2020)
9. Chen, C., Liu, M.Y., Tuzel, O., Xiao, J.: R-cnn for small object detection. In: ACCV. pp. 214–230 (2016)
10. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., et al.: Mmdetection: Open mmlab detection toolbox and benchmark. arxiv 2019. arXiv preprint arXiv:1906.07155 (2019)
11. Cheng, G., Yuan, X., Yao, X., Yan, K., Zeng, Q., Xie, X., Han, J.: Towards large-scale small object detection: Survey and benchmarks. *IEEE TPAMI* **45**(11), 13467–13488 (2023)
12. Coluccia, A., Fascista, A., Schumann, A., Sommer, L., Dimou, A., Zarpalas, D., Akyon, F.C., Eryuksel, O., Ozfuttu, K.A., Altinuc, S.O., et al.: Drone-vs-bird detection challenge at ieee avss2021. In: 2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS). pp. 1–8 (2021)
13. Dai, X., Chen, Y., Xiao, B., Chen, D., Liu, M., Yuan, L., Zhang, L.: Dynamic head: Unifying object detection heads with attentions. In: CVPR. pp. 7373–7382 (2021)
14. Deng, C., Wang, M., Liu, L., Liu, Y., Jiang, Y.: Extended feature pyramid network for small object detection. *IEEE TMM* **24**, 1968–1979 (2021)

15. Du, D., Zhu, P., Wen, L., Bian, X., Lin, H., Hu, Q., Peng, T., Zheng, J., Wang, X., Zhang, Y., et al.: Visdrone-det2019: The vision meets drone object detection in image challenge results. In: Proceedings of the IEEE/CVF international conference on computer vision workshops. pp. 213–226 (2019)
16. Gao, S.H., Cheng, M.M., Zhao, K., Zhang, X.Y., Yang, M.H., Torr, P.: Res2net: A new multi-scale backbone architecture. *IEEE TPAMI* **43**(2), 652–662 (2019)
17. Ghiasi, G., Lin, T.Y., Le, Q.V.: Nas-fpn: Learning scalable feature pyramid architecture for object detection. In: CVPR. pp. 7036–7045 (2019)
18. Girshick, R.: Fast r-cnn. In: ICCV. pp. 1440–1448 (2015)
19. Gong, Y., Yu, X., Ding, Y., Peng, X., Zhao, J., Han, Z.: Effective fusion factor in fpn for tiny object detection. In: WACV. pp. 1160–1168 (2021)
20. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NeurIPS. pp. 2672–2680 (2014)
21. Guo, G., Chen, P., Yu, X., Han, Z., Ye, Q., Gao, S.: Save the tiny, save the all: hierarchical activation network for tiny object detection. *IEEE TCSVT* (2023)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
23. Hu, X., Xu, X., Xiao, Y., Chen, H., He, S., Qin, J., Heng, P.A.: Sinet: A scale-insensitive convolutional neural network for fast vehicle detection. *IEEE transactions on intelligent transportation systems* **20**(3), 1010–1019 (2018)
24. Jiang, L., Dai, B., Wu, W., Loy, C.C.: Focal frequency loss for image reconstruction and synthesis. In: ICCV. pp. 13919–13929 (2021)
25. Kisantant, M., Wojna, Z., Murawski, J., Naruniec, J., Cho, K.: Augmentation for small object detection. arXiv preprint arXiv:1902.07296 (2019)
26. Kong, T., Sun, F., Liu, H., Jiang, Y., Li, L., Shi, J.: Foveabox: Beyond anchor-based object detection. *IEEE TIP* **29**, 7389–7398 (2020)
27. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: ECCV. pp. 734–750 (2018)
28. Li, J., Liang, X., Wei, Y., Xu, T., Feng, J., Yan, S.: Perceptual generative adversarial networks for small object detection. In: CVPR. pp. 1222–1230 (2017)
29. Li, Y., Chen, Y., Wang, N., Zhang, Z.: Scale-aware trident networks for object detection. In: ICCV. pp. 6054–6063 (2019)
30. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR. pp. 2117–2125 (2017)
31. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017)
32. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
33. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: CVPR. pp. 8759–8768 (2018)
34. Lu, X., Li, B., Yue, Y., Li, Q., Yan, J.: Grid r-cnn. In: CVPR. pp. 7363–7372 (2019)
35. Meethal, A., Granger, E., Pedersoli, M.: Cascaded zoom-in detector for high resolution aerial images. In: CVPRW. pp. 2045–2054 (2023)
36. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: ICML. pp. 807–814 (2010)
37. Noh, J., Bae, W., Lee, W., Seo, J., Kim, G.: Finding tiny faces in the wild with generative adversarial network. In: ICCV. pp. 9725–9734 (2019)
38. Qiao, S., Chen, L.C., Yuille, A.: Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In: CVPR. pp. 10213–10224 (2021)



39. Rabbi, J., Ray, N., Schubert, M., Chowdhury, S., Chao, D.: Small-object detection in remote sensing images with end-to-end edge-enhanced gan and object detector network. *Remote Sensing* **12**(9), 1432 (2020)
40. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: *CVPR*. pp. 7263–7271 (2017)
41. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018)
42. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS* **28**, 91–99 (2015)
43. Rodriguez-Ramos, A., Rodriguez-Vazquez, J., Sampedro, C., Campoy, P.: Adaptive inattentive framework for video object detection with reward-conditional training. *IEEE Access* **8**, 124451–124466 (2020)
44. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. pp. 234–241 (2015)
45. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *IJCV* **115**(3), 211–252 (2015)
46. Singh, B., Davis, L.S.: An analysis of scale invariance in object detection snip. In: *CVPR*. pp. 3578–3587 (2018)
47. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., et al.: Sparse r-cnn: End-to-end object detection with learnable proposals. In: *CVPR*. pp. 14454–14463 (2021)
48. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: *CVPR*. pp. 10781–10790 (2020)
49. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: *ICCV*. pp. 9627–9636 (2019)
50. Vu, T., Jang, H., Pham, T.X., Yoo, C.: Cascade rpn: Delving into high-quality region proposal network with adaptive convolution. *NeurIPS* **32**, 1432–1442 (2019)
51. Wang, J., Yang, W., Guo, H., Zhang, R., Xia, G.S.: Tiny object detection in aerial images. In: *ICPR*. pp. 3791–3798 (2021)
52. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Proceedings of the IEEE conference on computer vision and pattern recognition. In: *CVPR*. pp. 1492–1500 (2017)
53. Xu, C., Wang, J., Yang, W., Yu, H., Yu, L., Xia, G.S.: Rfla: Gaussian receptive field based label assignment for tiny object detection. In: *ECCV*. pp. 526–543 (2022)
54. Yang, C., Huang, Z., Wang, N.: Querydet: Cascaded sparse query for accelerating high-resolution small object detection. In: *CVPR*. pp. 13668–13677 (2022)
55. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: *ICCV*. pp. 9657–9666 (2019)
56. Yu, X., Gong, Y., Jiang, N., Ye, Q., Han, Z.: Scale match for tiny person detection. In: *WACV*. pp. 1257–1265 (2020)
57. Yuan, X., Cheng, G., Yan, K., Zeng, Q., Han, J.: Small object detection via coarse-to-fine proposal generation and imitation learning. In: *ICCV*. pp. 6317–6327 (2023)
58. Zeiler, M.D., Taylor, G.W., Fergus, R.: Adaptive deconvolutional networks for mid and high level feature learning. In: *ICCV*. pp. 2018–2025 (2011)
59. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In: *ICLR* (2022)

60. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: CVPR. pp. 9759–9768 (2020)
61. Zhao, J., Zhang, J., Li, D., Wang, D.: Vision-based anti-uav detection and tracking. *IEEE transactions on intelligent transportation systems* **23**(12), 25323–25334 (2022)
62. Zhou, Y., Deng, W., Tong, T., Gao, Q.: Guided frequency separation network for real-world super-resolution. In: CVPRW. pp. 428–429 (2020)
63. Zhu, P., Wen, L., Du, D., Bian, X., Ling, H., Hu, Q., Nie, Q., Cheng, H., Liu, C., Liu, X., et al.: Visdrone-det2018: The vision meets drone object detection in image challenge results. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 437–468 (2018)
64. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. ICLR (2021)