

Uncertainty-Aware Gradient Stabilization for Small Object Detection

Huixin Sun¹, Yanjing Li¹, Linlin Yang⁴, Xianbin Cao^{1†}, Baochang Zhang^{2,3†}

¹School of Electronic Information Engineering, Beihang University

²School of Artificial Intelligence, Beihang University

³Zhongguancun Laboratory, Beijing, China

⁴State Key Laboratory of Media Convergence and Communication, CUC

[†]Corresponding authors

Abstract

Despite advances in generic object detection, there remains a performance gap in detecting small objects compared to normal-scale objects. We reveal that conventional object localization methods suffer from gradient instability in small objects due to sharper loss curvature, leading to a convergence challenge. To address the issue, we propose Uncertainty-Aware Gradient Stabilization (UGS), a framework that reformulates object localization as a classification task to stabilize gradients. UGS quantizes continuous labels into interval non-uniform discrete representations. Under a classification-based objective, the localization branch generates bounded and confidence-driven gradients, mitigating instability. Furthermore, UGS integrates an uncertainty minimization (UM) loss that reduces prediction variance and an uncertainty-guided refinement (UR) module that identifies and refines high-uncertainty regions via perturbations. Evaluated on four benchmarks, UGS consistently improves anchor-based, anchor-free, and leading small object detectors. Notably, UGS enhances DINO-5scale by 2.6 AP on VisDrone, surpassing prior state-of-the-art performance.

1. Introduction

Recent advances in deep neural networks (DNNs) [16, 59] have substantially improved the object detection field [11, 33]. Nevertheless, detecting small objects remains a persistent challenge [55]. For instance, Cascade R-CNN [4], one of the most representative two-stage object detectors, achieves 45.5% and 55.2% AP on medium and large-sized objects but only 23.7% AP on small objects in the COCO test-dev set [33]. The significant performance gap between small and normal-sized objects limits the effectiveness of object detectors in real-world applications, such as driving assistance, traffic management, and anomaly detection.

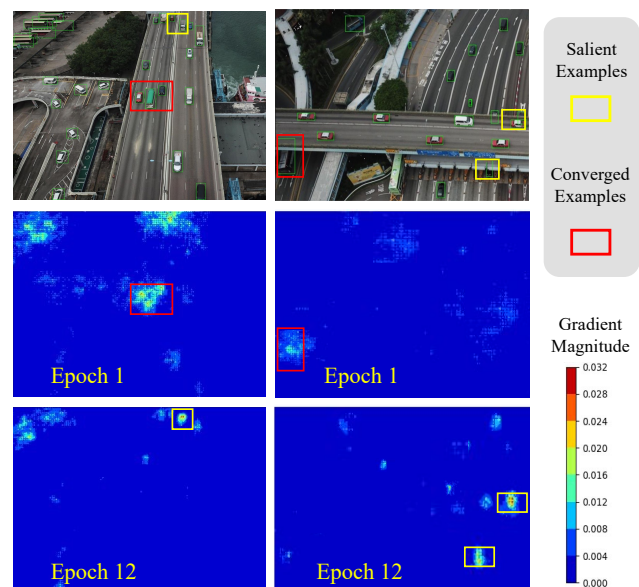


Figure 1. Gradient magnitude maps based on the ℓ_2 -Norm of the gradient in the regression features [15]. The maps are generated using a $1 \times$ Faster R-CNN [44] detector with ResNet-18 backbone. By epoch 12, most objects exhibit converged gradients (red areas) while multiple small object examples (yellow areas) exhibit salient gradients, indicating a convergence challenge.

Due to the limited pixel inputs (under 32×32 [33]), one of the primary challenges in small object detection is the extraction of discriminative foreground features [32]. The challenge is compounded in cluttered environments [5], where occlusions, background noise, and low signal-to-noise ratio conditions induce feature ambiguity. Consequently, generic detectors can develop a feature bias towards distinguishing the foreground from background regions that resemble it [54]. Recent efforts address these problems by increasing the resolution of feature maps [1, 34], integrating contextual information [10, 56], and employing auxiliary self-reconstruction branches [5, 54] to en-

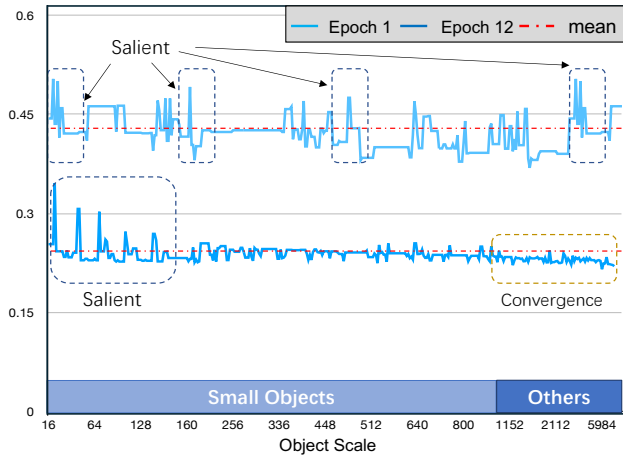


Figure 2. Statistics of the average object gradient. Object gradient is derived by summing the gradient magnitudes within each bounding box based on the gradient maps. The average gradient of each object scale is calculated as the mean across objects in the corresponding size category. The figure is generated using 100 randomly selected samples in the VisDrone test [62]. The bottom curve (epoch 12) reveals that gradients for medium and large objects have significantly reduced after training. In contrast, gradients for small objects remain pronounced, indicating the convergence challenge.

hance object features. In this work, we analyze the SOD challenge from a novel perspective of gradient stability.

Motivated by prior research on gradient anomalies [38, 49], we examine the Hessian matrix of the object localization losses and demonstrate that conventional methods [44, 57] suffer from gradient instability in small objects due to the sharper loss curvature, which leads to a convergence challenge. We illustrate this phenomenon using the gradient magnitudes generated from the localization features in Fig. 1. As shown, the detector converges well for most objects (red areas) but exhibits salient gradients in multiple small object examples at epoch 12 (yellow areas). We further quantified the average object gradient based on the gradient magnitudes. As shown in Fig. 2, the gradients of small objects remain pronounced at epoch 12, while those of normal-scale objects converge effectively.

In light of the analysis, we propose Uncertainty-Aware Gradient Stabilization (UGS), a framework that stabilizes gradients by reformulating regression as a classification task with uncertainty-aware mechanisms. The framework of UGS is illustrated in Fig. 4. Building on prior classification-based localization methods [29, 42], UGS quantizes continuous labels into interval non-uniform discrete representations to address the label imbalance in small objects. Under a classification-based objective, the localization branch generates bounded and confidence-driven gradients, mitigating instability. To further stabilize training, UGS introduces a dual uncertainty-aware mechanism consisting of an uncer-

tainty minimization (UM) loss and an uncertainty-guided refinement (UR) module. The UM loss explicitly models and minimizes prediction uncertainty via entropy, reducing the prediction variance. Additionally, the UR module leverages adversarial perturbations to identify and refine high-uncertainty regions, thereby improving feature robustness.

To summarize, our main contributions are three-fold:

1. A gradient analysis is conducted to investigate the Small Object Detection (SOD) challenge, which shows that conventional object localization suffers from unstable gradients on small objects and leads to a convergence challenge. We propose a novel Uncertainty-Aware Gradient Stabilization (UGS) method to improve gradient stability and foster better convergence.
2. UGS integrates three key components: a classification-based localization objective for generating bounded and confidence-driven gradients, an uncertainty minimization loss that explicitly models and minimizes prediction uncertainty in small objects, and an uncertainty-guided refinement module that leverages adversarial perturbations to identify and refine high-uncertainty regions.
3. UGS exhibits consistent increases on baseline detectors and state-of-the-art small-object detectors, demonstrating effectiveness in general object detection and high-resolution detection.

2. Related Work

2.1. Small Object Detection

Small object detection remains a major challenge in computer vision due to the inherent limitations of pixel-level inputs [54]. Key difficulties include weak feature representation [32, 34], information loss during down-sampling [24], and fewer assigned positives due to higher sensitivity in IoU calculations [52, 53, 58]. To address these issues, various approaches have been proposed, which can be categorized into four main strategies: feature enhancement, where methods like [6, 8, 34] improve feature representation through advanced architectures and fusion techniques; data augmentation and oversampling, where techniques like [36, 65] increase the diversity and quantity of training data for small objects; scale-aware training, where approaches like [31, 47] adapt models to handle objects at multiple scales; super-resolution-based method, where frameworks like [2, 27] enhance small object features by reconstructing high-resolution representations. Recent state-of-the-art methods focus on label assignment and proposal refinement to ensure both the quantity and quality of proposals for small objects [53, 58]. Additionally, auxiliary self-reconstruction [5, 54] and spectral enhancement [48] methods have been proposed to enhance the weak representations of objects. Our method is orthogonal to existing small object detection methods from a new perspective of solving the convergence challenge in localization.

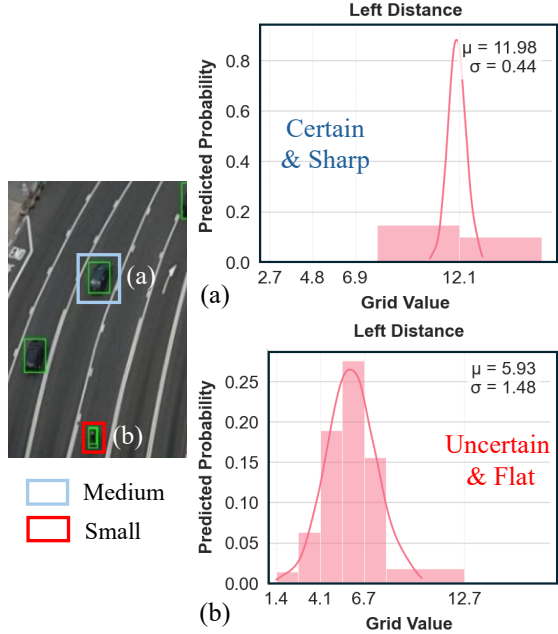


Figure 3. A qualitative example of localization uncertainty using FCOS [50] trained with \mathcal{L}_{CE} (Eq. (10)). (a) The medium object instance shows certain and sharp prediction distribution ($\sigma = 0.44$), while (b) the small object instance exhibits flatter distribution ($\sigma = 1.48$), reflecting higher prediction uncertainty.

2.2. Uncertainty Estimation in Object Detection

Recent advances in localization uncertainty estimation aim to quantify prediction uncertainty by modeling bounding box distributions. Early work by [22] directly predicts the Intersection-over-Union (IoU) between ground-truth and predicted boxes as a confidence measure. Subsequent approaches like [18] formulate uncertainty via KL divergence loss between Dirac delta ground-truths and Gaussian predictions, enabling variance-based refinement. The state-of-the-art GFL V1 [29] further generalizes this paradigm by modeling arbitrary distributions through distribution focal loss. However, the methods are based on anchor-free frameworks (e.g., FCOS [50]), limiting their compatibility with anchor-based methods (e.g., RetinaNet [60]), modern anchor-free methods (e.g., YOLO series [3, 43, 64]), two-stage detectors (e.g., Faster R-CNN [44]), and DETR-based detectors (e.g., DINO [59]). Moreover, while uncertainty estimation addresses challenges such as occlusion and noise [26], it remains unexplored for small object detection.

3. Method

In this section, we first analyze the gradient instability in norm-based and IoU-based localization methods, then introduce the Uncertainty-Aware Gradient Stabilization (UGS) design as a solution.

3.1. Gradient Instability in Small Objects

Norm-based Localization. We analyze gradient instability using the \mathcal{L}_2 loss in norm-based localization. Following previous detectors [13, 14, 44], we denote the localization targets and predictions as:

$$\begin{aligned} \{T_x, T_y, T_w, T_h\} &= \left\{ \frac{x - x_a}{w_a}, \frac{y - y_a}{h_a}, \log \frac{w}{w_a}, \log \frac{h}{h_a} \right\}, \\ \{\hat{T}_x, \hat{T}_y, \hat{T}_w, \hat{T}_h\} &= \left\{ \frac{\hat{x} - x_a}{w_a}, \frac{\hat{y} - y_a}{h_a}, \log \frac{\hat{w}}{w_a}, \log \frac{\hat{h}}{h_a} \right\}, \end{aligned} \quad (1)$$

where (x_a, y_a, w_a, h_a) denote the anchor coordinates, (x, y, w, h) the ground-truth coordinates, and $(\hat{x}, \hat{y}, \hat{w}, \hat{h})$ the predicted coordinates, respectively. The \mathcal{L}_2 loss can be formulated as:

$$\mathcal{L}_2(T_x, \hat{T}_x) = \|T_x - \hat{T}_x\|_2^2, \quad (2)$$

which applies to y , w , and h . For center coordinates (x, y) , the Hessian derives as:

$$\mathbf{H}_x = \frac{\partial^2 \mathcal{L}_2}{\partial \hat{x}^2} = \frac{2}{w_a^2}, \quad \mathbf{H}_y = \frac{\partial^2 \mathcal{L}_2}{\partial \hat{y}^2} = \frac{2}{h_a^2}, \quad (3)$$

where the Lipschitz constants K_x and K_y scale inversely with the square of anchor size. For small objects assigned with smaller anchors, the Lipschitz constants increase, leading to steeper loss curvature. For size regression, the Hessian is:

$$\mathbf{H}_w = \frac{\partial^2 \mathcal{L}_2}{\partial \hat{w}^2} = 2 \cdot \left(\frac{1}{\hat{w}} \right)^2, \quad (4)$$

where the Lipschitz constant K_w depends inversely on the predicted width \hat{w} , resulting in steeper loss curvature for small objects. Steep loss curvature can result in unstable updates [39] that are prone to oscillate or diverge near minima, causing a convergence challenge. As shown in Fig. 2, the gradients of small objects remain pronounced at epoch 12, while those of normal-scale objects converge effectively.

IoU-based Localization. We analyze the issue using the IoU loss [57], which is applicable to other IoU-based losses [45, 61]. Following [52], we consider the case where the ground truth (x, y, w, h) and the prediction $(\hat{x}, \hat{y}, \hat{w}, \hat{h})$ are aligned square boxes with a center displacement $d = |x - \hat{x}|$. The IoU loss is expressed as:

$$\mathcal{L}_{IoU} = -\ln \left(\frac{I}{U} \right), \quad (5)$$

where I and U represent the intersection and union. For overlapping boxes ($|x - \hat{x}| < w$), the gradient with respect to \hat{x} is derived as:

$$\frac{\partial \mathcal{L}_{IoU}}{\partial \hat{x}} = \frac{2}{w^2 - d^2} \cdot \text{sign}(x - \hat{x}), \quad (6)$$

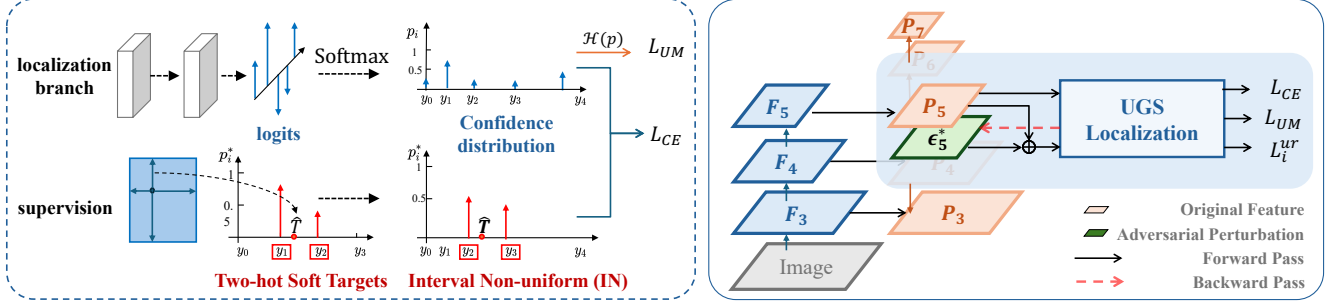


Figure 4. Overview of the proposed UGS localization method with FCOS [50] framework, consisting a classification-based localization objective (\mathcal{L}_{CE}), an uncertainty minimization loss (\mathcal{L}_{UM}), and an uncertainty-guided refinement module (\mathcal{L}^{ur}).

the Hessian can be derived as:

$$\frac{\partial^2 \mathcal{L}_{IoU}}{\partial \hat{x}^2} = \frac{4w}{(w^2 - d^2)^2}, \quad (7)$$

where the gradient scales inversely with w and the Hessian with w^3 . Consequently, smaller objects exhibit larger gradients and sharper loss curvature. This can lead to more significant fluctuations in the loss landscape and oscillations near the ground truth, resulting in instability during optimization. Derivations of Eq. (3), Eq. (4), Eq. (6), and Eq. (7) are provided in the supplementary material.

3.2. Uncertainty-Aware Gradient Stabilization

To address the gradient instability problem in small object localization, we propose the Uncertainty-Aware Gradient Stabilization (UGS) method, depicted in Fig. 4. UGS integrates three components: a classification-based localization objective that generates bounded and confidence-driven gradients, an uncertainty minimization loss to further stabilize optimization, and an uncertainty-guided refinement module that leverages adversarial perturbations to identify and refine regions of high uncertainty.

3.2.1. Classification-based Localization

Building upon GFL V1 [29], we propose a classification-based localization objective that quantizes continuous regression targets into interval non-uniform discrete grid representations. Specifically, the continuous regression range $[-\alpha, \alpha]$ for each target value T is partitioned into $n + 1$ uniformly spaced intervals, yielding the discrete grid set $Y = \{y_0, y_1, \dots, y_n\}$. The ground truth T is mapped to adjacent grids i_l and i_r using two-hot soft targets [29]:

$$\mathbf{p}_i^* = \begin{cases} |y_i - T| \cdot \frac{n+1}{2\alpha}, & \text{if } i = i_l \text{ or } i_r, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

For small objects, the regression targets are distributed in a limited range. When transformed into soft targets, the head range of the target distribution occupies few grids,

leading to class imbalance and ineffective training. Furthermore, in offset-based label systems [44, 60], target values decrease with training and are constantly assigned to the same grid interval after some iterations, further hindering optimization. To address the issues, we design interval non-uniform (IN) labels through exponential grid spacing:

$$\mathbf{y}_i^{\text{IN}} = \text{sign}(y_i) \cdot \frac{\alpha}{e^{\alpha\beta} - 1} \left(e^{\beta|y_i|} - 1 \right), \quad (9)$$

where $y_i = -\alpha + \frac{2\alpha i}{n}$ are uniformly spaced grids within $[-\alpha, \alpha]$, and β is a factor controlling the density of grid points. A larger β results in denser grids near zero, balancing the localization targets for small objects and facilitating better optimization.

For optimization, the framework minimizes the cross-entropy loss between the predicted distribution \mathbf{p} and the ground truth \mathbf{p}^* :

$$\mathcal{L}_{CE} = -\mathbf{p}_{i_l}^* \log \mathbf{p}_{i_l} - \mathbf{p}_{i_r}^* \log \mathbf{p}_{i_r}. \quad (10)$$

where the continuous coordinate prediction can be restored as $\hat{T} = \sum_{i=0}^n \mathbf{p}_i * \mathbf{y}_i$. We demonstrate that the gradient of the classification-based localization objective is both bounded and confidence-driven:

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}}{\partial \mathbf{p}_i} &= \mathbf{p}_i - \mathbf{p}_i^*, \\ &= \begin{cases} \mathbf{p}_i - |y_i - T| \cdot \frac{n+1}{2\alpha}, & i = i_l, i_r, \\ \mathbf{p}_i, & \text{otherwise,} \end{cases} \end{aligned} \quad (11)$$

which shows that the gradient magnitude of the cross-entropy loss with respect to logits $|\mathbf{p}_i - \mathbf{p}_i^*|$ is bounded within $[0, 1]$ across object scales, mitigating instability for small objects. Additionally, the gradient magnitude is proportional to the confidence difference $|\mathbf{p}_i - \mathbf{p}_i^*|$ and enables confidence-driven learning, where large differences lead to more aggressive updates and small differences result in fine-grained refinements, facilitating precise localization.

3.2.2. Uncertainty Minimization

To further stabilize training, we explicitly model and minimize the prediction uncertainty in the confidence distribution \mathbf{p} :

$$\mathcal{L}_{\text{UM}} = \mathcal{H}(\mathbf{p}) = - \sum_{i=0}^n \mathbf{p}_i \log \mathbf{p}_i, \quad (12)$$

where $\mathcal{H}(\mathbf{p})$ represents the entropy of the distribution \mathbf{p} [46]. By minimizing the entropy, the loss can suppress uncertain predictions that amplify gradient variance and reduce the coordinate prediction variance, stabilizing the optimization. Moreover, it penalizes high-confidence errors less aggressively than the cross-entropy objective through:

$$\frac{\partial \mathcal{L}_{\text{UM}}}{\partial \mathbf{l}_i} = -(\log \mathbf{p}_i + 1), \quad (13)$$

which benefits small objects as they tend to generate noisy and uncertain predictions due to insufficient feature representations [30], exhibited in Fig. 3 (b). The UM loss also maintains stable updates across object scales via bounded gradients.

3.2.3. Uncertainty-guided Refinement

Inspired by prior works in beneficial noise learning [21], we propose an Uncertainty-guided Refinement (UR) module that leverages perturbations derived from the uncertainty minimization loss (\mathcal{L}_{UM}) to identify and refine regions of high uncertainty. This approach also enhances feature robustness and stabilizes the optimization process for small objects.

The UR formulation for FPN [34] layer \mathbf{F}_i establishes a min-max objective:

$$\min_{\mathbf{F}_i, \theta_i} \left(\max_{\|\epsilon_i\| \leq \rho} \mathcal{L}_{\text{UM}}(\mathbf{F}_i + \epsilon_i) + \gamma \|\mathbf{F}_i\|_2^2 \right), \quad (14)$$

where \mathcal{L}_{UM} quantifies prediction uncertainty via entropy minimization. The inner optimization introduces adversarial perturbation ϵ_i into the feature space of \mathbf{F}_i , with ρ controlling the perturbation magnitude and γ regulating the regularization strength. Here, θ_i represents the model parameters at the i -th layer. Following [12], we derive the closed-form perturbation under L_2 -norm:

$$\epsilon_i^* \approx \rho \cdot \frac{\nabla_{\mathbf{F}_i} \mathcal{L}_{\text{UM}}(\mathbf{F}_i)}{\|\nabla_{\mathbf{F}_i} \mathcal{L}_{\text{UM}}(\mathbf{F}_i)\|_2}, \quad (15)$$

where the perturbation ϵ_i^* targets regions where \mathcal{L}_{UM} exhibits high sensitivity to activation changes, indicating areas of high uncertainty. The introduced perturbation can amplify the refinement [23] of these uncertain regions while preserving stable updates in regions with high confidence. Fig. 5 shows that the method can learn about occluded objects and noise that resembles objects. Additionally, the adversarial objective encourages the learning of robust feature

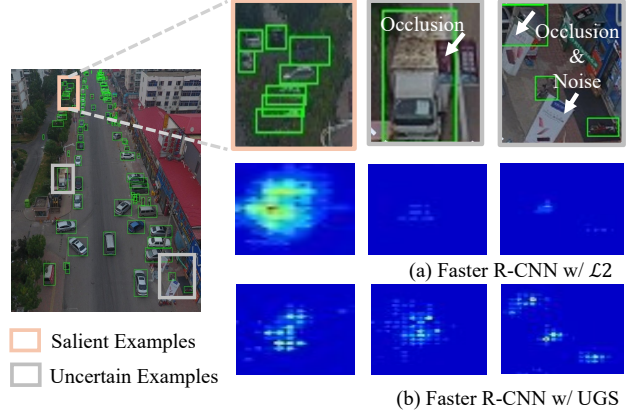


Figure 5. Comparing the gradient magnitudes maps [15] using a $1 \times$ Faster R-CNN detector with ResNet-18 backbone. In (a), normal-scale objects converge by epoch 12 while the gradients for some small objects remain salient, indicating a convergence challenge. In (b), the gradients in the salient object are less pronounced.

representations [19], further stabilizing the training process for small objects.

We formulate the overall training objective as:

$$\mathcal{L}_{\text{localization}} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{UM}} + \underbrace{\gamma \cdot \sum_{i=1}^N \mathcal{L}_i^{\text{ur}}(\mathbf{F}_i + \epsilon_i^*)}_{\text{Uncertainty-Guided Refinement Loss}}, \quad (16)$$

where λ, γ are balancing hyperparameters, and $\mathcal{L}_i^{\text{ur}}$ denotes the layer-wise uncertainty-guided refinement loss. Described in Sec. 6, our approach reduces gradient variance by $2.5 \times$ compared to Smooth- \mathcal{L}_1 loss for small objects.

4. Experiments

4.1. Datasets and Implementation Details

4.1.1. Datasets

We evaluate the proposed method on **VisDrone** [62], **SODA-A**[9], **COCO 2017** [33], and the **PASCAL VOC** dataset [11]. The main experiments were conducted using the VisDrone [62] dataset, which contains a high proportion of small object instances. The dataset consists of 10,209 drone-shot images divided into a training set (6,471 images), a validation set (548 images), and a test set (3,190 images). Meanwhile, we conduct experiments on SODA-A [9] dataset, which comprises 2513 high-resolution images of aerial scenes, which has 872069 instances annotated with *oriented* rectangle box annotations over 9 classes. We also perform experiments on two general object detection benchmarks: the COCO 2017 [33] dataset and the PASCAL VOC [11] dataset.

Table 1. Detection performance on the VisDrone val set [62]. +(\cdot) indicates the improvement over the baseline, **bold** denotes the best results. Baselines include general and small-object-oriented detectors.

Method	Year	AP	AP ₅₀	AP _s	AP _m	AP _l
QueryDet [55]	2022	28.3	48.1	-	-	-
DetectoRS [41]	2021	29.4	49.3	-	-	-
CZDet [37]	2023	33.2	58.3	-	-	-
DQ-DETR [20]	2024	37.0	60.9	-	-	-
FCOS [50]	2019	19.9	37.7	11.4	30.7	35.2
w/ UGS	-	22.4	39.7	13.0	33.2	37.8
Faster R-CNN [44]	2015	21.3	36.4	12.8	32.9	38.9
w/ UGS	-	24.2	41.3	15.8	37.6	42.5
Cascade R-CNN [4]	2018	22.4	37.0	13.1	34.1	42.1
w/ UGS	-	24.4	40.1	15.8	37.8	44.3
GFL V1 [29]	2020	28.4	50.0	15.9	39.7	42.6
w/ CEASC [10]	2023	28.7	50.7	-	-	-
w/ UGS	-	31.2	53.0	19.2	42.6	54.4
DINO-5scale [59]	2023	35.5	58.0	22.4	45.6	51.2
w/ UGS	-	38.1	61.9	24.2	47.3	56.6

Table 2. Detection performance of YOLO-based detectors on the VisDrone val [62], using CSP-D53 [51] as backbone. +(\cdot) indicates the performance improvement, and **bold** denotes the best results.

Method	Input Size	AP	AP _s	AP _m
TPH-YOLOv5-x [64]	(640, 640)	26.5	14.1	22.0
w/ UGS		29.0	16.7	23.9
TPH-YOLOv5-x [64]	(1536, 1536)	39.2	22.4	34.7
w/ UGS		41.7	24.2	36.2

4.1.2. Implementation Details.

We conduct the experiments on a computer with 4 NVIDIA RTX 3090 GPUs. The models are implemented using PyTorch [40], with core functionalities built upon the MMDetection framework [7]. For the VisDrone dataset [62], we use an input resolution of 1333×800 with a batch size of 16. For fair comparisons, we utilize default optimization parameters and $1 \times$ learning schedule without multi-scale training. We utilize the SGD optimizer with initial learning rate of $2e-2$, decayed by 0.1 at epoch 8 and 11, and weight decay of $1e-4$. We evaluate our method on two state-of-the-art small object detectors. For YOLO-based TPH-YOLOv5 [64], we train the model for 80 epochs using the default hyperparameters from [64]. We evaluate under two input resolutions— 640×640 and 1536×1536 , with batch sizes of 16 and 4, respectively. We also train a transformer-based detector, DINO [59], with 5-scale feature maps for 24 epochs as a baseline. Training employs the AdamW optimizer with initial learning rate $2e-4$, and adopts DETR-style augmentation [63] including random cropping and multi-scale resizing. At test time, the maximum number of predictions per image is increased to 1500 following DQ-DETR [20] to accommodate dense small-object scenarios.

4.2. Experiments on VisDrone

We evaluate the proposed UGS method with various baselines and two state-of-the-art small object detectors on the VisDrone val set [62]. As shown in Tab. 1, UGS enhances all baseline detectors by approximately 2% in AP, a notable improvement. Specifically, UGS boosts the performance of the anchor-free FCOS [50] by 2.5% in AP and 1.6% in AP_s. Integrated with two-stage detectors, UGS improves Faster R-CNN [44] by 2.6% AP and 1.5% AP_s, while enhancing Cascade R-CNN [4] by 2.0% AP and 2.7% AP_s, yielding consistent performance gains. Furthermore, we apply UGS to GFL V1 [29], a strong general object detector that employs a classification-based localization objective. UGS improves GFL V1 by a significant 3.5% AP, surpassing the state-of-the-art comparison CEASC [10]. In addition, we evaluate UGS’s compatibility with DETR-based architecture DINO [59]. UGS boosts the performance of the DINO-5scale baseline by 2.6% AP, surpassing the prior art DQ-DETR [20]. Tab. 2 highlights the compatibility of UGS with YOLO-based architectures. Specifically, UGS improves TPH-YOLOv5 by 2.5% AP at both 640^2 and 1536^2 resolutions. In particular, UGS achieves 41.7% AP in the VisDrone validation set without test-time augmentation.

4.3. Experiments on SODA-A

We further validate the UGS method on the aerial rotated small object detection dataset SODA-A [9]. As shown in Tab. 3, the proposed UGS method effectively improves the performance of both anchor-free and anchor-based rotating detectors. Specifically, for the single-stage anchor-based model, Rotated RetinaNet [35], UGS improved the AP by 4.5%. For the anchor-free model, Rotated FCOS, UGS enhanced the AP by 2.6%. For the proposal-based Oriented RCNN, UGS boosted the AP by 1.6%.

Table 3. Main results with various frameworks and FPN features on SODA-A. Note that models are trained on the SODA-A `train` and validated on the SODA-A `test`. We report APs (%) with different IoU threshold and APs (%) for objects in various sizes based on the SODA-A criterion. The **bold** denotes the best result.

Method	AP	AP _{0.5}	AP _{0.75}	AP _{eS}	AP _{rS}	AP _{gS}	AP _N
Rotated RetinaNet [35]	22.3	57.7	10.6	7.7	18.0	28.0	23.0
w/ UGS	26.8	63.2	15.9	8.9	21.7	34.9	27.5
Rotated FCOS [50]	32.6	69.4	24.8	11.5	30.0	42.7	41.6
w/ UGS	35.2	72.9	28.7	12.3	30.5	46.3	39.2
Oriented Reppoints [28]	25.4	60.0	16.8	8.6	21.9	29.6	27.8
w/ UGS	28.2	64.5	19.7	9.8	23.8	33.6	31.2
Oriented RCNN [28]	34.4	70.7	28.6	12.5	28.6	44.5	36.7
w/ UGS	36.0	73.1	30.3	13.7	30.2	47.8	38.1

Table 4. Experiments with Faster R-CNN on VOC [11] and COCO [33]. +(\cdot) indicates the performance improvement, and **bold** denotes the best results. ‘-’ indicates that the result is not reported or not publicly available.

Backbone	Method	Dataset	AP	AP _s
R-50	Faster R-CNN	VOC	78.4	-
	w/ UGS		82.2	-
R-50	Faster R-CNN	COCO	37.4	21.2
	w/ UGS		38.9	22.6
R-101	Faster R-CNN	VOC	79.9	-
	w/ UGS		81.7	-
R-101	Faster R-CNN	COCO	39.4	22.4
	w/ UGS		41.4	24.9

4.4. Experiments on COCO and VOC

To evaluate the generalizability of UGS, we conduct experiments on two general object detection benchmarks: PASCAL VOC [11] and COCO [33]. As shown in Tab. 4, UGS achieves significant improvements over the baselines on VOC. Under the ResNet-50 backbone, UGS improves Faster R-CNN by 3.8% AP, raising the performance from 78.4% to 82.2%. With the ResNet-101 backbone, UGS achieves a 1.8% improvement. On the COCO dataset, UGS demonstrates consistent improvements, particularly for small objects (AP_s). With the ResNet-50 backbone, UGS improves overall AP by 1.5% and AP_s by 1.4%. For ResNet-101, the improvements are more pronounced. These results underscore UGS’s effectiveness in improving small object detection while maintaining strong performance across all object sizes.

4.5. Training Cost of UGS

We analyze the training cost of our UGS on 1 NVIDIA RTX 3090 GPU with 24 GB memory. We test them using a batch size of 2 and input resolution of 1333 × 800 at each iteration. As shown in Tab. 5, UGS introduces 15%, 0.6%, 13% increase in training time, computational and memory cost, which is moderate.

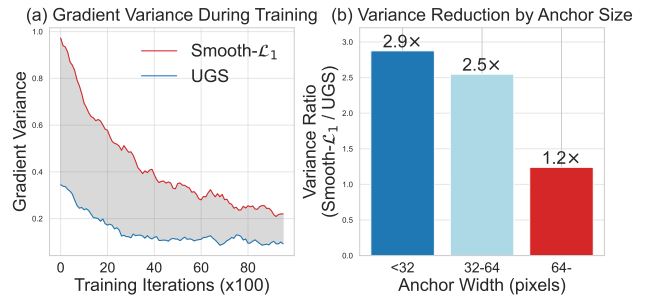


Figure 6. Gradient variance analysis for small objects. (a) Variance of localization loss gradients across training iterations. (b) Mean variance ratio (UGS vs. Smooth- \mathcal{L}_1) grouped by anchor size. Our method reduces variance by 2.9 \times for $w_a < 32$ pixels.

Table 5. Comparing the training cost of Baseline and our UGS using various detection frameworks using ResNet-50 backbone.

Method	AP	Time (s/batch)	GFLOPs	Params (M)
FCOS	19.9	0.151	127.25	42.40
w/ UGS	22.4	0.174	128.06	48.12

4.6. Gradient Variance Analysis

To quantify gradient stability, we measure the gradient variance during training for both our method (UGS) and the Smooth- \mathcal{L}_1 baseline. For small objects ($w_a < 32$), we compute the variance of localization loss gradients across 100 training iterations and report the mean variance ratio, defined as the ratio of the gradient variance of Smooth- \mathcal{L}_1 to that of UGS. Results in Fig. 6 demonstrate that UGS reduces gradient variance by up to 2.9 \times for small objects compared to Smooth- \mathcal{L}_1 . This validates our analysis in Sec. 3.2 that classification-based localization objective and the proposed uncertainty-aware designs lead to smoother optimization landscapes and more stable gradients.

4.7. Visualizations

We show in Fig. 5 that our design leads to a more stable training process, where the gradient magnitude of salient

Table 6. Ablation study of UGS components on VisDrone validation set using Faster R-CNN. Improvements over \mathcal{L}_2 baseline are marked in blue. UR denotes the uncertainty-guided refinement loss.

Method	λ	γ	ρ	IN	AP	AP ₅₀	AP _s
\mathcal{L}_2 (Baseline)	-	-	-	-	21.3	36.4	12.8
Smooth- \mathcal{L}_1	-	-	-	-	21.5	36.4	12.7
IoU-loss [57]	-	-	-	-	21.8	37.0	13.0
Gradient Clipping	-	-	-	-	21.6	36.7	12.9
Bayesian YOLO [25]	-	-	-	-	22.0	37.3	13.1
<i>Classification-based Localization</i>							
\mathcal{L}_{CE}	-	-	-	-	22.1	37.1	13.2
$\mathcal{L}_{CE} + \text{IN}$	-	-	-	✓	22.5	38.2	13.4
<i>CE + Uncertainty Minimization (UM)</i>							
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM}$	0.1	-	-	✓	22.7	38.2	13.6
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM}$	0.5	-	-	✓	22.9	38.4	13.6
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM}$	1.0	-	-	✓	22.6	37.2	13.7
<i>CE + UM + Uncertainty-guided Refinement (UR)</i>							
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM} + \gamma \mathcal{L}^{\text{ur}}$	0.5	0.1	0.5	✓	23.5	39.0	14.1
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM} + \gamma \mathcal{L}^{\text{ur}}$	0.5	0.5	0.5	✓	24.2	41.3	15.8
$\mathcal{L}_{CE} + \lambda \mathcal{L}_{UM} + \gamma \mathcal{L}^{\text{ur}}$	0.5	0.5	1	✓	23.2	38.7	13.9

small object examples becomes less pronounced compared to the \mathcal{L}_2 loss. Fig. 5 also highlights UGS’s ability to learn from uncertain examples, including occluded objects, objects in cluttered backgrounds, and noise-resembling objects.

4.8. Ablation Study

In the following experiments, we demonstrate that UGS leads to consistent performance increment when applied on baseline detectors and explore the best-performing structure by ablating and tuning each component. We use Faster R-CNN [44] with ResNet-50 [17] as the baseline and conduct all tests on the VisDrone val dataset [62].

Effectiveness of Components. According to Tab. 6, replacing \mathcal{L}_2 with the classification-based localization objective (\mathcal{L}_{CE}) significantly improves performance, achieving 22.1 AP. Utilizing IN labels further increases performance by 1.2% AP and 0.6% AP_s, demonstrating the importance of non-uniform quantization for ensuring balanced supervision for small object detection. Adding uncertainty minimization loss with $\lambda = 0.5$ increases performance to 22.9 AP, validating the effectiveness of entropy-based uncertainty modeling. The full UGS framework, incorporating CE, UM, and the uncertainty-guided refinement module (\mathcal{L}^{ur}), achieves the best performance with 24.2 AP, 41.3 AP₅₀ and 15.8 AP_s. This represents a 2.9 AP improvement over the \mathcal{L}_2 baseline, which is significant. We also evaluate existing gradient stabilization techniques for comparison, including normalized regression losses, gradient clipping ($\|\nabla\| \leq 1.0$), and uncertainty-aware method Bayesian YOLO [25]. While the methods effect in mitigating gradient instability, their improvements are relatively marginal

Table 7. Parameter analysis of UGS components on Faster R-CNN: Effects of range (α), IN modulator (β), and grid number (n). Best results per stage are in bold.

Stage	Parameter	α	β	n	AP	AP ₅₀	AP _s
RPN	α - β	2	1.0	10	21.9	37.4	13.0
		2	1.5	10	21.8	37.2	13.1
		3	1.0	10	21.8	37.2	13.1
		3	1.5	10	21.7	37.1	13.0
	Grid Number	2	1.0	2	21.8	36.9	12.9
				4	21.9	37.1	13.0
		20	1.0	10	21.9	37.4	13.0
				20	21.8	37.1	13.0
R-CNN	α - β	4	1.0	5	22.2	37.6	13.1
		4	1.5	5	21.9	37.2	13.0
		5	1.0	5	22.5	38.2	13.4
		5	1.5	5	22.0	37.0	13.1
	Grid Number	5	1.0	5	22.5	38.2	13.4
				10	22.3	37.6	13.2
		20	1.0	20	22.4	37.9	13.4
				20	22.4	37.9	13.4

compared to UGS.

Effect of α , β , and n in Interval Non-uniform (IN) Label Quantization. We investigate the impact of three key parameters in UGS’s interval non-uniform (IN) label quantization: the range parameter α , the IN modulator β , and the grid number n . Experiments show that a moderate range of $\alpha = 2$ for RPN and $\alpha = 5$ for R-CNN yields the best performance improvements, achieving 1.2% AP and 0.6% AP_s. Notably, the performance remains robust across different values of α , indicating that UGS is not overly sensitive to this parameter. A smaller β (e.g., 1.0) generally leads to better performance, as it creates a denser grid distribution around zero. This design is particularly effective for capturing finer details in small object instances, where precise localization is critical. UGS demonstrates flexibility in the choice of grid number n , with performance remaining stable across different configurations.

5. Conclusions

In this paper, we demonstrate that conventional object localization methods tend to produce unstable gradients on small objects and introduce an Uncertainty-Aware Gradient Stabilization (UGS) method to rectify gradients. UGS quantizes continuous labels into interval non-uniform discrete representations as supervision. In optimization, UGS employs a classification-based localization objective that generates bounded and confidence-driven gradients. Further, UGS integrates a dual uncertainty-aware mechanism to enhance the robustness of localization. UGS consistently improves baselines and state-of-the-art small object detectors, demonstrating compatibility across different detection architectures. Future work includes extending UGS to video or multi-modal small object detection scenarios.

Acknowledgments. This work was supported by the National Key Research and Development Program of China (No. 2023YFC3306401). This research was also supported by the National Natural Science Foundation of China (No. 61827901, 623B2016, 62406298), the Zhejiang Provincial Natural Science Foundation (No. LD24F020007), the Beijing Natural Science Foundation (No. L223024, L244043, Z241100001324017), the “One Thousand Plan” projects in Jiangxi Province (No. Jxsq2023102268), and the Fundamental Research Funds for the Central Universities (CUC25QT17).

References

- [1] Fatih Cagatay Akyon, Sinan Onur Altinuc, and Alptekin Temizel. Slicing aided hyper inference and fine-tuning for small object detection. In *ICIP*, pages 966–970. IEEE, 2022. 1
- [2] Syed Muhammad Arsalan Bashir and Yi Wang. Small object detection in remote sensing images with residual feature aggregation-based super-resolution and object detector network. *Remote Sensing*, 13(9):1854, 2021. 2
- [3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020. 3
- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, pages 6154–6162, 2018. 1, 6
- [5] Bing Cao, Haiyu Yao, Pengfei Zhu, and Qinghua Hu. Visible and clear: Finding tiny objects in difference map. In *ECCV*, 2024. 1, 2
- [6] Chenyi Chen, Ming-Yu Liu, Oncel Tuzel, and Jianxiong Xiao. R-cnn for small object detection. In *ACCV*, pages 214–230. Springer, 2017. 2
- [7] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 6
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2017. 2
- [9] Gong Cheng, Xiang Yuan, Xiwen Yao, Kebin Yan, Qinghua Zeng, Xingxing Xie, and Junwei Han. Towards large-scale small object detection: Survey and benchmarks. *IEEE TPAMI*, 2023. 5, 6
- [10] Bowei Du, Yecheng Huang, Jiabin Chen, and Di Huang. Adaptive sparse convolutional networks with global context enhancement for faster object detection on drone images. In *CVPR*, pages 13435–13444, 2023. 1, 6
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. 1, 5, 7
- [12] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. *ICLR*, 2021. 5
- [13] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 3
- [14] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014. 3
- [15] Jianyuan Guo, Kai Han, Yunhe Wang, Han Wu, Xinghao Chen, Chunjing Xu, and Chang Xu. Distilling object detectors via decoupled features. In *CVPR*, pages 2154–2164, 2021. 1, 5
- [16] Kaiming He and Jian Sun. Convolutional neural networks at constrained time cost. In *CVPR*, pages 5353–5360, 2015. 1
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 8
- [18] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *CVPR*, pages 2888–2897, 2019. 3
- [19] Yanxu Hu and Andy J Ma. Adversarial feature augmentation for cross-domain few-shot classification. In *ECCV*, pages 20–37. Springer, 2022. 5
- [20] Yi-Xin Huang, Hou-I Liu, Hong-Han Shuai, and Wen-Huang Cheng. Dq-detr: Detr with dynamic query for tiny object detection. In *ECCV*, pages 290–305. Springer, 2024. 6
- [21] Masato Ishii and Atsushi Sato. Training deep neural networks with adversarially augmented features for small-scale training datasets. In *IJCNN*, pages 1–8. IEEE, 2019. 5
- [22] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yunying Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV*, pages 784–799, 2018. 3
- [23] Woo Jae Kim, Yoonki Cho, Junsik Jung, and Sung-Eui Yoon. Feature separation and recalibration for adversarial robustness. In *CVPR*, pages 8183–8192, 2023. 5
- [24] Mate Kisantal, Zbigniew Wojna, Jakub Murawski, Jacek Naruniec, Kyunghyun Cho, et al. Augmentation for small object detection. In *CS & IT Conference Proceedings. CS & IT Conference Proceedings*, 2019. 2
- [25] Florian Kraus and Klaus Dietmayer. Uncertainty estimation in one-stage object detection. In *ITSC*, pages 53–60. IEEE, 2019. 8
- [26] Youngwan Lee, Joong-won Hwang, Hyung-Il Kim, Kimin Yun, Yongjin Kwon, Yuseok Bae, and Sung Ju Hwang. Localization uncertainty estimation for anchor-free object detection. In *ECCV*, pages 27–42. Springer, 2022. 3
- [27] Jianan Li, Xiaodan Liang, Yunchao Wei, Tingfa Xu, Jiashi Feng, and Shuicheng Yan. Perceptual generative adversarial networks for small object detection. In *CVPR*, pages 1222–1230, 2017. 2
- [28] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *CVPR*, pages 1829–1838, 2022. 7
- [29] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In *NeurIPS*, 2020. 2, 3, 4, 6

- [30] Xiang Li, Chengqi Lv, Wenhai Wang, Gang Li, Lingfeng Yang, and Jian Yang. Generalized focal loss: Towards efficient representation learning for dense object detection. *IEEE TPAMI*, 2022. 5
- [31] Yanghao Li, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Scale-aware trident networks for object detection. In *ICCV*, pages 6054–6063, 2019. 2
- [32] Xi Liang, Jing Zhang, Li Zhuo, Yuzhao Li, and Qi Tian. Small object detection in unmanned aerial vehicle images using feature fusion and scaling-based single shot detector with spatial context analysis. *IEEE TCSVT*, 30(6):1758–1770, 2019. 1, 2
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. 1, 5, 7
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 1, 2, 5
- [35] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 6, 7
- [36] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37, 2016. 2
- [37] Akhil Meethal, Eric Granger, and Marco Pedersoli. Cascaded zoom-in detector for high resolution aerial images. In *CVPR*, pages 2046–2055, 2023. 6
- [38] Qi Ming, Lingjuan Miao, Zhe Ma, Lin Zhao, Zhiqiang Zhou, Xuhui Huang, Yuanpei Chen, and Yufei Guo. Deep dive into gradients: Better optimization for 3d object detection with gradient-corrected iou supervision. In *CVPR*, pages 5136–5145, 2023. 2
- [39] Yurii Nesterov et al. *Lectures on convex optimization*. Springer, 2018. 3
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS Workshop*, 2017. 6
- [41] Siyuan Qiao, Liang-Chieh Chen, and Alan Yuille. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *CVPR*, pages 10213–10224, 2021. 6
- [42] Heqian Qiu, Hongliang Li, Qingbo Wu, and Hengcan Shi. Offset bin classification network for accurate object detection. In *CVPR*, pages 13188–13197, 2020. 2
- [43] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 3
- [44] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1, 2, 3, 4, 6, 8
- [45] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, pages 658–666, 2019. 3
- [46] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 5
- [47] Bharat Singh and Larry S Davis. An analysis of scale invariance in object detection snip. In *CVPR*, pages 3578–3587, 2018. 2
- [48] Huixin Sun, Runqi Wang, Yanjing Li, Linlin Yang, Shaohui Lin, Xianbin Cao, and Baochang Zhang. Set: Spectral enhancement for tiny object detection. In *CVPR*, pages 4713–4723, 2025. 2
- [49] Jingru Tan, Xin Lu, Gang Zhang, Changqing Yin, and Quanquan Li. Equalization loss v2: A new gradient balance approach for long-tailed object detection. In *CVPR*, pages 1685–1694, 2021. 2
- [50] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 3, 4, 6, 7
- [51] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Scaled-yolov4: Scaling cross stage partial network. In *CVPR*, pages 13029–13038, 2021. 6
- [52] Chang Xu, Jinwang Wang, Wen Yang, and Lei Yu. Dot distance for tiny object detection in aerial images. In *CVPR*, pages 1192–1201, 2021. 2, 3
- [53] Chang Xu, Jinwang Wang, Wen Yang, Huai Yu, Lei Yu, and Gui-Song Xia. Rfla: Gaussian receptive field based label assignment for tiny object detection. In *ECCV*, pages 526–543, 2022. 2
- [54] Sheng Xu, Mingze Wang, Yanjing Li, Mingbao Lin, Baochang Zhang, David Doermann, and Xiao Sun. Learning 1-bit tiny object detector with discriminative feature refinement. In *ICML*, 2024. 1, 2
- [55] Chenhongyi Yang, Zehao Huang, and Naiyan Wang. Querydet: Cascaded sparse query for accelerating high-resolution small object detection. In *CVPR*, pages 13668–13677, 2022. 1, 6
- [56] Jiangnan Yang, Shuangli Liu, Jingjun Wu, Xinyu Su, Nan Hai, and Xueli Huang. Pinwheel-shaped convolution and scale-based dynamic loss for infrared small target detection. In *AAAI*, 2024. 1
- [57] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 516–520, 2016. 2, 3, 8
- [58] Xiang Yuan, Gong Cheng, Kebing Yan, Qinghua Zeng, and Junwei Han. Small object detection via coarse-to-fine proposal generation and imitation learning. In *ICCV*, pages 6317–6327, 2023. 2
- [59] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *ICLR*, 2023. 1, 3, 6
- [60] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, pages 9759–9768, 2020. 3, 4
- [61] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and bet-

- ter learning for bounding box regression. In *AAAI*, pages 12993–13000, 2020. [3](#)
- [62] Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Haibin Ling, Qinghua Hu, Qinqin Nie, Hao Cheng, Chenfeng Liu, Xiaoyu Liu, et al. Visdrone-det2018: The vision meets drone object detection in image challenge results. In *ECCVW*, pages 0–0, 2018. [2](#), [5](#), [6](#), [8](#)
- [63] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. [6](#)
- [64] Xingkui Zhu, Shuchang Lyu, Xu Wang, and Qi Zhao. Tph-yolov5: Improved yolov5 based on transformer prediction head for object detection on drone-captured scenarios. In *ICCV*, pages 2778–2788, 2021. [3](#), [6](#)
- [65] Barret Zoph, Ekin D Cubuk, Golnaz Ghiasi, Tsung-Yi Lin, Jonathon Shlens, and Quoc V Le. Learning data augmentation strategies for object detection. In *ECCV*, pages 566–583. Springer, 2020. [2](#)