

Article

ADG-YOLO: A Lightweight and Efficient Framework for Real-Time UAV Target Detection and Ranging

Hongyu Wang *, Zheng Dang , Mingzhu Cui, Hanqi Shi , Yifeng Qu, Hongyuan Ye, Jingtao Zhao and Duosheng Wu

School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai 200240, China

* Correspondence: redrain@sjtu.edu.cn

Highlights

What are the main findings?

- ADG-YOLO achieves a lightweight architecture with only 1.77 M parameters and 5.7 GFLOPs, while maintaining a high detection accuracy of 98.4% mAP_{0.5} and 27 FPS on an edge computing device.
- The integrated monocular ranging method based on similar triangles achieves an average distance estimation error of 2.40–4.18% across 0.5–50 m for three different UAV models.

What is the implication of the main finding?

- This work provides a practical and efficient solution for real-time UAV detection and ranging on resource-constrained edge platforms, enabling onboard intelligence for autonomous UAV operations.
- The proposed framework demonstrates the feasibility of deploying advanced perception systems on low-power devices, paving the way for wider adoption of AI-driven UAVs in military and commercial applications.



Academic Editors: Peiyang Zhang and Athanasios V. Vasilakos

Received: 22 August 2025

Revised: 6 October 2025

Accepted: 7 October 2025

Published: 13 October 2025

Citation: Wang, H.; Dang, Z.; Cui, M.; Shi, H.; Qu, Y.; Ye, H.; Zhao, J.; Wu, D. ADG-YOLO: A Lightweight and Efficient Framework for Real-Time UAV Target Detection and Ranging. *Drones* **2025**, *9*, 707. <https://doi.org/10.3390/drones9100707>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract

The rapid evolution of UAV technology has increased the demand for lightweight airborne perception systems. This study introduces ADG-YOLO, an optimized model for real-time target detection and ranging on UAV platforms. Building on YOLOv11n, we integrate C3Ghost modules for efficient feature fusion and ADown layers for detail-preserving downsampling, reducing the model's parameters to 1.77 M and computation to 5.7 GFLOPs. The Extended Kalman Filter (EKF) tracking improves positional stability in dynamic environments. Monocular ranging is achieved using similarity triangle theory with known target widths. Evaluations on a custom dataset, consisting of 5343 images from three drone types in complex environments, show that ADG-YOLO achieves 98.4% mAP_{0.5} and 85.2% mAP_{0.5:0.95} at 27 FPS when deployed on Lubancat4 edge devices. Distance measurement tests indicate an average error of 4.18% in the 0.5–5 m range for the DJI NEO model, and an average error of 2.40% in the 2–50 m range for the DJI 3TD model. These results suggest that the proposed model provides a practical trade-off between detection accuracy and computational efficiency for resource-constrained UAV applications.

Keywords: UAV detection; monocular ranging; edge computing; YOLOv11n; real-time tracking

1. Introduction

With the rapid evolution of unmanned aerial vehicle (UAV) technology, its significance has been steadily rising across military, civilian, and commercial domains. In modern military operations, UAVs are increasingly used not only in auxiliary roles but also for direct operational tasks, influencing combat strategies and operational planning [1–3]. Breakthroughs in artificial intelligence (AI), low-cost manufacturing, and stealth flight technology have endowed various unmanned systems with enhanced perception, autonomous decision-making, and strike capabilities. A representative case is the “Operation Spider’s Web” conducted in June 2025, in which a coordinated deployment of over 100 FPV UAVs demonstrated the potential of UAV swarms for precision strike tasks. This operation highlights the growing complexity of UAV missions and the operational challenges for conventional defense systems, emphasizing the need for advancements in target detection, path planning, and low-latency response mechanisms.

Currently, the development of unmanned systems technology exhibits two prominent trends. First, AI-empowered intelligent systems are progressively enhancing autonomous decision-making and operational coordination, with UAV swarms demonstrating closed-loop capabilities in simulated and real-world missions. Second, counter-UAV strategies are evolving, including communication jamming, deception interference, and terrain-based evasion tactics [4,5]. Within this technological contest, improving the perception capabilities of unmanned systems—ensuring reliable detection, operational stability, and energy efficiency—remains a critical focus for tactical applications.

Compared to conventional ground-based fixed observation platforms, airborne perception systems integrate visual ranging functions directly onto UAV platforms, achieving mobile and real-time sensing. Airborne platforms allow closer proximity to targets, enable dynamic close-range tracking, and support complex tasks such as autonomous obstacle avoidance, formation coordination, and precision operations—thereby enhancing operational autonomy in diverse scenarios [2]. However, airborne deployment also imposes stringent demands on system lightweight design, low power consumption, and real-time responsiveness, motivating the development of visual ranging and target detection algorithms optimized for embedded edge computing platforms. Advancing efficient and reliable airborne visual ranging systems is not only a key technology for UAV intelligence upgrades but also an important enabler for improving future UAV operational effectiveness and mission reliability.

At present, mainstream ranging methods include radar [6], laser [7], ultrasonic [8], and visual ranging [9]. While radar and laser offer high precision and long detection ranges, their bulky size, high power consumption, and cost render them unsuitable for small UAV applications. Visual ranging, with its non-contact nature, low cost, and ease of integration, has become a research focus for lightweight perception systems. Based on camera configurations, visual ranging can be categorized into binocular and monocular systems. Binocular vision offers higher accuracy but requires extensive calibration and disparity computation resources [10,11]. In contrast, monocular vision—with its simple hardware structure and algorithmic flexibility—has emerged as the preferred solution for edge and embedded platforms [12,13].

Monocular visual ranging methods fall into two categories: depth map estimation and physical distance regression. The former relies on convolutional neural networks to generate relative depth maps, which are suitable for scene modeling but lack direct physical distance outputs [14–19]. The latter directly outputs target distances based on geometric modeling or end-to-end regression, offering faster and more practical responses. Geometric modeling approaches, such as perspective transformation, inverse perspective mapping (IPM), and similar triangle modeling, estimate depth by mapping pixel data to physical

parameters [20–25]. In recent years, deep learning techniques have been introduced to monocular ranging models to enhance adaptability in unstructured environments [26]. Traditional two-stage detection algorithms such as RCNN [27], Fast R-CNN [28], and Faster R-CNN [29] excel in detection accuracy but are limited by their complex architectures and high computational demands—making them less suitable for edge platforms that require low power and real-time performance, thus restricting their practical applications in UAV scenarios.

In the field of target detection, the YOLO (You Only Look Once) family of algorithms has gained widespread use in unmanned systems due to its fast detection speed and compact architecture, ideal for edge deployment [30–41]. To enhance small-object detection and long-range recognition on embedded platforms, researchers have proposed various improvements based on YOLOv5, YOLOv8, and YOLOv11, incorporating lightweight convolutions, attention mechanisms, and feature fusion modules [42–45]. However, most existing studies focus on ground-based or general-purpose deployments and lack systematic research on lightweight airborne deployment and integrated perception-ranging systems. Particularly on ultra-low-power, computation-constrained mobile platforms, balancing detection accuracy, ranging stability, and frame rate remains a critical unresolved challenge. Reference [46] explored a low-power platform based on Raspberry Pi combined with YOLOv5 for real-time UAV target detection. Although it did not deeply address deployment efficiency and response latency, it provided a valuable reference for lightweight applications. In addition, several recent studies have emphasized the importance of dataset augmentation strategies to enhance small-drone detection robustness and improve generalization across diverse environments [47–51].

To this end, this study expands the research perspective and application scope of UAV visual perception. Transitioning from traditional ground-based observation models to onboard autonomous sensing, we propose a lightweight airborne perception system mounted directly on UAV platforms. This system is designed to meet the integrated requirements of target detection and distance measurement, achieving real-time, in-flight optimization of both functionalities. The main innovations and contributions of this study are summarized as follows:

1. **Lightweight Detection Architecture Design:** Based on the YOLOv11n model, this study introduces the C3GHOST and ADown modules to construct an efficient detection architecture tailored for edge computing platforms. The C3GHOST module reduces computational overhead through lightweight feature fusion while enhancing feature representation capability. The ADown module employs an efficient down-sampling strategy that lowers computational cost without compromising detection accuracy. Systematic evaluation on a custom-built dataset demonstrates the model's capability for joint optimization in terms of frame rate and ranging precision.
2. **Target Tracking Optimization:** To further improve the stability of UAV target tracking, this study incorporates the Extended Kalman Filter (EKF) approach. EKF performs target position estimation and trajectory prediction in dynamic environments, significantly reducing position jitter and sporadic false detections during the tracking process, thereby enhancing robustness and consistency.
3. **Dataset Expansion:** Based on a publicly available dataset from CSDN, this study conducts further expansion by constructing a comprehensive dataset that covers a wide range of UAV models and complex environments. The dataset includes image samples captured under varying flight altitudes, viewing angles, and lighting conditions. This expansion enables the proposed model to not only adapt to different types of UAV targets, but also maintain high detection accuracy and stability in complex flight environments.

4. **Model Conversion and Deployment on Edge Devices:** To facilitate practical deployment, the trained model was converted from its standard format to a format compatible with edge computing devices based on the RK3588S chipset. The converted model was successfully deployed onto the edge platform, ensuring efficient operation on resource-constrained hardware.

2. Methodology

2.1. ADG-YOLO

In this section, we present a lightweight UAV visual perception framework based on the YOLOv11n architecture for edge devices, integrating C3Ghost and ADown modules. The framework achieves a balanced trade-off between detection accuracy and real-time performance, enabling reliable target detection.

2.1.1. C3Ghost

In this study, the standard C3k2 modules are replaced with C3Ghost modules to achieve a more lightweight design and improved computational efficiency in feature fusion. The C3Ghost module consists of a series of GhostConv layers integrated within a Cross-Stage Partial (CSP) architecture, combining efficient information flow with compact network structure [52]. As shown in Figure 1, GhostConv divides the input feature map into two parts: the first part generates primary features using standard convolution, while the second part produces complementary “ghost” features through low-cost linear transformations such as depthwise separable convolutions. These two parts are then concatenated to form the final output. By exploiting the inherent redundancy in feature representations, this design significantly reduces the number of parameters and floating-point operations, while retaining a representational capacity comparable to conventional convolutions.




Figure 1. The Ghost module. The figure is redrawn by the authors for clarity, based on the original design of Han et al. [52].

On this basis, C3Ghost integrates multiple GhostConv layers into the CSP structure to form a lightweight feature fusion unit. As illustrated in Figure 2, the input is split into two parallel branches. One branch extracts higher-level features through a series of stacked GhostBottleneck layers, and the other directly passes the input via a shortcut connection to preserve original information. The outputs from both branches are then concatenated and fused using a 1×1 convolution. This design enhances feature representation while keeping computational cost and model complexity to a minimum [53].




Figure 2. The C3Ghost module. The figure is redrawn by the authors for clarity, based on the original design of Ji et al. [53].

2.1.2. ADown

In this study, an ADown module is introduced to replace conventional convolution-based downsampling operations, aiming to improve efficiency while preserving fine-grained feature information. As illustrated in Figure 3, the core design of ADown consists of the following stages: the input feature map is first processed through average pooling and downsampled to half the spatial resolution. It is then split into two branches along the channel dimension. The first branch applies a 3×3 convolution to extract local detail features, while the second branch undergoes max pooling for downsampling, followed by a 1×1 convolution for channel compression and nonlinear transformation. Finally, the outputs from both branches are concatenated along the channel axis to form the downsampled output [54].




Figure 3. The ADown module. The figure is redrawn by the authors for clarity, based on the original design of Fang et al. [54].

Compared to standard convolution or conventional pooling-based downsampling, ADown offers several distinct advantages. Its multi-path structure allows for the integration of both global and local information, mitigating the severe loss of detail often caused by traditional downsampling methods. Additionally, by reducing spatial resolution through pooling before applying lightweight convolutions, ADown achieves efficient feature extraction with significantly fewer parameters and lower FLOPs, without sacrificing representational power.

Moreover, ADown can be seamlessly integrated with existing multi-scale feature fusion modules, such as SPPF or FPN, enhancing the capacity of both the backbone and neck components to retain fine-grained features. This is particularly beneficial for small object detection tasks. In UAV imagery, where targets tend to be small and highly resolution-sensitive, the hierarchical detail preserved by ADown proves critical for accurately detecting small-scale objects. Its design not only improves fine-feature retention but also enhances the overall robustness of the model in multi-scale environments.

2.1.3. Proposed ADG-YOLO

To enhance the detection capability of lightweight networks for low-altitude, low-speed, and small-sized UAVs—commonly referred to as “low-slow-small” targets—under resource-constrained environments, this study proposes a novel architecture named ADG-YOLO (ADown + Ghost modules YOLO), based on the original YOLOv11n framework. While maintaining detection accuracy, ADG-YOLO significantly reduces model parameter size and computational complexity. The architecture incorporates systematic structural optimizations in three key areas: feature extraction, downsampling strategy, and multi-scale feature fusion. These improvements collectively strengthen the model’s perception capability for low-altitude UAV targets in ground-based scenarios, thereby better meeting the practical demands of UAV detection from aerial perspectives. The overall network architecture of ADG-YOLO is illustrated in Figure 4.

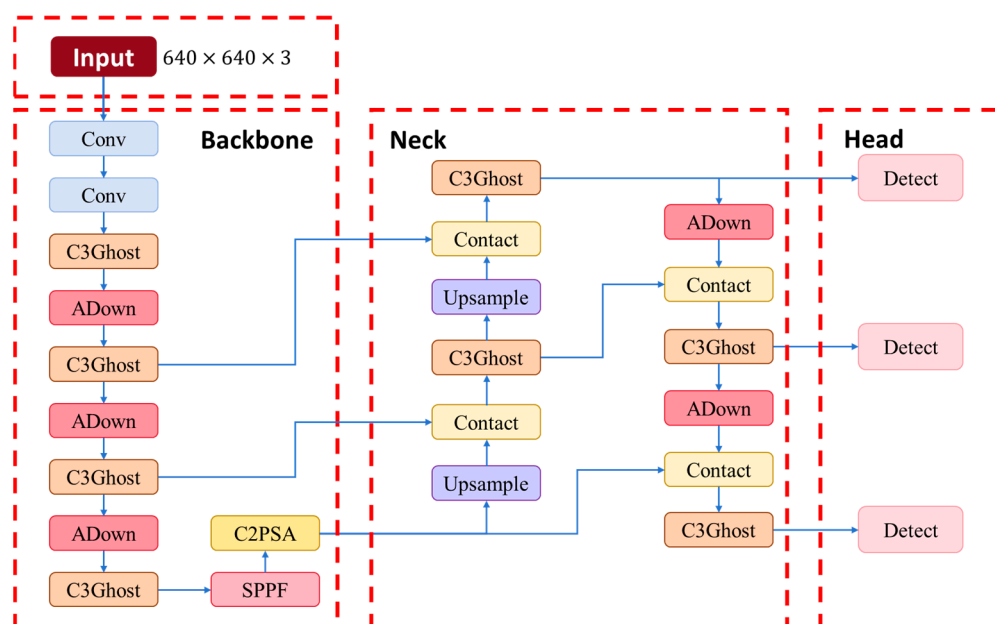


Figure 4. ADG-YOLO algorithm structure.

Firstly, in both the backbone and the neck, the original C3k2 modules were systematically replaced with C3Ghost modules. Specifically, in the backbone, the C3k2 blocks at the 256-, 512-, and 1024-channel stages were replaced, while in the neck, all C3k2 layers immediately following each feature upsampling–concatenation operation were also substituted with C3Ghost. This design ensures consistent lightweight representation across the entire network without altering the overall FPN–PAN topology. This ensures consistent lightweight representation across both backbone and neck without altering the overall FPN–PAN topology. C3Ghost is a lightweight residual module constructed using GhostConv, initially introduced in GhostNet, and incorporates a Cross Stage Partial (CSP) structure to enable cross-stage feature fusion. Its core design concept lies in generating primary features through standard convolution, while reusing redundant information by producing additional “ghost” features through low-cost linear operations such as depthwise separable convolutions. This approach significantly reduces the number of parameters and the overall computational cost (FLOPs), making it especially suitable for deployment on edge devices with limited computing resources. Additionally, the stacked structure of GhostBottleneck layers further enhances the network’s ability to represent features across multiple semantic levels.

Secondly, all stride = 2 downsampling operations in the network are replaced with the ADown module. Instead of conventional 3×3 strided convolutions, ADown adopts a dual-path structure composed of average pooling and max pooling for spatial compression. Each path extracts features at different scales through lightweight 3×3 and 1×1 convolutions, and the outputs are concatenated along the channel dimension. This asymmetric parallel design allows ADown to preserve more rich texture and edge information while reducing feature map resolution. Such a design is particularly beneficial in UAV-based detection scenarios where objects are small and captured from high-altitude viewpoints against complex backgrounds. Compared to standard convolutions, ADown effectively reduces computational burden without compromising detection accuracy, while improving the flexibility and robustness of the downsampling process.

Thirdly, the Spatial Pyramid Pooling—Fast (SPPF) module is retained at the end of the backbone to enhance the modeling of long-range contextual information. Meanwhile, in the neck, a series of alternating operations—upsampling, feature concatenation, and

downsampling via ADown—are introduced for feature fusion. This design facilitates the precise supplementation of low-level detail with high-level semantic information and improves alignment and interaction across multi-scale feature maps. As a result, the model's ability to detect small objects and capture boundary-level details is significantly enhanced. Combined with the lightweight feature extraction capability of the C3Ghost modules at various stages, the entire network achieves high detection accuracy while substantially reducing deployment cost and system latency.

In summary, the improvements of ADG-YOLO presented in this study are threefold: the C3Ghost modules enable efficient lightweight feature representation; the ADown module reconstructs a more effective downsampling pathway; and the SPPF module, together with multi-scale path interactions, strengthens fine-grained feature aggregation, particularly for small object detection. The complete network architecture of ADG-YOLO is shown in Figure 4, where the overall design seamlessly integrates lightweight structure with multi-scale feature enhancement. This model achieves a well-balanced trade-off among accuracy, inference speed, and computational resource consumption, offering high adaptability and practical value for real-world deployment.

2.2. Model Conversion and Edge Deployment

Considering factors such as device size, weight, scalability, and cost, this study selects the LubanCat 4 development board as the deployment platform for the ADG-YOLO model. The board is equipped with the Rockchip RK3588S processor and integrates an AI acceleration NPU capable of INT4, INT8, and INT16 mixed-precision computing, with a peak performance of up to 6 TOPS. It includes 4 GB of onboard memory and supports peripheral interfaces such as mini HDMI output and USB camera input, with an overall weight of approximately 62 g.

Given the limited computing capacity of the CPU, it is necessary to maximize inference efficiency by converting the model from its original .pt format (trained with PyTorch 1.8.1) to the .rknn format compatible with the NPU. The conversion pipeline proceeds as follows: first, the trained model is exported to the ONNX format using the `torch.onnx.export` interface in PyTorch; then, the RKNN Toolkit is used to convert the ONNX model into RKNN format. The overall conversion process is illustrated in Figure 5.




Figure 5. ADG-YOLO Model Conversion Process Diagram.

After conversion, the model is deployed onto the development board to enable real-time detection of UAV targets from live video input via the connected USB camera. With the aid of hardware acceleration provided by the NPU, the system is capable of maintaining a high frame rate and fast response speed while ensuring detection accuracy, thereby fulfilling the dual demands of real-time performance and lightweight deployment in practical application scenarios.

2.3. Target Monitoring Based on ADG-YOLO Detection and EKF Tracking

In this study, we propose a method that integrates the ADG-YOLO object detection algorithm with the Extended Kalman Filter (EKF) for target monitoring in dynamic scenarios. The YOLO model is employed to extract bounding box information from consecutive image frames in real time, including the center position and size parameters of the detected targets. To enable temporal filtering and motion trajectory prediction of the detected objects,

the target state is modeled as a six-dimensional vector $x = [c_x, c_y, v_x, v_y, w, h]^T$, consisting of the center coordinates (c_x, c_y) , the horizontal and vertical velocity components (v_x, v_y) , and the width w and height h of the bounding box. Considering that targets typically follow a constant velocity linear motion within short time intervals and that their size changes are relatively stable, a state transition model is formulated under this assumption. The corresponding state transition matrix is defined as follows:

$$F = \begin{bmatrix} 1 & 0 & \Delta t & 0 & 0 & 0 \\ 0 & 1 & 0 & \Delta t & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (1)$$

The observations provided by YOLO are the bounding box parameters of the detected target in the image, represented as $[c_x, c_y, w, h]^T$. The correspondence between these observations and the system state vector is modeled through an observation matrix, which is expressed as follows:

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad (2)$$

The execution process of the Extended Kalman Filter (EKF) consists of two stages: prediction and update. In the prediction stage, the target state and its covariance are estimated based on the current state and the state transition model, as expressed by:

$$\begin{cases} x' = F \cdot x, \\ P' = F \cdot P \cdot F^T + Q \end{cases} \quad (3)$$

Here, Q denotes the process noise covariance matrix, and P represents the observation noise covariance matrix. Upon receiving a new observation z from the YOLO algorithm, the update stage is performed as follows:

Residual computation:

$$y = z - H \cdot x' \quad (4)$$

Kalman gain computation:

$$K = P' \cdot H^T \cdot (H \cdot P' \cdot H^T + R)^{-1} \quad (5)$$

State update:

$$x = x' + K \cdot y \quad (6)$$

Covariance update:

$$P = (I - K \cdot H) \cdot P' \quad (7)$$

Here, I denotes the identity matrix.

The integration of the EKF module helps to mitigate the potential localization fluctuations and occasional false detections that may occur in YOLO's single-frame inference. This facilitates smoother target position estimation and further enhances the tracking consistency and robustness of the system in multi-frame processing scenarios.

2.4. Monocular Ranging for UAVs Using Similar Triangles

Figure 6 illustrates the UAV target detection results based on the YOLO model, where the red bounding boxes accurately locate and outline the position and size of the targets in the monocular images. The pixel width of the bounding box is denoted as p , representing the projected size of the target in the image, which serves as a key parameter for subsequent distance estimation. Neglecting lens distortion, and based on the principle of similar triangles, when the actual width of the target is W , the camera focal length is f , and the physical size of a single pixel on the image sensor is s , the actual projected width w of the target on the imaging plane can be expressed as:

$$w = p \cdot s \quad (8)$$




Figure 6. Drone Projection Width w Diagram.

As shown in Figure 7, when the target plane Ω_1 is perpendicular to the optical axis of the camera, the imaging process can be abstracted as two similar triangles, which satisfy the following proportional relationship:

$$\frac{W}{D} = \frac{w}{f} \quad (9)$$

Here, D denotes the distance from the target to the camera along the optical axis. Based on this relationship, the formula for computing the target distance is derived as follows:

$$D = \frac{W \cdot f}{p \cdot s} \quad (10)$$

In this study, the training dataset comprises three different types of UAVs, each associated with a known physical width W_n . The YOLO model not only outputs the bounding box coordinates but also possesses target classification capability, enabling precise identification of the specific UAV type. Once the target type is detected, the corresponding W_n is automatically selected and substituted into the distance estimation Formula (10), thereby enhancing the accuracy and generalizability of the distance measurement.




Figure 7. UAV Projection Width w Diagram.

3. Model Analysis

3.1. Dataset

A comprehensive UAV detection dataset was constructed for this study, comprising a total of 5343 high-resolution images. This dataset integrates two subsets: a custom-target subset with 2670 images and a generalization subset with 2664 images. The custom subset focuses on three specific UAV models: DJI 3TD (DJI, Shenzhen, China), with 943 training and 254 testing images (labeled as drone1); DJI NEO (DJI, China), with 739 training and 170 testing images (drone2); and DWI-S811 (DWI, China), with 454 training and 110 testing images (drone3). All images in this subset were captured under strictly controlled conditions, with target distances ranging from 5 to 30 m and 360-degree coverage, to reflect variations in object appearance under different perspectives and distances.

The generalization subset was collected from publicly available multirotor UAV image resources published on the CSDN object detection platform. It contains quadrotor and hexarotor UAVs from popular brands, appearing in diverse environments including urban buildings, rural landscapes, highways, and industrial areas. Additionally, the images cover challenging weather conditions such as bright sunlight, fog, and rainfall. All images were annotated using LabelImg, a widely used open-source image annotation tool, ensuring consistency and efficiency. To ensure annotation consistency, a subset of the labeled images was cross-checked by multiple annotators, and discrepancies were resolved through consensus. The annotations strictly follow the YOLOv11 format, including normalized center coordinates (x, y) and relative width w and height h of each bounding box.

All images in the custom-target subset are collected and annotated by our team. The dataset is made publicly available to facilitate reproducibility and further studies. For the generalization subset, any publicly sourced images that might raise personal privacy or portrait rights concerns have been removed, and all data comply with applicable usage and licensing regulations.

Figure 8 shows representative images of the three specific UAV models from the custom subset, highlighting multi-angle and multi-distance variations. Figure 9 presents sample images from the generalization subset, illustrating environmental and visual diversity. Table 1 provides an overview of the dataset composition, including the number of training and testing images for each subset and the corresponding label formats.

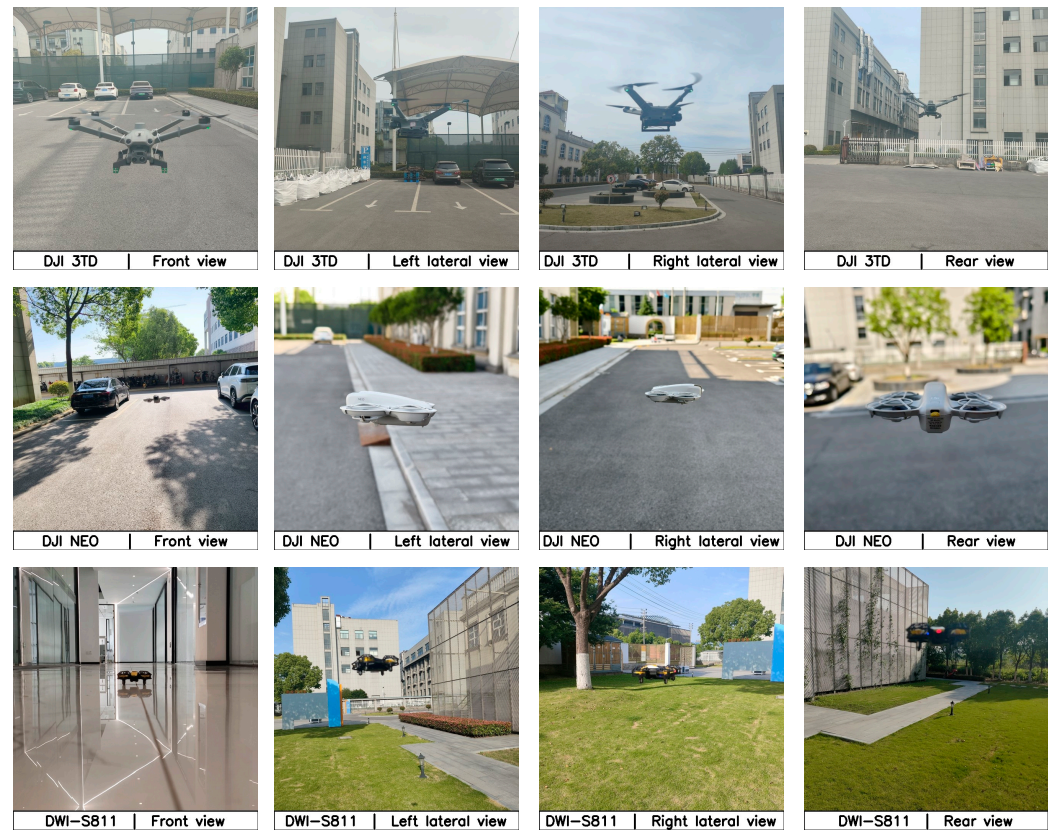


Figure 8. Typical Samples of Custom Subset UAVs.

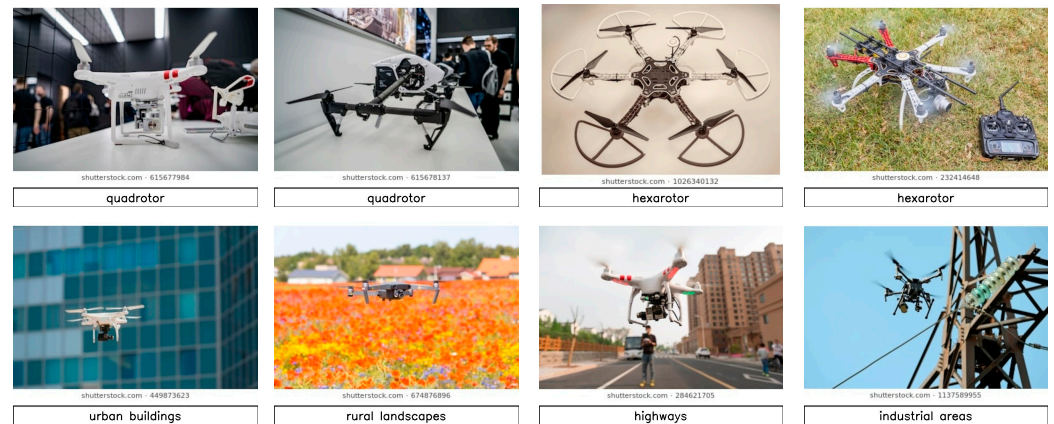


Figure 9. Samples of Multirotor UAVs in Generalization Subset.

Table 1. Overview of the UAV Detection Dataset Used in This Study.

Subset	Total Images	Training Set	Testing Set	Drone Models	Annotation Standard
Custom-Target (three kinds)	2670	2136	534	DJI 3TD/NEO/DWI-S811	drone1/drone2/dro-ne3 labels
Generalization	2664	2363	301	Multi-brand quad/hexa-rotor	drone label
Total	5334	4499	835	-	YOLO format

A differentiated sampling strategy was used to partition the dataset. The custom subset contains 2136 training images and 534 testing images, while the generalization subset includes 2363 training and 301 testing images. To enhance detection performance on

specific UAV types, the test set was intentionally supplemented with additional samples of DJI 3TD, DJI NEO, and DWI-S811. This allows the model to better learn and evaluate fine-grained appearance features of these UAVs, contributing to improved detection accuracy and robustness.

3.2. Experimental Environment and Experimental Parameters

To evaluate the performance of the proposed ADG-YOLO model in UAV object detection tasks, the model was trained on the custom dataset described in Section 3.1. Comparative experiments were conducted against several representative algorithms under identical training configurations. To ensure the reproducibility and fairness of the results, the experimental environment settings and training parameters are summarized in Tables 2 and 3, respectively.

Table 2. Configuration experimental environment.

Environment	Parameters
CPU	Intel (R) Xeon (R) Platinum 8358P
GPU	RTX 3090
GPU memory size	90 GB
Operating system	ubuntu18.04
Language	Python 3.8
Frame	PyTorch 1.8.1
CUDA version	CUDA 11.1

Table 3. Training parameters setting.

Parameters	Setup
Epochs	500
Input image size	640×640
Batch size	16
Optimizer	SGD
Initial learning rate	0.01

3.3. Evaluation Metrics

In target detection tasks, the Mean Average Precision (mAP) is widely employed to evaluate the detection performance of a model [55]. Based on the model's prediction results, two key metrics can be further computed: Precision (P) and Recall (R). Precision measures the proportion of correctly predicted targets among all samples identified as targets by the model, whereas Recall reflects the model's ability to detect actual targets, defined as the ratio of correctly detected targets to all true targets. Typically, there exists a trade-off between Precision and Recall, where improving one may lead to a reduction in the other. Therefore, a Precision–Recall (PR) curve is plotted to comprehensively analyze the detection performance of the model. For a single category, the Average Precision (AP) is defined as the area under the PR curve, which is calculated as follows:

$$AP = \int_0^1 P(R) dR \quad (11)$$

In practical computations, a discrete approximation method is typically employed:

$$AP = \frac{1}{m} \sum_{i=1}^m P(R_i) \quad (12)$$

Here, R_i represents the sampled recall values, and $P(R_i)$ denotes the corresponding precision at each recall point. The calculation of AP varies slightly across different datasets. For instance, the PASCAL VOC dataset adopts an interpolation method based on 11 fixed recall points, whereas the COCO evaluation protocol computes the mean over all recall points. For multi-class object detection, the Mean Average Precision (mAP) is defined as the mean AP across all categories:

$$mAP = \frac{1}{N} \sum_{i=1}^N AP_i \quad (13)$$

Here, N denotes the total number of categories, and AP_i represents the Average Precision of the i -th target category.

In practical object detection scenarios, in addition to model accuracy, the actual runtime speed of the model is also of significant concern. Frames Per Second (FPS) is a key metric for evaluating the runtime efficiency of a model, representing the number of image frames the model can process per second [56]. The FPS can be calculated as follows:

$$FPS = \frac{N_s}{T} \quad (14)$$

Here, N_s denotes the total number of processed frames, and T represents the total processing time in seconds. A higher FPS indicates that the model can process input images more rapidly, thereby enhancing its capability for real-time detection.

3.4. Ablation Study

To evaluate the effectiveness of the proposed lightweight modules, we conducted an ablation study based on YOLOv11n using the dataset described in Section 3.1, sequentially incorporating the C3Ghost and ADown structures, with the original network serving as the baseline. As shown in Table 4, the baseline model contains 2.58 M parameters and 6.3 GFLOPs, achieving 98.2% mAP with a power consumption of 5.71 W and an energy per frame of 0.2284 J. Introducing the C3Ghost module reduces the parameter size to 2.25 M with a slight increase in GFLOPs to 6.7, while slightly lowering power consumption to 5.67 W and energy per frame to 0.2181 J, indicating that C3Ghost can enhance efficiency without compromising accuracy. Incorporating the ADown module alone decreases both parameters (2.06 M) and computational cost (4.8 GFLOPs), yielding comparable accuracy (98.3% mAP) with reduced power consumption of 5.51 W and energy per frame of 0.2119 J. When both C3Ghost and ADown are integrated, the model achieves the best trade-off, reducing parameters to 1.77 M and GFLOPs to 5.7, while maintaining 98.4% mAP with 5.39 W power consumption and 0.1996 J energy per frame.

Table 4. Ablation experiment.

Baseline	C3Ghost	ADown	Params (M)	GFLOPs (G)	mAP _{0.5}	mAP _{0.5:0.95}	FPS	Power Consumption (W)	Energy Per Frame (J)
✓			2.58	6.3	98.2%	85.3%	25	5.71	0.2284
✓	✓		2.25	6.7	98.2%	85.2%	26	5.67	0.2181
✓		✓	2.06	4.8	98.3%	84.5%	26	5.51	0.2119
✓	✓	✓	1.77	5.7	98.4%	85.2%	27	5.39	0.1996

To further visualize the trade-offs between accuracy, computational cost, and power consumption, two Pareto plots are presented in Figure 10a,b; the vertical axis shows detection accuracy (mAP), and the color denotes power consumption in watts. Figure 10b

illustrates $mAP_{50:95}$ vs. GFLOPs, where the horizontal axis represents computational complexity (GFLOPs). Figure 10a presents $mAP_{50:95}$ vs. Params, with model parameter count on the horizontal axis, providing a direct view of the efficiency–accuracy trade-off in terms of model size. Both plots clearly demonstrate that the integration of C3Ghost and ADown substantially reduces model complexity and computational burden, increases efficiency, and maintains high detection accuracy, highlighting the superior performance–efficiency–power trade-off of the proposed ADG-YOLO.

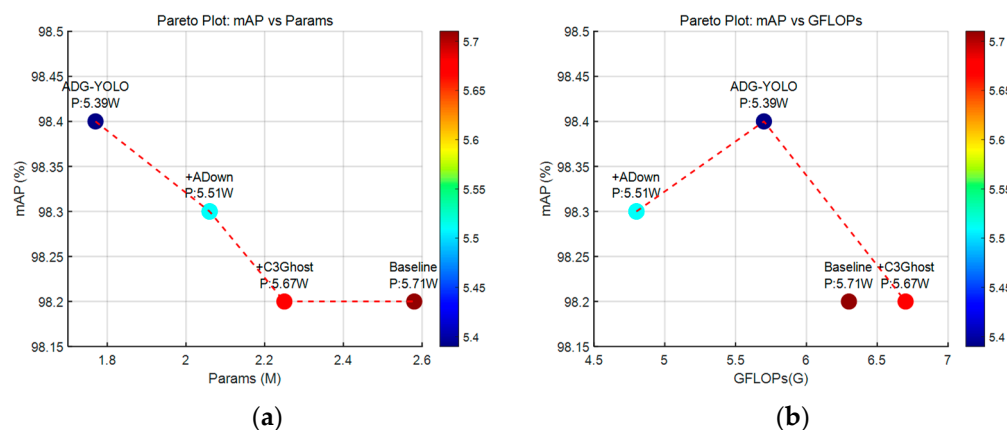


Figure 10. Detection Results of the ADG-YOLO Model. (a) $mAP_{0.5:0.95}$ vs. Params. (b) $mAP_{0.5:0.95}$ vs. GFLOPs.

In summary, the ablation study demonstrates that integrating C3Ghost and ADown effectively reduces model complexity and computational cost, lowers power consumption, minimizes energy per frame, and maintains high mAP, confirming the efficiency and robustness of ADG-YOLO.

3.5. Compration Exprimtent

To evaluate the overall performance of the ADG-YOLO model in UAV target detection, three mainstream lightweight models—YOLOv5s, YOLOv8n, and YOLOv11n—were selected as baseline comparisons. All models were converted to RKNN INT8 models using the RKNN Toolkit v2(Rockchip, Fuzhou, China), with quantization performed on a representative calibration dataset of 20 images randomly sampled from the validation set to ensure accurate weight scaling. During real-time inference on the Lubancat 4 development board, the input resolution was set to 640×640 . The power consumption measurements were conducted using the same KWS-X1 Type-C USB power meter (TGEINHVDU, Shenzhen, China) as described in the previous section, with the specific measurement scenario illustrated in Figure 11. The FPS, power consumption, and energy per frame reported in this study were measured under these quantized runtime conditions. These YOLO variants are chosen because they are currently among the most mature and high-performing single-stage detection methods, which facilitates a fair and consistent comparison across model parameters, computational complexity (GFLOPs), FPS, detection accuracy (mAP), power consumption (W), and energy per frame (J). The experimental results are summarized in Table 5.

ADG-YOLO contains only 1.77 M parameters and requires 5.7 GFLOPs, representing a substantial simplification compared to YOLOv5s, which has 7.02 M parameters and 15.8 GFLOPs. It is also more lightweight than YOLOv8n (3.00 M parameters, 8.1 GFLOPs) and YOLOv11n (2.58 M parameters, 6.3 GFLOPs), making it well-suited for deployment on resource-constrained edge platforms. In terms of detection accuracy, ADG-YOLO achieves 98.4% $mAP_{0.5}$ and 85.2% $mAP_{0.5:0.95}$, slightly outperforming the other models. Notably,

its $mAP_{0.5:0.95}$ is significantly higher than that of YOLOv5s and YOLOv8n (both 84.2%), and comparable to YOLOv11n (85.3%), demonstrating strong robustness. In addition, ADG-YOLO achieves a competitive inference speed of 27 FPS, with a power consumption of 5.39 W and an energy per frame of 0.1996 J, indicating an optimal balance between model compactness, computational efficiency, and energy efficiency for edge deployment.




Figure 11. Experimental setup for power consumption measurement.

Considering the full inference pipeline, including image capture, pre-processing, NPU inference, post-processing/Non-Maximum Suppression (NMS), and EKF-based monocular ranging, the end-to-end latency is expected to be within the frame interval of the camera (approximately 25–28 ms per frame for ADG-YOLO at 27 FPS). Given the camera’s actual frame rate of 35–40 FPS, the pipeline can process images in near real-time, indicating that the proposed system is suitable for practical UAV target detection on edge platforms.

To further evaluate the generalization capability of the ADG-YOLO model, we conducted a separate performance analysis on the generalization subset. The results are summarized in Table 6. Compared with the baseline models, ADG-YOLO consistently achieves higher mAP values (95.1% $mAP_{0.5}$ and 66.5% $mAP_{0.5:0.95}$), outperforming YOLOv5s (94.1% $mAP_{0.5}$, 62.8% $mAP_{0.5:0.95}$), YOLOv8n (94.5% $mAP_{0.5}$, 64.4% $mAP_{0.5:0.95}$), and YOLOv11n (94.5% $mAP_{0.5}$, 66.0% $mAP_{0.5:0.95}$). These results demonstrate that ADG-YOLO not only excels on the controlled subset but also generalizes effectively to unseen UAV types and diverse environments, confirming its robustness and practical applicability beyond the three specific UAV models. Moreover, ADG-YOLO maintains the most compact model size (1.77 M parameters) and the lowest computational complexity (5.7 GFLOPs) among the compared models, highlighting its suitability for deployment on resource-constrained platforms.


Table 5. Comparison of Lightweight Detection Models in Terms of Model Size, Accuracy, and Inference Speed (all kinds).

Model	Params (M)	GFLOPs (G)	$mAP_{0.5}$	$mAP_{0.5:0.95}$	FPS	Power Consumption (W)	Energy Per Frame (J)
YOLOv5s	7.02	15.8	98.2%	84.2%	15	3.96	0.264
YOLOv8n	3.00	8.1	98.2%	84.2%	28	5.78	0.206
YOLOv11n	2.58	6.3	98.2%	85.3%	25	5.71	0.2284
ADG-YOLO	1.77	5.7	98.4%	85.2%	27	5.39	0.1996

Table 6. Comparison of Lightweight Detection Models in Terms of Model Size, mAP (generalization subset).

Model	Params (M)	GFLOPs (G)	mAP _{0.5}	mAP _{0.5:0.95}
YOLOv5s	7.02	15.8	94.1%	62.8%
YOLOv8n	3.00	8.1	94.5%	64.4%
YOLOv11n	2.58	6.3	94.5%	66.0%
ADG-YOLO	1.77	5.7	95.1%	66.5%

To further compare the ADG-YOLO model with other mainstream lightweight models, we conducted additional experiments on the VisDrone dataset, considering Params (M), GFLOPs, FPS, power consumption (W), and energy per frame (J). The corresponding results are summarized in Table 7. As shown in the table, ADG-YOLO achieves the smallest model size with only 1.77 million parameters and the lowest computational cost of 5.7 GFLOPs, significantly lower than YOLOv5s (7.02 M, 15.8 G) and YOLOv8n (3.00 M, 8.1 G). In terms of inference speed, ADG-YOLO reaches 20 FPS, outperforming all compared models, including YOLOv11n (16 FPS). Moreover, ADG-YOLO maintains competitive detection accuracy with an mAP of 33.2% and a recall of 19.0%. Notably, ADG-YOLO exhibits superior energy efficiency, consuming only 0.1600 joules per frame—the lowest among all models—while operating at a power consumption of 3.2 watts. This highlights its advantage in resource-constrained UAV applications. In addition, PR curves for the ADG-YOLO model were generated, as shown in Figure 12, providing a quantitative evaluation of its detection performance and demonstrating its strong precision and recall. These results further confirm the robustness and practical applicability of ADG-YOLO in real-world UAV detection scenarios.

**Figure 12.** PR curves of ADG-YOLO on Visdrone dataset.**Table 7.** Comparison of Lightweight Detection Models in Terms of Model Size, mAP and power. (Visdrone dataset).

Model	Params (M)	GFLOPs (G)	mAP _{0.5}	mAP _{0.5:0.95}	FPS	Power Consumption (W)	Energy Per Frame (J)
YOLOv5s	7.02	15.8	34.3%	18.9%	12	3.4	0.2833
YOLOv8n	3.00	8.1	32.6%	19.0%	10	3.2	0.3200
YOLOv11n	2.58	6.3	34.2%	20.0%	16	3.4	0.2125
ADG-YOLO	1.77	5.7	33.2%	19.0%	20	3.2	0.1600

In summary, ADG-YOLO achieves an optimal trade-off among accuracy, model size, and resource consumption, making it particularly suitable for real-time UAV detection tasks in computationally constrained environments. The model also exhibits strong engineering adaptability for practical deployment.

4. Target Distance Estimation Experiment

4.1. Experimental Setup and Deployment Overview

To verify the accuracy of UAV altitude measurement, the ADG-YOLO model was converted into the .rknn format and deployed on the Lubancat 4 development board (EmbedFire, Dongguan, China). The target UAVs used in the distance measurement experiments were DJI 3TD and DJI NEO, both known for their flight stability. Visual data were captured using a Raspberry Pi USB camera (Zhongwei Aoke, Shenzhen, China) module, which was connected to the development board via a USB cable. Five lenses with focal lengths of 12 mm, 16 mm, 25 mm, 35 mm, and 50 mm were selected for distance measurement experiments at various ranges. The captured images and corresponding distance information were displayed on a YCXSQ-10 display screen (ZINCTUNG, Shenzhen, China), which features a 10-inch size and a resolution of 1920×1080 pixels. The camera was mounted on a tripod, while the display screen was connected to the development board via an HDMI cable for real-time visualization of detection results and distance measurements.

Figure 13 illustrates the experimental setup: (a) the distance measurement platform, (b) the lenses used in this experiment, and (c) the UAVs employed for the experiment.




Figure 13. Experimental Setup. (a) Distance Measurement Platform. (b) The lens used in this experiment. (c) UAVs used for Experiment.

To further validate the model beyond laboratory conditions, experiments were conducted both indoors and outdoors. The DJI 3TD, with strong wind resistance, was used for outdoor experiments in an open area with buildings, trees, and pavement as the background, under sufficient natural lighting. During testing, the UAV was manually flown along a straight path aligned with the camera's optical axis, maintaining a level attitude to ensure stable visual features and reduce interference from yaw or pitch. The DJI NEO, due to its smaller size and lower wind tolerance, was tested indoors in a closed room with stable lighting and minimal airflow, providing favorable conditions for high-precision distance calibration.

The Lubancat 4 development board runs the ADG-YOLO model with an estimated processing time of ~ 37 ms per frame (corresponding to ~ 27 FPS) and a power consumption of approximately 5 W. This setup demonstrates the feasibility of deploying ADG-YOLO in real-world scenarios with both indoor and outdoor UAV tests, while providing high-speed and energy-efficient performance.

4.2. Distance Measurement for UAV Targets

In our UAV distance measurement experiments, two drone models were selected as test targets: the DJI 3TD and the DJI NEO, with rotor spans of 62 cm and 15 cm, respectively. To assess the adaptability of the proposed measurement method across different UAV sizes and operational environments, each model was tested under distinct conditions.

The DJI 3TD, featuring strong wind resistance, was used for outdoor experiments. During testing, the UAV was manually flown along a straight path aligned with the camera's optical axis, maintaining a level attitude to ensure stable visual features and reduce interference from yaw or pitch. As shown in Figure 14a, the test was conducted in an open outdoor area, where a standard measuring tape was laid along the flight path to mark reference distance points. The camera system was fixed on a stationary tripod, and images were captured at each distance for subsequent evaluation.

In contrast, the DJI NEO, due to its smaller size and lower wind tolerance, was tested indoors to ensure stable hovering. As shown in Figure 14b, the indoor experiment was conducted in a closed room, with the measuring tape placed along a straight line. The UAV hovered at various predefined points to collect image samples at known distances. The controlled indoor environment—with stable lighting and minimal airflow—provided favorable conditions for high-precision distance calibration.

(a) DJI 3TD



(b) DJI NEO



Figure 14. Distance Measurement Experimental Scene.

In this study, distance estimation was performed using the principle of similar triangles. Based on the known physical width of the UAV and the width of the corresponding bounding box output by the detection model, the distance between the UAV and the camera was calculated. All estimated distances were compared with ground-truth values obtained from physical measurements using a tape measure. This comparison was conducted to evaluate the effectiveness and stability of the proposed distance estimation method under real-world conditions.

To further quantify the accuracy of the system in UAV distance estimation, relative error was introduced as a performance evaluation metric. By calculating the ratio between the prediction error and the ground-truth distance, this metric reflects the overall precision

of the proposed distance measurement method. The formula for computing the relative error e_{mea} is given in Equation (11) [57].

$$e_{mea} = \frac{|d_{real} - d_{mea}|}{d_{real}} \times 100\% \quad (15)$$

The distance estimation results for the DJI NEO are summarized in Table 8, obtained using a 12 mm lens. The experimental results for the DJI 3TD are presented in Table 9, based on tests conducted with five different lens focal lengths. In addition, representative experimental images are provided to visually demonstrate the distance measurement process and outcomes—Figure 15 shows the measurement setup for the DJI NEO, while Figure 16 presents the measurement setup for the DJI 3TD.




Figure 15. Distance Measurement Scene of DJI NEO.

The UAV distance estimation method based on the principle of similar triangles demonstrated excellent performance in real-world scenarios. As shown in Tables 8 and 9, the DJI NEO achieved an average relative error of 4.18% across 10 test cases within the range of 0.5 to 5 m. For the DJI 3TD, a total of 45 measurements across various focal lengths and distances ranging from 2 to 50 m resulted in a combined average relative error of only 2.40%. High accuracy was maintained across all test distances, with the maximum error not exceeding 12.33%. Notably, even at a distance of 50 m, the method achieved a minimal error of 0.26%, further validating the effectiveness and stability of the proposed approach.

Table 8. Distance Measurement of DJI NEO ($f = 12$ mm).

Experiment ID	d_{real} (m)	d_{mea} (m)	e_{mea}
1	0.5	0.52	4%
2	1	0.95	5%
3	1.5	1.64	9.3%
4	2	2.13	6.5%

Table 8. Cont.

Experiment ID	$d_{\text{real}}(\text{m})$	$d_{\text{mea}}(\text{m})$	e_{mea}
5	2.5	2.36	5.6%
6	3	3.05	1.7%
7	3.5	3.61	3.1%
8	4	3.98	0.5%
9	4.5	4.36	3.1%
10	5	4.90	2%



Figure 16. Distance Measurement Scene of DJI 3TD.

Table 9. Distance Measurement of DJI 3TD.

Experiment ID	Focal Length (mm)	$d_{\text{real}}(\text{m})$	$d_{\text{mea}}(\text{m})$	e_{mea}
1	12	2	2.15	7.50%
2	12	3	3.37	12.33%
3	12	4	4.28	7.00%
4	12	5	4.99	0.20%
5	12	6	5.85	2.50%
6	12	7	7.58	8.29%

Table 9. Cont.

Experiment ID	Focal Length (mm)	d_{real} (m)	d_{mea} (m)	e_{mea}
7	12	8	8.07	0.88%
8	12	9	8.81	2.11%
9	16	11	10.59	3.73%
10	16	12	11.62	3.17%
11	16	13	13.01	0.08%
12	16	14	13.73	1.93%
13	16	15	14.79	1.40%
14	16	16	16.35	2.19%
15	16	17	16.87	0.76%
16	16	18	18.02	0.11%
17	16	19	18.33	3.53%
18	25	21	21.06	0.29%
19	25	22	22.36	1.64%
20	25	23	22.59	1.78%
21	25	24	23.66	1.42%
22	25	25	25.39	1.56%
23	25	26	26.78	3.00%
24	25	27	27.66	2.44%
25	25	28	28.21	0.75%
26	25	29	29.33	1.14%
27	35	31	30.52	1.55%
28	35	32	31.78	0.69%
29	35	33	33.26	0.79%
30	35	34	33.70	0.88%
31	35	35	34.96	0.11%
32	35	36	35.78	0.61%
33	35	37	38.21	3.27%
34	35	38	38.62	1.63%
35	35	39	39.45	1.15%
36	50	41	40.73	0.66%
37	50	42	42.56	1.33%
38	50	43	43.34	0.79%
39	50	44	43.97	0.07%
40	50	45	45.26	0.58%
41	50	46	45.69	0.67%
42	50	47	46.91	0.19%

5. Discussion

The proposed ADG-YOLO framework demonstrates significant advancements in real-time UAV detection and distance estimation on edge devices. Nonetheless, several challenges remain that warrant further investigation to enhance its scalability and real-world applicability.

First, although the current custom dataset (5343 images) includes three UAV models across diverse backgrounds, its limited scope constrains the generalization of the proposed framework. The dataset size was inherently restricted by practical limitations, including the availability of UAV models, time, and funding, which prevented the acquisition of a larger and more diverse collection. Despite these constraints, the dataset maintains significant diversity through multi-angle, multi-distance, and multi-condition image acquisition, as well as supplementation from publicly available UAV resources. Experimental results demonstrate that the model achieves robust detection performance across both the custom-target and generalization subsets, indicating that the dataset is sufficient to support the objectives of the current study. Nevertheless, future work should focus on building

a large-scale, open-source UAV dataset covering various drone types (e.g., quadrotors, hexarotors, fixed-wing), sizes (micro to commercial), and environmental conditions (e.g., night, adverse weather, swarm operations), particularly under low-SNR settings prone to false positives. Collaborative data collection across platforms may further accelerate this expansion and improve the model's generalization capability. Compared with large-scale UAV benchmarks such as UAVDT, which contain over 100,000 images, the current dataset is relatively limited. Future efforts will focus on collaborative data collection across institutions and open-source release to further enhance the dataset's diversity and support broader research reproducibility.

Second, current distance estimation depends on known UAV dimensions (e.g., DJI 3TD: 62 cm, DJI NEO: 15 cm), which limits its flexibility in handling unknown models. Future research should explore multi-model support through an onboard UAV identification module containing pre-calibrated physical parameters. The feasibility of the pre-calibrated identification module is supported by the possibility to store physical parameters of known UAV models onboard, enabling rapid retrieval during detection without significant computational overhead. Additionally, geometry-independent approaches, such as monocular depth estimation fused with detection outputs, offer promising alternatives that remove dependency on prior shape knowledge. For instance, monocular depth estimation can be combined with multi-scale feature fusion and uncertainty-aware refinement to reduce errors arising from perspective distortion, particularly in high-altitude scenarios. Although the current controlled dataset covers target distances from 5 to 30 m, real-world UAV applications such as aerial surveillance or infrastructure monitoring often involve higher flight height exceeding 50 m. Under such scenarios, monocular distance estimation may experience increased errors due to reduced image resolution and perspective distortion. To mitigate these challenges, adaptive focal length calibration, multi-scale training strategies, or fusion with onboard UAV identification modules and monocular depth estimation could be adopted. These enhancements are expected to improve model robustness and extend its applicability to high-altitude UAV operations.

Third, while ADG-YOLO achieves 27 FPS on the Lubancat4 edge device, its practical deployment on UAVs introduces additional challenges. These include optimizing the model for ultra-low-power processors, ensuring efficient thermal dissipation during extended operation, and compensating for dynamic motion via IMU and EKF integration to stabilize detection during rapid pitch or yaw movements. Moreover, expanding the system to air-to-air detection, such as in drone swarm environments, requires altitude-invariant ranging models and training strategies that are robust to occlusion. Furthermore, evaluating energy consumption and latency across different edge processors can guide optimization for practical UAV deployment. Strategies such as occlusion-aware training and data augmentation are expected to improve swarm detection robustness under challenging scenarios. Finally, achieving sub-20 ms end-to-end latency is essential for enabling closed-loop tasks such as autonomous interception and cooperative formation flight.

Building on the above discussion of ADG-YOLO's performance, limitations, and potential improvements, it is also informative to consider alternative UAV detection modalities, such as LiDAR- and radar-based systems. While LiDAR and radar provide advantages in ranging accuracy, wider field-of-view (FoV), and robustness to environmental conditions, they typically incur higher power consumption, increased weight, and elevated cost, which may limit their deployment on small UAV platforms. In contrast, computer vision (CV)-based approaches, exemplified by ADG-YOLO, offer lightweight, low-power solutions well-suited for onboard UAV integration, albeit with trade-offs in maximum detection range and sensitivity under low-SNR or occluded scenarios. Importantly, these approaches are not mutually exclusive; rather, they can be complementary. Hybrid sys-

tems that integrate CV, LiDAR, and radar could leverage the strengths of each modality, enabling more robust, flexible, and energy-efficient UAV perception. Our work focuses on CV for edge-deployed, real-time UAV detection, which aligns with rather than conflicts with LiDAR- or radar-focused studies, and provides a foundation for future multi-sensor collaborative frameworks.

6. Conclusions

This study proposes ADG-YOLO, a lightweight and efficient framework for real-time UAV target detection and distance estimation on edge devices. The framework integrates multiple key innovations: (1) a computationally optimized architecture that incorporates C3Ghost modules and ADown layers, reducing model parameters to 1.77 M and GFLOPs to 5.7, while maintaining high detection accuracy with 98.4% mAP_{0.5}; (2) an EKF-based tracking mechanism that significantly improves detection stability in dynamic environments; (3) a monocular distance estimation method based on similarity triangle theory, which achieves average relative errors ranging from 2.40% to 4.18% over distances of 0.5–50 m; (4) successful real-time deployment on the Lubancat4 edge platform (RK3588S NPU) at 27 FPS, demonstrating its practical applicability in resource-constrained settings.

Overall, ADG-YOLO effectively balances detection accuracy and computational efficiency, bridging the gap between advanced perception and edge deployment for UAV-based applications. Future work will focus on expanding large-scale UAV datasets, enabling generalized ranging for unknown UAV models, and facilitating deployment in autonomous aerial systems to support next-generation capabilities in both military and commercial UAV operations. In addition, the proposed framework holds strong potential for broader application domains, including environmental monitoring, precision agriculture, infrastructure inspection, urban air traffic management, and search and rescue, as well as security and defense operations. These directions will further enhance the scalability and generalization of ADG-YOLO in real-world scenarios.

Author Contributions: Conceptualization, H.W. and Z.D.; methodology, H.W.; software, M.C.; validation, H.S., Y.Q., Z.D. and H.Y.; formal analysis, J.Z.; investigation, D.W.; resources, H.W.; data curation, Z.D.; writing—original draft preparation, Z.D.; writing—review and editing, H.W.; visualization, M.C.; supervision, H.Y.; project administration, H.Y.; funding acquisition, H.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All experimental data analyzed in this study are included in this published article. The neural framework for model training and the datasets used can be accessed via the GitHub repository: <https://github.com/bigdang123/YOLOV11-UAV-detection.git> (accessed on 6 October 2025). No additional external datasets were utilized.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Lin, Y.-H.; Joubert, D.A.; Kaeser, S.; Dowd, C.; Germann, J.; Khalid, A.; Denton, J.A.; Retski, K.; Tavui, A.; Simmons, C.P.; et al. Field Deployment of Wolbachia-Infected Aedes Aegypti Using Uncrewed Aerial Vehicle. *Sci. Robot.* **2024**, *9*, eadk7913. [[CrossRef](#)] [[PubMed](#)]
2. Zhou, X.; Wen, X.; Wang, Z.; Gao, Y.; Li, H.; Wang, Q.; Yang, T.; Lu, H.; Cao, Y.; Xu, C.; et al. Swarm of Micro Flying Robots in the Wild. *Sci. Robot.* **2022**, *7*, eabm5954. [[CrossRef](#)]
3. Han, J.; Yan, Y.; Zhang, B. Towards Efficient Multi-UAV Air Combat: An Intention Inference and Sparse Transmission Based Multi-Agent Reinforcement Learning Algorithm. *IEEE Trans. Artif. Intell.* **2025**, *1*, 1–12. [[CrossRef](#)]
4. Karimov, C.Y. The Role of Unmanned Aircraft Vehicles in the Russian-Ukrainian War. *Endless Light Sci.* **2025**, 83–89. Available online: <https://cyberleninka.ru/article/n/the-role-of-unmanned-aircraft-vehicles-in-the-russian-ukrainian-war> (accessed on 13 August 2025).

5. Wennerholm, D. Above the Trenches: Russian Military Lessons Learned about Drone Warfare from Ukraine. Master's Thesis, Uppsala University, Uppsala, Sweden, 2025.
6. Tang, Z.; Ma, H.; Qu, Y.; Mao, X. UAV Detection with Passive Radar: Algorithms, Applications, and Challenges. *Drones* **2025**, *9*, 76. [\[CrossRef\]](#)
7. Abir, T.A.; Kuantama, E.; Han, R.; Dawes, J.; Mildren, R.; Nguyen, P. Towards Robust Lidar-Based 3D Detection and Tracking of UAVs. In Proceedings of the Ninth Workshop on Micro Aerial Vehicle Networks, Systems, and Applications, Helsinki, Finland, 18 June 2023; pp. 1–7.
8. Qiu, Z.; Lu, Y.; Qiu, Z. Review of Ultrasonic Ranging Methods and Their Current Challenges. *Micromachines* **2022**, *13*, 520. [\[CrossRef\]](#)
9. Rahmaniar, W.; Wang, W.-J.; Caesarendra, W.; Glowacz, A.; Oprzędkiewicz, K.; Sułowicz, M.; Irfan, M. Distance Measurement of Unmanned Aerial Vehicles Using Vision-Based Systems in Unknown Environments. *Electronics* **2021**, *10*, 1647. [\[CrossRef\]](#)
10. Tian, X.; Liu, R.; Wang, Z.; Ma, J. High Quality 3D Reconstruction Based on Fusion of Polarization Imaging and Binocular Stereo Vision. *Inf. Fusion* **2022**, *77*, 19–28. [\[CrossRef\]](#)
11. Tang, Y.; Zhou, H.; Wang, H.; Zhang, Y. Fruit Detection and Positioning Technology for a Camellia Oleifera C. Abel Orchard Based on Improved YOLOv4-Tiny Model and Binocular Stereo Vision. *Expert Syst. Appl.* **2023**, *211*, 118573. [\[CrossRef\]](#)
12. Bao, D.; Wang, P. Vehicle Distance Detection Based on Monocular Vision. In Proceedings of the 2016 IEEE International Conference on Progress in Informatics and Computing (PIC), Shanghai, China, 23–25 December 2016; pp. 187–191.
13. Ali, A.A.; Hussein, H.A. Distance Estimation and Vehicle Position Detection Based on Monocular Camera. In Proceedings of the 2016 IEEE Al-Sadeq International Conference on Multidisciplinary in IT and Communication Science and Applications (AIC-MITCSA), Baghdad, Iraq, 9–10 May 2016; pp. 1–4.
14. Liu, F.; Shen, C.; Lin, G. Deep Convolutional Neural Fields for Depth Estimation from a Single Image. In Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, 7–12 June 2015; pp. 5162–5170.
15. Li, J.; Klein, R.; Yao, A. A Two-Streamed Network for Estimating Fine-Scaled Depth Maps from Single RGB Images. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 3372–3380.
16. Eigen, D.; Fergus, R. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. In Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 7–13 December 2015; pp. 2650–2658.
17. Eigen, D.; Puhrsch, C.; Fergus, R. Depth Map Prediction from a Single Image Using a Multi-Scale Deep Network. *Adv. Neural Inf. Process. Syst.* **2014**, *27*.
18. Jiao, J.; Cao, Y.; Song, Y.; Lau, R. Look Deeper into Depth: Monocular Depth Estimation with Semantic Booster and Attention-Driven Loss. In *Computer Vision—ECCV 2018*; Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2018; Volume 11219, pp. 55–71, ISBN 978-3-030-01266-3.
19. Zhe, T.; Huang, L.; Wu, Q.; Zhang, J.; Pei, C.; Li, L. Inter-Vehicle Distance Estimation Method Based on Monocular Vision Using 3D Detection. *IEEE Trans. Veh. Technol.* **2020**, *69*, 4907–4919. [\[CrossRef\]](#)
20. Mallot, H.A.; Bühlhoff, H.H.; Little, J.J.; Bohrer, S. Inverse Perspective Mapping Simplifies Optical Flow Computation and Obstacle Detection. *Biol. Cybern.* **1991**, *64*, 177–185. [\[CrossRef\]](#)
21. Tuohy, S.; O'Cualain, D.; Jones, E.; Glavin, M. Distance Determination for an Automobile Environment Using Inverse Perspective Mapping in OpenCV. In Proceedings of the IET Irish Signals and Systems Conference (ISSC 2010), Cork, Ireland, 23–24 June 2010; pp. 100–105.
22. Wongsaree, P.; Sinchai, S.; Wardkein, P.; Koseeyaporn, J. Distance Detection Technique Using Enhancing Inverse Perspective Mapping. In Proceedings of the 2018 IEEE 3rd International Conference on Computer and Communication Systems (ICCCS), Nagoya, Japan, 27–30 April 2018; pp. 217–221.
23. Huang, L.; Zhe, T.; Wu, J.; Wu, Q.; Pei, C.; Chen, D. Robust Inter-Vehicle Distance Estimation Method Based on Monocular Vision. *IEEE Access* **2019**, *7*, 46059–46070. [\[CrossRef\]](#)
24. Qi, S.H.; Li, J.; Sun, Z.P.; Zhang, J.T.; Sun, Y. Distance Estimation of Monocular Based on Vehicle Pose Information. *J. Phys. Conf. Ser.* **2019**, *1168*, 032040. [\[CrossRef\]](#)
25. Jiafa, M.; Wei, H.; Weiguo, S. Target Distance Measurement Method Using Monocular Vision. *IET Image Process.* **2020**, *14*, 3181–3187. [\[CrossRef\]](#)
26. Yang, R.; Yu, S.; Yao, Q.; Huang, J.; Ya, F. Vehicle Distance Measurement Method of Two-Way Two-Lane Roads Based on Monocular Vision. *Appl. Sci.* **2023**, *13*, 3468. [\[CrossRef\]](#)
27. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, 23–28 June 2014; pp. 580–587.
28. Girshick, R. Fast R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 7–13 December 2015; pp. 1440–1448.

29. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *39*, 1137–1149. [[CrossRef](#)]
30. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. In Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 779–788.
31. Redmon, J.; Farhadi, A. YOLO9000: Better, Faster, Stronger. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 6517–6525.
32. Redmon, J.; Farhadi, A. YOLOv3: An Incremental Improvement. *arXiv* **2018**, arXiv:1804.02767. [[CrossRef](#)]
33. Bochkovskiy, A.; Wang, C.Y.; Liao, H.Y.M. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv* **2020**, arXiv:2004.10934. [[CrossRef](#)]
34. Zhang, Y.; Guo, Z.; Wu, J.; Tian, Y.; Tang, H.; Guo, X. Real-Time Vehicle Detection Based on Improved YOLO v5. *Sustainability* **2022**, *14*, 12274. [[CrossRef](#)]
35. Li, C.; Li, L.; Jiang, H.; Weng, K.; Geng, Y.; Li, L.; Ke, Z.; Li, Q.; Cheng, M.; Nie, W.; et al. YOLOv6: A Single-Stage Object Detection Framework for Industrial Applications. *arXiv* **2022**, arXiv:2209.02976. [[CrossRef](#)]
36. Wang, C.Y.; Bochkovskiy, A.; Liao, H.Y.M. YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 17–24 June 2023; pp. 7464–7475. [[CrossRef](#)]
37. Jacob, I.J.; Piramuthu, S.; Falkowski-Gilski, P. (Eds.) *Data Intelligence and Cognitive Informatics: Proceedings of ICDICI 2023; Algorithms for Intelligent Systems*; Springer Nature: Singapore, 2024; ISBN 978-981-99-7999-8.
38. Wang, C.Y.; Yeh, I.H.; Liao, H.Y.M. YOLOv9: Learning What You Want to Learn Using Programmable Gradient Information. In *Computer Vision—ECCV 2024, Proceedings of the European Conference on Computer Vision, Milan, Italy, 29 September–4 October 2024*; Springer Nature: Cham, Switzerland, 2024; pp. 1–21.
39. Wang, A.; Chen, H.; Liu, L.; Chen, K.; Lin, Z.; Han, J.; Ding, G. YOLOv10: Real-Time End-to-End Object Detection. *Adv. Neural Inf. Process. Syst.* **2024**, *37*, 107984–108011.
40. Khanam, R.; Hussain, M. YOLOv11: An Overview of the Key Architectural Enhancements. *arXiv* **2024**, arXiv:2410.17725. [[CrossRef](#)]
41. Tian, Y.; Ye, Q.; Doermann, D. YOLOv12: Attention-Centric Real-Time Object Detectors. *arXiv* **2025**, arXiv:2502.12524.
42. Cheng, Q.; Wang, Y.; He, W.; Bai, Y. Lightweight Air-to-Air Unmanned Aerial Vehicle Target Detection Model. *Sci. Rep.* **2024**, *14*, 2609. [[CrossRef](#)] [[PubMed](#)]
43. Su, J.; Qin, Y.; Jia, Z.; Liang, B. MPE-YOLO: Enhanced Small Target Detection in Aerial Imaging. *Sci. Rep.* **2024**, *14*, 17799. [[CrossRef](#)]
44. Wang, C.; Han, Y.; Yang, C.; Wu, M.; Chen, Z.; Yun, L.; Jin, X. CF-YOLO for Small Target Detection in Drone Imagery Based on YOLOv11 Algorithm. *Sci. Rep.* **2025**, *15*, 16741. [[CrossRef](#)]
45. Zhou, S.; Yang, L.; Liu, H.; Zhou, C.; Liu, J.; Wang, Y.; Zhao, S.; Wang, K. Improved YOLO for Long Range Detection of Small Drones. *Sci. Rep.* **2025**, *15*, 12280. [[CrossRef](#)]
46. Kanjalkar, P.; Kinhikar, S.; Zagade, A.; Rane, S.; Kanjalkar, J. Intelligent Surveillance Tower for Detection of the Drone from the Other Aerial Objects Using Deep Learning. In *Intelligent Systems for Smart Cities, Proceedings of the International Conference on Information Science and Applications (ICISA 2023), Singapore, 23–25 May 2023*; Springer Nature: Singapore, 2023; pp. 39–51.
47. Laroca, R.; Santos, M.D.; Menotti, D. Improving Small Drone Detection through Multi-Scale Processing and Data Augmentation. *arXiv* **2025**, arXiv:2504.19347. [[CrossRef](#)]
48. Dadboud, F.; Patel, V.; Mehta, V.; Bolic, M.; Mantegh, I. Single-Stage UAV Detection and Classification with YOLOV5: Mosaic Data Augmentation and PANet. In Proceedings of the 2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), Washington, DC, USA, 16 November 2021; pp. 1–8.
49. Arezoomandan, S.; Klohoker, J.; Han, D.K. Data Augmentation Pipeline for Enhanced UAV Surveillance. In *Pattern Recognition, Proceedings of the International Conference on Pattern Recognition (ICPR 2025), Kolkata, India, 1–5 December 2024*; Springer: Cham, Switzerland, 2025; pp. 366–380.
50. Gishyan, K. Improving UAV Object Detection through Image Augmentation. *Math. Probl. Comput. Sci.* **2020**, *54*, 53–68. [[CrossRef](#)]
51. Kim, B.H.; Khan, D.; Bohak, C.; Choi, W.; Lee, H.J.; Kim, M.Y. V-RBNN Based Small Drone Detection in Augmented Datasets for 3D LADAR System. *Sensors* **2018**, *18*, 3825. [[CrossRef](#)]
52. Han, K.; Wang, Y.; Xu, C.; Guo, J.; Xu, C.; Wu, E.; Tian, Q. GhostNets on Heterogeneous Devices via Cheap Operations. *Int. J. Comput. Vis.* **2022**, *130*, 1050–1069. [[CrossRef](#)]
53. Ji, C.L.; Yu, T.; Gao, P.; Wang, F.; Yuan, R.Y. YOLO-TLA: An Efficient and Lightweight Small Object Detection Model Based on YOLOv5. *J. Real-Time Image Process.* **2024**, *21*, 141. [[CrossRef](#)]
54. Fang, S.; Chen, C.; Li, Z.; Zhou, M.; Wei, R. YOLO-ADual: A Lightweight Traffic Sign Detection Model for a Mobile Driving System. *World Electr. Veh. J.* **2024**, *15*, 323. [[CrossRef](#)]

55. Tian, Z.; Chu, X.; Wang, X.; Wei, X.; Shen, C. Fully Convolutional One-Stage 3D Object Detection on LiDAR Range Images. In *Advances in Neural Information Processing Systems (NeurIPS 2022), Proceedings of the 36th International Conference on Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022*; Neural Information Processing Systems Foundation Inc.: San Diego, CA, USA, 2022; Volume 35, pp. 34899–34911.
56. Liu, S.; Huang, D.; Wang, Y. Learning Spatial Fusion for Single-Shot Object Detection. *arXiv* **2019**, arXiv:1911.09516. [[CrossRef](#)]
57. Ky, H.H. Notes on the Use of Propagation of Error Formulas. *J. Res. Natl. Bur. Stand. C Eng. Instrum.* **1966**, *70*, 263.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.