

MODA: The First Challenging Benchmark for Multispectral Object Detection in Aerial Images

Shuaihao Han, Tingfa Xu*, Peifu Liu, Jianan Li*

Beijing Institute of Technology

{3120230561, ciom_xtf1}@bit.edu.cn, bitlpf@163.com, lijianan15@gmail.com

Abstract

Aerial object detection faces significant challenges in real-world scenarios, such as small objects and extensive background interference, which limit the performance of RGB-based detectors with insufficient discriminative information. Multispectral images (MSIs) capture additional spectral cues across multiple bands, offering a promising alternative. However, the lack of training data has been the primary bottleneck to exploiting the potential of MSIs. To address this gap, we introduce the first large-scale dataset for Multispectral Object Detection in Aerial images (MODA), which comprises 14,041 MSIs and 330,191 annotations across diverse, challenging scenarios, providing a comprehensive data foundation for this field. Furthermore, to overcome challenges inherent to aerial object detection using MSIs, we propose OSSDet, a framework that integrates spectral and spatial information with object-aware cues. OSSDet employs a cascaded spectral-spatial modulation structure to optimize target perception, aggregates spectrally related features by exploiting spectral similarities to reinforce intra-object correlations, and suppresses irrelevant background via object-aware masking. Moreover, cross-spectral attention further refines object-related representations under explicit object-aware guidance. Extensive experiments demonstrate that OSSDet outperforms existing methods with comparable parameters and efficiency.

Datasets — <https://github.com/shuaihao-han/MODA>

Introduction

Aerial object detection has become a prominent research focus in recent years (Zhang et al. 2022b; Xie et al. 2021), yet remains challenging in practice due to small targets and strong background noise (Leng et al. 2024), which obscure target features and degrade detection performance. This issue is particularly pronounced for RGB-based detectors that rely heavily on spatial cues (Fig. 1).

Multispectral images (MSIs) capture additional spectral cues that characterize the target’s intrinsic reflectance properties (Fang et al. 2023b), offering valuable insights in cases where spatial features are limited by aforementioned challenges (Fang et al. 2023a). For example, distinct spectral curves reliably differentiate targets from background, even

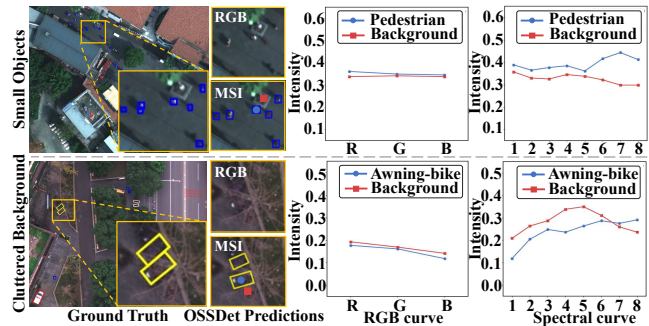


Figure 1: In challenging scenarios, limited spatial information in RGB data hampers effective detection. In contrast, multispectral images offer additional spectral cues that significantly enhance target discrimination for robust detection.

in cases involving small objects or cluttered scenes (Fig. 1). Moreover, critical spectral features generally remain stable despite variations in object appearance (Qin et al. 2024), offering robust cues for detection. Consequently, MSIs represent a promising approach for object detection in challenging aerial images. However, the limited availability of large-scale datasets hampers further advancement.

To bridge this gap, we present MODA, the first large-scale challenging dataset for multispectral aerial object detection, containing 14,041 MSIs with 330,191 oriented annotations across 8 categories. Each MSI features a large image size of 1200×900 and covers 8 spectral bands (395~950nm). The data is captured in diverse urban scenes over various time periods. Additionally, MODA includes 8 challenging attributes commonly encountered in real-world scenarios (e.g., small objects, low illumination, truncation, occlusion). This comprehensive design renders MODA well-suited for practical applications and establishes it as a valuable resource for advancing research in multispectral aerial object detection.

In terms of model design, leveraging spectral-spatial information of MSI for aerial object detection encounters three key challenges: (i) prior multispectral object detection (MOD) methods decouple spectral and spatial information via dimensionality reduction such as PCA (Pearson 1901) and band selection (Zhang et al. 2024) along with two-stream network for independent processing, incur high

*Correspondence to: Tingfa Xu and Jianan Li.

computational costs and spectral information loss; (ii) existing aerial object detection methods primarily exploiting spatial features (Xiao et al. 2024), and spectral cues is under-utilized; (iii) extensive background interference in complex scenarios dilute distinctive target features (Zhu et al. 2021), impairs model attention and object spectral-spatial feature learning, degrade detector performance.

To address these issues, we propose OSSDet, an object-aware spectral-spatial learning framework for multispectral aerial object detection. OSSDet integrates spectral and spatial cues in a unified object-aware design, eliminating explicit decoupling. It employs a cascaded spectral-spatial awareness interactive modulation structure to optimize target perception and a Spectral-Guided Adaptive Cross-Layer Fusion module that adaptively aggregates spectral-related features to reinforce intra-object correlations, resisting noise interference while preserving spatial texture details to minimize information loss during cross-layer fusion.

Moreover, to mitigate significant background noise that dilutes target features, especially for small objects, we introduce an object-aware masking branch. This branch generates an object activation mask from low-level features using an object activation loss as guidance. As an object-aware prior, the mask refines these features by preserving object-specific information while suppressing irrelevant noise. Subsequently, features with explicit object-aware cues are integrated with cross-spectral attention to bolster object-related representations in both spatial and spectral domains.

Extensive experiments on MODA and other MOD datasets demonstrate that OSSDet achieves state-of-the-art performance while preserving computational efficiency. The primary contributions of this work are summarized below:

- We introduce the first large-scale dataset for multispectral aerial object detection, featuring diverse challenges to support research advancements in this domain.
- We propose OSSDet, a novel framework that integrates spectral and spatial information with object-aware cues for robust detection in aerial scenarios using MSI.
- We develop cascaded spectral-spatial awareness modulation to enhance target perception, spectral aggregation to reinforce intra-object correlations, and object-aware masking branch to reduce noise interference, optimizing MSI utilization and improving detection accuracy.

Related Work

Multispectral Object Detection Datasets. Growing interest in multispectral object detection has led to the introduction of several datasets. HOD-1 (Yan et al. 2021) is the first bounding box-annotated MSI dataset, followed by HOD3K (He et al. 2023), which captures MSIs in natural scenes. However, HOD-1 includes manually placed objects with simple fore-backgrounds, and HOD3K relies on fixed camera positions, resulting in static backgrounds that lack real-world dynamics. To address these limitations, we propose MODA, a large-scale dataset with diverse real-world challenges for advancing research in multispectral aerial object detection. To our best knowledge, MODA is the largest multispectral object detection dataset available.

Aerial Object Detection. Many existing methods address the boundary discontinuity issue in oriented object angle prediction (Yang et al. 2021b; Xu et al. 2024), improve sample quality and allocation (Li et al. 2022), or enhance object representations (Han et al. 2021; Li et al. 2024; Yuan et al. 2025). However, these approaches are confined to RGB images, overlooking the rich spectral cues in MSIs. To bridge this gap, we propose a novel framework that exploits spectral information to advance multispectral aerial object detection.

Multispectral Object Detection. Despite advances in general object detection (Carion et al. 2020; Tian, Ye, and Dorrermann 2025), MOD remains underexplored due to limited data. (Yan et al. 2021) pioneered this area with the first dataset and a 3D CNN for spatial-spectral feature extraction, achieving promising results. More recently, (He et al. 2023) proposed S2ADet, a two-stream network separating spatial and spectral branches for independent feature extraction and fusion, improving accuracy while incurring high computational cost due to the decoupling and two-stream network complexity. In contrast, our method adopts a single-stream design, integrating spectral and spatial information directly, achieves superior performance while maintaining computational efficiency, and avoids the limitations of prior methods.

MODA Dataset

Overview

Data Collection. Multispectral sensors trade spatial resolution for spectral bands. MODA targets aerial scenes with widespread small objects requiring high spatial resolution; thus, we select a professional drone-mounted multispectral camera (1280×960 image size; 8 spectral bands range of 395~950 nm; 4.5cm/pixel at 100m height) to record MSIs across diverse scenes, times, and weather, yielding 14,041 MSIs (9,156 training; 4,885 testing) across 50 urban areas.

High Quality Annotation. For high annotation quality, we adopted a three-stage protocol: (i) drafting detailed guidelines (object definitions, annotation tool usage) and training annotators via trial tasks; (ii) trained annotators labeled the raw data accordingly; (iii) two verification passes by the author team to correct errors and ensure accuracy.

Dataset Comparison. Existing MOD datasets suffer from limited scale and artificial ground scenarios (Table 1). HOD-1 targets manually placed targets, while HOD3K features fixed scenes with static backgrounds, limiting their generalization (Fig. 2). In contrast, MODA offers 14,041 MSIs with 330,191 targets across 8 categories using oriented bounding boxes. Collected from diverse aerial views, MODA greatly expands data scale and complexity, better reflecting real-world challenges to support broader practical applications.

Dataset Analysis

Challenge Attributes. As Fig. 2 shows, HOD-1 targets manually placed objects captured from real scenes or screens, and HOD3K focuses on occlusion and small objects in natural scenes. MODA targets to tackle more complex aerial scenarios by integrating spectral and spatial information. We define and record eight key challenge attributes

Dataset	Scene	Images	Image size	Bands	Categories	Annotations	Avg.labels/image	Spectral bands	Type	Year
HOD-1 (Yan et al. 2021)	Manual	454	467×336	96	8	1657	3.57	400~1000nm	HBB	2021
HOD3K (He et al. 2023)	Natural	3242	512×256	16	3	15149	4.37	470~620nm	HBB	2023
MODA (Ours)	Aerial	14041	1200×900	8	8	330191	23.52	390~950nm	OBB	2025

Table 1: Comparison of multispectral object detection datasets. HBB: horizontal bounding box; OBB: oriented bounding box.

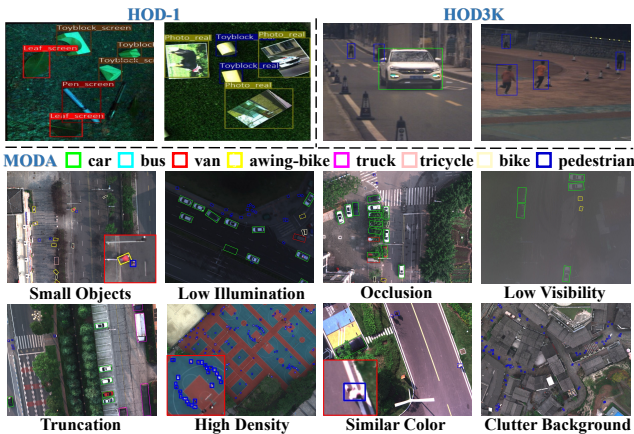


Figure 2: Comparison with other multispectral object detection datasets and examples of challenge attributes in MODA.

of aerial object detection, including small objects, occlusion, and low visibility (Fig. 2), which hinder detectors that rely solely on RGB data, underscoring the need for spectral cues. **Statistical Analysis.** MODA comprises 50 scenes, divided into 103 non-overlapping sub-scenes for training (70%) and testing (30%), while preserving an approximate category-wise annotation ratio of 7:3, which ensures balanced and distinct challenge attributes and annotations across both subsets (Fig. 3 (a)). Fig. 3 (b) depicts the instance distribution per MSI, revealing that over 3.5% of MSIs contain more than 100 instances, contributing to high-density challenges. Furthermore, Fig. 3 (c) presents the relative and absolute instance size distributions, indicating that 95% of instances occupy less than 1% of the image area. Notably, “Pedestrian” and “bike” are particularly small, with an average pixel size below 20, and exhibit significant scale variations.

Method

We present OSSDet, an object-aware spectral-spatial learning framework for multispectral aerial object detection, as shown in Fig. 4. OSSDet employs Cascaded Spectral-Spatial joint Perception (CSSP) to optimize target region perception through cascaded spectral-spatial awareness interactive modulation. Subsequently, Spectral-guided Adaptive Cross-layer Fusion (SACF) modules build a top-down pathway to propagate target-aware cues across layers, reinforcing intra-object correlations and texture details. Object-aware masking branch filters background noise while retaining discriminative object representations. Finally, Cross-spectral Attention Feature Refinement (CAFR) refines object-related fea-

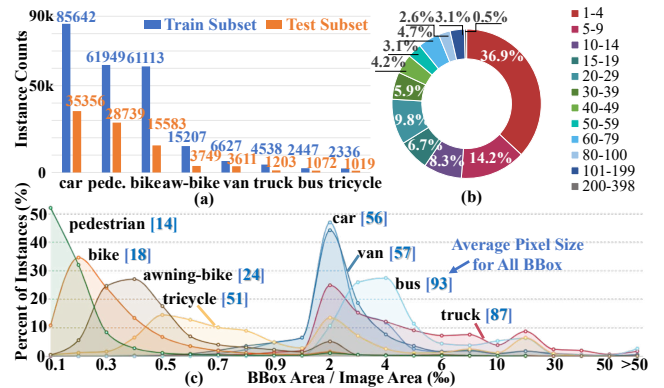


Figure 3: Statistical analysis of MODA. Distribution of instance counts across 8 categories (a) and per MSI (b). (c) Relative and absolute distributions of instance sizes.

tures using object-aware cues and cross-spectral attention.

Cascaded Spectral-Spatial Joint Perception

Successive downsampling during top-level feature extraction introduces spatial aliasing, allowing background clutter to interfere with target regions and degrade model attention. To address this and leverage spectral information, we propose CSSP, which integrates spectral and spatial awareness.

Let $\{F_i \mid i \in 1, \dots, 5\}$ denote multiscale features extracted from MSI input $I \in \mathbb{R}^{h \times w \times b}$, b is the bands number. As Fig. 4 (b) shows, CSSP first applies spectral and spatial attention to top-level feature $F_5 \in \mathbb{R}^{C \times H \times W}$ via two dedicated sub-networks, extracting spectral and spatial aware features. For spectral-wise attention, F_5 is compacted using global average pooling (GAP), then scaled with learnable weight w_e and bias vector b_e along the spectral dimension:

$$F_e = f_s(w_e \cdot \text{GAP}(F_5) + b_e) \odot F_5, \quad (1)$$

where $f_s(\cdot)$ is sigmoid, \odot is element-wise multiplication. Similarly, for spatial-wise attention, F_5 is averaged along the spectral dimension, and a learnable weight map w_a and a bias map b_a is applied to amplify spatially salient regions:

$$F_a = f_s(w_a \cdot \text{AVG}(F_5) + b_a) \odot F_5, \quad (2)$$

We employ a cascaded interaction modulation network to achieve joint spectral-spatial perception. Specifically, spectral and spatial aware features are reshaped to $F_e, F_a \in \mathbb{R}^{C \times HW}$, from which the spectral-spatial cross-correlation matrix $M_{e,a}$ is computed to project spectral-aware feature into spatial space, modulating the spatially aware feature:

$$M_{e,a} = \text{Tanh}\left(\frac{F_e^T F_a}{\|F_e^T F_a\|_2}\right), \hat{F}_a = F_e M_{e,a} + F_a, \quad (3)$$

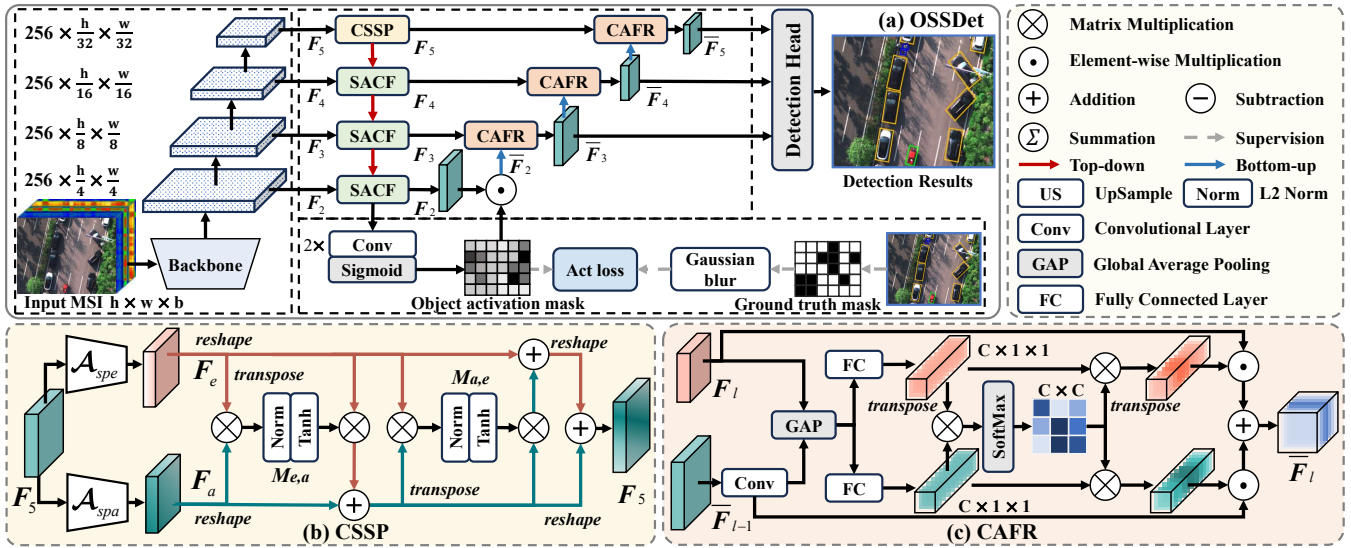


Figure 4: (a) Overall OSSDet framework. SACF fuses aggregated spectral features with spatially enhanced details to reinforce intra-object correlations and spatial texture details; (b) CSSP integrates spectral and spatial awareness to improve target perception; (c) CAFR refines object-related representations with explicit object-aware cues and cross-spectral attention.

where $\|\cdot\|_2$ denotes the L2 norm. Then, a similar operation is applied to reconstruct the spectrally aware feature F_e using the modulated spatially aware feature \hat{F}_a :

$$M_{a,e} = f_t \left(\frac{\hat{F}_a^T F_e}{\|\hat{F}_a^T F_e\|_2} \right), \hat{F}_e = \hat{F}_a M_{a,e} + F_e, \quad (4)$$

Finally, the modulated spectral and spatial features are reshaped back and fused to form the reconstructed feature \hat{F}_5 .

Spectral-guided Adaptive Cross-layer Fusion

As shown in Fig. 5, SACF incorporates Spectral Feature Aggregator (SFA) to adaptively aggregate spectral information and Spatial Detail Enhancer (SDE) to improve spatial details in high-resolution features. A cross-layer fusion unit subsequently embeds the aggregated spectral features and enhanced spatial textures into the low-level feature.

Spectral Feature Aggregator. SFA adaptively aggregates spectral features by examining the relationship between a central spectral feature vector $\mathbf{p}_{i,j} \in \mathbb{R}^{C \times 1 \times 1}$ and its neighbors within a $k \times k$ spatial patch $\mathbf{P}_{i,j}$ centered at (i, j) in the high-level feature $\hat{F}_l \in \mathbb{R}^{C \times H \times W}$, where $\mathbf{P}_{i,j} = [\mathbf{p}_{i+\delta_m, j+\delta_n}]_{\delta_m, \delta_n}$, and $-\frac{k-1}{2} \leq \delta_m, \delta_n \leq \frac{k-1}{2}$ denotes the spatial offsets. To quantify the similarity between the central spectral feature vector $\mathbf{p}_{i,j}$ and its neighbors, the similarity weight $e_{i+\delta_m, j+\delta_n}$ is computed for each neighboring vector:

$$e_{i+\delta_m, j+\delta_n} = \frac{\exp(-\|\mathbf{p}_{i,j} - \mathbf{p}_{i+\delta_m, j+\delta_n}\|_2)}{\sum_{\delta_u, \delta_v} \exp(-\|\mathbf{p}_{i,j} - \mathbf{p}_{i+\delta_u, j+\delta_v}\|_2)}, \quad (5)$$

where δ_u, δ_v share same spatial offsets with δ_m, δ_n . Then the central spectral feature vector $\mathbf{p}_{i,j}$ is updated by weighted aggregation of the spectral feature vectors within $\mathbf{P}_{i,j}$:

$$\mathbf{p}'_{i,j} = \sum_{\delta_m, \delta_n} e_{i+\delta_m, j+\delta_n} \cdot \mathbf{p}_{i+\delta_m, j+\delta_n} + \mathbf{p}_{i,j}, \quad (6)$$

Finally, \hat{F}_l is updated to F'_l by adaptively aggregating neighboring spectral information for each spectral vector, which enhances intra-object feature correlations, thereby improving robustness against noise interference.

Spatial Detail Enhancer. Downsampling during feature extraction can induce spectral aliasing, which destabilizes feature aggregation—particularly at edges with prevalent mixed pixels—and leads to a loss of spatial texture details. To mitigate this, the SDE is introduced to enhance spatial details, minimize information loss during cross-layer transmission.

For a low-level feature $F_{l-1} \in \mathbb{R}^{C \times 2H \times 2W}$, SDE first decomposes it into low-frequency $F_{l-1}^{lf} \in \mathbb{R}^{C \times H \times W}$ and high-frequency $F_{l-1}^{hf} \in \mathbb{R}^{C \times 2H \times 2W}$ components:

$$F_{l-1}^{lf} = \text{AP}(F_{l-1}), F_{l-1}^{hf} = F_{l-1} - \text{US}(F_{l-1}^{lf}), \quad (7)$$

where AP is 3×3 average pooling. Next, the high-frequency component, which stores spatial texture details (Fig. 5), is enhanced via a lightweight detail enhancement block:

$$\hat{F}_{l-1}^{hf} = (1 + \mathbf{f}_s(\phi_{1 \times 1}(F_{l-1}^{hf}))) \odot F_{l-1}^{hf}, \quad (8)$$

where $\phi_{k \times k}$ defines a $k \times k$ convolution. Finally, the detail-enhanced feature is obtained by fusing the enhanced high-frequency textures with low-frequency structural context:

$$F'_{l-1} = \phi_{1 \times 1}([\hat{F}_{l-1}^{hf}, \text{US}(\phi_{3 \times 3}(F_{l-1}^{lf}))]) + F_{l-1}, \quad (9)$$

where $[\cdot]$ denotes channel-wise concatenation.

Adaptive Cross-layer Fusion. To adaptively fuse aggregated spectral features with enhanced spatial details, we first align the spatial dimensions of F'_l and F'_{l-1} . Global confidence are then computed by applying GAP to them separately, which are concatenated and passed a 1×1 convolution followed by activation to generate adaptive fusion weights:

$$\mathbf{W} = \mathbf{f}_s(\phi_{1 \times 1}([\text{AVG}(F'_l), \text{AVG}(F'_{l-1})])), \quad (10)$$

$$\mathbf{W}_l, \mathbf{W}_{l-1} = \mathbf{f}_{\text{split}}(\mathbf{W}) \in \mathbb{R}^{2H \times 2W},$$

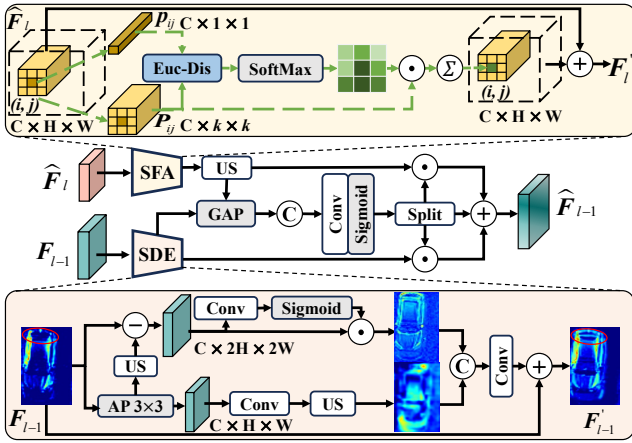


Figure 5: Illustration of SACF. The Euc-Dist denotes Euclidean Distance for spectral vector similarity measurement.

where $f_{split}(\cdot)$ denotes channel-wise split. Finally, we embed F'_l and F'_{l-1} to low-level features via adaptive weighted summation: $\hat{F}_{l-1} = F'_l \odot W_l + F'_{l-1} \odot W_{l-1}$, thereby increasing foreground-background contrast.

Object-aware Masking

SACFs propagate target-aware information across network stages. However, preserving distinctive features for instances prone to background interference, such as small objects, remains challenging. To address this, we introduce the Object Activation Loss \mathcal{L}_{act} , which explicitly enforces object-specific activations. As shown in Fig. 4 (a), the spectral-spatial features of objects in \hat{F}_2 is projected to an activation representations: $M_p = f_s(\phi_{3 \times 3}(\phi_{3 \times 3}(\hat{F}_2)))$, where $M_p \in \mathbb{R}^{H \times W}$ is the predicted object activation mask.

Simultaneously, the ground truth instance mask $M_g \in \mathbb{R}^{H \times W}$ is generated from bounding boxes using a Gaussian blur to model the instance as 2D Gaussian distribution.

\mathcal{L}_{act} comprises an intersection loss \mathcal{L}_I promotes the activation of target-related features, and a difference loss \mathcal{L}_D penalizes the activation of irrelevant background features.

$$\mathcal{L}_I = 1 - \frac{\sum M_p M_g}{\sum M_g}, \quad \mathcal{L}_D = \frac{\sum M_p (1 - \mathcal{H}(M_g))}{\sum M_p}, \quad (11)$$

where $\mathcal{H}(\cdot)$ maps values greater than 0 to 1. Finally, the object activation loss is formulated as: $\mathcal{L}_{act} = \mathcal{L}_I + \gamma \mathcal{L}_D$, with γ controlling foreground-background balance.

Supervision by \mathcal{L}_{act} , M_p highlights target regions, enabling \hat{F}_2 to be masked as $\bar{F}_2 = \hat{F}_2 \odot M_p$ to retain object-specific features while suppressing irrelevant response.

Cross-spectral Attention Feature Refinement

To further refine object-related representations in multiscale features using explicit object-aware cues from \bar{F}_2 , we construct a bottom-up flow with CAFRs. As shown in Fig. 4 (c), given $\hat{F}_l \in \mathbb{R}^{C \times H \times W}$ and its lower-level feature $\bar{F}_{l-1} \in \mathbb{R}^{C \times 2H \times 2W}$, CAFR first aligns them via a 3×3 convolution. Global average pooling followed by a fully connected layer

yields global spectral vectors $V_l, V_{l-1} \in \mathbb{R}^{C \times 1 \times 1}$, which are used to calculate fusion weights W_l, W_{l-1} through cross attention, mitigating spectral bias across layers:

$$A = \text{softmax}(V_l^T V_{l-1}), \quad (12)$$

$$W_l, W_{l-1} = A^T V_l, A V_{l-1},$$

Finally, \hat{F}_l is refined by embedding spatial object-aware and aligned spectral cues: $\bar{F}_l = W_l \odot \hat{F}_l + W_{l-1} \odot \bar{F}_{l-1}$.

Learning Objective

OSSDet incorporates detection loss \mathcal{L}_{det} from (Li et al. 2022) and \mathcal{L}_{act} introduced above, the total model optimize function is:

$$\mathcal{L} = \mathcal{L}_{det} + \alpha \mathcal{L}_{act}. \quad (13)$$

where α is a weighting factor.

Experiments

We evaluate OSSDet on MODA and HOD3K. For fairness, all methods are evaluated in identical settings. Further results and details are available in the supplementary material.

Results on the MODA Dataset

Quantitative Results. Table 2 reports results on MODA. Overall, OSSDet achieves superior performance, outperforming the suboptimal method by 2.5%, 1.8%, and 1.7% in mAP50, mAP75, and mAP, respectively. The latest MSI-oriented detector S2ADet struggles with small objects like “bike” and “pedestrian”. OSSDet achieves top results in all categories except “awning-bike”, with notable gains of 2.3% and 2.6% in “tricycle” and “pedestrian”, underscoring its effectiveness in preserving small objects’ distinctive features.

Efficiency Analysis. As Table 2 shows, two-stage methods offer high accuracy but complex architectures, while single-stage ones like LSKNet are lightweight yet less accurate. The MSI-oriented S2ADet adopts a two-stream design to decouple spatial and spectral cues, incurs substantial preprocessing and computational cost. In contrast, OSSDet achieves best accuracy with 36.5M parameters and 263.1G FLOPs, offering a favorable accuracy-efficiency balance.

Qualitative Results. As shown in Fig. 6, OSSDet enhances the network’s focus on target regions while suppressing irrelevant noise, even for challenging scenarios, such as small objects, low visibility, or cluttered backgrounds, resulting in a substantial reduction of false and missed detections.

Visualization of Feature Distribution. Fig. 7 shows T-SNE (Van der Maaten and Hinton 2008) visualization, where MSI inputs enlarge inter-class distances for easily confused (car & van), small (pedestrian & bike), and few-sample (tricycle) categories, yielding clearer feature separation. This confirms that spectral cues improve object discriminability and validate OSSDet’s effectiveness in feature learning.

Effect of Components. Fig. 8 visualizes features involved in key components. (a) CSSP optimizes target perception via cascaded spectral-spatial modulation; (b) SACF integrates spectral aggregation with spatial detail enhancement to strengthen intra-object correlations and clear boundaries; (c) object-aware masking backpropagates \mathcal{L}_{act} to preserve

Method		car	bus	van	aw-bike	truck	tricycle	bike	pedestrian	mAP ₅₀	mAP ₇₅	mAP	FLOPs	Params
Two-stage	Gliding Vertex (Xu et al. 2020)	90.3	89.1	73.8	69.5	66.0	46.7	41.4	22.6	62.4	35.7	34.7	230.8G	41.4M
	Roi Transformer (Ding et al. 2019)	90.5	89.3	75.2	73.3	68.7	51.8	44.5	30.0	65.4	43.4	40.7	244.7G	55.3M
	StripRCNN (Yuan et al. 2025)	90.5	89.0	76.1	73.6	67.1	56.5	45.2	30.7	66.1	44.0	41.0	231.8G	45.2M
One-stage	GWD (Yang et al. 2021b)	90.4	79.9	70.7	59.7	49.3	30.8	26.7	29.7	54.7	37.8	34.6	233.5G	36.5M
	R3Det (Yang et al. 2021a)	90.4	88.7	74.0	63.8	57.8	45.1	29.4	32.2	60.2	36.7	34.8	362.2G	42.0M
	S ² ANet (Han et al. 2021)	90.4	88.7	74.4	69.1	62.7	48.9	30.1	40.7	63.1	38.7	36.7	216.5G	38.8M
	R3Det-KLD (Yang et al. 2021c)	90.4	88.7	74.8	66.6	60.9	42.6	29.8	35.5	61.1	39.4	36.8	306.7G	39.6M
	LSKNet-S ² ANet (Li et al. 2024)	90.5	89.5	76.1	72.5	68.6	51.8	38.9	43.1	66.4	41.3	38.6	194.9G	29.9M
	S2ADet [†] (He et al. 2023)	90.3	86.6	72.1	71.3	57.2	54.1	35.0	40.8	63.5	41.1	38.9	406.0G	65.2M
	Rot. FCOS (Tian et al. 2019)	90.3	86.3	75.2	71.2	59.0	53.8	34.7	37.4	63.5	41.8	39.1	226.9G	32.2M
	CFA (Guo et al. 2021)	90.4	89.0	76.7	69.7	64.4	55.1	43.1	41.5	66.2	43.2	40.6	213.6G	36.9M
	Ori. RepPoints (Li et al. 2022)	90.5	89.2	77.7	71.2	66.2	53.1	43.0	41.1	66.5	44.1	40.9	213.6G	36.9M
	OSSDet (Ours)	90.5	89.9	79.2	72.7	69.7	58.8	45.3	45.7	69.0	45.9	42.7	263.1G	36.5M

Table 2: Comparison with other methods on MODA. [†] indicates methods originally designed for multispectral object detection.

Method	mAP ₅₀	mAP ₇₅	mAP	FLOPs	Params
CO-DETR (Zong et al. 2023)	77.9	57.3	51.7	200.7G	64.5M
DINO (Zhang et al. 2022a)	89.1	61.6	56.3	114.7G	47.6M
Fovea (Kong et al. 2020)	92.2	62.2	56.8	123.9G	38.0M
TOOD (Feng et al. 2021)	90.3	65.6	57.8	108.5G	32.1M
VFNet (Zhang et al. 2021)	92.3	66.5	59.0	104.5G	32.8M
C-RCNN (Cai et al. 2019)	90.9	66.9	59.6	149.5G	69.2M
S2ADet [†] (He et al. 2023)	93.4	-	59.8	169.2G	48.6M
OSSDet (Ours)	93.4	68.8	60.9	131.2G	36.6M

Table 3: Comparison with other detectors on HOD3K.

Method	Input	mAP ₅₀	mAP ₇₅	mAP	FLOPs	Params
Strip	RGB	65.2	41.9	39.6	227.4G	45.13M
RCNN	MSI	66.1 (+0.9)	44.0 (+2.1)	41.0 (+1.4)	231.8G	45.15M
Oriented	RGB	66.2	41.6	39.4	209.2G	36.83M
RepPoints	MSI	66.5 (+0.3)	44.1 (+2.5)	40.9 (+1.5)	213.6G	36.85M
OSSDet (Ours)	RGB	67.1	41.3	39.7	258.7G	36.46M
	MSI	69.0 (+1.9)	45.9 (+4.6)	42.7 (+3.0)	263.1G	36.48M

Table 4: Comparison with different input.

object features, even for small objects-reinforces object-specific features while suppressing noise.

Results on the HOD3K Dataset

Quantitative Results. OSSDet achieves the best overall performance, surpassing the suboptimal method by 1.9% in mAP₇₅ and 1.1% in mAP, as shown in Table 3. While S2ADet matches OSSDet in mAP₅₀, its two-stream design causes high preprocessing and computation overhead. In contrast, OSSDet directly processes MSI inputs, delivering better accuracy with fewer parameters and FLOPs.

Qualitative Results. Other methods often miss occluded or background-blended objects (Fig. 9). In contrast, OSSDet exploits spectral cues to enhance target-background separation, yielding predictions closer to the ground truth.

CSSP	SACF	Obj-aware	CAFR	mAP ₅₀	mAP ₇₅	mAP
✗	✗	✗	✗	66.5	44.1	40.9
✓	✗	✗	✗	67.6	44.1	41.4
✗	✓	✗	✗	67.4	44.7	41.6
✓	✓	✗	✗	68.1	44.9	41.9
✓	✓	✓	✗	68.7	45.6	42.5
✓	✓	✓	✓	69.0	45.9	42.7

Table 5: Ablation studies on the key components of OSSDet.

SDE	SFA	mAP ₅₀	mAP ₇₅	mAP
✓	✗	67.8	44.5	41.5
✗	✓	68.2	44.7	41.7
✓	✓	69.0	45.9	42.7

Table 6: Ablation study of the key components of SACF.

Ablation Studies

Input Ablation Study. As shown in Table 4, MSI inputs significantly improve performance over RGB due to richer spectral information. OSSDet achieves the largest gains by effectively integrating spectral-spatial features with object-aware cues, while adding only marginal cost (4.4 GFLOPs, 0.02 M params), making this enhancement nearly cost-free.

Component Ablation Study. Each component improves detection, with full integration performing best (Table 5). Notably, adding SACF and object-aware branch offers the most significant gains across metrics, underscoring the value of integrating spectral aggregation, spatial detail enhancement, and object-aware cues for spectral-spatial learning.

Ablation Study on SACF. As shown in Table 6, removing either SACF component causes clear performance drops, confirming their complementarity. SACF adaptively aggregates spectral features within a spatial patch of size k , overly large patches introduce irrelevant features that weaken target specificity and degrade performance (Table 7).

Ablation Study on CSSP. CSSP enhances target perception through cascaded spectral-spatial interaction. To assess its

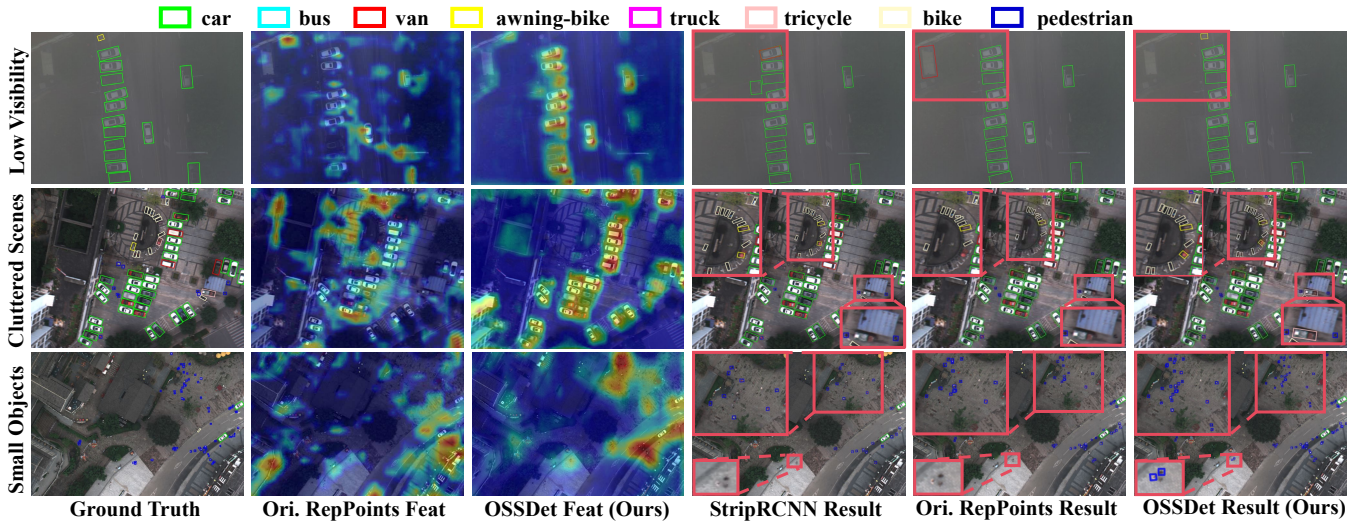


Figure 6: Visualization comparison of detection results and feature maps obtained by different methods on MODA.




Figure 7: 2D t-SNE visualization of feature distributions learned by OSSDet on MODA with RGB and MSI inputs.

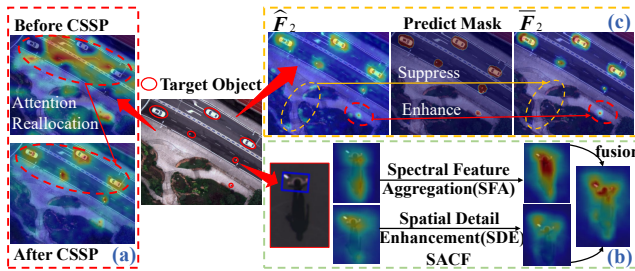


Figure 8: Effectiveness visualization of key components.



Figure 9: Detection results on HOD3K. **Red** boxes indicate the ground truth, while **green** boxes denote the predictions.

effectiveness, we compare it with addition, concatenation, and a Softmax-based variant (Table 7), CSSP outperforms

Method	mAP ₅₀	mAP ₇₅	mAP	k	mAP ₅₀	mAP ₇₅	mAP
SoftMax	66.7	42.8	40.5	3	69.0	45.9	42.7
Concat	67.9	44.7	41.9	5	68.7	45.8	42.4
Addition	68.5	45.4	42.4	7	61.6	40.4	37.6
CSSP	69.0	45.9	42.7	9	61.4	38.5	36.8

Table 7: Ablation study of CSSP and patch size k in SACF.

γ	mAP ₅₀	mAP ₇₅	mAP	α	mAP ₅₀	mAP ₇₅	mAP
0.5	67.6	44.8	41.8	0.4	68.2	45.2	41.8
0.25	68.2	45.4	42.1	0.6	69.0	45.9	42.7
0.1	69.0	45.9	42.7	0.8	68.0	44.6	41.7
0.05	67.4	44.6	41.5	1.0	68.3	45.1	42.0

Table 8: Ablation study of weighting factors.

all, while the Softmax version shows unstable convergence. **Ablation Study on Weighting Factors.** We evaluated the effects of γ and loss weight α (Table 8). A large γ suppresses early activations, whereas a small one fails to filter noise. OSSDet performs best with $\gamma = 0.1$ and $\alpha = 0.6$.

Conclusion

We introduce MODA, the first large-scale challenging dataset for multispectral aerial object detection, which addresses training sample scarcity with high-resolution MSIs featuring diverse scenarios, challenging attributes, and high-quality oriented annotations. We also propose OSSDet, a novel baseline integrating spectral-spatial information with object-aware cues. OSSDet enhances target perception via cascaded spectral-spatial interactive modulation, reinforces intra-object correlations via spectral aggregation, and suppresses background noise through explicit object-aware guidance, boosting detection accuracy. We believe MODA and OSSDet will catalyze future research in this domain.

Acknowledgments

This work was financially supported by the National Natural Science Foundation of China (No. 62571031).

References

- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*, 213–229. Springer.
- Ding, J.; Xue, N.; Long, Y.; Xia, G.-S.; and Lu, Q. 2019. Learning RoI transformer for oriented object detection in aerial images. In *CVPR*, 2849–2858.
- Fang, L.; Jiang, Y.; Yan, Y.; Yue, J.; and Deng, Y. 2023a. Hyperspectral image instance segmentation using spectral-spatial feature pyramid network. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13.
- Fang, L.; Yan, Y.; Yue, J.; and Deng, Y. 2023b. Toward the vectorization of hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–14.
- Feng, C.; Zhong, Y.; Gao, Y.; Scott, M. R.; and Huang, W. 2021. TOOD: Task-aligned One-stage Object Detection. In *ICCV*.
- Guo, Z.; Liu, C.; Zhang, X.; Jiao, J.; Ji, X.; and Ye, Q. 2021. Beyond bounding-box: Convex-hull feature adaptation for oriented and densely packed object detection. In *CVPR*, 8792–8801.
- Han, J.; Ding, J.; Li, J.; and Xia, G.-S. 2021. Align deep features for oriented object detection. *IEEE transactions on geoscience and remote sensing*, 60: 1–11.
- He, X.; Tang, C.; Liu, X.; Zhang, W.; Sun, K.; and Xu, J. 2023. Object detection in hyperspectral image via unified spectral-spatial feature aggregation. *IEEE Transactions on Geoscience and Remote Sensing*.
- Kong, T.; Sun, F.; Liu, H.; Jiang, Y.; Li, L.; and Shi, J. 2020. Foveabox: Beyond anchor-based object detection. *IEEE Transactions on Image Processing*, 29: 7389–7398.
- Leng, J.; Ye, Y.; Mo, M.; Gao, C.; Gan, J.; Xiao, B.; and Gao, X. 2024. Recent Advances for Aerial Object Detection: A Survey. *ACM Computing Surveys*, 56(12): 1–36.
- Li, W.; Chen, Y.; Hu, K.; and Zhu, J. 2022. Oriented re-points for aerial object detection. In *CVPR*, 1829–1838.
- Li, Y.; Li, X.; Dai, Y.; Hou, Q.; Liu, L.; Liu, Y.; Cheng, M.-M.; and Yang, J. 2024. LSKNet: A Foundation Lightweight Backbone for Remote Sensing. *International Journal of Computer Vision*.
- Pearson, K. 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science*, 2(11): 559–572.
- Qin, H.; Xu, T.; Liu, P.; Xu, J.; and Li, J. 2024. DMSSN: Distilled mixed spectral-spatial network for hyperspectral salient object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–18.
- Tian, Y.; Ye, Q.; and Doermann, D. 2025. YOLOv12: Attention-Centric Real-Time Object Detectors. *arXiv preprint arXiv:2502.12524*.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. FCOS: Fully Convolutional One-Stage Object Detection. In *ICCV*, 9626–9635. IEEE Computer Society.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Xiao, Y.; Xu, T.; Yu, X.; Fang, Y.; and Li, J. 2024. A Lightweight Fusion Strategy with Enhanced Inter-layer Feature Correlation for Small Object Detection. *IEEE Transactions on Geoscience and Remote Sensing*.
- Xie, X.; Cheng, G.; Wang, J.; Yao, X.; and Han, J. 2021. Oriented R-CNN for object detection. In *ICCV*, 3520–3529.
- Xu, H.; Liu, X.; Xu, H.; Ma, Y.; Zhu, Z.; Yan, C.; and Dai, F. 2024. Rethinking boundary discontinuity problem for oriented object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17406–17415.
- Xu, Y.; Fu, M.; Wang, Q.; Wang, Y.; Chen, K.; Xia, G.-S.; and Bai, X. 2020. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE transactions on pattern analysis and machine intelligence*, 43(4): 1452–1459.
- Yan, L.; Zhao, M.; Wang, X.; Zhang, Y.; and Chen, J. 2021. Object detection in hyperspectral images. *IEEE Signal Processing Letters*, 28: 508–512.
- Yang, X.; Yan, J.; Feng, Z.; and He, T. 2021a. R3det: Refined single-stage detector with feature refinement for rotating object. In *AAAI*, volume 35, 3163–3171.
- Yang, X.; Yan, J.; Ming, Q.; Wang, W.; Zhang, X.; and Tian, Q. 2021b. Rethinking rotated object detection with gaussian wasserstein distance loss. In *International conference on machine learning*, 11830–11841. PMLR.
- Yang, X.; Yang, X.; Yang, J.; Ming, Q.; Wang, W.; Tian, Q.; and Yan, J. 2021c. Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. *Advances in Neural Information Processing Systems*, 34: 18381–18394.
- Yuan, X.; Zheng, Z.; Li, Y.; Liu, X.; Liu, L.; Li, X.; Hou, Q.; and Cheng, M.-M. 2025. Strip R-CNN: Large Strip Convolution for Remote Sensing Object Detection. *arXiv preprint arXiv:2501.03775*.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2022a. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. *arXiv:2203.03605*.
- Zhang, H.; Wang, Y.; Dayoub, F.; and Sunderhauf, N. 2021. Varifocalnet: An iou-aware dense object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8514–8523.
- Zhang, P.; Zhao, J.; Wang, D.; Lu, H.; and Ruan, X. 2022b. Visible-Thermal UAV Tracking: A Large-Scale Benchmark and New Baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8886–8895.
- Zhang, Y.; Qi, J.; Wang, X.; Cai, Z.; Peng, J.; and Zhou, Y. 2024. Tensorial global-local graph self-representation for hyperspectral band selection. *IEEE Transactions on Circuits and Systems for Video Technology*.

Zhu, P.; Wen, L.; Du, D.; Bian, X.; Fan, H.; Hu, Q.; and Ling, H. 2021. Detection and Tracking Meet Drones Challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.