

Fully Rotation-Equivariant Spectral-Spatial Learning for Multispectral Object Detection

Peng Zhang[✉], Tingfa Xu[†][✉], Shuaihao Han[✉], and Jianan Li[†][✉]

Beijing Institute of Technology, Beijing, China

Abstract. Existing multispectral detectors are limited by discrete spectral processing, a scale-dependent shift in the relative reliability of spectral and spatial cues across pyramid levels, and the lack of explicit rotation-equivariant geometric priors for arbitrarily oriented objects. To tackle these limitations, we propose FressDet, a fully rotation-equivariant spectral-spatial learning framework for multispectral object detection, capable of capturing the continuous, ordered nature of spectral structure and enabling reliable spectral-spatial fusion across pyramid levels under arbitrary in-plane rotations. FressDet integrates three complementary components. Spectral Implicit Warp (SpeIW) enables query-based spectral resampling via a coordinate-conditioned implicit field, yielding a monotone, order-preserving warp. Rotation-Equivariant Consistency Weighting (ReCoW) adaptively fuses spectral and spatial branches based on branch reliability, reinforcing informative cues while suppressing noise across pyramid levels. The oriented-aware head exploits group-indexed features to stably predict oriented objects without parameter replication. Taken together, FressDet learns more discriminative and robust spectral-spatial representations even under rotational perturbations. By achieving state-of-the-art performance with 93% fewer parameters on five public benchmarks, FressDet demonstrates its effectiveness and generalizability. Code is available at <https://github.com/Riiluo/FressDet>.

Keywords: Multispectral object detection · Equivariant neural networks · Implicit spectral warping · Spectral-spatial learning

1 Introduction

Multispectral object detection (MOD) improves localization and recognition by jointly exploiting spatial structure and spectral cues. Compared with RGB imagery, multispectral images (MSIs) measure spectral responses at sampled wavelengths, yielding per-pixel spectral signatures with strong inter-band correlation. Prior methods either decouple spectral and spatial information through dimensionality reduction (e.g., PCA [37] and band selection [42, 74]) or rely on attention-based modeling to capture spectral-spatial interactions [15]. Although these methods achieve promising performance, two challenges remain in MOD: (i) discrete spectral processing that neglects the continuous and ordered

[†] Correspondence to: Tingfa Xu and Jianan Li.




Fig. 1: (a) Transformation consistency in vanilla (V-CNN) and equivariant (Eq-CNN) networks. T_e and T_r denote the identity and a rotation, f is a feature extractor, and Z_0 , Z_1 are features of $I_0 = T_e(I_0)$ and $I_1 = T_r(I_0)$. Ideally $T_r^{-1}(Z_1) \approx Z_0$; the error map confirms Eq-CNN yields lower transformation error. However, existing methods have not yet unified continuous spectral modeling with rotation-equivariant geometric priors. (b) Detection robustness under rotational perturbations at 1° intervals on MODA. FressDet exhibits the best robustness. (c) mAP50 vs. parameter count on MODA, where circle size is proportional to FLOPs. FressDet achieves the best performance with the fewest parameters and lowest FLOPs.

structure of the spectral dimension; (ii) a scale-dependent shift in the relative reliability between spectral and spatial cues across pyramid levels, which can degrade fusion performance in complex scenes.

Rotation equivariance is a desirable geometric inductive bias that exploits the intrinsic rotational symmetry and arbitrary orientations of objects, improving representation stability and generalization across diverse visual tasks [6, 52], including ground-level [54, 55] and aerial [14, 53] object detection as well as semantic segmentation [61], yet many multispectral detectors still rely on heavy rotation-based data augmentation in practice. However, extending this inductive bias to MOD requires rethinking the modeling paradigm. In particular, the continuous, ordered spectral structure of MSIs is mismatched with most equivariant designs built on discrete sampling (Fig. 1 (a)). Moreover, reliability-aware pyramid fusion under equivariant architectures is non-trivial in complex scenes. Furthermore, group-indexed interactions increase the computational burden. Therefore, a key challenge in MOD is to explicitly model spectral continuity and enable reliable spectral-spatial fusion across scales, while enforcing rotation-equivariant geometric priors with manageable overhead.

To address the challenges of spectral discreteness, scale-dependent fusion reliability, and the lack of rotation-equivariant geometric priors, we propose Fress-

Det, the first fully rotation-equivariant spectral-spatial learning framework for multispectral object detection. FressDet captures continuous spectral structure, enables reliability-aware pyramid fusion across levels, and stably predicts oriented objects without parameter replication.

To overcome the discreteness of spectral processing, we draw on the universal approximation theorem [7, 17]: implicit neural representations (INRs) parameterize signals as coordinate-conditioned neural functions, enabling approximation of continuous functions with arbitrarily small error. Accordingly, we introduce Spectral Implicit Warp (SpeIW), an efficient spectral operator that models multispectral representations as a coordinate-conditioned function queried along the spectral dimension with Fourier-encoded spectral coordinates [38, 45]. This encoding mitigates the spectral bias of coordinate-based networks, enabling the representation of higher-frequency spectral variations. Specifically, we parameterize spectral warping over rotation-equivariant features via a factorization into coordinate-dependent bases and feature-conditioned coefficients. The bases encode regularities along the spectral dimension, while the coefficients provide pixel-wise modulation of the warping offsets. By predicting strictly positive spectral steps and integrating them cumulatively, we obtain an order-preserving mapping by construction.

Addressing scale-dependent reliability shifts across pyramid levels, we propose Rotation-Equivariant Consistency Weighting (ReCoW), a lightweight refinement module with a dual prototype-driven routing mechanism that adaptively selects the more reliable branch. On the spectral branch, ReCoW performs aggregation with soft assignment to capture discriminative spectral patterns. On the spatial branch, it improves geometric coherence, reducing background clutter while preserving fine-grained texture details. A learned branch-selection gate finally produces a convex combination of the two branches. Overall, ReCoW improves spectral-spatial fusion across pyramid levels in complex scenes.

For arbitrarily oriented objects, we propose an oriented-aware head that exploits feature maps with explicit orientation indices, enabling stable prediction across orientations without parameter replication. This design enhances feature utilization and substantially reduces the parameter count.

Taken together, these components enable FressDet to learn more discriminative and robust spectral-spatial representations. Extensive experiments on five public benchmarks demonstrate that FressDet achieves state-of-the-art accuracy with 93% fewer parameters than the previous best method (Fig. 1 (c)) and the best robustness even under input rotational perturbations (Fig. 1 (b)). Our contributions are summarized as follows:

- Present FressDet, the first fully rotation-equivariant spectral-spatial learning framework unifying continuous spectral modeling, reliability-aware pyramid fusion, and stable orientation-aware prediction.
- Introduce Spectral Implicit Warp for continuous, order-preserving spectral resampling via a coordinate-conditioned implicit field.
- Propose Rotation-Equivariant Consistency Weighting, which performs adaptive, reliability-aware spectral-spatial fusion across pyramid levels.

- Design an oriented-aware head for stably predicting arbitrarily oriented objects without parameter replication.

2 Related Work

Multispectral Object Detection. Multispectral imagery provides spectral cues beyond RGB data, facilitating more discriminative object representations. Existing methods either separate spectral and spatial streams with cross-branch fusion [16, 63] or integrate them in a unified design [15]. However, most treat spectral features independently, leaving the continuous and ordered nature of spectral structure insufficiently modeled, while scale-dependent shifts in spectral-spatial reliability across pyramid levels remain unaddressed. In contrast, our method explicitly models spectral continuity and enables reliability-aware spectral-spatial fusion across pyramid levels.

Rotation-equivariant Networks. Since Cohen *et al.* [6] introduced G-CNNs, rotation-equivariant networks have shown strong performance across diverse vision tasks. In object detection, ReDet [14] pioneered rotation-equivariant aerial detectors, and subsequent works [53, 54] advanced equivariant architectures. However, these methods are only approximately equivariant. Lee *et al.* [26] present FRED, the first fully rotation-equivariant detector with deformable convolutions, though its group-indexed interactions incur considerable computational overhead. We develop the first fully rotation-equivariant spectral-spatial learning framework for multispectral object detection, integrating spatial geometry with discriminative spectral cues while maintaining manageable overhead.

Implicit Neural Representations. Implicit neural representations (INRs) parameterize signals as coordinate-to-value neural fields, achieving success in 3D scene modeling [36] and continuous image representation [4, 60]. In spectral imaging, INRs have been applied to spectral rendering [28] and multispectral-hyperspectral fusion [31]. However, prior INR methods primarily address reconstruction or fusion tasks, whereas multispectral detection requires modeling spectral continuity under a strict order-preserving constraint. Our method addresses this by introducing an order-preserving spectral warping mechanism that maintains rotation equivariance.

3 Method

In this section, we introduce the preliminaries on rotation equivariance and describe how the components of FressDet are assembled into a fully rotation-equivariant pipeline (Fig. 2). A complete equivariance proof for the entire network is provided in the supplementary material.

3.1 Preliminaries

Symmetry Groups and Equivariance. We consider equivariance to discrete in-plane rotations within the Euclidean group $E(2) = \mathbb{R}^2 \rtimes O(2)$, the semi-direct




Fig. 2: The overall framework of FressDet. (a) SpeIW resamples spectral features via a continuous, order-preserving warp. (b) ReCoW adaptively fuses spectral and spatial branches based on branch reliability. (c) The oriented-aware head stably predicts oriented objects without parameter replication.

product of translations and orthogonal transformations. By restricting to the cyclic subgroup $C_N = \{r_k \mid k = 0, \dots, N-1\}$ of N equally spaced rotations, we sample features at N discrete orientations. We denote this group by G ($|G| = N$) in the sequel. A function $\Phi: X \rightarrow Y$ is equivariant with respect to G if:

$$\Phi(\rho_g^X(x)) = \rho_g^Y(\Phi(x)), \quad \forall g \in G, x \in X, \quad (1)$$

where ρ_g^X and ρ_g^Y denote group actions on X and Y , respectively. Intuitively, equivariance ensures that transforming the input produces a correspondingly transformed output, enabling the network to generalize across orientations without relying on explicit data augmentation. When ρ_g^Y is the identity for all g , equivariance reduces to invariance.

Lifting and Regular Representation. Under an input rotation r_k , a group feature map $F: \mathbb{R}^2 \times G \rightarrow \mathbb{R}^C$ transforms as follows:

$$[\rho_{r_k}(F)](x, g) = F(r_k^{-1}x, r_k^{-1}g), \quad (2)$$

concretely, a cyclic permutation by k along the group axis composed with a rotation by r_k^{-1} . Following this representation, the planar input $\mathbf{x}_0 \in \mathbb{R}^{c \times h \times w}$ is lifted to a group feature map as follows:

$$\mathbf{y}(g) = \sum_{i=1}^c \mathbf{x}_{0,i} * \mathcal{R}_g(\mathbf{W}_i), \quad g \in \{0, 1, \dots, N-1\}, \quad (3)$$

where $\mathbf{W} \in \mathbb{R}^{c' \times c \times s \times s}$ is a learnable weight matrix with kernel size s and $\mathcal{R}_g$ denotes spatial rotation by $2\pi g/N$. For subsequent layers, group-to-group con-

volution take the form:

$$\mathbf{y}'(g) = \sum_{i=1}^c \sum_{g'=0}^{N-1} \mathbf{y}_i(g') * \mathcal{R}_g(\mathbf{W}_{i,g^{-1}g'}), \quad g \in \{0, 1, \dots, N-1\}, \quad (4)$$

where the relative index $g^{-1}g'$ induces a cyclic shift along the group axis. Group convolutions preserve this transformation law, so equivariance propagates through successive layers. The $\mathcal{O}(c^2Ns^2)$ cost of Eq. (4) is prohibitive for real-time detection. We decompose it into two stages that jointly preserve Eq. (2).

Stage 1: Depthwise rotated spatial filtering. Each channel i is convolved with a single spatial kernel $\mathbf{W}_i^{\text{dw}} \in \mathbb{R}^{s \times s}$, rotated per group element:

$$\tilde{F}_i(x, g) = F_i(x, g) * \mathcal{R}_g(\mathbf{W}_i^{\text{dw}}), \quad (5)$$

preserving the cyclic-shift law of Eq. (2) channel-wise at $\mathcal{O}(cNs^2)$ cost.

Stage 2: Pointwise group mixing. A 1×1 convolution mixes channels and group indices via a weight $\mathbf{W}^{\text{pw}} \in \mathbb{R}^{c' \times c \times N}$ whose group axis is cyclically shifted:

$$F'_j(x, g) = \sum_{i=1}^c \sum_{g'=0}^{N-1} W_{j,i,g^{-1}g'}^{\text{pw}} \tilde{F}_i(x, g'), \quad (6)$$

at $\mathcal{O}(c^2N)$ per spatial location. The cyclic indexing $g^{-1}g'$ guarantees that the composition of Eqs. (5)–(6) satisfies Eq. (2) end-to-end with cost $\mathcal{O}(cNs^2 + c^2N)$.

3.2 Spectral Implicit Warp

While the lifting layer endows features with rotational structure, the spectral dimension remains discretized into independent channels. SpeIW addresses this by predicting a monotone warp index ϕ for each spectral channel and resampling via differentiable interpolation (Fig. 2 (a)), enabling adaptive spectral modeling while preserving equivariance.

Coordinate-Conditioned Factorization. Let $F \in \mathbb{R}^{C \times |G| \times H \times W}$ denote a group feature map with C spectral channels and $|G|$ orientation slots. Each spectral channel is assigned a normalized spectral coordinate $\lambda_c \in [-1, 1]$, lifted via Fourier positional encoding $\psi(\lambda) = [\sin(2^i\pi\lambda), \cos(2^i\pi\lambda)]_{i=0}^{L-1} \in \mathbb{R}^{2L}$ to mitigate the frequency bias of coordinate-based networks [38, 45]. We decompose the raw warp field into a spatial latent code $\mathbf{z}(x, g) \in \mathbb{R}^R$, derived from F via a learned equivariant projection, and a spectral basis decoded by a coordinate-conditioned implicit function $f_\theta: \mathbb{R}^{2L} \rightarrow \mathbb{R}^R$:

$$\delta_c(x, g) = \langle \mathbf{z}(x, g), f_\theta(\psi(\lambda_c)) \rangle, \quad c = 1, \dots, C. \quad (7)$$

Because f_θ is shared across all C spectral channels and takes continuous spectral coordinates as input, it parameterizes a spectral basis that can be queried at arbitrary coordinates, with parameter cost that does not scale with C .

Order-Preserving Mapping. Directly using $\delta_c(x, g)$ as resampling indices may cause index crossing (i.e., $\phi_c(x, g) \geq \phi_{c+1}(x, g)$), resulting in non-monotonic

spectral mappings and invalid resampling. We therefore enforce monotonically ordered indices $\{\phi_1(x, g) < \phi_2(x, g) < \dots < \phi_C(x, g)\}$ to preserve the channel order. We obtain nonnegative increments via softplus, and enforce strict ordering by cumulative summation followed by normalization:

$$\Delta_c(x, g) \triangleq \sigma_+(\delta_c(x, g)), \quad (8a)$$

$$\phi_c(x, g) = (C-1) \cdot \frac{\sum_{j=2}^c \Delta_j(x, g)}{\sum_{j=2}^C \Delta_j(x, g)}, \quad (8b)$$

where σ_+ denotes the softplus function. Each $\Delta_c(x, g)$ serves as a positive increment along the spectral dimension, so $\phi_c(x, g)$ is strictly increasing w.r.t. c by construction. This yields a valid, non-crossing resampling function and prevents degenerate channel permutations.

Continuous Spectral Resampling. Given the monotonically ordered indices $\{\phi_c(x, g)\}_{c=1}^C$, $\phi_c(x, g) \in [0, C-1]$, we obtain the warped feature via differentiable linear interpolation along the spectral dimension:

$$\hat{F}_c(x, g) = \mathcal{I}[F(x, g, \cdot); \phi_c(x, g)], \quad (9)$$

where $\mathcal{I}[\cdot; t]$ denotes a differentiable linear interpolation operator evaluated at continuous index t . Since $\mathbf{z}(x, g)$ is rotation-equivariant and f_θ is shared across orientations, the predicted indices $\phi_c(x, g)$ are equivariant as well. Since the interpolation operates only along the spectral dimension and is applied independently for each (x, g) , it commutes with the group action.

3.3 Rotation-Equivariant Consistency Weighting

Multi-scale fusion in feature pyramids aggregates representations across pyramid levels, which is essential for scale robustness. However, repeated cross-level aggregation can wash out discriminative spectral evidence and reduce class separability for spectrally similar categories. ReCoW addresses this issue (Fig. 2 (b)) by combining a spectral branch that preserves discriminative spectral patterns with a spatial branch that enforces geometric coherence, and fusing them with a learned branch-selection gate. This design maintains this balance during pyramid fusion with provable equivariance guarantees.

Dual-Branch Prototype Routing. Both branches operate on local windows of size $W \times W$. A rotation-invariant descriptor $\bar{F}(x) = \frac{1}{|G|} \sum_g F(x, g)$ is first obtained by group averaging. Within each window, we form $K = q^2$ prototype centers $\{\mathbf{c}_k\}_{k=1}^K$ by adaptive average pooling of $\bar{F}$ to a $q \times q$ spatial grid. The grid preserves local spatial layout and provides region-specific centers, where x indexes locations within the window.

The two branches adopt complementary routing rules. Spectral signatures often exhibit continuous mixing; thus, we employ cosine-based soft assignment

to preserve spectral diversity:

$$s_k^{\text{spec}}(x) = \tau \cdot \cos(\mathbf{c}_k, \bar{F}(x)), \quad (10a)$$

$$p_k(x) = \frac{\exp(s_k^{\text{spec}}(x))}{\sum_{k'} \exp(s_{k'}^{\text{spec}}(x))}, \quad (10b)$$

where τ is a learnable temperature. The routed spectral representation is obtained by weighted aggregation of prototype centers:

$$\hat{F}_{\text{spec}}(x) = \sum_{k=1}^K p_k(x) \mathbf{c}_k. \quad (11)$$

We share $\hat{F}_{\text{spec}}(x)$ across g to obtain the refinement $F_{\text{spec}}(x, g)$ for all orientations, preserving equivariance.

Spatial grouping benefits from crisp region assignments; accordingly, we perform hard routing on a rotation-invariant spatial score map. We compute a scalar score $s(x) = \frac{1}{C|G|} \sum_{c,g} F_c(x, g)$ and obtain scalar prototype centers $\{\mu_k\}_{k=1}^K$ by adaptive average pooling of s onto the same $q \times q$ grid. Hard assignment is obtained via a Gaussian affinity:

$$a_k(x) = \exp(-\gamma \|\mu_k - s(x)\|_2^2), \quad (12a)$$

$$k^*(x) = \arg \max_k a_k(x), \quad (12b)$$

where γ is a learnable inverse bandwidth that controls the sharpness of the spatial affinity. Given $k^*(x)$, the spatial branch aggregates the original group features by cluster mean within each window:

$$\hat{F}_{\text{spat}}(x, g) = \frac{1}{|\Omega_{k^*(x)}|} \sum_{x' \in \Omega_{k^*(x)}} F(x', g), \quad (13)$$

where Ω_k denotes the set of locations assigned to prototype k in the window. The assignment is shared across group indices $g \in G$, while aggregation is performed separately per group element to obtain the refinement $F_{\text{spat}}(x, g)$.

Branch-Selection Gate. Let F_{spec} and F_{spat} denote the routed outputs of the spectral and spatial branches. Rather than fixed-weight fusion, we learn a position-wise gate that adaptively balances their contributions. We obtain $\bar{F}_{\text{spec}}$ and $\bar{F}_{\text{spat}}$ by group averaging and project them into a shared d -dimensional space using lightweight projections π_s and π_t . The branch-selection gate is:

$$\beta(x) = \sigma(\langle \pi_s(\bar{F}_{\text{spec}}(x)), \pi_t(\bar{F}_{\text{spat}}(x)) \rangle) \in [0, 1], \quad (14)$$

where σ is the sigmoid function and $\langle \cdot, \cdot \rangle$ denotes the channel-wise dot product. Intuitively, $\beta(x)$ quantifies the relative reliability of the spectral branch at each spatial location. The fused output is a convex combination:

$$F'(x, g) = \beta(x) \cdot F_{\text{spec}}(x, g) + (1 - \beta(x)) \cdot F_{\text{spat}}(x, g). \quad (15)$$

This dot product measures agreement between the two branches in a learned subspace. Since β is computed from descriptors averaged over the group and applied consistently across $g \in G$, the fusion operation is equivariance preserving.

3.4 Oriented-Aware Head Network

Detection requires category-score fields that are invariant along the orientation axis while remaining spatially equivariant, whereas box and angle predictions must transform consistently with the input. Let $F_n \in \mathbb{R}^{C \times |G| \times H \times W}$ denote the group feature from the neck and ρ_g the group action defined in Eq. (2). We design the three branches of the head to satisfy (Fig. 2 (c)):

$$f_{\text{cls}}(\rho_g F_n)(x) = f_{\text{cls}}(F_n)(g^{-1}x), \quad (16a)$$

$$f_{\text{box}}(\rho_g F_n) = \rho_g f_{\text{box}}(F_n), \quad (16b)$$

$$f_{\text{ang}}(\rho_g F_n)(x) = \text{wrap}_\pi \left(f_{\text{ang}}(F_n)(g^{-1}x) + \frac{2\pi}{|G|} \text{idx}(g) \right), \quad (16c)$$

where $\text{idx}(g)$ returns the integer index of g in G , and wrap_π denotes the π -periodic OBB angle normalization.

All three branches are built from the depthwise-pointwise equivariant convolutions of Eqs. (5)–(6) with shared parameters across orientations, keeping the head compact. The classification branch achieves invariance via group mean pooling at the readout stage:

$$\mathbf{z}_{\text{inv}} = \frac{1}{|G|} \sum_{g \in G} \mathbf{z}_g, \quad (17)$$

where $\mathbf{z}_g$ is the feature response at orientation g ; since group pooling removes only the orientation index while preserving the transformed spatial coordinate, Eq. (16a) follows directly. The box branch keeps group-indexed features and uses a cyclic-tied DFL readout whose weights depend on relative group offsets, while the angle branch predicts group-indexed residual candidates and decodes them by soft circular aggregation. These readouts preserve the equivariant box-side transformation and the π -periodic angle shift law in Eqs. (16b)–(16c).

4 Experiments

We evaluate FressDet on five multispectral object detection benchmarks: MODA [15], HOD3K [16], and DroneVehicle [44] in the main paper, with LLVIP [18] and VEDAI [39] in the supplementary material. These datasets span diverse spectral configurations, object scales, and imaging conditions, forming a challenging evaluation suite. All results are reported with the standard COCO AP metric [33]. Further dataset details are provided in the supplementary material.

4.1 Implementation Details

The backbone employs four stages, each equipped with Fourier positional encoding of $L=8$ frequencies and an implicit MLP of depth 5 with ReLU activations. The neck applies ReCoW at each fusion level with local window size $W=4$

Table 1: Comparison with other methods on MODA. R-50: ResNet-50, ReR-50: ReResNet-50, ReR-N-50: ReResNet-N-50, CSP-D: CSPDarknet, CSP-N: CSPNeXt, RE-CSP-N: RE-CSPNeXt. #P (M): trainable parameters. #F (G): inference FLOPs.

Methods	Backbone	Car	Bus	Van	Aw.	Tru.	Tri.	Bike	Ped.	mAP ₅₀	mAP ₇₅	mAP	#F	#P
one-stage methods														
R Retina. [32]	R-50	90.2	81.9	72.4	59.0	45.5	30.8	21.7	25.4	53.4	37.3	34.1	233.5	36.5
GWD [65]	R-50	90.4	79.9	70.7	59.7	49.3	30.8	26.7	29.7	54.7	37.8	34.6	233.5	36.5
RepPts-R [67]	R-50	89.9	70.2	67.0	66.1	50.2	41.5	33.3	40.5	57.3	34.9	33.0	213.6	36.9
R3Det [64]	R-50	90.4	88.7	74.0	63.8	57.8	45.1	29.4	32.2	60.2	36.7	34.8	362.2	42.0
R3Det-KLD [66]	R-50	90.4	88.7	74.8	66.6	60.9	42.6	29.8	35.5	61.2	39.4	36.8	306.7	39.6
YOLOv8n [20]	CSP-D	95.8	85.2	63.7	67.4	33.6	35.9	49.6	65.1	62.0	48.0	43.3	12.6	3.2
S ² ANet [13]	R-50	90.4	88.7	74.4	69.1	62.7	48.9	30.1	40.7	63.1	38.7	36.7	216.5	38.8
S2ADet [16]	R-50	90.3	86.6	72.1	71.3	57.2	54.1	35.0	40.8	63.5	41.1	38.9	406.0	65.2
FCOS-R [48]	R-50	90.3	86.3	75.2	71.2	59.0	53.8	34.7	37.4	63.5	41.8	39.1	226.9	32.2
YOLO26m [41]	CSP-D	95.5	85.0	67.3	74.4	35.7	33.8	56.9	66.6	64.4	51.2	46.5	167.2	31.8
RTMDet [34]	CSP-N	90.0	88.4	75.3	71.1	62.8	52.5	42.9	41.0	65.5	41.5	40.1	258.5	52.3
YOLOv8m [20]	CSP-D	96.2	87.1	64.7	73.3	40.4	36.3	57.0	69.4	65.6	53.5	47.9	112.0	26.4
YOLOv12m [47]	CSP-D	96.3	85.3	68.8	74.3	40.3	36.3	57.4	69.5	66.0	52.4	46.8	115.0	24.7
CFA [12]	R-50	90.4	89.0	76.7	69.7	64.4	55.1	43.1	41.5	66.2	43.2	40.6	213.6	36.9
O-RepPts [29]	R-50	90.5	89.2	77.7	71.2	66.2	53.1	43.0	41.1	66.5	44.1	40.9	213.6	36.9
YOLO11m [22]	CSP-D	96.4	88.0	68.2	75.7	44.4	38.0	58.1	68.7	67.2	54.6	48.9	126.2	25.1
DCFL [58]	ReR-50	90.6	89.7	77.5	71.3	67.2	55.5	44.2	44.0	67.5	44.8	41.5	241.0	37.0
MessDet [53]	RE-CSP-N	90.0	89.1	75.4	71.4	70.2	54.9	48.4	45.5	68.1	42.4	40.1	616.0	15.1
OSSDet [15]	R-50	90.5	89.9	79.2	72.7	69.7	58.8	45.3	45.7	69.0	45.9	42.7	263.1	36.5
two-stage methods														
GV [62]	R-50	90.3	89.1	73.8	69.5	66.0	46.7	41.4	22.6	62.4	35.7	34.7	230.8	41.4
RoI-T [10]	R-50	90.5	89.3	75.2	73.3	68.7	51.8	44.5	30.0	65.4	43.4	40.7	244.7	55.3
O-RCNN [56]	R-50	90.5	89.7	74.6	72.5	66.6	54.7	45.1	30.4	65.5	44.0	40.9	244.7	55.3
LSKNet [30]	LSKNet-S	90.4	88.5	74.5	72.5	66.0	52.5	41.0	28.2	64.2	42.8	39.4	221.9	32.6
StripRCNN [69]	R-50	90.5	89.0	76.1	73.6	67.1	56.5	45.2	30.7	66.1	44.0	41.0	231.8	45.2
ReDet [14]	ReR-N-50	90.6	89.6	76.8	72.9	69.6	52.4	41.5	41.0	66.8	44.1	41.3	245.2	31.6
DETR based methods														
AO ² -DETR [9]	R-50	89.4	88.1	76.0	69.5	65.7	54.8	41.9	41.0	65.8	43.0	40.0	240.0	42.0
ARS-DETR [70]	Swin-T	90.5	89.2	77.2	70.6	64.8	53.6	43.2	39.7	66.1	43.4	40.5	254.0	44.2
RHINO [27]	Swin-T	90.7	89.1	77.3	70.8	65.6	54.0	42.9	41.0	66.4	44.0	40.8	319.5	50.8
O-DETR [75]	Swin-T	90.8	89.3	78.8	72.8	68.0	56.0	46.6	40.1	67.8	45.1	41.9	492.0	57.6
FressDet (Ours)	SpelWNet	97.7	90.3	75.4	77.8	59.5	49.4	62.0	72.5	73.1	60.2	54.3	31.7	2.3

and $K = 16$ prototypes arranged on a 4×4 spatial grid. The rotation group is constructed on C_4 in all experiments unless otherwise specified.

FressDet is trained from scratch without data augmentation, as the built-in rotation equivariance provides sufficient geometric robustness. The framework is built on Ultralytics YOLO [20] with default settings, expanding the first convolution to match the number of MSI bands. All experiments are conducted in PyTorch on two NVIDIA RTX 3090 GPUs. Additional details and results are provided in the supplementary material.

4.2 Comparison with State-of-the-Art Methods

Results on MODA. Tab. 1 reports results on MODA. FressDet outperforms OSSDet by 4.1% mAP₅₀, 14.3% mAP₇₅, and 11.6% mAP, with only 6.3% of OSSDet’s parameters and 12.0% of its FLOPs, and achieves large gains on small classes (bike +16.7%, pedestrian +26.8%). Consistently, Fig. 3 demonstrates




Fig. 3: Visualization comparison of detection results and feature maps obtained by different methods on MODA. Please zoom in to see details.

Table 2: Performance comparisons on HOD3K.

Methods	People	Bike	Car	mAP ₅₀	mAP ₇₅	mAP	FLOPs	#P
YOLOF [3]	60.8	67.8	67.6	65.4	31.7	28.6	98.2G	42.1M
CO-DETR [79]	70.2	75.3	88.2	77.9	57.3	51.7	200.7G	64.5M
RetinaNet [32]	85.8	95.0	92.8	91.2	59.4	53.3	205.7G	36.2M
YOLOv5 [21]	79.3	94.0	91.2	88.1	59.8	54.4	48.3G	20.9M
DINO [71]	82.5	92.3	92.5	89.1	61.6	56.3	114.7G	47.6M
Fovea [24]	88.1	95.8	92.7	92.2	62.2	56.8	123.9G	38.0M
Faster RCNN [40]	81.8	94.5	91.7	89.4	62.7	56.9	206.7G	41.1M
ATSS [73]	87.2	94.0	92.5	91.2	64.3	57.4	110.6G	32.2M
TOOD [11]	88.1	90.1	92.7	90.3	65.6	57.8	108.5G	32.1M
MethaneMapper [25]	80.2	95.1	90.9	88.7	63.8	57.9	267.6G	80.3M
VFNet [72]	88.5	96.2	92.2	92.3	66.5	59.0	104.5G	32.8M
S2ADet [16]	87.2	97.7	95.3	93.4	66.2	59.8	169.2G	48.6M
OSSDet [15]	88.3	96.5	95.4	93.4	68.8	60.9	131.2G	36.6M
FressDet (Ours)	88.9	97.1	95.3	93.8	69.1	61.1	35.2G	2.3M

fewer misses and false positives under clutter and low visibility. Fig. 5 (a) further verifies component effectiveness: SpeIW yields more localized activations on target regions through continuous spectral resampling, ReCoW adaptively enhances the reliable branch and suppresses the weaker one, reinforcing target responses while attenuating background noise. Fig. 5 (b) shows that continuous spectral modeling enlarges inter-class margins (e.g., car vs. van) and forms tighter clusters for small classes (bike, pedestrian), improving discriminability.

Results on HOD3K. Tab. 2 shows FressDet surpasses the previous best OSSDet by 0.2% in mAP while achieving the best mAP₅₀ with only 6.3% of OSSDet’s parameters and 26.8% of its FLOPs. Fig. 4 shows qualitative results on HOD3K. Competing methods often miss occluded or background-blended objects, whereas FressDet enhances target-background separation, yielding predictions closer to the ground truth.



Fig. 4: Detection results on HOD3K. Red boxes indicate the ground truth, while green boxes denote the predictions.

Table 3: Performance comparisons on DroneVehicle.

Methods	Car	Bus	Truck	Freight-car	Van	mAP ₅₀	mAP
C^2 Former- S^2 ANet [68]	90.2	89.8	68.3	64.4	58.5	74.2	47.3
SLBAF-Net [5]	90.2	89.9	76.0	68.2	59.9	76.8	49.5
CSOM-ODAF [2]	90.1	89.8	75.6	68.2	61.8	77.1	50.1
Multimodal DINO [43]	89.5	88.8	75.4	54.3	54.3	72.5	50.3
CMA [19]	95.8	93.1	75.9	59.8	59.4	76.8	50.4
CRSIOD [50]	95.6	92.2	71.7	50.5	55.8	73.2	50.8
YOLOFIV [49]	95.9	91.6	64.2	34.6	37.3	64.7	53.1
GLFNet [23]	90.3	88.0	72.7	53.6	52.6	71.4	54.8
Dual-YOLO [1]	95.9	91.6	69.7	55.9	46.6	71.9	55.2
DMM [76]	90.4	88.7	77.8	63.0	66.0	77.2	55.8
CAFN-IA [59]	89.1	90.8	62.0	57.3	47.1	69.3	56.1
IV-YOLO [46]	97.2	94.3	65.4	63.1	53.0	74.6	56.8
WaveMamba [77]	95.0	90.6	80.4	68.5	64.5	79.8	60.5
FressDet (Ours)	98.6	95.6	79.0	65.9	62.4	80.3	61.6

Results on DroneVehicle. As shown in Tab. 3, our method surpasses the best [77] by 1.1% mAP and 0.5% mAP₅₀ under identical settings. FressDet achieves the best Car and Bus mAP₅₀, leading by 1.4% and 1.3% over the best competing methods in each category. FressDet benefits from the representational continuity that naturally emerges along the spectral dimension.

4.3 Ablation Studies and Analysis

All ablations are conducted on MODA under identical settings unless noted.

Effectiveness of FressDet’s Key Components. As shown in Tab. 4, introducing rotation-equivariant group convolutions alone yields a +4.4% mAP gain over the baseline, confirming that geometric priors are as critical as spectral modeling for multispectral detection. Each subsequent module provides improvements, and the full model achieves the best result across all metrics. Among the three proposed modules, SpeIW contributes the most, underscoring the value of continuous spectral modeling in multispectral object detection.




Fig. 5: (a) Effectiveness visualization of key components. (b) 2D t-SNE [35] visualization of feature distributions on MODA.

Table 4: Ablation studies on the key components of FressDet. SIW: *SpeIW*, RCW: *ReCoW*, OH: *Oriented-Aware Head*.

Equiv.	SIW	RCW	OH	mAP ₅₀	mAP ₇₅	mAP
✗	✗	✗	✗	62.0	48.0	43.3
✓	✗	✗	✗	66.8	54.1	47.7
✓	✓	✗	✗	69.9	55.8	50.3
✓	✗	✓	✗	69.0	55.1	49.6
✓	✗	✗	✓	68.6	55.5	49.3
✓	✓	✗	✓	71.7	58.8	52.5
✓	✓	✗	✓	71.6	58.4	52.2
✓	✓	✓	✓	73.1	60.2	54.3

Table 5: Comparison of spectral operators in the *SpeIWNet* backbone.

Method	mAP ₅₀	mAP ₇₅	mAP
C ₄ -CNN	71.4	58.6	52.1
RE-DCN	71.6	58.7	52.9
RE-DCNv2	72.0	59.3	53.2
RE-DCNv3	71.7	59.0	53.1
RE-DCNv4	71.9	59.2	53.0
SpeIW	73.1	60.2	54.3

Ablation on Spectral Operator. We ablate the spectral operator by replacing *SpeIW* with rotation-equivariant DCN variants [8, 51, 57, 78] (implementation details in the supplementary material). Tab. 5 shows that the plain C₄-CNN performs worst, underscoring the value of adaptive spectral resampling. *SpeIW* surpasses all DCN-based alternatives, with the closest trailing by 1.1 mAP, indicating that order-preserving continuous resampling along the spectral dimension provides a stronger inductive bias than discrete deformable offsets.

Transferability Across Detectors. To evaluate *SpeIWNet*’s generalization capability, we replace the backbones of five detectors while keeping other components and training unchanged. Fig. 6 shows consistent performance gains across single- and two-stage architectures, along with reduced model size, demonstrating *SpeIWNet*’s effectiveness as a strong plug-and-play feature extractor.

Rotation Equivariance Error. To quantify how well rotation equivariance is preserved, we measure the equivariance error at each backbone stage, defined as the RMS difference between rotating the input before and the output after the




Fig. 6: SpeIWNNet backbone generalization across detectors on MODA.

Table 6: Ablation on group order in the FressDet backbone.

Group	mAP ₅₀	mAP ₇₅	mAP	#F(G)	#P(M)
C_2	71.2	58.3	51.8	15.4	1.1
C_4	73.1	60.2	54.3	31.7	2.3
C_8	73.2	60.3	54.3	65.1	4.8




Fig. 7: Equivariance error across backbone stages (S0: stem, S1–S4: stages 1–4).

Table 7: Fourier frequency L ablation in the SpeIW module.

L	mAP ₅₀	mAP ₇₅	mAP
0	71.5	58.7	52.6
2	72.3	59.4	53.4
4	72.8	59.9	54.0
8	73.1	60.2	54.3
16	72.9	60.0	54.1

network [53]: Fig. 7 shows FressDet maintains low equivariance error across all stages, while OSSDet exhibits consistently higher errors.

Rotation Group Order. We compare C_2 , C_4 , and C_8 groups in Tab. 6. Higher group order improves accuracy at the cost of increased parameters and FLOPs. C_4 achieves the best balance between accuracy and efficiency: C_8 incurs more than $2\times$ FLOPs for only marginal improvement, suggesting that 90° discretization already captures the dominant orientation variations.

Fourier Frequency in SpeIW. Tab. 7 ablates the Fourier positional encoding frequency L in SpeIW. Increasing L improves high-frequency spectral modeling, whereas excessive frequencies may overfit; $L=8$ gives the best performance.

Spectral Sensitivity. Tab. 8 compares RGB/MSI inputs across MODA detectors. MSI improves every method, with FressDet showing the largest gain (+2.7 mAP₅₀ and +3.7 mAP); this supports continuous spectral modeling as a better way to exploit the available bands more fully.

Spectral-Spatial Routing. Tab. 9 shows that routing should match the branch property. Soft spectral routing preserves diversity, while hard spatial routing keeps sharper boundaries; the mismatched alternatives lose one of these cues.

Prototype Number in ReCoW. Tab. 10 shows that the 4×4 grid ($K=16$) best balances spatial resolution and redundancy; smaller grids miss local patterns, while larger grids introduce redundant prototypes and tend to overfit.

Branch Fusion Strategy in ReCoW. We ablate the fusion strategy for combining the spectral and spatial branches in ReCoW (Tab. 11). Our learned gate

Table 8: Spectral sensitivity on MODA. Entries report RGB/MSI results with gains in parentheses.

Method	mAP ₅₀	mAP	#F(G)	#P(M)	Spe. Spa.	mAP ₅₀	mAP
S-RCNN	65.2/66.1(+0.9)	39.6/41.0(+1.4)	227.4/231.8	45.13/45.15	Hard Hard	70.4	51.2
O-RepPoints	66.2/66.5(+0.3)	39.4/40.9(+1.5)	209.2/213.6	36.83/36.85	Soft Soft	71.6	52.4
OSSDet	67.1/69.0(+1.9)	39.7/42.7(+3.0)	258.7/263.1	36.46/36.48	Hard Soft	69.8	50.6
FressDet	70.4/73.1(+2.7)	50.6/54.3(+3.7)	31.0/31.7	2.29/2.30	Soft Hard	73.1	54.3

Table 9: Routing ablation in ReCoW.**Table 10:** Prototype grid ablation in ReCoW.

K	Grid	mAP ₅₀	mAP ₇₅	mAP
4	2×2	71.8	59.0	52.9
9	3×3	72.5	59.6	53.6
16	4×4	73.1	60.2	54.3
25	5×5	72.8	59.9	54.0

Table 11: Fusion strategy ablation in ReCoW.

Fusion	mAP ₅₀	mAP ₇₅	mAP
Addition	71.9	59.0	52.8
Concat.	72.4	59.5	53.4
Softmax	72.7	59.8	53.8
β gate	73.1	60.2	54.3

Table 12: Dual-branch ablation in ReCoW.

Spe.	Spa.	mAP ₅₀	mAP ₇₅	mAP
✓	✗	72.3	59.5	53.6
✗	✓	71.8	59.1	53.2
✓	✓	73.1	60.2	54.3

outperforms all alternatives, confirming the benefit of adaptive branch selection over fixed or implicit fusion.

Dual-Branch Design in ReCoW. We ablate each ReCoW branch (Tab. 12). The spectral and spatial branches achieve 72.3% and 71.8% mAP₅₀, respectively. Their combination reaches 73.1%, confirming that the two branches capture complementary information and yield consistent improvements when combined.

5 Conclusion

We present FressDet, the first fully rotation-equivariant spectral-spatial learning framework for multispectral object detection. FressDet addresses discrete spectral processing, spectral-spatial reliability shifts, and arbitrary orientations through three components: SpeIW for continuous, monotonic spectral resampling, ReCoW for reliable spectral-spatial fusion across pyramid levels, and an oriented-aware head for stable orientation prediction without parameter replication. Extensive experiments on five public benchmarks demonstrate that FressDet achieves state-of-the-art accuracy with the fewest parameters among compared multispectral detectors and exhibits the best robustness even under input rotational perturbations.

Acknowledgements

This work was financially supported by the National Natural Science Foundation of China (No. 62101032), the Chongqing Excellent Young Scientists Fund (No. CSTB2025NSCQ-JQX0017), and the High-Quality Development Special Project of the Ministry of Industry and Information Technology (TC240HAJ9-35).

References

1. Bao, C., Cao, J., Hao, Q., Cheng, Y., Ning, Y., Zhao, T.: Dual-YOLO architecture from infrared and visible images for object detection. *Sensors* **23**(6), 2934 (2023)
2. Chen, C., Qi, J., Liu, X., Bin, K., Fu, R., Hu, X., Zhong, P.: Weakly misalignment-free adaptive feature alignment for UAVs-based multimodal object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 26836–26845 (2024)
3. Chen, Q., Wang, Y., Yang, T., Zhang, X., Cheng, J., Sun, J.: You only look one-level feature. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 13039–13048 (2021)
4. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8628–8638 (June 2021)
5. Cheng, X., Geng, K., Wang, Z., Wang, J., Sun, Y., Ding, P.: SLBAF-Net: Super-lightweight bimodal adaptive fusion network for UAV detection in low recognition environment. *Multimedia Tools and Applications* **82**(30), 47773–47792 (2023)
6. Cohen, T.S., Welling, M.: Group equivariant convolutional networks. In: *ICML* (2016)
7. Cybenko, G.: Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems* **2**(4), 303–314 (1989). <https://doi.org/10.1007/BF02551274>
8. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. pp. 764–773 (2017). <https://doi.org/10.1109/ICCV.2017.89>
9. Dai, L., Liu, H., Tang, H., Wu, Z., Song, P.: AO2-DETR: Arbitrary-oriented object detection transformer. *IEEE transactions on circuits and systems for video technology* **33**(5), 2342–2356 (2022)
10. Ding, J., Xue, N., Long, Y., Xia, G.S., Lu, Q.: Learning RoI transformer for oriented object detection in aerial images. In: *CVPR*. pp. 2849–2858 (2019)
11. Feng, C., Zhong, Y., Gao, Y., Scott, M.R., Huang, W.: TOOD: Task-aligned one-stage object detection. In: *ICCV* (2021)
12. Guo, Z., Liu, C., Zhang, X., Jiao, J., Ji, X., Ye, Q.: Beyond bounding-box: Convex-hull feature adaptation for oriented and densely packed object detection. In: *CVPR*. pp. 8792–8801 (2021)
13. Han, J., Ding, J., Li, J., Xia, G.S.: Align deep features for oriented object detection. *IEEE transactions on geoscience and remote sensing* **60**, 1–11 (2021)
14. Han, J., Ding, J., Xue, N., Xia, G.S.: ReDet: A rotation-equivariant detector for aerial object detection. In: *CVPR*. pp. 2786–2795 (2021)
15. Han, S., Xu, T., Liu, P., Li, J.: MODA: The first challenging benchmark for multi-spectral object detection in aerial images. *arXiv preprint arXiv:2512.09489* (2025)
16. He, X., Tang, C., Liu, X., Zhang, W., Sun, K., Xu, J.: Object detection in hyperspectral image via unified spectral-spatial feature aggregation. *IEEE Transactions on Geoscience and Remote Sensing* **61**, 1–13 (2023). <https://doi.org/10.1109/TGRS.2023.3307288>
17. Hornik, K.: Approximation capabilities of multilayer feedforward networks. *Neural Networks* **4**(2), 251–257 (1991). [https://doi.org/10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T)
18. Jia, X., Zhu, C., Li, M., Tang, W., Zhou, W.: LLVIP: A visible-infrared paired dataset for low-light vision. In: *ICCV*. pp. 3496–3504 (2021)

19. Jiang, C., Ren, H., Yang, H., Huo, H., Zhu, P., Yao, Z., Li, J., Sun, M., Yang, S.: M2FNet: Multi-modal fusion network for object detection from visible and thermal infrared images. *International Journal of Applied Earth Observation and Geoinformation* **130**, 103918 (2024). <https://doi.org/10.1016/j.jag.2024.103918>
20. Jocher, G., Chaurasia, A., Qiu, J.: YOLOv8. Ultralytics (2023), available at <https://github.com/ultralytics/ultralytics>. Accessed 5 March 2026
21. Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., Kwon, Y., Michael, K., Fang, J., Yifu, Z., Wong, C., Montes, D., et al.: YOLOv5. GitHub repository (2020), available at <https://github.com/ultralytics/yolov5>. Accessed 5 March 2026
22. Jocher, G., Qiu, J.: Ultralytics YOLO11 (2024), available at <https://github.com/ultralytics/ultralytics>. Accessed 5 March 2026
23. Kang, X., Yin, H., Duan, P.: Global-local feature fusion network for visible-infrared vehicle detection. *IEEE Geoscience and Remote Sensing Letters* **21**, 1–5 (2024). <https://doi.org/10.1109/LGRS.2024.3375634>
24. Kong, T., Sun, F., Liu, H., Jiang, Y., Li, L., Shi, J.: Foveabox: Beyond anchor-based object detection. *IEEE Transactions on Image Processing* **29**, 7389–7398 (2020)
25. Kumar, S., Arevalo, I., Iftekhhar, A., Manjunath, B.: MethaneMapper: Spectral absorption aware hyperspectral transformer for methane detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17609–17618 (2023)
26. Lee, C., Son, J., Shon, H., Jeon, Y., Kim, J.: FRED: Towards a full rotation-equivariance in aerial image object detection. In: *AAAI*. pp. 2883–2891 (2024). <https://doi.org/10.1609/aaai.v38i4.28069>
27. Lee, H., Song, M., Koo, J., Seo, J.: Hausdorff distance matching with adaptive query denoising for rotated detection transformer. In: *Proceedings of the Winter Conference on Applications of Computer Vision (WACV)*. pp. 1872–1882 (February 2025)
28. Li, R., Liu, J., Liu, G., Zhang, S., Zeng, B., Liu, S.: SpectralNeRF: Physically based spectral rendering with neural radiance field. In: Wooldridge, M.J., Dy, J.G., Natarajan, S. (eds.) *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Vancouver, Canada, February 20-27, 2024*. pp. 3154–3162. AAAI Press (2024). <https://doi.org/10.1609/AAAI.V38I4.28099>
29. Li, W., Chen, Y., Hu, K., Zhu, J.: Oriented RepPoints for aerial object detection. In: *CVPR*. pp. 1829–1838 (2022)
30. Li, Y., Li, X., Dai, Y., Hou, Q., Liu, L., Liu, Y., Cheng, M.M., Yang, J.: LSKNet: A foundation lightweight backbone for remote sensing. *International Journal of Computer Vision* (2024). <https://doi.org/10.1007/s11263-024-02247-9>
31. Liang, Y., Cao, Z., Deng, S., Dou, H., Deng, L.: Fourier-enhanced implicit neural fusion network for multispectral and hyperspectral image fusion. In: Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J.M., Zhang, C. (eds.) *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024* (2024)
32. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. In: *ICCV* (2017)
33. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: *Computer Vision – ECCV 2014*. pp. 740–755 (2014). https://doi.org/10.1007/978-3-319-10602-1_48

34. Lyu, C., Zhang, W., Huang, H., Zhou, Y., Wang, Y., Liu, Y., Zhang, S., Chen, K.: RTMDet: An empirical study of designing real-time object detectors. arXiv preprint arXiv:2212.07784 (2022)
35. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *Journal of machine learning research* **9**(11) (2008)
36. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 12346, pp. 405–421. Springer (2020). https://doi.org/10.1007/978-3-030-58452-8_24
37. Pearson, K.: 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 (1901)
38. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F.A., Bengio, Y., Courville, A.: On the spectral bias of neural networks (2019)
39. Razakarivony, S., Jurie, F.: Vehicle detection in aerial imagery: A small target detection benchmark. *Journal of Visual Communication and Image Representation* **34**, 187–203 (2016)
40. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: *NeurIPS* (2015)
41. Sapkota, R., Cheppally, R.H., Sharda, A., Karkee, M.: YOLO26: Key architectural enhancements and performance benchmarking for real-time object detection. arXiv preprint arXiv:2509.25164 (2025)
42. Sun, W., Du, Q.: Hyperspectral band selection: A review. *IEEE Geoscience and Remote Sensing Magazine* **7**(2), 118–139 (2019)
43. Sun, X., Yu, Y., Cheng, Q.: Adaptive multimodal feature fusion with frequency domain gate for remote sensing object detection. *Remote Sensing Letters* **15**(2), 133–144 (2024). <https://doi.org/10.1080/2150704X.2024.2305177>
44. Sun, Y., Cao, B., Zhu, P., Hu, Q.: Drone-based RGB-infrared cross-modality vehicle detection via uncertainty-aware learning. *IEEE TCSVT* **32**(10), 6700–6713 (2022). <https://doi.org/10.1109/TCSVT.2022.3168279>
45. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains (2020)
46. Tian, D., Yan, X., Zhou, D., Wang, C., Zhang, W.: IV-YOLO: A lightweight dual-branch object detection network. *Sensors* **24**(19), 6181 (2024). <https://doi.org/10.3390/s24196181>
47. Tian, Y., Ye, Q., Doermann, D.: YOLOv12: Attention-centric real-time object detectors. arXiv preprint arXiv:2502.12524 (2025)
48. Tian, Z., Shen, C., Chen, H., He, T.: FCOS: Fully convolutional one-stage object detection. In: *ICCV*. pp. 9626–9635 (2019)
49. Wang, H., Wang, C., Fu, Q., Si, B., Zhang, D., Kou, R., Yu, Y., Feng, C.: YOLOFIV: Object detection algorithm for around-the-clock aerial remote sensing images by fusing infrared and visible features. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **17**, 15269–15287 (2024)
50. Wang, H., Wang, C., Fu, Q., Zhang, D., Kou, R., Yu, Y., Song, J.: Cross-modal oriented object detection of UAV aerial images based on image feature. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–21 (2024)

51. Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., Wang, X., Qiao, Y.: InternImage: Exploring large-scale vision foundation models with deformable convolutions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14408–14419 (2023). <https://doi.org/10.1109/CVPR52729.2023.01385>
52. Weiler, M., Cesa, G.: General E(2)-equivariant steerable CNNs. In: NeurIPS (2019)
53. Wu, X., Wang, X., Zhu, X., Yang, L., Liu, J., Hu, X.: Measuring the impact of rotation equivariance on aerial object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2025)
54. Wu, Z., Liu, Y., Dong, H., Tang, X., Yang, J., Jin, B., Chen, M., Wei, X.: R2Det: Exploring relaxed rotation equivariance in 2D object detection. In: International Conference on Learning Representations (ICLR) (2025)
55. Wu, Z., Liu, Y., Sun, L., Yang, J., Dong, H., Lin, S.J., Tang, X., Mi, J., Jin, B., Wei, X.: Relaxed rotational equivariance via G-biases in vision. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). pp. 8541–8549 (2025). <https://doi.org/10.1609/AAAI.V39I8.32922>, arXiv:2408.12454
56. Xie, X., Cheng, G., Wang, J., Yao, X., Han, J.: Oriented R-CNN for object detection. In: ICCV. pp. 3520–3529 (2021)
57. Xiong, Y., Li, Z., Chen, Y., Wang, W., Zhu, X., Zhang, L., Lu, T., Dai, J., Qiao, Y.: Efficient deformable convnets: Rethinking dynamic and sparse operator for vision applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5765–5775 (2024). <https://doi.org/10.1109/CVPR52733.2024.00551>
58. Xu, C., Ding, J., Wang, J., Yang, W., Yu, H., Yu, L., Xia, G.S.: Dynamic coarse-to-fine learning for oriented tiny object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7318–7328 (June 2023)
59. Xu, J., Mo, B., Zhao, J., Zhao, C., Tao, Y., Han, S.: Cross-modal adaptive fusion object detection based on illumination-awareness. In: 2024 39th Youth Academic Annual Conference of Chinese Association of Automation (YAC). pp. 931–938. IEEE (2024)
60. Xu, X., Wang, Z., Shi, H.: UltraSR: Spatial encoding is a missing key for implicit image function-based arbitrary-scale super-resolution (2022)
61. Xu, X., Liu, H., Zhang, T., Xiong, H., Yu, W.: PreCM: The padding-based rotation equivariant convolution mode for semantic segmentation. IEEE Transactions on Image Processing **34**, 2781–2795 (2025). <https://doi.org/10.1109/TIP.2025.3558425>, arXiv:2411.01624
62. Xu, Y., Fu, M., Wang, Q., Wang, Y., Chen, K., Xia, G.S., Bai, X.: Gliding vertex on the horizontal bounding box for multi-oriented object detection. IEEE transactions on pattern analysis and machine intelligence **43**(4), 1452–1459 (2020)
63. Yan, L., Zhao, M., Wang, X., Zhang, Y., Chen, J.: Object detection in hyperspectral images. IEEE Signal Processing Letters **28**, 508–512 (2021)
64. Yang, X., Yan, J., Feng, Z., He, T.: R3Det: Refined single-stage detector with feature refinement for rotating object. In: AAAI. vol. 35, pp. 3163–3171 (2021)
65. Yang, X., Yan, J., Ming, Q., Wang, W., Zhang, X., Tian, Q.: Rethinking rotated object detection with gaussian wasserstein distance loss. In: International conference on machine learning. pp. 11830–11841. PMLR (2021)
66. Yang, X., Yang, X., Yang, J., Ming, Q., Wang, W., Tian, Q., Yan, J.: Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. Advances in Neural Information Processing Systems **34**, 18381–18394 (2021)

67. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: ICCV. pp. 9657–9666 (2019)
68. Yuan, M., Wei, X.: C²Former: Calibrated and complementary transformer for RGB-infrared object detection. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–12 (2024)
69. Yuan, X., Zheng, Z., Li, Y., Liu, X., Liu, L., Li, X., Hou, Q., Cheng, M.M.: Strip R-CNN: Large strip convolution for remote sensing object detection. *arXiv preprint arXiv:2501.03775* (2025)
70. Zeng, Y., Chen, Y., Yang, X., Li, Q., Yan, J.: ARS-DETR: Aspect ratio-sensitive detection transformer for aerial oriented object detection. *IEEE transactions on geoscience and remote sensing* **62**, 1–15 (2024)
71. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: DINO: DETR with improved denoising anchor boxes for end-to-end object detection. In: ICLR (2023)
72. Zhang, H., Wang, Y., Dayoub, F., Sunderhauf, N.: VarifocalNet: An IoU-aware dense object detector. In: CVPR. pp. 8514–8523 (2021)
73. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: CVPR (2020)
74. Zhang, Y., Qi, J., Wang, X., Cai, Z., Peng, J., Zhou, Y.: Tensorial global-local graph self-representation for hyperspectral band selection. *IEEE Transactions on Circuits and Systems for Video Technology* (2024)
75. Zhao, Z., Xue, Q., He, Y., Bai, Y., Wei, X., Gong, Y.: Projecting points to axes: Oriented object detection via point-axis representation. In: ECCV. pp. 161–179 (2024). https://doi.org/10.1007/978-3-031-73390-1_10
76. Zhou, M., Li, T., Qiao, C., Xie, D., Wang, G., Ruan, N., Mei, L., Yang, Y.: DMM: Disparity-guided multispectral mamba for oriented object detection in remote sensing (2024)
77. Zhu, H., Dong, W., Yang, L., Li, H., Yang, Y., Ren, Y., Zhu, Q., Feng, Z., Li, C., Lin, S., et al.: WaveMamba: Wavelet-driven mamba fusion for RGB-infrared object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11219–11229 (2025)
78. Zhu, X., Hu, H., Lin, S., Dai, J.: Deformable convnets v2: More deformable, better results. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
79. Zong, Z., Song, G., Liu, Y.: DETRs with collaborative hybrid assignments training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6748–6758 (2023)