

RiO-DETR: DETR for Real-time Oriented Object Detection

Zhangchi Hu¹, Yifan Zhao¹, Yansong Peng¹, Wenzhang Sun³, Jie Chen¹, Peixi Wu¹, Xinghao Wang², Xiangchen Yin¹, Dongsheng Jiang², Hebei Li^{1†}, and Xiaoyan Sun^{1,4}

¹ University of Science and Technology of China

² Huawei Technologies Co., Ltd.

³ Tsinghua University

⁴ Institute of Artificial Intelligence, Hefei Comprehensive National Science Center
{huzhangchi, lihebei}@mail.ustc.edu.cn, sunxiaoyan@ustc.edu.cn

[†] Corresponding author




Fig. 1: Comparisons with other detectors in terms of model size (left), latency (mid), and computational cost (right) on DOTA-1.0 under single-scale training and testing protocol. * denotes a community implemented version.

Abstract. We present **RiO-DETR: DETR for Real-time Oriented Object Detection**, the first real-time oriented detection transformer to the best of our knowledge. Adapting DETR to oriented bounding boxes (OBBs) poses three challenges: semantics-dependent orientation, angle periodicity that breaks standard Euclidean refinement, and an enlarged search space that slows convergence. RiO-DETR resolves these issues with task-native designs while preserving real-time efficiency. First, we propose Content-Driven Angle Estimation by decoupling angle from positional queries, together with Rotation-Rectified Orthogonal Attention to capture complementary cues for reliable orientation. Second, Decoupled Periodic Refinement combines bounded coarse-to-fine updates with a Shortest-Path Periodic Loss for stable learning across angular seams. Third, Oriented Dense O2O injects angular diversity into dense supervision to speed up angle convergence at no extra cost. Extensive experiments on DOTA-1.0, DIOR-R, and FAIR-1M-2.0 demonstrate RiO-DETR establishes a new speed-accuracy trade-off for real-time oriented detection. Code is available at <https://github.com/RicePasteM/RiO-DETR>.

Keywords: Oriented Object Detection · Detection Transformer · Aerial Object Detection · Real-time Object Detection

1 Introduction

Oriented object detection extends horizontal bounding boxes (HBBs) to localize objects with arbitrary rotations, which is crucial for aerial imagery, remote sensing, and scene text understanding [6, 9, 12, 34, 35, 39, 40, 42, 46, 63, 65]. With the rise of edge computing, the bottleneck has evolved from mere detection accuracy to the speed-accuracy trade-off. This raises a critical question: how can we maintain high-quality oriented bounding boxes (OBBs) detection while meeting rigorous real-time performance requirements?

In the CNN paradigm, real-time oriented detectors have established robust baselines, especially oriented variants of YOLO and RTMDet-style frameworks [11, 13, 20, 31, 45, 47, 49]. Meanwhile, DETR-style detectors have recently achieved real-time efficiency in horizontal detection by leveraging lightweight designs and efficient training strategies [4, 16, 19, 28, 33, 36, 60]. However, real-time oriented DETRs remain under-explored. Existing oriented DETR variants are often burdened by heavy attention designs and complex sampling modules in pursuit of peak accuracy [7, 17, 22, 57, 59, 61], which inherently limits their ability to match the high throughput of CNN-based counterparts.

This gap is not primarily due to insufficient model capacity; rather, it reflects inherent architectural bottlenecks that emerge when adapting DETR from HBB to OBB. We identify three key bottlenecks that systematically limit real-time oriented DETRs: **(1) Semantic–Geometric Coupling and Feature Collapse.** Unlike the standard (cx, cy, w, h) format, the orientation θ is strongly driven by semantic appearance cues such as texture flow and dominant axes. Encoding θ solely as a geometric prior in positional queries can introduce noise and misguide the attention mechanism. Furthermore, naively aligning attention with the major axis risks feature collapse, where lateral structures are insufficiently attended. **(2) Periodicity Mismatch in Angle Refinement.** Standard DETR decoders typically refine bounding boxes through Euclidean additive updates, such as the inverse-sigmoid formulation. Applying such updates to the cyclic angular domain introduces discontinuities at periodic boundaries, leading to unstable gradients and unreliable refinement [5, 51–54]. **(3) Slow Convergence in the Expanded Search Space.** OBBs introduce additional degrees of freedom, which significantly expands the search space for bipartite matching and slows convergence. While techniques like dense supervision or one-to-many training are effective for HBB detection, they often lack sufficient angular diversity to accelerate orientation learning in the OBB setting.

These observations motivate a fundamental shift in design: achieving real-time oriented detection requires moving beyond simply appending an angle branch or increasing computational overhead. Instead, we focus on reformulating core components to natively handle oriented geometry. To this end, we present **RiO-DETR**, a **R**eal-time **O**riented detection transformer that effec-

tively bridges the gap between accuracy and latency. Built upon a delicately implemented baseline, RiO-DETR introduces three task-specific designs to resolve the aforementioned bottlenecks:

- To decouple semantic and geometric cues, we introduce **Content-Driven Angle Estimation**. This approach employs a **Geometry-Decoupled Query Encoding** that facilitates angle prediction by leveraging semantic context rather than relying on rigid geometric priors. To address feature collapse with minimal overhead, we further propose **Rotation-Rectified Orthogonal Attention**, which captures both axial and lateral cues to enhance orientation inference.
- To address periodicity issues, we introduce **Decoupled Periodic Refinement**. Specifically, we replace standard Euclidean updates with a bounded coarse-to-fine periodic mechanism and a **Shortest-path Periodic L_1 Loss**. This approach enables stable optimization across angular boundaries and resolves refinement artifacts caused by periodic discontinuities.
- To address slow convergence, we develop **Oriented Dense O2O**, a training strategy that injects angular diversity by applying independent random rotations to stitched image quadrants, significantly accelerating the convergence of angle predictions.

RiO-DETR achieves a superior speed–accuracy balance for real-time oriented detection (Fig. 1). On DOTA-1.0, RiO-DETR-n attains **78.4 AP₅₀** with only **2.7 ms** end-to-end latency (TensorRT FP16 on NVIDIA T4). Meanwhile, RiO-DETR-x reaches **81.8 AP₅₀** at **29.9 ms**. Both models outperform state-of-the-art real-time detectors at similar speeds. Our approach also shows consistent improvements on DIOR-R and FAIR-1M-2.0 across various scales. These results demonstrate that end-to-end transformers can be highly efficient for oriented object detection. We hope our design provides a robust framework and inspires further exploration in this field.

2 Related Works

CNN-based Oriented Object Detection. Early oriented detectors mainly extended horizontal CNN detectors with angle regression. Two-stage methods [27, 47, 50] (e.g., RoI Transformer [8]) achieve high precision via rotation-invariant RoI operations, while single-stage/anchor-free counterparts (e.g., S2ANet [56], FCOS-O [42]) simplify the pipeline for higher efficiency. Recently, oriented YOLO variants [11, 20, 21, 37, 38, 41, 49] and RTMDet-R [31] have advanced real-time oriented detection. However, most still depend on dense prediction and heuristic anchors, typically requiring NMS, which complicates deployment and adds hyperparameter sensitivity. Moreover, lightweight designs for oriented detection [2, 25, 55] often introduce complex structures that raise compute/memory access costs, leading to noticeable inference latency.

Real-time Object Detection. While the YOLO [20] ecosystem long dominated real-time detection, the emergence of DETR [3] introduced a NMS-free, end-to-end paradigm. Early DETR variants [23, 29, 32, 58, 64] struggled with computational overhead, but recent models like RT-DETR [4, 60] have bridged this gap using efficient hybrid encoders and uncertainty-minimal query selection. Further advancements [16, 26, 30, 36, 44], such as D-FINE [33] and DEIM [18, 19], have enhanced precision and accelerated convergence through refined optimization and dense supervision. These breakthroughs prove the feasibility of real-time DETRs, yet their success remains largely confined to horizontal bounding boxes.

DETR-based Oriented Object Detection. Motivated by the success of DETRs [3], several works have attempted to adapt the Transformer architecture for oriented bounding boxes [59]. Methods like AO2-DETR [7] and EMO2-DETR [17] introduce oriented proposal generation and rotation-aware attention mechanisms to align features with arbitrary object orientations. Recent state-of-the-art models, such as ARS-DETR [57], Oriented-DETR [61] and RHINO [22], further address the matching and representation problem of OBBs. Despite their strong performance, these methods predominantly build upon heavy architectures and prioritize precision over inference speed, excluding them from real-time applications. To the best of our knowledge, designing a DETR-based oriented detector that maintains high accuracy while meeting strict real-time constraints remains an unresolved challenge.

3 Methods

3.1 Best of Both Worlds: Building a Strong Real-time Baseline

While numerous studies have migrated oriented object detection to the DETR family, most of them rely on DINO [58] or Deformable-DETR [64], which is not fully optimized for inference efficiency. RT-DETRv2 is originally designed for horizontal object detection and does not come with an oriented detection (OBB) implementation. Therefore, we first develop an oriented RT-DETRv2 baseline by introducing an OBB regression head and adapting the angle representation, losses, and denoising scheme following RHINO-DETR [22]. To make the baseline strong and fair under real-time constraints, we further consolidate two complementary training designs that have proven effective for DETR-style detectors: the universal matching strategy from D-FINE and Dense O2O from DEIM [19]. Additional implementation details are provided in Appendix A.

Table 6 shows the comparison between our implemented baseline and previous works. Our implementation exhibits competitive performance, which forms a solid foundation for our method. However, there remains a noticeable gap compared with state-of-the-art counterparts (1.18 AP₅₀), indicating that simply porting existing DETR techniques is insufficient for closing the accuracy-efficiency gap in oriented detection. It is our task-oriented improvements that solve the rooted issues in DETR-based oriented object detection, and elevate RiO-DETR




Fig. 2: The main architecture of our proposed RiO-DETR. The framework highlights three key components: (A) Content-Driven Angle Estimation, (B) Rotation-Rectified Orthogonal Attention, and (C) Decoupled Periodic Refinement.

to state-of-the-art accuracyefficiency performance while keeping the parameter count, FLOPs, and inference latency virtually unchanged.

3.2 Content-Driven Angle Estimation

Existing DETR-based oriented object detectors treat the orientation θ as a geometric component symmetric to box coordinates (c_x, c_y, w, h) . Consequently, these methods [7, 17, 22, 57, 59, 61] intuitively embed the full 5-dimensional tuple (c_x, c_y, w, h, θ) into the query embeddings. This joint design implicitly assumes that spatial localization and angular rotation can be optimized in a homogeneous Euclidean manner, letting the decoder refine the entire geometric state as a unified representation.

However, unlike (c_x, c_y) , the OBB orientation θ is not uniquely determined by geometric correspondence. It is periodic and admits equivalent parameterizations; for instance, (w, h, θ) and $(h, w, (\theta + \pi/2) \bmod \pi)$ describe the same physical rectangle. As a result, θ is effectively a canonical choice induced by annotation conventions (e.g., long-side direction or heading) rather than a purely geometric quantity, and is often disambiguated only by appearance cues such as dominant axes, part layouts, texture flows, and semantic definitions. Injecting θ into positional embeddings therefore imposes a rigid geometric prior that can be noisy during early training and potentially non-smooth around periodic boundaries, which may misguide attention sampling and interfere with content-driven refinement in the decoder. We provide a more formal discussion from the viewpoint of quotient-consistency in Appendix B.

Motivated by this asymmetry, we propose **Content-Driven Angle Estimation** (illustrated in Fig. 2(A)(B)). Specifically, we decouple orientation from

the positional prior via **Geometry-Decoupled Query Encoding**, encouraging the model to regress angles primarily from semantic context learned through content queries. Then, we design **Rotation-Rectified Orthogonal Attention** to extract rotation-aligned object features for more robust angle estimation.

Geometry-Decoupled Query Encoding We formulate the object query Q as a composition of a content part $Q_{content}$ and a positional part Q_{pos} . Unlike previous works, we strictly limit the positional embedding to the 4-dimensional spatial domain, explicitly excluding angular information.

Let $\mathbf{p}_{ref} \in \mathbb{R}^{N \times 5}$ denote the reference points containing (cx, cy, w, h, θ) , where N is the number of object queries. The positional embedding is generated via a coordinate encoder $\phi(\cdot)$, which consists of a sinusoidal positional encoding (PE) and a multi-layer perceptron (MLP):

$$Q_{pos} = \phi(\mathbf{p}_{ref}[:, :4]) = \text{MLP}(\text{PE}(\mathbf{p}_{ref}[:, :4])) \quad (1)$$

where $\mathbf{p}_{ref}[:, :4]$ represents only the spatial coordinates (cx, cy, w, h) . By masking the θ dimension during embedding generation, we ensure the positional queries remain rotation-invariant.

Simultaneously, the orientation information is latently modeled within the learnable content embeddings $Q_{content}$. This forces the decoder to extract rotation cues from the image features such as texture direction and object heading.

Rotation-Rectified Orthogonal Attention While decoupling angle from positional embedding aids convergence, precise feature extraction requires sampling locations that align with the object’s geometry. Previous oriented DETRs uniformly align all attention heads with the object’s major axis, which leads to a feature collapse where the model disproportionately focuses on longitudinal details while neglecting lateral structural information.

To address this, we impose an orthogonal constraint on the multi-head attention mechanism to ensure feature coverage along both the major and minor axes of the oriented bounding box. For a query with predicted orientation θ , we divide the total attention heads H (assuming H is an even integer) into two distinct groups. The sampling rotation angle $\theta^{(h)}$ for the h -th head, where $h \in \{1, 2, \dots, H\}$, is defined as:

$$\theta^{(h)} = \begin{cases} \theta & \text{if } h \leq \frac{H}{2} \\ \theta + \frac{\pi}{2} & \text{if } h > \frac{H}{2} \end{cases} \quad (2)$$

This formulation forces the first half of the heads to sample features aligned with the object’s predicted heading, while the second half samples orthogonally. Notably, this design introduces merely no additional parameters or GFlops, achieving performance gains without extra computational overhead.

Let $\Delta \mathbf{p}_{qk} \in \mathbb{R}^2$ be the learned offset for the k -th sampling point of the q -th query, and $\mathbf{s}_q \in \mathbb{R}^2$ be its corresponding scale factor (width and height). The




Fig. 3: An intuitive illustration of (D) Decoupled Periodic Refinement and (E) Oriented Dense O2O.

spatial center of the query is denoted as $\mathbf{p}_q = \mathbf{p}_{ref}[q, : 2]$. The final sampling location S is rectified by the head-specific rotation matrix $\mathbf{R}(\theta^{(h)})$:

$$S(\mathbf{p}_q, \Delta \mathbf{p}_{qk}) = \mathbf{p}_q + \mathbf{R}(\theta^{(h)})(\Delta \mathbf{p}_{qk} \odot \mathbf{s}_q) \quad (3)$$

where $\odot$ denotes the element-wise multiplication, and $\mathbf{R}(\theta^{(h)})$ is the rotation matrix derived from the head-specific angle:

$$\mathbf{R}(\theta^{(h)}) = \begin{bmatrix} \cos(\theta^{(h)}) & -\sin(\theta^{(h)}) \\ \sin(\theta^{(h)}) & \cos(\theta^{(h)}) \end{bmatrix} \quad (4)$$

This orthogonal splitting strategy allows the model to explicitly disentangle features along the object’s length and width, improving robustness for aspect-ratio and angle combined prediction.

3.3 Decoupled Periodic Refinement

Standard DETR-style decoders refine boxes by iteratively adding predicted offsets in the inverse-sigmoid space. This update rule implicitly assumes that each regressed variable lies in a Euclidean and globally continuous domain, which holds well for spatial coordinates (c_x, c_y, w, h) . However, the orientation θ in OBB resides on a periodic space. Treating θ as an unconstrained Euclidean scalar, i.e., directly applying additive updates and standard L_1 regression induces a mismatched geometric assumption in decoder refinement: numerically distant angles can be geometrically adjacent across the boundary (e.g., $0 \leftrightarrow \pi$), leading to discontinuous gradients and unstable refinement near periodic seams.

To resolve this geometry mismatch, we introduce **Decoupled Periodic Refinement**, which jointly redesigns both the refinement update and the optimization metric for θ . In other words, we make the decoder refinement consistent with the periodic topology of orientation, rather than merely adding an auxiliary trick. The overall workflow is illustrated in Fig. 2 (C) and 3 (D).

For spatial dimensions, we retain the standard inverse-sigmoid update. Given the predicted spatial offset $\Delta \mathbf{b}$ and the reference box spatial parameters $\mathbf{b}_{ref}$,

the refined spatial coordinates are computed as $\sigma(\Delta \mathbf{b} + \sigma^{-1}(\mathbf{b}_{ref}))$, where $\sigma(\cdot)$ denotes the sigmoid function. For orientation, we adopt a bounded coarse-to-fine periodic update that explicitly respects the cyclic domain. Let $\Delta\theta_i$ be the raw angle offset predicted by the i -th decoder layer ($i \in \{1, \dots, L\}$) and θ_{ref} be the reference angle. We first bound the update magnitude with $\tanh(\cdot)$ and scale it using a layer-wise decaying factor α_i :

$$\alpha_i = \alpha_0^{-i}, \quad (5)$$

where $\alpha_0 > 1$ controls the decay rate. This design enforces a coarse-to-fine refinement schedule: early layers perform larger corrective rotations, while later layers are restricted to fine-grained tuning, improving stability under periodicity. The bounded update is then applied as

$$\theta_{raw} = \theta_{ref} + \tanh(\Delta\theta_i) \cdot \alpha_i. \quad (6)$$

Finally, we map the result back onto the canonical angular domain $[0, \pi)$ via periodic normalization:

$$\theta_{new} = \begin{cases} (\theta_{raw} \bmod \pi) + \pi & \text{if } (\theta_{raw} \bmod \pi) < 0 \\ (\theta_{raw} \bmod \pi) & \text{otherwise.} \end{cases} \quad (7)$$

Crucially, a periodic update alone is insufficient if the optimization objective still measures distance in a Euclidean sense. A standard L_1 loss over-penalizes boundary-adjacent cases, producing gradients that point along the longer arc and thus contradict the periodic refinement behavior. Therefore, we additionally replace the angular regression metric with a **Shortest-path Periodic L_1 Loss**:

$$\mathcal{L}_{angle} = \min(|\theta_{pred} - \theta_{tgt}|, \pi - |\theta_{pred} - \theta_{tgt}|), \quad (8)$$

which guarantees that gradients always follow the shortest angular displacement on the circle. Appendix C provides further analysis and discussion on this point.

By simultaneously adopting a periodic, bounded coarse-to-fine update and a shortest-path periodic loss, our decoder refinement for θ becomes consistent with the orientation setting. The overall box regression loss concatenates the standard L_1 loss for (c_x, c_y, w, h) with $\mathcal{L}_{angle}$ for θ .

3.4 Oriented Dense O2O

We introduce **Oriented Dense O2O** to provide dense supervision while explicitly accelerating the convergence of angle predictions. As shown in Fig. 3 (E), building upon the Dense O2O [19] which combines four replicated images into a single composite grid to increase GT counts, we apply an independent random rotation $\theta_{rot} \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ to each individual quadrant prior to stitching them together. By artificially enriching the angular diversity within a single training image, this computation-free mechanism forces the model to simultaneously process semantic features at various orientations. Consequently, Oriented Dense O2O seamlessly integrates rotational variance into the dense supervision framework, accelerating the model’s convergence on angle predictions and enhancing its robustness against rotated objects.

4 Experiments

4.1 Evaluation Datasets

DOTA-1.0. DOTA-1.0 [46] is a large-scale oriented object detection dataset containing 2,806 images and 188,282 instances across 15 categories, with significant variations in orientation, shape, and scale. Two evaluation protocols are provided: (A) single-scale and (B) multi-scale training and testing. In single-scale settings, images are cropped into 1024×1024 patches with a stride of 824. In multi-scale settings, images are resized to 0.5, 1.0, and 1.5 scales before being cropped into 1024×1024 patches with a stride of 524.

DIOR-R. DIOR-R [6] is constructed from DIOR by adding annotations for rotated bounding boxes. It comprises 23,463 images of size 800×800 , totaling 190,288 oriented instances across 20 categories. The dataset is split into training, validation, and testing sets with a ratio of 1:1:2.

FAIR-1M-2.0. FAIR-1M-2.0 [1] is a recently published remote sensing dataset that consists of 15,266 high-resolution images and more than 1 million instances. It contains 5 categories and 37 subcategories. For multi-scale training and testing, images are processed in the same way as in DOTA.

4.2 Implementation Details

For DOTA-1.0 and DIOR-R, models are trained on the combined train and val sets and evaluated on the test set. For FAIR-1M-2.0, models are trained on the train set and evaluated on the val set. The results on DOTA are reported using the official evaluation server. The training schedules for different RiO-DETR variants and the corresponding hyperparameter configurations are detailed in Appendix A. All models are trained on a single NVIDIA A800 or H200 GPU. All latencies are measured equally on a single NVIDIA T4 GPU using TensorRT 10 with FP16. For end-to-end models, the latency excludes pre-processing and post-processing time, while for others the `BatchedNMSPlugin` is employed to get the actual latency under production scenes. Unless otherwise specified, all ablation studies are conducted on the DIOR-R dataset with RiO-DETR-m.

4.3 Comparison with State-of-the-Art Methods

Results on DOTA-1.0. Under single-scale training and testing protocol (Table 1), RiO-DETR performs strongly across all scales. RiO-DETR-n achieves 78.4% AP_{50} with 2.7 ms latency and 4.0M parameters, surpassing YOLO26n-obb. At the high end, RiO-DETR-x reaches 81.8% AP_{50} at 29.9 ms, outperforming YOLO26x-obb (80.4%, 30.5 ms) and heavy DETR variants such as RHINO-DETR (79.4%, 242.6 ms). Under multi-scale evaluation (Table 2), RiO-DETR-x attains 81.76% AP_{50} with the lowest latency among high-precision models. Per-class AP_{50} under multi-scale protocol, results under high precision metrics and results on different backbones are provided in Appendix D.

Results on DIOR-R. RiO-DETR maintains a clear speedaccuracy advantage on DIOR-R (Table 3). RiO-DETR-s achieves 74.44% AP₅₀ at 3.01 ms, while RiO-DETR-x reaches 77.43% at 17.31 ms, exceeding YOLO26x-obb (76.48%) under similar latency and outperforming prior state-of-the-art models.

Results on FAIR-1M-2.0. On FAIR-1M-2.0 (Table 4), RiO-DETR-x achieves a new state of the art with 47.4 AP₅₀ under multi-scale training and testing, surpassing YOLO26x-obb (46.7%), LSKNet-S (46.3%), and ReDet (43.2%), demonstrating strong scalability to large-scale remote sensing benchmarks. Per-class AP₅₀ are provided in Appendix D.

4.4 Efficiency Analysis of RiO-DETR

Our RiO-DETR series operates at the same latency level as the YOLO26 family across all model scales (nx). Specifically, RiO-DETR-n/s/m/l/x achieve 2.7/5.2/8.8/13.4/29.9 ms, closely matching the corresponding YOLO26 variants (2.8/4.9/10.2/13.0/30.5 ms), demonstrating that RiO-DETR maintains a fully real-time regime comparable to state-of-the-art CNN-based real-time detectors.

Several CNN-based counterparts [2, 25, 55] that claim efficiency exhibit substantially higher latency (200360 ms). This gap stems largely from computationally expensive components such as large-kernel convolutions. By explicitly avoiding such high-latency operators, RiO-DETR achieves a clear advantage over these methods. Notably, RiO-DETR is the first oriented DETR to achieve end-to-end real-time inference for a single image, closing the long-standing efficiency gap between transformer-based and CNN-based oriented detectors.

4.5 Ablation Studies

The Roadmap to RiO-DETR. Table 5 showcases the stepwise progression from the baseline model to our proposed RiO-DETR framework. Starting from RT-DETRv2 adding a simple OBB head, we first introduce a stronger bipartite matching and optimization objective, which yields a notable improvement in AP₅₀. Specifically, we formulate the matching cost $\mathcal{C}_{match}$ as a combination of focal loss, KLD loss, and Hausdorff distance:

$$\mathcal{C}_{match} = \lambda_{cls}\mathcal{C}_{focal} + \lambda_{kld}\mathcal{C}_{kld} + \lambda_{hausdorff}\mathcal{C}_{hausdorff}$$

Correspondingly, the overall training loss $\mathcal{L}$ is optimized using a shortest-path periodic L1 loss ($\mathcal{L}_{angle}$) alongside the focal and KLD losses:

$$\mathcal{L} = \lambda_{cls}\mathcal{L}_{focal} + \lambda_{kld}\mathcal{L}_{kld} + \lambda_{angle}\mathcal{L}_{angle}$$

By integrating the matching strategy from D-FINE and the training scheme from DEIM, the baseline increases to 73.47 AP₅₀. On this basis, Content-Driven Angle Estimation yields steady gains: Geometry-Decoupled Query Encoding raises AP₅₀ to 74.18 with negligible cost, and Rotation-Rectified Orthogonal Attention further improves it to 74.74 with slight latency overhead. Finally, Decoupled Periodic Refinement and Oriented Dense O2O push performance to 75.73 AP₅₀ while keeping parameters, FLOPs, and latency nearly unchanged.

Table 1: Comparison with state-of-the-art oriented object detectors on DOTA-1.0 under single-scale training and testing protocol.

Methods	Backbone	#P	Flops	Lat.	AP ₅₀	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC
CNN-based Non-Real-time Oriented Object Detectors																				
LSKNet-S [25]	LSKNet-S	31.0M	161G	203.5	77.5	89.7	85.5	57.7	75.7	75.0	78.7	88.2	90.9	90.9	86.4	66.9	63.8	77.8	74.5	64.8
PKINet-S [2]	PKINet-S	30.8M	190G	359.7	78.4	89.7	84.2	55.8	77.6	80.2	84.5	88.1	90.9	87.6	86.1	66.9	70.2	77.5	73.6	62.9
PSD-SQ [10]	R-50	-	-	-	78.5	89.7	85.4	57.3	75.2	80.0	81.2	88.3	89.9	88.0	86.3	69.6	68.5	75.3	69.5	72.7
RVSA [43]	ViTAE-B	114.4M	414G	-	79.0	89.4	84.3	59.4	73.2	80.0	85.4	88.1	90.9	88.5	86.5	58.9	72.2	77.3	79.6	71.2
Strip R-CNN-S [55]	StripNet-S	30.5M	159G	241.9	80.1	88.9	86.4	57.4	76.4	79.7	84.4	88.2	90.9	86.7	87.5	69.9	66.8	79.2	82.9	75.6
DETR-based Non-Real-time Oriented Object Detectors																				
R. D-DETR [57]	R-50	41.1M	409G	216.9	69.5	84.8	70.7	46.0	61.9	73.9	78.8	87.7	90.0	77.9	78.4	47.0	54.4	66.8	67.6	55.6
R. D-DETR [57]	R-50 w/ CSL	41.4M	411G	221.7	72.2	86.2	76.6	46.6	65.2	76.8	76.3	87.7	90.7	79.3	82.3	54.0	61.7	66.0	70.4	61.9
EMO2-DETR [17]	Swin-T	-	-	-	72.3	89.0	79.6	48.7	60.2	77.3	76.4	84.5	90.8	84.8	85.7	48.9	67.6	66.3	71.5	53.5
ARS-DETR [57]	R-50	41.4M	411G	221.7	74.2	86.9	75.5	48.3	69.2	77.9	77.9	87.6	90.5	77.3	82.8	60.2	64.5	74.8	71.7	66.6
ARS-DETR [57]	Swin-T	41.9M	431G	303.6	75.5	87.7	76.5	50.6	69.9	79.8	83.9	87.9	90.3	86.2	85.1	54.6	67.0	75.6	73.7	63.4
AO2-DETR [7]	R-50	74.3M	304G	-	77.7	89.3	85.0	56.7	74.9	78.9	82.7	87.3	90.5	84.7	85.4	62.0	70.0	74.7	72.4	71.6
RHINO-DETR [22]	R-50	47.6M	566G	187.9	78.7	88.2	85.1	55.8	72.7	80.2	83.1	89.0	90.8	87.1	86.8	65.3	71.6	77.7	81.2	64.7
RHINO-DETR [22]	Swin-T	50.8M	609G	242.6	79.4	88.2	84.8	58.5	77.7	81.1	85.6	89.2	90.9	87.0	86.4	65.5	71.3	78.2	82.8	64.3
Oriented-DETR [61]	R-50	57.2M	302G	217.3	79.1	89.2	86.4	57.7	75.3	81.1	84.7	89.1	90.9	86.1	87.0	59.5	70.3	79.3	81.5	68.8
Oriented-DETR [61]	Swin-T	57.7M	309G	235.5	79.8	89.4	85.1	57.8	75.0	81.2	86.1	89.1	90.9	88.7	87.0	62.9	69.1	80.7	82.8	71.0
Real-time Oriented Object Detectors																				
PPYOLOE-R-l [49]	CRN-l	52.2M	141G	33.3	78.1	89.2	81.0	54.0	70.2	81.8	85.2	88.8	90.8	87.0	88.0	62.9	67.9	76.6	79.1	69.7
PPYOLOE-R-x [49]	CRN-x	98.4M	264G	51.0	78.3	89.5	79.7	55.0	75.6	82.4	85.2	88.3	90.8	85.7	87.7	63.2	69.5	77.1	75.1	69.4
RTMDet-R-m [31]	CSPNext-m	24.7M	100G	22.4	78.2	89.2	84.7	53.9	74.7	81.5	84.0	88.7	90.8	87.4	87.2	59.4	66.7	77.7	82.4	65.3
RTMDet-R-l [31]	CSPNext-l	52.3M	205G	30.2	78.8	89.4	84.2	55.2	75.1	80.8	84.5	89.0	90.9	87.4	87.2	63.1	67.9	78.1	80.8	69.1
YOLO26n-obb [37]	YOLO26n	2.5M	14G	<u>2.8</u>	<u>77.7</u>	89.7	84.7	52.1	71.1	83.1	78.9	88.7	91.1	87.3	86.9	58.5	69.1	76.6	82.8	65.1
YOLO26s-obb [37]	YOLO26s	9.8M	55G	<u>4.9</u>	<u>79.7</u>	89.7	86.3	55.5	71.8	82.5	81.9	88.9	91.1	88.2	88.1	62.9	72.5	78.0	83.9	74.5
YOLO26m-obb [37]	YOLO26m	21.2M	183G	10.2	<u>80.0</u>	89.2	86.8	56.4	73.0	82.6	78.3	89.1	91.2	88.3	87.9	64.5	78.0	78.3	83.1	79.7
YOLO26l-obb [37]	YOLO26l	25.6M	230G	<u>13.0</u>	<u>80.2</u>	89.5	86.3	56.2	74.2	82.8	78.8	89.3	91.1	88.6	88.1	67.8	72.2	78.2	83.5	76.5
YOLO26x-obb [37]	YOLO26x	57.6M	517G	<u>30.5</u>	<u>80.4</u>	89.7	86.6	57.3	73.8	82.8	79.2	89.2	91.1	89.1	88.1	66.7	70.6	77.9	83.3	80.7
RiO-DETR-n	HGNet-B0	4.0M	17G	2.7	78.4	87.3	86.1	54.9	73.6	80.5	85.1	88.0	90.8	87.4	86.9	60.9	72.0	76.8	73.4	71.8
RiO-DETR-s	HGNet-B2	8.2M	53G	<u>5.2</u>	80.3	85.6	84.8	57.5	75.5	80.7	86.2	89.0	90.8	88.4	88.2	65.2	74.2	78.3	80.1	79.6
RiO-DETR-m	HGNet-B2	18.6M	158G	8.8	80.9	85.9	86.4	61.7	78.1	82.0	86.8	89.2	90.8	88.9	87.7	67.4	73.3	78.8	77.4	78.8
RiO-DETR-l	HGNet-B4	27.5M	230G	<u>13.4</u>	81.7	86.7	86.8	60.8	79.4	82.0	86.4	89.0	90.8	88.7	88.5	70.0	76.0	79.2	83.4	76.9
RiO-DETR-x	HGNet-B5	62.5M	527G	29.9	81.8	88.3	87.5	61.9	79.2	83.2	86.5	89.2	90.8	88.4	87.8	71.2	75.1	78.5	82.2	77.1

Table 2: Comparison with state-of-the-art oriented object detectors on DOTA-1.0 under multi-scale training and testing protocol.

Model	PKINet-S [2]	LSKNet-S [25]	Strip R-CNN-S [55]	PPYOLOE-R-x [49]	RTMDet-R-l [31]	YOLO26m-obb [37]	Ours-m	YOLO26x-obb [37]	Ours-x
Latency	359.7	203.5	241.9	51.0	30.2	10.2	8.8	30.5	29.9
AP₅₀	81.06	81.64	82.28	80.73	81.33	81.00	81.49	81.70	81.76

Ablation on Geometry-Decoupled Query Encoding. Table 9 analyzes different geometry components used in query encoding. Encoding the center coordinates alone achieves limited performance (72.56 AP₅₀), and additionally encoding the angle further degrades accuracy (72.34), suggesting that directly injecting θ into the query can introduce optimization ambiguity. Incorporating size information provides a clear gain (73.57), and the best result is obtained by encoding center and size while leaving the angle to be inferred from content features, reaching 73.81 AP₅₀. These results validate our design choice of geometry decoupling: use (cx, cy, w, h) to guide localization while predicting θ in a content-driven manner.

Ablation on Decoupled Periodic Refinement Components. Table 8 analyzes the contribution of each component in Decoupled Periodic Refinement. Using SP- L_1 alone slightly improves AP₅₀ from 74.18 to 74.32. The periodic update mechanism alone leads to a small drop in AP₅₀ (74.18 $\rightarrow$ 74.05), sug-

Table 3: Performance comparison with state-of-the-art oriented object detectors on DIOR-R.

Methods	Backbone	#P	Flops	Lat.	AP ₅₀
Non-Real-time Oriented Object Detectors					
LSKNet-S [25]	LSKNet-S	31.0M	111G	138.75	65.90
PKINet-S [2]	PKINet-S	30.8M	118G	217.33	67.03
Strip R-CNN-S [55]	StripNet-S	30.5M	157G	236.98	68.70
RHINO-DETR [22]	Swin-T	50.8M	383G	164.51	72.67
Real-time Oriented Object Detectors					
YOLO26n-obb [37]	YOLO26n	2.5M	10G	<u>2.23</u>	<u>71.88</u>
YOLO26s-obb [37]	YOLO26s	9.8M	39G	2.84	<u>74.07</u>
YOLO26m-obb [37]	YOLO26m	21.2M	129G	<u>5.91</u>	<u>74.65</u>
YOLO26l-obb [37]	YOLO26l	25.6M	158G	7.53	<u>75.31</u>
YOLO26x-obb [37]	YOLO26x	57.6M	353G	<u>17.66</u>	<u>76.48</u>
RiO-DETR-n	HGNet-B0	4.0M	11G	2.20	71.92
RiO-DETR-s	HGNet-B2	8.2M	16G	<u>3.01</u>	74.44
RiO-DETR-m	HGNet-B2	18.6M	97G	5.10	75.73
RiO-DETR-l	HGNet-B4	27.5M	141G	<u>7.76</u>	76.11
RiO-DETR-x	HGNet-B5	62.7M	324G	17.31	77.43

Table 4: Performance comparison with SoTA methods on FAIR-1M-2.0 under multi-scale training and testing protocol.

Methods	Backbone	#P	Flops	Lat.	AP ₅₀
SASM RepPoints [15]	R-101	55.8M	542G	144.5	30.9
R-FCOS [42]	R-101	50.9M	284G	156.9	36.1
S2A-Net [56]	R-50	31.6M	588G	-	37.4
R-Faster RCNN [50]	R-101	60.1M	289G	163.4	37.5
O-RepPoints [24]	Swin-T	37.3M	200G	-	38.9
O-RCNN [47]	R-101	60.3M	289G	160.8	40.4
RoI Trans. [8]	R-101	67.8M	607G	-	40.2
LOOD (RT) [62]	R-50	-	-	-	42.6
ReDet [14]	ReR-50	31.8M	225G	448.2	43.2
PKINet-S [2]	PKINet-S	30.8M	190G	359.7	44.5
LOOD (RD) [62]	R-101	-	-	-	44.9
LSKNet-S [25]	LSKNet-S	31.0M	161G	203.5	46.3
Strip-RCNN-S [55]	StripNet-S	30.5M	159G	241.9	46.8
YOLO26m-obb [37]	YOLO26m	21.2M	183G	<u>10.2</u>	<u>42.5</u>
YOLO26x-obb [37]	YOLO26x	57.6M	517G	<u>30.5</u>	<u>46.7</u>
RiO-DETR-m	HGNet-B2	18.6M	158G	8.8	43.6
RiO-DETR-x	HGNet-B5	62.5M	527G	29.9	47.4

Table 5: Step-by-step modifications from baseline model to RiO-DETR-m on DIOR-R. Each step shows changes in AP₅₀, the number of parameters, latency, and Flops.

Methods	#P	Flops	Latency	AP ₅₀
Oriented RT-DETRv2 (Ours Scratch Implementation)	18.61M	97.06G	5.04	70.35
+ Hausdorff Matching from RHINO-DETR [22] & KLD Loss [53]				72.86
+ Universal Matching Strategy from D-FINE [33]				73.33
+ Dense O2O from DEIM [19] (Our Implemented Baseline)	18.61M	97.06G	5.04	73.47
+ <i>Content-Driven Angle Estimation</i>				
<i>Geometry-Decoupled Query Encoding</i>	18.59M	97.01G	5.04	74.18
<i>Rotation-Rectified Orthogonal Attention</i>	18.67M	97.18G	5.10	74.74
+ <i>Decoupled Periodic Refinement</i>				75.46
+ <i>Oriented Dense O2O</i> (Our RiO-DETR-m)	18.67M	97.18G	5.10	75.73

gesting that constraining angles within $[0, \pi)$ without a consistent periodic loss is insufficient and may introduce instability. Combining both components yields the largest gain, reaching 74.74 AP₅₀, demonstrating that consistent periodic modeling in both the loss and update steps is crucial for stable OBB refinement.

Ablation on Oriented Dense O2O. Compared to Regular Augmentation (73.33 AP₅₀ at 86 epochs), Dense O2O achieves 73.47 AP₅₀ in 68 epochs, proving that denser supervision mitigates sparse matching. Oriented Dense O2O reaches the highest AP₅₀ (73.88) with the fastest convergence, demonstrating the value of rotation diversity. While random rotation yields a smaller gain (73.66), we attribute this to noisy angle supervision from orientation ambiguity in near-square objects. In contrast, the discrete rotations in Oriented Dense O2O enrich patterns while reducing ambiguity, resulting in more stable training and superior performance. Oriented Dense O2O is essentially equivalent to performing angle-stratified dense supervision within a single composite image, which improves the angular separability of targets in each Hungarian matching step. Moreover,

Table 6: Comparison of our implemented baseline with SoTA methods on DIOR-R.

Methods	#P	Flops	Latency	AP ₅₀
RoI Trans. [8]	68.1M	373G	-	63.87
RVSA-ViTAE-B [43]	114.4M	-	-	71.05
DCFL-ReR-101 [48]	-	-	-	71.03
LSKNet-S [25]	31M	111G	138.75	65.90
PKINet-S [2]	30.8M	118G	217.33	67.03
Strip R-CNN-S [55]	30.5M	157G	236.98	68.70
RHINO-DETR-Swin-T [22]	50.8M	383G	164.51	72.67
YOLO26m-obb [37]	21.2M	129G	<u>5.91</u>	74.65
Baseline	18.61M	97G	5.04	<u>73.47</u>

Table 8: Ablation study on the components of Decoupled Periodic Refinement.

SP- L_1	Periodic	AP ₅₀
		74.18
✓		74.32
	✓	74.05
✓	✓	74.74

Table 7: Ablation on Oriented Dense O2O. Epoch notes the one with highest AP₅₀ on Stage 1.

Methods	Epoch	AP ₅₀
No Augmentation	94	72.18
Regular Augmentation	86	73.33
Dense O2O [19]	68	73.47
Oriented Dense O2O (Ours)	60	73.88
Dense O2O w/ random rotation	62	73.66
Pure Angle Augmentation	84	73.53

Table 9: Ablation study on Geometry-Decoupled Query Encoding.

Center (c_x, c_y)	Size (w, h)	Angle (θ)	AP ₅₀
✓			72.56
✓		✓	72.34
✓	✓		73.57
✓	✓	✓	73.81

Table 10: Ablation on injecting angle into the positional query.

Positional Query Representation	Angle in Positional Branch	Periodicity-aware Encoding	AP ₅₀
(c_x, c_y, w, h, θ)	Yes	No	73.47
($c_x, c_y, w, h, \sin \theta, \cos \theta$)	Yes	Yes	72.60
($c_x, c_y, w, h, \sin 2\theta, \cos 2\theta$)	Yes	Yes	73.52
(c_x, c_y, w, h) (Ours)	No	N/A	74.18

using discrete rotations maximizes angular coverage while minimizing equivalent-parameterization noise of near-square objects; as a result, the angle branch enters a stable refinement regime earlier, with lower gradient variance and reduced jitter near periodic boundaries. Further analysis is provided in Appendix E.

Ablation on the Design of Rotation-Rectified Orthogonal Attention. Table 10 analyzes head splitting in Rotation-Rectified Orthogonal Attention. Aligning all heads to θ yields 73.81 AP₅₀, while an asymmetric 6:2 split slightly improves performance to 74.02. The symmetric orthogonal design (4:4 for θ and $\theta + \pi/2$) performs best at 74.18, showing that modeling orthogonal directions better captures longitudinal and lateral structures. Distributing heads across four angles further reduces performance (73.95), due to diluted capacity per direction.

4.6 Visualization Analysis

Visualization Analysis of Sampling Patterns. To intuitively understand how Rotation-Rectified Orthogonal Attention improves feature extraction, we visualize the deformable attention sampling points in Fig. 4. Our proposed




Fig. 4: Visualization of deformable attention sampling points on DOTA.




Fig. 5: Visual illustration of layer-wise angular refinement for a specific instance.

Fig. 6: t-SNE visualization of angular features, demonstrating Geometry-Decoupled Query Encoding’s clustering.

Rotation-Rectified Orthogonal Attention enforces an orthogonal sampling pattern, effectively capturing both longitudinal and lateral structural features of the oriented objects.

Visualization Analysis of Angular Prediction. Figure 5 illustrates this refinement on a challenging instance. Starting from a coarse estimate, the predicted OBB progressively aligns with object boundaries: Layer 1 corrects the global heading, while Layers 24 remove minor angular deviations without overshooting, further validating the stability and precision of the proposed strategy. More cases are provided in Appendix E.

Visualizing Angular Feature Distribution. To further validate the effectiveness of Geometry-Decoupled Query Encoding, we employ t-SNE to visualize the query angular features on FAIR-1M-2.0 in Fig. 6. Our decoupled formulation yields content-query embeddings that organize into clear and separated clusters across orientation intervals, indicating that the model encodes orientation-sensitive semantic evidence in the content space rather than relying on a rigid

geometric prior, which is consistent with our design choice to avoid injecting θ as positional geometry. Further t-SNE analyses are provided in Appendix E.

5 Conclusion

In this paper, we present **RiO-DETR**, the first real-time oriented detection transformer that offers a meaningful step forward in oriented object detection. Content-Driven Angle Estimation, Decoupled Periodic Refinement, and Oriented Dense O2O jointly address angle-content entanglement, periodicity-induced optimization instability, and slow convergence of existing methods. Extensive experiments on DOTA-1.0, DIOR-R, and FAIR-1M-2.0 demonstrate a new accuracy-latency frontier compared to all state-of-the-art CNN and DETR counterparts.

Limitation and Future Work: While RiO-DETR significantly optimizes the detection pipeline, designing a *real-time* yet high-performance backbone specifically tailored for the complexities of oriented/remote sensing object detection remains an open challenge. Future work may focus on specialized feature extraction to further enhance the real-time trade-off. We hope RiO-DETR serves as a strong, practical baseline for the community and inspires further exploration into real-time end-to-end perception for oriented objects.

Acknowledgements

This work was in part supported by the National Natural Science Foundation of China under grants 62472399 and Open Fund of APKL of BIIP, IAI, Hefei Comprehensive National Science Center under grants 24YGXT003.

References

1. Fair1m: A benchmark dataset for fine-grained object recognition in high-resolution remote sensing imagery. ISPRS Journal of Photogrammetry and Remote Sensing **184**, 116–130 (2022). <https://doi.org/https://doi.org/10.1016/j.isprsjprs.2021.12.004>
2. Cai, X., Lai, Q., Wang, Y., Wang, W., Sun, Z., Yao, Y.: Poly kernel inception network for remote sensing detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 27706–27716 (2024)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
4. Chen, Q., Su, X., Zhang, X., Wang, J., Chen, J., Shen, Y., Han, C., Chen, Z., Xu, W., Li, F., et al.: Lw-detr: A transformer replacement to yolo for real-time detection. arXiv preprint arXiv:2406.03459 (2024)
5. Chen, Z., Chen, K., Lin, W., See, J., Yu, H., Ke, Y., Yang, C.: Piou loss: Towards accurate oriented object detection in complex environments. In: European conference on computer vision. pp. 195–211. Springer (2020)

6. Cheng, G., Wang, J., Li, K., Xie, X., Lang, C., Yao, Y., Han, J.: Anchor-free oriented proposal generator for object detection. *IEEE Transactions on Geoscience and Remote Sensing* **60**, 1–11 (2022)
7. 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)
8. Ding, J., Xue, N., Long, Y., Xia, G.S., Lu, Q.: Learning roi transformer for oriented object detection in aerial images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2849–2858 (2019)
9. Ding, J., Xue, N., Xia, G.S., Bai, X., Yang, W., Yang, M.Y., Belongie, S., Luo, J., Datcu, M., Pelillo, M., et al.: Object detection in aerial images: A large-scale benchmark and challenges. *IEEE transactions on pattern analysis and machine intelligence* **44**(11), 7778–7796 (2021)
10. Feng, S., Wang, B.: Psd-sq: Point set decoding based on semantic query for object detection in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–12 (2024)
11. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430* (2021)
12. Girshick, R.: Fast r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1440–1448 (2015)
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: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2786–2795 (2021)
15. Hou, L., Lu, K., Xue, J., Li, Y.: Shape-adaptive selection and measurement for oriented object detection. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 36, pp. 923–932 (2022)
16. Hu, Z., Wu, P., Chen, J., Zhu, H., Wang, Y., Peng, Y., Li, H., Sun, X.: Dome-detr: Detr with density-oriented feature-query manipulation for efficient tiny object detection. In: *Proceedings of the 33rd ACM International Conference on Multimedia*. pp. 101–110 (2025)
17. Hu, Z., Gao, K., Zhang, X., Wang, J., Wang, H., Yang, Z., Li, C., Li, W.: Emo2-detr: Efficient-matching oriented object detection with transformers. *IEEE transactions on geoscience and remote sensing* **61**, 1–14 (2023)
18. Huang, S., Hou, Y., Liu, L., Yu, X., Shen, X.: Real-time object detection meets dinov3. *arXiv preprint arXiv:2509.20787* (2025)
19. Huang, S., Lu, Z., Cun, X., Yu, Y., Zhou, X., Shen, X.: Deim: Detr with improved matching for fast convergence. In: *Proceedings of the computer vision and pattern recognition conference*. pp. 15162–15171 (2025)
20. Jocher, G., Chaurasia, A., Stoken, A., Borovec, J., Kwon, Y., Michael, K., Fang, J., Wong, C., Yifu, Z., Montes, D., et al.: ultralytics/yolov5: v6. 2-yolov5 classification models, apple m1, reproducibility, clearml and deci. ai integrations. *Zenodo* (2022)
21. Khanam, R., Hussain, M.: Yolov11: An overview of the key architectural enhancements. *arXiv preprint arXiv:2410.17725* (2024)
22. Lee, H., Song, M., Koo, J., Seo, J.: Hausdorff distance matching with adaptive query denoising for rotated detection transformer. In: *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 1872–1882. *IEEE* (2025)

23. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13619–13627 (2022)
24. Li, W., Chen, Y., Hu, K., Zhu, J.: Oriented reppoints for aerial object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1829–1838 (2022)
25. 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: Y. li et al. *International Journal of Computer Vision* **133**(3), 1410–1431 (2025)
26. Liao, Z., Zhao, Y., Shan, X., Yan, Y., Liu, C., Lu, L., Ji, X., Chen, J.: Rt-detr4: Painlessly furthering real-time object detection with vision foundation models. arXiv preprint arXiv:2510.25257 (2025)
27. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
28. Liu, H., Sun, W., Zhang, Q., Di, D., Gong, B., Li, H., Wei, C., Zou, C.: Hi-vae: Efficient video autoencoding with global and detailed motion. arXiv preprint arXiv:2506.07136 (2025)
29. Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: Dab-detr: Dynamic anchor boxes are better queries for detr. arXiv preprint arXiv:2201.12329 (2022)
30. Lv, W., Zhao, Y., Chang, Q., Huang, K., Wang, G., Liu, Y.: Rt-detr2: Improved baseline with bag-of-freebies for real-time detection transformer. arXiv preprint arXiv:2407.17140 (2024)
31. 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)
32. Meng, D., Chen, X., Fan, Z., Zeng, G., Li, H., Yuan, Y., Sun, L., Wang, J.: Conditional detr for fast training convergence. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3651–3660 (2021)
33. Peng, Y., Li, H., Wu, P., Zhang, Y., Sun, X., Wu, F.: D-fine: Redefine regression task in detr as fine-grained distribution refinement. arXiv preprint arXiv:2410.13842 (2024)
34. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
35. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence* **39**(6), 1137–1149 (2016)
36. Robinson, I., Robicheaux, P., Popov, M., Ramanan, D., Peri, N.: Rf-detr: neural architecture search for real-time detection transformers. arXiv preprint arXiv:2511.09554 (2025)
37. 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)
38. Sohan, M., Sai Ram, T., Rami Reddy, C.V.: A review on yolov8 and its advancements. In: International conference on data intelligence and cognitive informatics. pp. 529–545. Springer (2024)
39. Sun, W., Che, Y., Huang, H., Guo, Y.: Neural reconstruction of relightable human model from monocular video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 397–407 (2023)

40. Sun, W., Wang, Z., Hu, Z., Wang, C., Li, H., Chen, W.: Muse: A multi-agent framework for unconstrained story envisioning via closed-loop cognitive orchestration. arXiv preprint arXiv:2602.03028 (2026)
41. Tian, Y., Ye, Q., Doermann, D.: Yolov12: Attention-centric real-time object detectors. arXiv preprint arXiv:2502.12524 (2025)
42. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9627–9636 (2019)
43. Wang, D., Zhang, Q., Xu, Y., Zhang, J., Du, B., Tao, D., Zhang, L.: Advancing plain vision transformer toward remote sensing foundation model. IEEE transactions on geoscience and remote sensing **61**, 1–15 (2022)
44. Wang, S., Xia, C., Lv, F., Shi, Y.: Rt-detr3: Real-time end-to-end object detection with hierarchical dense positive supervision. In: WACV. pp. 1628–1636 (2025)
45. Wen, L., Cheng, Y., Fang, Y., Li, X.: A comprehensive survey of oriented object detection in remote sensing images. Expert Systems with Applications **224**, 119960 (2023)
46. Xia, G.S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L.: Dota: A large-scale dataset for object detection in aerial images. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3974–3983 (2018)
47. Xie, X., Cheng, G., Wang, J., Yao, X., Han, J.: Oriented r-cnn for object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3520–3529 (2021)
48. 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. pp. 7318–7328 (2023)
49. Xu, S., Wang, X., Lv, W., Chang, Q., Cui, C., Deng, K., Wang, G., Dang, Q., Wei, S., Du, Y., et al.: Pp-yoloe: An evolved version of yolo. arXiv preprint arXiv:2203.16250 (2022)
50. Yang, S., Pei, Z., Zhou, F., Wang, G.: Rotated faster r-cnn for oriented object detection in aerial images. In: Proceedings of the 2020 3rd International Conference on Robot Systems and Applications. pp. 35–39 (2020)
51. Yang, X., Yan, J.: Arbitrary-oriented object detection with circular smooth label. In: European conference on computer vision. pp. 677–694. Springer (2020)
52. 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)
53. 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)
54. Yang, X., Zhou, Y., Zhang, G., Yang, J., Wang, W., Yan, J., Zhang, X., Tian, Q.: The kfiou loss for rotated object detection. arXiv preprint arXiv:2201.12558 (2022)
55. 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)
56. Yujie, L., Xiaorui, S., Wenbin, S., Yafu, Y.: S2anet: Combining local spectral and spatial point grouping for point cloud processing. Virtual Reality & Intelligent Hardware **6**(4), 267–279 (2024)
57. 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)

58. 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. arXiv preprint arXiv:2203.03605 (2022)
59. Zhao, J., Ding, Z., Zhou, Y., Zhu, H., Du, W.L., Yao, R., El Saddik, A.: Orientedformer: An end-to-end transformer-based oriented object detector in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–16 (2024)
60. Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., Chen, J.: Detrs beat yolos on real-time object detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16965–16974 (2024)
61. Zhao, Z., Xue, Q., He, Y., Bai, Y., Wei, X., Gong, Y.: Projecting points to axes: Oriented object detection via point-axis representation. In: *European conference on computer vision*. pp. 161–179. Springer (2024)
62. Zhou, B., Bi, Q., Ding, J., Xia, G.S.: Boosting fine-grained oriented object detection via text features. In: *International Conference on Pattern Recognition*. pp. 109–125. Springer (2024)
63. Zhou, Y., Yang, X., Zhang, G., Wang, J., Liu, Y., Hou, L., Jiang, X., Liu, X., Yan, J., Lyu, C., et al.: Mmrotate: A rotated object detection benchmark using pytorch. In: *Proceedings of the 30th ACM international conference on multimedia*. pp. 7331–7334 (2022)
64. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)
65. Zou, Z., Chen, K., Shi, Z., Guo, Y., Ye, J.: Object detection in 20 years: A survey. *Proceedings of the IEEE* **111**(3), 257–276 (2023)