

Integrating Diverse Assignment Strategies into DETRs

Yiwei Zhang^{1,2,5*}, Jin Gao^{1,2,5†}, Hanshi Wang^{1,2,5}, Fudong Ge^{1,2,5},
Guan Luo^{1,2,5}, Weiming Hu^{1,2,5,6}, Zhipeng Zhang^{3,4†}

¹State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS), CASIA

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³AutoLab, School of Artificial Intelligence, Shanghai Jiao Tong University,

⁴Anyverse Robotics

⁵Beijing Key Laboratory of Super Intelligent Security of Multi-Modal Information

⁶School of Information Science and Technology, ShanghaiTech University

zhipeng.zhang.cv@outlook.com, jin.gao@nlpr.ia.ac.cn, ywzhang.ai@gmail.com

Abstract

Label assignment is a critical component in object detectors, particularly within DETR-style frameworks where the one-to-one matching strategy, despite its end-to-end elegance, suffers from slow convergence due to sparse supervision. While recent works have explored one-to-many assignments to enrich supervisory signals, they often introduce complex, architecture-specific modifications and typically focus on a single auxiliary strategy, lacking a unified and scalable design. In this paper, we first systematically investigate the effects of “one-to-many” supervision and reveal a surprising insight that *performance gains are driven not by the sheer quantity of supervision, but by the diversity of the assignment strategies employed*. This finding suggests that a more elegant, parameter-efficient approach is attainable. Building on this insight, we propose LoRA-DETR, a flexible and lightweight framework that seamlessly integrates diverse assignment strategies into any DETR-style detector. Our method augments the primary network with multiple Low-Rank Adaptation (LoRA) branches during training, each instantiating a different one-to-many assignment rule. These branches act as auxiliary modules that inject rich, varied supervisory gradients into the main model and are discarded during inference, thus incurring no additional computational cost. This design promotes robust joint optimization while maintaining the architectural simplicity of the original detector. Extensive experiments on different baselines validate the effectiveness of our approach. Our work presents a new paradigm for enhancing detectors, demonstrating that diverse “one-to-many” supervision can be integrated to achieve state-of-the-art results without compromising model elegance.

Code — <https://github.com/Z1zyw/LoRA-DETR>

Introduction

Label assignment, which involves matching predicted objects to ground-truth annotations during training, serves as a foundational component in object detection methods. It critically determines which predictions are treated as positive

*Work done co-mentored by Prof. Zhipeng Zhang.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

or negative samples, thereby shaping the optimization process and ultimately influencing the detector’s performance. Traditional detectors, such as Faster R-CNN (Ren et al. 2015) and RetinaNet (Lin et al. 2017), typically employ hand-crafted rules based on anchors and IoU thresholds for this assignment, followed by Non-Maximum Suppression (NMS) to eliminate redundant predictions from overlapping anchors. In contrast, recent transformer-based frameworks like DETR (Carion et al. 2020) have introduced more principled mechanisms of one-to-one matching via the Hungarian algorithm, which assigns each object query to a unique ground-truth target. This approach effectively obviates the need for NMS during inference.

While the one-to-one assignment strategy in DETR-style detection frameworks enables a clean, end-to-end formulation, it inherently limits the number and diversity of positive samples during training, resulting in slow convergence. To address this limitation, recent studies have explored “one-to-many” assignment strategies to enhance supervision. However, incorporating these strategies often entails non-trivial modifications to the decoder pipeline. Such architectural modifications, as observed in certain methods (Zhao et al. 2024), introduce unnecessary complexity that departs from the elegant simplicity of the original DETR design, making them potentially avoidable via more efficient approaches that maintain the core structure while delivering comparable or superior supervision gains. More interestingly, as we will demonstrate in the later section, the benefits of some modification appear to be less significant than previously assumed.

Moreover, in addition to this limitation, another issue is that these methods are typically tailored to a *single* auxiliary assignment strategy, with their integration mechanisms tightly coupled to specific architectures or training procedures. This absence of a unified design hampers their scalability and flexibility, particularly when incorporating multiple assignment strategies is desired.

Then, *how can we devise a more elegant, flexible, and effective one-to-many assignment design for object detection?*

To investigate this, we begin with a straightforward design inspired by MS-DETR (Fig. 1(b)), adapting the vanilla DETR decoder layer (Fig. 1(a)) into a multi-branch struc-

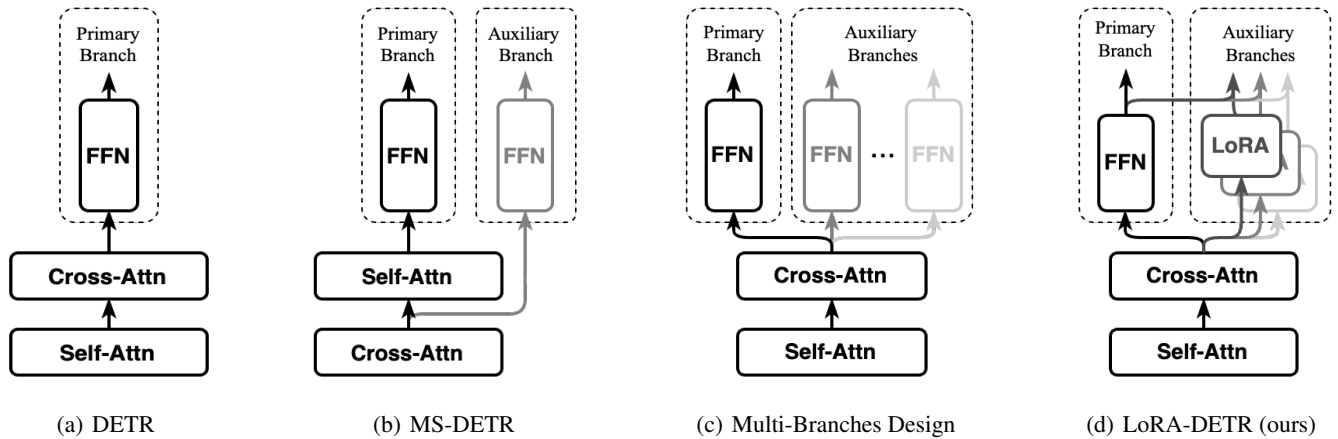


Figure 1: Different Decoder Architectures. (a) Standard DETR decoder layer, supervised solely by the one-to-one Hungarian matching strategy. (b) MS-DETR design that exchanges the positions of the cross-attention and self-attention modules, enabling supervision from both one-to-one and one-to-many assignment strategies, where cross-attention is guided by the latter. (c) A naive multi-branch extension of (a), where multiple FFNs are independently supervised by different assignment strategies. (d) Our proposed LoRA-DETR, which integrates lightweight LoRA branches to support multiple assignment strategies with minimal structural overhead. For clarity, the figure simplifies the merge between LoRA and FFN; see Eq. (3) for details. In all designs, only the primary branch output is fed to the next decoder layer.

ture (Fig. 1(c)). In this setup, we integrate parallel feedforward networks (FFNs) into the decoder layers and supervise each with either identical or distinct assignment strategies. The main branch and auxiliary branches respectively use one-to-one and one-to-many assignment strategies. Experimental results, detailed in Tab. 1, offer valuable insights through targeted comparisons. **(1)** Simplifying MS-DETR to a DETR-like structure with an independent auxiliary FFN (#2 vs. #3) results in only a negligible performance decline ($< 0.1\%$), underscoring that the benefits of decoder modifications for one-to-many training have likely been exaggerated. **(2)** Building on this, experiments with additional auxiliary branches using identical one-to-many strategies (#3, #4, #6) demonstrate no further gains and sometimes even degradation, revealing that sheer increases in supervision quantity can introduce counterproductive noise rather than enhancement. **(3)** In contrast, incorporating diverse assignment strategies (#3, #5, #7, #8) consistently delivers modest yet reliable improvements, suggesting a deeper mechanism at play where varied supervisory signals enrich the model’s learning dynamics. Collectively, these observations emphasize supervision diversity as a pivotal driver for advancing DETR performance, often surpassing the value of mere architectural expansions in fostering robust generalization. Nevertheless, assigning dedicated FFNs to each strategy escalates parameter counts and fosters optimization redundancies, which can subtly undermine the primary branch’s convergence efficiency and overall training stability.

Then, *how can we introduce diverse assignment supervision in a lightweight yet more effective manner?*

To affirmatively answer the question posed by our motivation, we introduce **LoRA-DETR**, a parameter-efficient framework designed to integrate diverse assign-

#	Model	N_{aux}	Param.	Div. Aux.	mAP \uparrow
1	(a)	0	47.3		43.7
2	(b)	1	53.6		47.6
3	(c)	1	53.6		47.5
4	(c)	2	59.9		47.3
5	(c)	2	59.9	✓	47.7
6	(c)	3	66.2		47.5
7	(c)	3	66.2	✓	48.0
8	(c)	5	78.8	✓	48.2

Table 1: Results with different decoder architectures and numbers of auxiliary branches. Model correspond to Fig. 1. N_{aux} denotes the number of auxiliary branches. Param. refers to the number of parameters during training, Div. Aux. indicates the use of diverse auxiliary assignment strategies, and mAP stands for mean Average Precision.

ment strategies by leveraging lightweight Low-Rank Adaptation (LoRA) (Hu et al. 2022) branches (see Fig. 1(d)). Specifically, we augment the primary FFN with distinct LoRA modules, each tailored to a specific assignment strategy, and jointly optimize them with the primary “one-to-one” branch via unified backpropagation. During the training phase, these LoRA branches serve exclusively as auxiliary paths, injecting diverse supervisory signals while maintaining maximum parameter efficiency. Crucially, they are detached during inference, ensuring that our method introduces no additional computational overhead and the inference process remains identical to the standard architecture. The LoRA mechanism facilitates the efficient modulation of the primary FFN using low-rank residual parameters, which

allows the main “one-to-one” branch to receive gradients directly from all auxiliary tasks, thereby promoting robust parameter sharing and joint optimization. Furthermore, by adjusting the rank of each LoRA module, we can control the representational capacity of these residual parameters, thus fine-tuning the contribution of each auxiliary supervision signal to the optimization of the primary branch.

As detailed above, our paper explores the “one-to-many” design paradigm in object detection algorithms from a fresh perspective. Experiments reveal that simply incorporating additional branches equipped with diverse “one-to-many” supervision consistently enhances performance without requiring complex structural modifications. Our proposed LoRA-DETR builds on these gains and outperforms strong DETR-style baselines in various settings. For instance, with Deformable DETR (Zhu et al. 2020) as the baseline, our method improves its performance from 43.7 AP to 49.0 AP on COCO (Lin et al. 2014) using three “one-to-many” LoRA-based auxiliary branches. Even with a strong baseline like Relation-DETR (Hou et al. 2024b), our method boosts its performance from 51.7 AP to 52.5 AP. Within this framework, LoRA serves as an intermediary that enables more effective interactions between auxiliary tasks and the primary end-to-end detection task. Compared to simply adding auxiliary FFN networks, LoRA-DETR not only boosts parameter efficiency but also scales more effectively as the number of auxiliary branches increases.

In summary, our contributions include: **(1)** We offer a fresh perspective on the “one-to-many” design paradigm in object detection algorithms. Empirical results show that simply adding branches with **diverse** “one-to-many” supervision consistently boosts performance without complex structural changes. **(2)** We propose LoRA-DETR as an extremely simple and scalable framework that integrates diverse assignment strategies into DETR-style detectors through lightweight LoRA branches. **(3)** LoRA-DETR delivers competitive or state-of-the-art results on various baselines, highlighting the effectiveness and broad applicability of our design.

Related Work

Assignment Strategies in Object Detection. Assignment strategies are critical in object detection. Classic two-stage detectors like Faster-RCNN (Ren et al. 2015) use max-IoU for positive/negative assignment. Subsequent improvements, such as ATSS (Zhang et al. 2020) introduce adaptive IoU thresholds, while one-stage detectors, such as YOLOX (Ge et al. 2021) and YOLOF (Chen et al. 2021), adopt SimOTA (Ge et al. 2021) and top-k assignment to improve matching quality. NMS is required in these methods to resolve overlapping detection boxes.

The advent of DETR reframed detection as a set prediction task using one-to-one bipartite matching, which simplifies the pipeline but suffers from slow convergence due to sparse positive samples. To address this, recent works propose one-to-many assignment strategies to improve training efficiency. We categorize existing one-to-many supervision strategies into two types: **(1) Static Label-Query Construction:** These methods explicitly construct query-label pairs,

where the supervision signal is predetermined and independent of the model’s behavior. For example, DN-DETR (Li et al. 2022) and DINO (Zhang et al. 2022) generate positive and negative queries by perturbing ground-truth boxes, while Co-DETR (Zong, Song, and Liu 2023) introduces auxiliary detection heads to define additional query-label pairs. **(2) Dynamic Label Assignment:** These methods generate auxiliary predictions within the network and assign supervision dynamically based on matching strategies. For instance, H-DETR (Jia et al. 2023) and Align-DETR (Cai et al. 2023) apply one-to-many assignment strategies to intermediate decoder layers. Group-DETR (Chen et al. 2023b) constructs multiple groups of queries, each of which independently performs one-to-one supervision. Similarly, DAC-DETR (Hu et al. 2023), MS-DETR (Zhao et al. 2024), and Mr. DETR (Zhang, Zhong, and Han 2024) introduce parallel prediction branches and apply one-to-many supervision alongside the primary one-to-one branch. Our method falls into the second category, but instead of relying on a *single* one-to-many assignment strategy, we introduce lightweight auxiliary LoRA branches trained with diverse strategies.

LoRA and Parameter-Efficient Adaptation. Low-Rank Adaptation (LoRA) (Hu et al. 2022) was initially introduced to efficiently fine-tune large-scale language models by injecting task-specific low-rank matrices into the weight updates. This approach enables adaptation with lightweight parameters overhead and has since been extended to various vision tasks, such as segmentation (Zhang and Liu 2023) and object tracking (Lin et al. 2024) with transformers. A key advantage of LoRA is its ability to support multiple learning objectives within a shared backbone through modular and lightweight branches. In our work, we leverage LoRA not only for parameter efficiency but also to encourage output diversity. Specifically, LoRA modules are used to decouple the learning signals from different assignment strategies, enabling the detector to produce diverse predictions.

Methodology

Preliminary

DETR Architecture. The DETR detector comprises three main components: a backbone, a transformer encoder, and a transformer decoder. It first encodes the input image \mathbf{I} into a feature sequence $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$ using the backbone and encoder. A set of object queries $\mathbf{Q} = \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n\}$, along with the image features \mathbf{X} , are then passed to the decoder to predict classification scores $\mathbf{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_n\}$ and bounding boxes $\mathbf{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n\}$. The decoder consists of multiple layers, each composed sequentially of a self-attention module, a cross-attention module, and a FFN.

Low-Rank Adaptation. LoRA is a lightweight fine-tuning technique that improves parameter efficiency by introducing trainable low-rank matrices into existing weight matrices of pre-trained models. Instead of updating the full weight matrix $W \in \mathbb{R}^{d \times k}$, LoRA freezes W and learns a low-rank residual in the form:

$$W' = W + \Delta W = W + BA, \quad (1)$$

where $A \in \mathbb{R}^{r \times k}$, $B \in \mathbb{R}^{d \times r}$ and $r \ll \min(d, k)$. The

matrices A and B are the only trainable parameters, while W remains frozen during training.

Quality-Aware Loss. In traditional object detection, classification is treated as a binary task with hard labels (0 for negatives, 1 for positives). However, this formulation does not reflect the varying quality of positive samples. Quality-aware classification loss addresses this issue by replacing hard labels with a continuous quality score (Pu et al. 2023; Cai et al. 2023; Zhang et al. 2021; Huang et al. 2024), allowing the model to focus more on high-quality predictions. A simple quality-aware loss, Assignment-Quality Loss, is included in experiments but is not a core part of our method.

LoRA Branches in Decoder Layer

As showed in Fig. 1(d), we modified the FFN module in the decoder layers. FFN typically consists of two linear layers W_1, b_1, W_2, b_2 with an activation function σ (e.g. ReLU or GELU) between. Formally, given an input feature $x \in \mathbb{R}^d$, the FFN computes

$$\text{FFN}(x) = W_2\sigma(W_1x + b_1) + b_2. \quad (2)$$

For each auxiliary branch i , we insert a LoRA module into both W_1 and W_2 , resulting in:

$$\text{FFN}^{\text{+LoRA}_i}(x) = W_2^i\sigma(W_1^ix + b_1) + b_2, \quad (3)$$

$$W_j^i = W_j + B_j^iA_j^i. \quad (4)$$

Unlike conventional LoRA usage where the base model is frozen to ensure parameter efficiency, we adopt LoRA as a lightweight mechanism to construct auxiliary branches. As such, both the original FFN and the LoRA modules are jointly optimized during training, allowing more effective gradient flow and interaction between the main and auxiliary objectives. For clarity, we simplify the illustration of the merged FFN and LoRA components in Fig. 1(d).

Given the input object queries \mathbf{Q}_j at the j -th layer, we compute the main branch output and simultaneously compute the auxiliary LoRA branch as

$$\mathbf{Q}_{j+1} = \text{FFN} \circ \text{CA} \circ \text{SA}(\mathbf{Q}_j), \quad (5)$$

$$\mathbf{Q}_{j+1}^i = \text{FFN}^{\text{+LoRA}_i} \circ \text{CA} \circ \text{SA}(\mathbf{Q}_j), \quad (6)$$

where \circ denotes the sequential composition of self-attention (SA), cross-attention (CA), and FFN. Only the main output \mathbf{Q}_{j+1} is forwarded to the next layer, while the auxiliary output \mathbf{Q}_{j+1}^i is used solely for training. All branches within the same layer share classification head and bounding box regression head to produce detection results.

We view the LoRA module as an intermediary that absorbs the supervision signals from one-to-many assignment, while preserving the end-to-end formulation of the main one-to-one branch. This design allows us to decouple the learning objectives, enabling diverse training signals to be injected without modifying the shared query representation.

Assignment Strategies

In LoRA-DETR, we apply distinct assignment strategies across decoder branches to enhance supervision diversity.

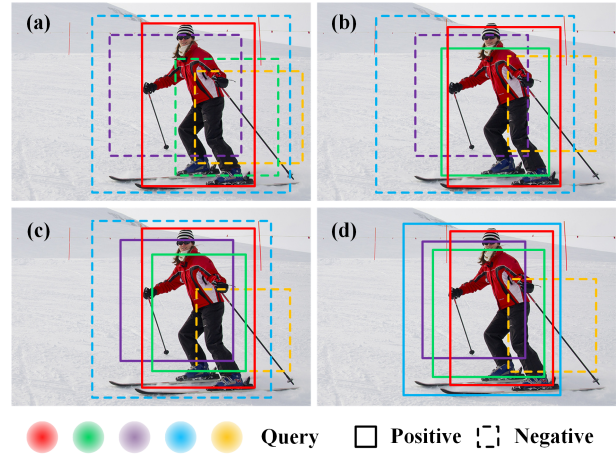


Figure 2: A set of queries simultaneously participates in diverse matching strategies. (a) illustrates the primary branch using one-to-one Hungarian matching, while (b), (c), and (d) depict auxiliary branches employing one-to-many matching with different hyperparameter of k .

The primary branch adopts the standard one-to-one matching strategy via the Hungarian algorithm, ensuring that each ground-truth box is assigned to a unique query prediction.

For LoRA branches, we adopt one-to-many matching strategies inspired by prior works (Ouyang-Zhang et al. 2022; Zhao et al. 2024; Hu et al. 2023; Zhang, Zhong, and Han 2024). Unlike the primary one-to-one assignment, each ground-truth box can be matched to multiple queries, providing denser supervisory signals and accelerating convergence. Each query prediction $\mathbf{p} = \langle \mathbf{c}, \mathbf{b} \rangle$ consists of a classification score \mathbf{c} and a predicted bounding box \mathbf{b} , while a ground-truth object is denoted as $\mathbf{y} = \langle \hat{\mathbf{c}}, \hat{\mathbf{b}} \rangle$. A matching score between \mathbf{p} and \mathbf{y} is defined as:

$$\mathbf{M}(\mathbf{p}, \mathbf{y}) = \alpha \cdot \mathbf{c}_{\hat{\mathbf{c}}} + (1 - \alpha) \cdot \text{IoU}(\mathbf{b}, \hat{\mathbf{b}}). \quad (7)$$

Given a threshold τ and a maximum match count k , we select up to k queries per ground-truth whose scores exceed τ . These queries are treated as positives for training. To introduce diverse supervisory signals across auxiliary branches, we vary the value of k to define different one-to-many tasks. Branches with smaller k focus on high-quality matches, while those with larger k allow more relaxed assignments.

Assignment Quality Loss

For the classification loss, we adopt a simplified variant of Varifocal Loss (Zhang et al. 2021) as our strong baseline. Following the design of DEIM (Huang et al. 2024), we apply the standard binary cross-entropy loss to positive samples using the quality score $s \in [0, 1]$ as the classification target, while omitting the original weighting factor introduced in Varifocal Loss. For negative samples, we retain the focal loss formulation to effectively mitigate class imbalance. The

Model	Num.	Param.	Epoches	Queries	mAP	AP _{50/75}	AP _{S/M/L}
Conditional-DETR (Meng et al. 2021)	-	-	50	300	40.9	61.8/43.3	20.8/44.6/60.2
Anchor-DETR (Wang et al. 2022)	-	-	50	300	42.1	63.1/44.9	22.3/46.2/60.0
DAB-DETR (Liu et al. 2022)	-	-	12	300	44.2	62.5/47.3	27.5/47.1/58.6
SQR-DETR (Chen et al. 2023a)	-	-	50	300	45.9	64.7/50.2	27.7/49.2/60.5
#Deformable DETR	0	47.3	12	300	43.7	62.2/46.9	26.4/46.4/57.9
DAC-DETR	1	47.3	12	300	47.1	64.8/51.1	29.2/50.6/62.4
MS-DETR	1	53.6	12	300	47.6	64.9/51.7	29.6/50.9/63.3
LoRA-DETR	1	48.1	12	300	48.8	65.6/52.8	30.4/52.6/64.5
LoRA-DETR	3	49.9	12	300	49.0	66.1/53.2	31.0/52.8/64.5
#Deformable DETR++	0	47.5	12	900	47.6	65.8/51.8	31.2/50.6/62.6
H-DETR	1	47.9	12	900	49.0	66.6/53.5	32.0/52.3/63.0
DAC-DETR	1	47.5	12	900	49.3	66.5/53.8	31.4/52.4/64.1
MS-DETR	1	53.8	12	900	50.0	67.3/54.4	31.6/53.2/64.0
LoRA-DETR	1	48.4	12	900	50.8	68.1/55.2	34.3/54.4/66.0
LoRA-DETR	3	51.4	12	900	51.0	68.5/55.6	34.1/54.7/66.2
#Deformable DETR++	0	47.5	24	900	49.8	67.0/54.2	31.4/52.8/64.1
MS-DETR	1	53.8	24	900	50.9	68.4/56.1	34.7/54.3/65.1
LoRA-DETR	1	48.4	24	900	52.1	69.4/56.8	34.5/55.6/67.0
#Relation-DETR	1	48.6	12	900	51.7	69.1/56.3	36.1/55.6/66.1
Relation-DETR	1	48.6	24	900	52.1	69.7/56.6	36.1/56.0/66.5
MI-DETR (Nan et al. 2025)	1	>75.0	12	900	52.4	69.8/57.0	35.6/56.1/67.2
LoRA-DETR	1+1	49.5	12	900	52.3	69.4/56.9	36.1/56.4/66.8
LoRA-DETR	1+3	51.3	12	900	52.5	69.8/57.3	37.2/56.5/66.8

Table 2: The performance on the COCO 2017 validation set. All models use ResNet-50 (He et al. 2016) as backbone. Num. is number of auxiliary branches and Param. is training parameter. We also consider the Hybrid branch proposed in H-DETR as an auxiliary branch. # stands for baseline.

resulting loss function is defined as:

$$\text{VFL}^+(p, s, y) = \begin{cases} -s \cdot \log(p) - (1-s) \cdot \log(1-p) & \text{if } y = 1 \\ -p^\gamma \cdot \log(1-p) & \text{if } y = 0. \end{cases} \quad (8)$$

where $p \in [0, 1]$ denotes the predicted classification score, and $y \in \{0, 1\}$ indicates the binary label.

In our framework, the definition of the quality score s varies across branches, reflecting their respective assignment strategies. Specifically, the one-to-one branch uses $s = \text{IoU}(\mathbf{b}, \hat{\mathbf{b}})$ between predicted and ground-truth boxes, while the one-to-many branches adopt the matching score $s = M(\mathbf{p}, \mathbf{y})$ introduced in Eq. (7), which combines classification confidence and localization quality. While the Quality-Aware Loss is employed in our experiments, it is not the primary focus of this paper.

Experiments

Implementation Details

Dataset. We conduct extensive experiments on the COCO dataset (Lin et al. 2014), training LoRA-DETR using the `train2017` split and evaluating on `val2017`. Performance is measured using standard COCO metrics, including AP (averaged over IoU thresholds from 0.50 to 0.95 in increments of 0.05), as well as AP₅₀, AP₇₅, and scale-specific AP scores: AP_S, AP_M, and AP_L.

Training. Our training setup follows that of MS-DETR, including the same set of hyperparameters and encoder proposal matching strategy. For $N_{aux} > 1$, we average the losses across all auxiliary branches. Additionally, when denoising queries (Zhang et al. 2022) are incorporated into training, we down-weight the auxiliary losses by a factor of 0.5 to balance the overall training objectives. For the baseline trained with focal loss for classification, we adopt the assignment-quality loss proposed in Eq. (8) with $\gamma = 1.5$. For other baselines that utilize alternative quality-aware classification losses, we follow the loss functions originally used in their respective implementations, ensuring a fair comparison. Moreover, we observed abnormally large gradient norms in the bounding box head during training and thus reduced its learning rate by a factor of 0.1.

Assignment Strategies. The primary branch adopts one-to-one matching, ensuring end-to-end inference. The LoRA branches employ the one-to-many matching strategy introduced by Eq. (7). For different numbers of auxiliary branches, we use different hyperparameters k in Tab. 4 to encourage diverse behaviors across the branches, with the corresponding results reported in Tab. 1 and Tab. 7.

Main Results

We integrate our approach into Deformable-DETR (Zhu et al. 2020) and Relation-DETR (Hou et al. 2024b). The

Model	Epoches	Queries	mAP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
DINO	12	900	56.8	75.4	62.3	41.1	60.6	73.5
Saliency-DETR (Hou et al. 2024a)	12	900	56.5	75.0	61.5	40.2	61.2	72.8
Rank-DETR (Pu et al. 2023)	12	900	57.6	76.0	63.4	41.6	61.4	73.8
EASE-DETR (Gao et al. 2024)	12	900	57.8	76.7	63.3	40.7	61.9	73.7
Stable-DETR	12	900	57.7	75.7	63.5	39.8	62.0	74.7
#Relation-DETR	12	900	57.8	76.1	62.9	41.2	62.1	74.4
Relation-DETR	24	900	58.1	76.3	63.2	42.0	62.7	73.7
MI-DETR (Nan et al. 2025)	12	900	58.2	76.5	63.4	42.5	62.8	74.6
LoRA-DETR	12	900	58.2	76.2	63.4	41.8	63.0	74.4

Table 3: The performance on the COCO 2017 validation set. All models use Swin-L (Liu et al. 2021) as backbone. # denotes the baseline model. We train our LoRA-DETR using bfloat16 mixed-precision for improved efficiency.

N_{aux}	1	2	3	5
Identical k	{6}	{6,6}	{6,6,6}	-
Diverse k	{6}	{3,6}	{2,4,6}	{2,3,4,5,6}

Table 4: Hyperparameter k settings under different numbers of auxiliary branches (N_{aux}).

experimental results are presented in Tab. 2 and Tab. 3. Our method consistently yields stable performance improvements across different baselines. Compared to other one-to-many strategies, our approach demonstrates a clear advantage. Notably, Relation-DETR already incorporates the denoising training (introduced in DINO) and contrastive hybrid one-to-many supervision, providing a pretty strong baseline. But surprisingly, even when directly applied to its primary branch without the contrastive setting in Relation-DETR, our method achieves a significant gain of 0.8%. Moreover, with only 12 epochs of training, our approach outperforms Relation-DETR trained for 24 epochs, highlighting superior training efficiency. In contrast to MI-DETR, whose decoder introduces repeated modules and increases total inference cost (GFlops) by over 20% (Nan et al. 2025), our method adds no extra computation.

Ablation Study

Number of Auxiliary Branches. We conduct an ablation study to investigate the impact of the number of auxiliary branches on model performance. As shown in Tabs. 1 and 7, increasing the number of auxiliary branches generally leads to consistent performance improvements, demonstrating the benefit of diverse supervision. As illustrated in Fig. 3(a), the decrease in loss becomes relatively smaller when the number increases from 2 to 3/5. Given the trade-off between performance and computation, we choose 3 as the maximum number of auxiliary branches.

Diverse Assignment Strategies. Besides varying k , diversity can also be introduced by perturbing other parameters in Eq. (7). Under the 3 auxiliary branches setting, Diverse k still delivers the most substantial performance gain, with $48.0 > 47.7 > 47.3$ for Diverse k , Diverse τ , and Diverse α , respectively. Adjusting k thus appears to be the most effective strategy at this stage.

Rank r	16	32	64	w/o
Prim.	48.5	48.8	48.2	48.2
Aux. w. NMS	49.0	49.1	48.8	48.7

Table 5: Performance (mAP) comparison between the primary (Prim.) and auxiliary (Aux.) branches under different values of the hyperparameter rank r , where the auxiliary branch incorporates NMS. w/o means using naive FFN branch instead of LoRA.

Method	Queries	Aux.	Time ↓	Param.
MS-DETR	300	1 × 300	44	53.8
DAC-DETR	300	1 × 300	43	47.3
H-DETR	300	1 × 300	46	47.3
LoRA-DETR	300	1 × 300	40	48.1
LoRA-DETR	300	2 × 300	46	49.0
LoRA-DETR	300	3 × 300	53	49.9

Table 6: Training cost comparison with other one-to-many methods. Aux. indicates auxiliary supervision (auxiliary branches × queries per branch). Time (minute) is GPU time per epoch on 8×A100, and Param. is training parameters.

LoRA Rank. The representational capacity of the LoRA module is controlled by the hyperparameter rank r . A higher rank offers greater flexibility to fit the auxiliary supervision; however, if the capacity is too high, the LoRA module may overfit to the one-to-many targets or introduce conflicting gradients to the shared parameters. This can potentially slow down the convergence of the main branch or harm generalization. Conversely, too small rank may underutilize the auxiliary signals and limit the benefits of diverse supervision. We conduct an ablation study on the choice of LoRA rank in Tab. 5 and the best rank $r = 32$.

Progressive Ablation. In Tab. 7, we present a step-by-step transition from MS-DETR to our proposed LoRA-DETR. We begin by reproducing the results of MS-DETR, obtaining a comparable mAP to the reported value, with minor variations in other metrics. We then reorder the cross-attention and self-attention modules to follow the standard sequence (Architecture (c)), which results in a slight perfor-

Model	Param.	mAP	AP _{50/75}	AP _{S/M/L}
MS-DETR* (Zhao et al. 2024)	53.6	47.6	65.0/51.9	29.9/50.4/62.9
+ Architecture (b) → (c)	53.6	47.5	64.9/51.5	28.4/50.3/63.7
+ Box Head Learning Rate × 0.1	53.6	47.7	65.4/52.0	29.7/50.8/63.0
+ Assignment-Quality Loss	53.6	48.2	65.1/52.3	30.2/51.7/64.0
+ Architecture (c) → (d)	48.1	48.8	65.6/52.8	30.4/52.6/64.5
+ $N_{aux} \rightarrow 3$	49.9	49.0	66.1/53.2	31.0/52.8/64.5
+...+ Assignment-Quality Loss †	53.8	50.4	67.5/54.8	33.2/53.7/65.2
+ Architecture (c) → (d) †	48.4	50.8	68.1/55.2	34.3/54.4/66.0
+ $N_{aux} \rightarrow 3$ †	51.4	51.0	68.5/55.6	34.1/54.7/66.2

Table 7: Progressive Ablation. The table reports the model performance (AP) and corresponding parameter counts (Param.) at each stage. * indicates our reproduced results. † marks results based on Deformable DETR++.




Figure 3: One-to-one training loss of the primary branch. (a) Varying the number of auxiliary branches. (b) Comparing different architectural variants, as defined in Fig. 1. The x-axis denotes epochs and the y-axis denotes the loss value.

mance drop of 0.1%. Building on this revised architecture, we first adjust the learning rate of the bounding box head and subsequently introduce our assignment-quality loss, leading to a cumulative performance gain of 0.7% to 48.2%. Then, we incorporate our proposed LoRA-DETR design (Architecture (d)), which brings an additional improvement of 0.6% on this strong baseline. Finally, as the number of auxiliary branches increases, the performance continues to improve consistently.

Training Loss. We present in Fig. 3(a) the convergence behavior of the one-to-one training loss on the primary branch under varying numbers of auxiliary branches. As the number of auxiliary branches increases, the training loss consistently decreases across the same number of epochs, suggesting that more diverse auxiliary tasks can better facilitate the one-to-one matching process. In Fig. 3(b), comparing the convergence behavior of Arch.(c) (Fig. 1(c)) and LoRA-DETR (Fig. 1(d)) under assignment-quality loss setting, both models perform similarly in the early training stage, but LoRA-DETR shows a consistently lower loss as training progresses. This suggests that the LoRA branch contributes more beneficial gradients to the primary branch during the later stages of training, thereby facilitating better convergence. As shown in Tab. 2, under the Deformable-DETR++ setting, our method outperforms MS-DETR by 0.8% mAP with only 12 epochs of training (50.8% vs.

50.0%), and achieves a larger margin of 1.2% mAP with 24 epochs (52.1% vs. 50.9%), reaching performance on par with Relation-DETR (52.1%).

Computation Cost Analysis. Compared to other one-to-many approaches, our method requires only a *single* forward pass through the decoder to obtain outputs from all branches, incurring only a small amount of additional computation from the LoRA modules and partially repeated FFN operations. In contrast, methods such as H-DETR and Group-DETR need to recompute the entire decoder layer for each branch. The Tab. 6 provides a comparison of the training overhead across different methods.

Discussion. MS-DETR (Zhao et al. 2024) includes an ablation study with a design similar to our naive architecture (see Fig. 1(c)), but without introducing task-specific FFNs. In this setting, the one-to-one and one-to-many branches completely share the same FFN, with separation applied only at the output heads. This shallow decoupling results in a lower mAP (47.0%) compared to our naive design with separate FFNs (47.5%). The performance gap highlights that, despite both branches targeting object detection, their distinct label assignment strategies lead to different optimization behaviors and should be treated as heterogeneous tasks. Completely sharing intermediate computation in such cases may cause conflicts and hinder learning. In contrast, LoRA-DETR introduces lightweight, task-specific branches using LoRA modules, enabling a controllable degree of task decoupling through rank r adjustment, while maintaining minimal computational overhead.

Conclusion

This paper presents LoRA-DETR, a simple and scalable framework that introduces diverse supervision into DETR-style object detectors through lightweight LoRA-based auxiliary branches. The proposed method supports diverse assignment strategies with minimal computational and parameter overhead. In this framework, LoRA serves as an intermediary that balances parameter sharing and task-specific specialization. LoRA-DETR achieves competitive or state-of-the-art results under various settings, validating the effectiveness of our design. We believe this framework provides a strong and extensible baseline for future research on efficient transformer-based detector.

Acknowledgements

This work was supported in part by the Beijing Natural Science Foundation (Grant No. JQ22014, L223003), the Natural Science Foundation of China (Grant No. 62422317, U22B2056, 62036011, 62192782, U2441241, 62503323).

References

- Cai, Z.; Liu, S.; Wang, G.; Ge, Z.; Zhang, X.; and Huang, D. 2023. Align-detr: Improving detr with simple iou-aware bce loss. *arXiv preprint arXiv:2304.07527*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*.
- Chen, F.; Zhang, H.; Hu, K.; Huang, Y.-K.; Zhu, C.; and Savvides, M. 2023a. Enhanced training of query-based object detection via selective query recollection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 23756–23765.
- Chen, Q.; Chen, X.; Wang, J.; Zhang, S.; Yao, K.; Feng, H.; Han, J.; Ding, E.; Zeng, G.; and Wang, J. 2023b. Group detr: Fast detr training with group-wise one-to-many assignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6633–6642.
- Chen, Q.; Wang, Y.; Yang, T.; Zhang, X.; Cheng, J.; and Sun, J. 2021. You only look one-level feature. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13039–13048.
- Gao, Y.; Sun, Y.; Ding, X.; Zhao, C.; and Liu, S. 2024. Ease-detr: Easing the competition among object queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17282–17291.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. YoloX: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Hou, X.; Liu, M.; Zhang, S.; Wei, P.; and Chen, B. 2024a. Saliency detr: Enhancing detection transformer with hierarchical saliency filtering refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17574–17583.
- Hou, X.; Liu, M.; Zhang, S.; Wei, P.; Chen, B.; and Lan, X. 2024b. Relation detr: Exploring explicit position relation prior for object detection. In *European Conference on Computer Vision*, 89–105. Springer.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, Z.; Sun, Y.; Wang, J.; and Yang, Y. 2023. DAC-DETR: Divide the attention layers and conquer. *Advances in Neural Information Processing Systems*, 36: 75189–75200.
- Huang, S.; Lu, Z.; Cun, X.; Yu, Y.; Zhou, X.; and Shen, X. 2024. DEIM: DETR with Improved Matching for Fast Convergence. *arXiv preprint arXiv:2412.04234*.
- Jia, D.; Yuan, Y.; He, H.; Wu, X.; Yu, H.; Lin, W.; Sun, L.; Zhang, C.; and Hu, H. 2023. Detr with hybrid matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19702–19712.
- Li, F.; Zhang, H.; Liu, S.; Guo, J.; Ni, L. M.; and Zhang, L. 2022. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13619–13627.
- Lin, L.; Fan, H.; Zhang, Z.; Wang, Y.; Xu, Y.; and Ling, H. 2024. Tracking meets lora: Faster training, larger model, stronger performance. In *European Conference on Computer Vision*, 300–318. Springer.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *ICCV*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part v 13*, 740–755. Springer.
- Liu, S.; Li, F.; Zhang, H.; Yang, X.; Qi, X.; Su, H.; Zhu, J.; and Zhang, L. 2022. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*.
- Meng, D.; Chen, X.; Fan, Z.; Zeng, G.; Li, H.; Yuan, Y.; Sun, L.; and Wang, J. 2021. Conditional detr for fast training convergence. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3651–3660.
- Nan, Z.; Li, X.; Dai, J.; and Xiang, T. 2025. MI-DETR: An Object Detection Model with Multi-time Inquiries Mechanism. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 4703–4712.
- Ouyang-Zhang, J.; Cho, J. H.; Zhou, X.; and Krähenbühl, P. 2022. Nms strikes back. *arXiv preprint arXiv:2212.06137*.
- Pu, Y.; Liang, W.; Hao, Y.; Yuan, Y.; Yang, Y.; Zhang, C.; Hu, H.; and Huang, G. 2023. Rank-DETR for high quality object detection. *Advances in Neural Information Processing Systems*, 36: 16100–16113.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Wang, Y.; Zhang, X.; Yang, T.; and Sun, J. 2022. Anchor detr: Query design for transformer-based detector. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 2567–2575.
- Zhang, C.-B.; Zhong, Y.; and Han, K. 2024. Mr. DETR: Instructive Multi-Route Training for Detection Transformers. *arXiv preprint arXiv:2412.10028*.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H.-Y. 2022. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*.

- Zhang, H.; Wang, Y.; Dayoub, F.; and Sunderhauf, N. 2021. Varifocalnet: An iou-aware dense object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8514–8523.
- Zhang, K.; and Liu, D. 2023. Customized segment anything model for medical image segmentation. *arXiv preprint arXiv:2304.13785*.
- Zhang, S.; Chi, C.; Yao, Y.; Lei, Z.; and Li, S. Z. 2020. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9759–9768.
- Zhao, C.; Sun, Y.; Wang, W.; Chen, Q.; Ding, E.; Yang, Y.; and Wang, J. 2024. Ms-detr: Efficient detr training with mixed supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17027–17036.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv:2010.04159*.
- Zong, Z.; Song, G.; and Liu, Y. 2023. Dets with collaborative hybrid assignments training. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6748–6758.