

RecursiveDet: End-to-End Region-based Recursive Object Detection

Jing Zhao¹, Li Sun^{1,2*}, Qingli Li¹

¹Shanghai Key Laboratory of Multidimensional Information Processing,

²Key Laboratory of Advanced Theory and Application in Statistics and Data Science,
 East China Normal University, Shanghai, China

Abstract

End-to-end region-based object detectors like Sparse R-CNN usually have multiple cascade bounding box decoding stages, which refine the current predictions according to their previous results. Model parameters within each stage are independent, evolving a huge cost. In this paper, we find the general setting of decoding stages is actually redundant. By simply sharing parameters and making a recursive decoder, the detector already obtains a significant improvement. The recursive decoder can be further enhanced by positional encoding (PE) of the proposal box, which makes it aware of the exact locations and sizes of input bounding boxes, thus becoming adaptive to proposals from different stages during the recursion. Moreover, we also design centerness-based PE to distinguish the RoI feature element and dynamic convolution kernels at different positions within the bounding box. To validate the effectiveness of the proposed method, we conduct intensive ablations and build the full model on three recent mainstream region-based detectors. The RecursiveDet is able to achieve obvious performance boosts with even fewer model parameters and slightly increased computation cost. Codes are available at <https://github.com/bravezzzzz/RecursiveDet>.

1. Introduction

Object detection has been intensively investigated by computer vision community for decades. Traditional detectors built by deep convolutional neural network (CNN) are either anchor-based [13, 32, 26] or anchor-free [30, 37, 46]. The former performs classification and regression based on pre-defined densely tiled bounding boxes, while the latter only assumes grid points in the 2D image plane. On the other hand, detection can be completed in a single stage, two stages or even multiple cascade stages. The single-stage

*Corresponding author, email: sunli@ee.ecnu.edu.cn. This work is supported by the Science and Technology Commission of Shanghai Municipality under Grant No. 22511105800, 19511120800 and 22DZ2229004.

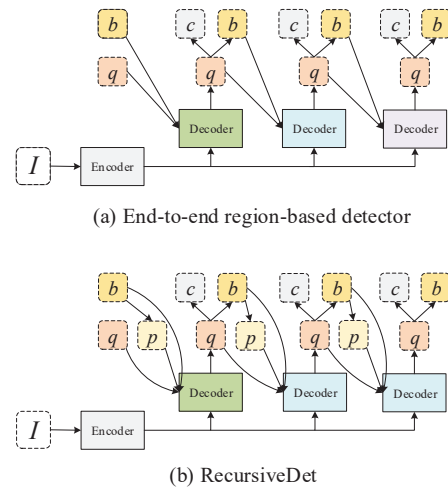


Figure 1. Comparison between end-to-end region-based detector and RecursiveDet. I is the input image feature. q and b are proposal feature and proposal box at input. c and b indicate the class and bounding box of the predictions at output. In RecursiveDet, we encode b into PE vector p . (a) Detectors like Sparse R-CNN, AdaMixer and DiffusionDet have cascade stages, which bring a huge amount of parameters. (b) Our proposed RecursiveDet simply shares the decoder parameters and makes the structure recursive, which reduces the model size.

method directly gives predictions without further modifications, which is usually simple and efficient. Two- or multi-stage methods repeatedly make corrections based on previous results, which offer better results but cost more model parameters and calculations. Except for the first stage, later stages are usually region-based, focusing on the local region within the bounding box, which is often realized by RoI Align [15].

Although these detectors have been widely used, they are often complained for complex pre- and post-designs, e.g., anchor sizes, standards for positive and negative samples, and non-maximum suppression (NMS) on results. DETection TRansformer (DETR) simplifies the framework. It relies on multi-stage transformer decoder layers to update learnable queries into object features, which can be decoded

into bounding box predictions at each stage. During training, DETR utilizes bipartite matching to determine positive and negative query samples in a dynamic way. Since each ground truth (GT) box is assigned with only one query as positive, NMS is no longer necessary. However, DETR has extremely low efficiency, mainly due to the lack of inductive bias in the transformer decoder. Later works incorporate either bounding box positional encoding (PE), iterative refinement or local deformation into the decoder, hence greatly improve the training efficiency.

Similar to DETR, region-based methods [35, 11, 4] also treat object detection as a set prediction task between predictions and GTs, and they all adopt a multi-stage strategy. Sparse R-CNN is a representative work depending on RoI Align. It employs a set of learnable proposal features to generate dynamic conv filters for image RoI features within corresponding proposal boxes. Compared to DETR, region-based methods obtain accurate results and high data efficiency. However, the model sizes are obviously larger than DETR-series, which not only limit their applications but also cause degradation for the small-scale dataset. We find that some specific layers in region-based methods account for most of the parameter number, so they can potentially be improved. Moreover, since proposal features are equivalent to object queries in DETR, it should be practical to introduce PE of bounding boxes like [25], so the self attention and dynamic conv know the exact global and local positions of image features.

Based on the above considerations, this paper proposes the RecursiveDet, a common cascade decoder structure built on dynamic conv for region-based detectors. Specifically, we share the decoder parameters and make the same module recursive in different stages, allowing training it based on proposal boxes of different qualities. By simply sharing all decoder parameters among different stages, the model size is apparently reduced while its performance only suffers a slight drop. Moreover, since the dynamic layer for generating conv kernels and out layer after dynamic conv are huge, we intend to make full use of them and apply the dynamic conv for more layers. In practice, we update the proposal feature after the first group of dynamic convs. Then the updated feature is given to the same dynamic layer, specifying conv kernels for the second group. In other words, the dynamic layer and out layer are reused to increase the model depth within a stage without extra parameter cost.

Due to sharing parameters of decoder stages, its input becomes indispensable to distinguish each decoding stage. Therefore, more information about the previous predictions needs to be utilized as input. Inspired by [28, 39, 25], we encode the positions and shapes of the bounding boxes into PE vectors, which are then employed for both position-aware self attention and dynamic conv. Since the PE of bound-

ing box is only computed from the global image coordinate and shape size, we design a centerness-based PE to represent different locations within the RoI as compensation. By combining the PE of bounding box and centerness within dynamic conv, our RecursiveDet model is able to fully exploit both proposal features and boxes in the end-to-end region-based detector.

To validate the effectiveness of the proposed RecursiveDet, intensive experiments and ablation studies are conducted on MS-COCO dataset. We implement RecursiveDet in three mainstream detectors, including Sparse R-CNN, AdaMixer and DiffusionDet, and achieve consistent performance boosts. We point that Sparse R-CNN and AdaMixer are end-to-end, while DiffusionDet is not. In summary, the contribution of this paper lies in the following aspects.

- We propose to share the decoder parameters among different stages in the region-based detector, which significantly reduces the model size without severely degrading the performance. Moreover, we repeatedly use the dynamic layer to generate conv kernels for more layers in a single stage, increasing the model depth for further performance gain.
- We design the bounding box PE according to the geometry information from predictions. Besides, the local coordinate is represented based on centerness PE, discriminating each feature within RoI. We make these PEs to participate in self attention and dynamic conv in the region-based detector.
- The proposed RecursiveDet is implemented in different end-to-end region-based detectors, showing that it enhances all of them and reduces the model size at the same time.

2. Related Work

CNN-based detector. Earlier CNN-based object detectors hold assumptions on densely located anchor boxes [13, 26, 22, 1] or grid points [19, 47, 37, 40] in 2D plane, and they compute the classification and bounding box regression loss for each candidate sample. Anchor-based methods heavily rely on hyper-parameters of anchor sizes. Moreover, they need heuristic rules, like the IoU threshold, to determine positive and negative anchors. The NMS post operation is also cumbersome. Grid point-based methods make detectors free from anchor boxes. They directly locate the GT box from the grid, however, positive and negative assignment is still a critical issue, which is often realized in a simple way. *E.g.*, FCOS [37] assumes the points inside a GT box to be positive, and RepPoints [40] labels a grid point nearest to the GT box as the positive. Giving the importance of sample assignment, researchers propose different schemes [43, 18, 12] to prevent heuristic rules. Another

way to characterize different detectors is based on the number of detection stages. Single stage detectors [26, 30, 31] are more efficient than two- [13, 32, 7] or multi-stage [1] competitors. They directly give results from whole image, while two- or multi-stage methods bring region-based feature from RoI Align to later stages to increase the accuracy.

Detection transformer. DETR [2] opens a new era for object detection, since it eliminates most handcrafted designs. It starts from a set of learnable queries and employs multiple decoder attention blocks to update them into object features for the detection head. DETR adopts dynamic bipartite matching to build one-to-one relation between the GTs and predictions. However, the slow convergence and large training set requirement become its obstacle. Later works [49, 10, 41, 28, 39, 25, 20, 42, 5] focus on incorporating spatial prior or extra query group into model for faster training and better results. Deformable DETR [49] utilizes deformable attention [8, 48] and iterative refinement. SMCA [10] predicts a spatial mask from query to modulate attention matrix. REGO [5] employs RoI Align in DETR-series. Conditional [28] and Anchor [39] DETR encode the center of bounding box as PE to assist the decoder, and DAB-DETR [25] further extends PE to bounding box coordinates including width and height. DN-DETR [20] introduces denoising query group to accelerate training. DINO [42] improves it through contrastive learning and query selection.

End-to-end region-based detector. Similar to DETR-series, end-to-end region-based detectors [35, 11, 45, 44, 4] start from learnable proposal features and update them into object features for final predictions in a cascade manner. They also borrow the dynamic bipartite matching in DETR to prevent NMS. However, instead of cross attention, dynamic conv is used to connect image with proposal features. Moreover, such detectors sample RoI feature within the bounding box from the previous stage for dynamic conv. The first representative work is Sparse R-CNN [35]. To further speed up the convergence, AdaMixer [11] replaces RoI Align by sampling points directly from the 3D feature pyramid and borrows the idea of MLP-Mixer [38] which changes axis of the second time dynamic conv. Diffusion-Det [4] first utilizes training loss in diffusion model, and builds a multi-step multi-stage detector based on Sparse R-CNN. Our RecursiveDet belongs to this type, and it can be applied in most of them to improve the cascade decoding stages.

Recursive module in deep model. The typical recursive model in neural networks is RNN [16, 6], which is unrolled in multi-step during training. Many models [17, 21, 36] with recursive module have been proposed to address different tasks. Recent work [34] applies recursive transformer block to increase the model depth. In object detection, recursive backbone [27] and FPN [29] have been proposed, however, none of the work deals with recursive decoder at-

attention, particularly for object detection task.

3. Method

3.1. Preliminaries

This paper focuses on the end-to-end region-based detectors, among which Sparse R-CNN is the first typical model. It has a learnable proposal feature set $\mathcal{Q} = \{q_i | q_i \in \mathbb{R}^c\}$, in which each q_i is a c dimension vector to represent an object in image. Here $i = 1, 2, \dots, N$ and N is the total number of proposal features. Correspondingly, it keeps a set of proposal boxes $\mathcal{B} = \{b_i | b_i = (x, y, w, h)\}$ with the same number of \mathcal{Q} , where b_i is a bounding box centered at (x, y) with its width and height representing by w and h . According to \mathcal{B} , the image RoI feature set $\mathcal{F} = \{f_i | f_i \in \mathbb{R}^{7 \times 7 \times c}\}$ is obtained from a multi-resolution feature pyramid by RoI Align operation, and each element is of the spatial size 7×7 . Sparse R-CNN has multiple cascade stages that gradually refine \mathcal{B} to approach the GT box set. In each stage, the self attention within \mathcal{Q} set is first computed. Then, the proposal feature q_i is given to a huge dynamic layer, providing a pair of dynamic conv kernels $k_i \in \mathbb{R}^{c \times d}$ and $v_i \in \mathbb{R}^{d \times c}$ for corresponding RoI feature f_i . Here k_i and v_i form a bottleneck, which can be regarded as two successive layers consisting of 1×1 conv whose parameters are specified as k_i and v_i . c and d are the channel of input RoI feature and hidden dimension respectively, with $c \gg d$. After the dynamic convs, f_i is updated into f'_i , which is then mapped into object feature o_i through an out layer, so that it can be utilized by the detection head to give results.

Sparse R-CNN performs the one-to-one dynamic interactions between \mathcal{Q} and \mathcal{B} . In each stage, it not only predicts the bounding boxes, but also updates \mathcal{Q} and \mathcal{B} . Particularly, object feature o_i updates proposal feature q_i , preparing for the next stage, and bounding box prediction also updates b_i as the next stage input. During training, Sparse R-CNN dynamically computes a cost matrix between the predictions and GTs. It carries out bipartite matching based on it, so that the GT box is only assigned with one prediction. The one-to-one sample assignment strategy prevents it from NMS post processing. The cost and training loss are the same and they can be formulated as $\mathcal{L} = \lambda_{cls} \mathcal{L}_{cls} + \lambda_{L1} \mathcal{L}_1 + \lambda_{giou} \mathcal{L}_{giou}$, where \mathcal{L}_{cls} is the focal BCE loss [23] for classification, \mathcal{L}_1 and \mathcal{L}_{giou} are L1 and generalized IoU loss [33] for regression.

Besides Sparse R-CNN, AdaMixer and DiffusionDet are also region-based detectors, the former is end-to-end while the latter is not. They have similar cascade structures consisting of dynamic convs whose kernels are provided from dynamic layer based on proposal feature set \mathcal{Q} . Our RecursiveDet shares the basic computation pipeline with all of them, which is shown in Fig. 2 (a). And it is compatible with all of them, with a significant reduction on the model

size and increase on the result accuracy.

3.2. Recursive Structure for Decoder

Compared with DETR-series, end-to-end region-based detectors usually have a large decoder, which obviously expands their model size. The decoder Dec^t at stage t , parameterized by θ^t , takes the current proposal set \mathcal{Q}^t and box set \mathcal{B}^t as input, and updates them into \mathcal{Q}^{t+1} and \mathcal{B}^{t+1} for the next stage. We formalize the decoder in Eq. (1). Note that RoI feature set \mathcal{F}^t is also utilized by Dec^t .

$$[\mathcal{Q}^{t+1}, \mathcal{B}^{t+1}] = Dec^t(\mathcal{Q}^t, \mathcal{B}^t, \mathcal{F}^t; \theta^t) \quad (1)$$

The key module in Dec is the dynamic conv, and it builds the connection between proposal feature q_i and image RoI feature f_i . However, we find one distinction between cross attention and dynamic conv. The former is usually lightweight, while the latter depends on a heavy dynamic layer Dyn , which maps q_i into two sets of kernels k_i and v_i , as is shown in Eq. (2).

$$[k_i, v_i] = Dyn(q_i) = FC_{c \rightarrow 2c \times d}(q_i) \quad (2)$$

Actually, Dyn is a single FC layer with its input dimension c the same as q_i , and output dimension equals $2c \times d$, which accommodates both k_i and v_i . It accounts for a large percentage of model parameters in Dec . Another module, which costs many params and does not even exist in DETR-series, is the out layer Out . It is also an FC layer for reducing the spatial dimension from 7×7 to 1, as is illustrated in Eq. (3). The input of Out is the updated RoI feature f'_i and the output is object feature o_i , which is utilized by the detection head for prediction and the next stage Dec^{t+1} as its input proposal feature.

$$o_i = Out(f'_i) + q_i = FC_{7 \times 7 \times c \rightarrow c}(f'_i) + q_i \quad q_i^{t+1} = o_i \quad (3)$$

The module details within Dec , particularly, the connection between different modules are illustrated in Fig. 2 (a).

In following Tab. 1, the number of parameters in a single stage of the decoder is listed. We can see that a single Dyn has 8.4M parameters, therefore, its parameter cost in six different stages already surpasses the full model of DETR. The second largest module is Out , and it costs 3.2M parameters in Sparse R-CNN and DiffusionDet, and 8.4M in AdaMixer. Note that the multi-head cross or deformable attention layer in DETR-series is much smaller than Dyn .

Based on the above analysis, we make two simple but important modifications to the decoder. The principle is to fully exploit it, particularly the Dyn and Out layers. Considering the large size of Dec and its cascade structures in multiple stages, we share their params θ among stages to reduce the model size, thus Dec becomes recursive with the same module repeatedly appearing in different stages.

Method	MSA	<i>Dyn</i> /MCA	<i>Out</i>	Head	Others	Total
DETR	0.3M	0.3M	-	26.0K*	1.0M	1.6M
Deformable DETR	0.3M	0.2M	-	0.2M	0.5M	1.2M
DAB-DETR	0.4M	0.4M	-	26.0K*	1.2M	2.0M
Sparse R-CNN	0.3M	8.4M	3.2M	0.3M	1.0M	13.2M
AdaMixer	0.3M	8.4M	8.4M	0.1M	1.2M	18.4M
DiffusionDet	0.3M	8.4M	3.2M	0.3M	1.6M	13.8M

Table 1. Analysis of parameters in a single decoder layer. 'MSA', '*Dyn*', 'MCA', '*Out*' and 'Head' indicate multi-head self attention, dynamic layer, multi-head cross attention, output layer and detection head. 'Others' includes ffn, layer norm, etc. Note that, the detection head in DETR and DAB-DETR is shared among all stages, thus results with '*' are divided by total stage number.

However, the recursion may result in unstable gradients during training, which is commonly encountered by RNN. Supposing the loss in stage t is \mathcal{L}^t , its gradient $\frac{\partial \mathcal{L}^t}{\partial q^t}$ takes its effect into stage $(t-1)$ since $o^{t-1} = q^t$, hence the Jacobian matrix $\frac{\partial q^t}{\partial q^{t-1}}$ keeps multiplying until the first stage, causing the possible gradient vanishing or exploding. We argue that the total recursion number is not big. Moreover, the detection loss is also computed in the intermediate stages, so the recursion of Dec does not bring instabilities during training. In practice, considering the difficulty of the first decoding stage, detecting the object from the initial whole image or randomly located bounding box, we keep it out of the recursion loop with unique parameters to boost the performance. This causes a slightly increased model size.

Besides the recursion in different stages, we additionally make full use of Dyn and Out layers by making a short recursion from object feature o_i to the input of the dynamic layer Dyn within the same stage. In other words, o_i is given to Dyn layer again for generating new kernels k'_i and v'_i . And they are used to perform dynamic conv on RoI f_i again and update it into f''_i , which is then given to Out to specify o'_i . We emphasize that the in-stage recursion does not increase the model size, since all trainable modules in the recursive loop, such as Dyn and Out , share the same parameters. It only increases the amount of computation.

3.3. Bounding Box Positional Encoding

By sharing Dec among stages, the model gets compact and already achieves a better result. However, due to the common model parameters, Dec treats input q_i and b_i from different stages in the same way, which could reduce its adaptability and be a limitation for the performance. We intend to introduce more information from the previous stage and make a compensation, so that Dec has a sense of decoding stages and adjusts itself accordingly. Note that for DETR-series, it is common to give bounding box PEs to decoder. They help better model the relation between queries and keys during self and cross attention, therefore accurately representing objects. However, PE is usually ignored by end-to-end region-based detectors, mainly due to the de-

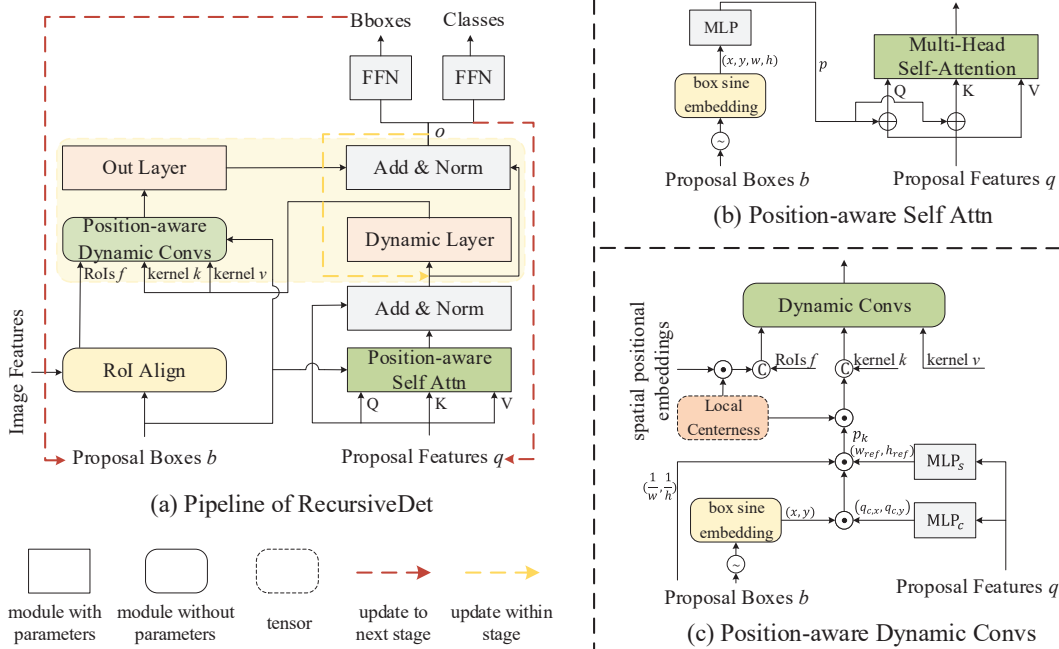


Figure 2. The details about decoder structure in RecursiveDet. (a) An overview of our method. We share the decoder parameters among different stages and make it recursive. The Dynamic layer and Out layer repeat twice for increasing the depth of decoder. The object feature o and predicted bounding boxes are updated to the next stage as proposal features q and proposal boxes b . The bounding boxes are encoded to PE vectors p and p_k to participate in position-aware self attention and dynamic conv in (b) and (c), respectively. As shown in (c), the local coordinate within RoI is calculated into a centerness mask to assist the bounding boxes PE in dynamic conv.

struction of global position by RoI Align.

We argue that PE of the bounding box can still be valid even for region-based detectors. First, detectors like Sparse R-CNN still need self attention to model the pairwise similarities among $q_i \in \mathcal{Q}$, which can be effectively measured between the spatial positions of two bounding boxes. Second, the dynamic kernels can be more adaptive and directly related to location and shape. So we build the position-aware self attention and dynamic conv modules, as is shown in Fig. 2 (b) and (c). We adopt a similar strategy for self attention with [25]. It first maps the 4D vector (x, y, w, h) to sinusoidal embedding. Then it is further projected into a PE vector $p \in \mathbb{R}^c$ by an MLP, and p is finally added onto the query and key tokens before self attention. For dynamic conv, the center (x, y) and box shape (w, h) are encoded in a separate manner. Utilizing two MLPs, denoted by MLP_c and MLP_s , proposal feature q is mapped into two vectors q_c and q_s in the geometry space, reflecting the center and shape of bounding box, respectively. The computations of q_c and q_s are provided in Eq. (4). Here q_s , consisting of two scalars w_{ref} and h_{ref} , is predicted from q by MLP_s . $q_{c,x}$ and $q_{c,y}$ together form q_c , and each contributes half of the full dimension.

$$\begin{aligned} q_c &= [q_{c,x}, q_{c,y}] = \text{MLP}_c(q) \in \mathbb{R}^c \\ q_s &= [w_{ref}, h_{ref}] = \text{MLP}_s(q) \in \mathbb{R}^2 \end{aligned} \quad (4)$$

Then, they are modulated by geometry features from (x, y) and (w, h) , specifying p_k for kernels. We formalize this in Eq. (5). Here \odot means the Hadamard product. $p_k(x, w)$ and $p_k(y, h)$ are finally concatenated into p_k .

$$\begin{aligned} p_k(x, w) &= \frac{w_{ref}}{w} \text{Sinusoidal}(x) \odot q_{c,x} \\ p_k(y, h) &= \frac{h_{ref}}{h} \text{Sinusoidal}(y) \odot q_{c,y} \end{aligned} \quad (5)$$

At the same time, image feature at each coordinate is also assigned with a PE vector according to the Sinusoidal function. As a result, each element in RoI feature f_i has its unique positional representation, denoted by p_f . Note that p_f and p_k have the same dimensions, and they are concatenated with f and k , respectively, completing the position-aware dynamic conv.

3.4. Centerness-based Positional Encoding

The bounding box PE for self attention and dynamic conv described in the previous section encodes (x, y, w, h) in the global image coordinate. Thus, for each RoI feature f , the generated kernels k and v are shared by all elements within it. We intend to further enhance the adaptability of dynamic conv by expanding kernels k to $k_e \in \mathbb{R}^{7 \times 7 \times c \times d}$, therefore, equipping each element in f with unique conv kernels. Based on this motivation, we propose a centerness-based PE, which encodes the local coordinate in bounding

Method	Backbone	Epoch	Params	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
DETR [2]	ResNet-50	500	41M	42.0	62.4	44.2	20.5	45.8	61.1
Deformable DETR [49]	ResNet-50	50	40M	45.9	64.9	49.3	28.2	48.9	61.7
Conditional DETR [28]	ResNet-50	50	44M	40.9	61.8	43.3	20.8	44.6	59.2
DAB-DETR [25]	ResNet-50	50	44M	42.2	63.1	44.7	21.5	45.7	60.3
DN-DETR [20]	ResNet-50	50	44M	44.1	64.4	46.7	22.9	48.0	63.4
RecursiveDet(Deformable DETR)	ResNet-50	50	36M	46.7	65.0	50.6	28.8	49.7	62.2
Sparse R-CNN [35]	ResNet-50	36	106M	45.0	63.4	48.2	26.9	47.2	59.5
AdaMixer [11]	ResNet-50	36	135M	47.0	66.0	51.1	30.1	50.2	61.8
DiffusionDet [4]	ResNet-50	60	111M	45.8	65.6	49.2	29.7	48.6	61.1
RecursiveDet (Sparse R-CNN)	ResNet-50	36	55M	46.5	65.4	50.9	29.7	49.1	59.9
RecursiveDet (AdaMixer)	ResNet-50	36	43M	47.9	66.7	52.2	32.5	50.7	61.9
RecursiveDet (DiffusionDet)	ResNet-50	60	57M	47.1	66.8	51.0	30.9	49.3	62.3
DETR [2]	ResNet-101	500	60M	43.5	63.8	46.4	21.9	48.0	61.8
Deformable DETR‡[49]	ResNet-101	50	59M	47.2	66.6	51.1	28.5	50.9	62.4
Conditional DETR [28]	ResNet-101	50	63M	42.8	63.7	46.0	21.7	46.6	60.9
DAB-DETR [25]	ResNet-101	50	63M	43.5	63.9	46.6	23.6	47.3	61.5
DN-DETR [20]	ResNet-101	50	63M	45.2	65.5	48.3	24.1	49.1	65.1
Sparse R-CNN [35]	ResNet-101	36	125M	46.4	64.6	49.5	28.3	48.3	61.6
AdaMixer [11]	ResNet-101	36	154M	48.0	67.0	52.4	30.0	51.2	63.7
DiffusionDet [4]	ResNet-101	60	130M	46.6	66.3	50.0	0.0	49.3	62.8
RecursiveDet (Sparse R-CNN)	ResNet-101	36	74M	47.1	65.7	51.8	29.3	50.6	61.5
RecursiveDet (AdaMixer)	ResNet-101	36	62M	48.9	67.8	53.1	32.1	52.2	63.8
RecursiveDet (DiffusionDet)	ResNet-101	60	76M	46.9	66.7	50.3	29.5	49.8	62.7
Sparse R-CNN [35]	Swin-T	36	110M	47.9	67.3	52.3	-	-	-
AdaMixer [11]	Swin-S	36	160M	51.3	71.2	55.7	34.2	54.6	67.3
Sparse R-CNN [35]	Swin-B	36	169M	52.0	72.2	57.0	35.8	55.1	68.2
RecursiveDet (Sparse R-CNN)	Swin-T	36	58M	48.9	68.2	53.6	33.1	51.3	63.2
RecursiveDet (AdaMixer)	Swin-S	36	68M	52.1	71.6	56.9	36.2	55.1	68.3
RecursiveDet (Sparse R-CNN)	Swin-B	36	118M	53.1	73.0	58.4	37.0	56.6	69.5

Table 2. Main results and comparisons with other object detectors on COCO 2017 val set. DETR uses 100 object queries, DiffusionDet and RecursiveDet built on it use 500 proposal features, and all other detectors use 300 of that. The results are from the original paper, mmdetection[3] and Detectron2 [14]. "‡" is the enhanced version with iterative box refinement and two-stage processing.

box b for feature f and kernels k . Particularly, a single channel centerness [37] mask $m \in \mathbb{R}^{7 \times 7}$ defined by Eq. (6) is first calculated. Here (x^*, y^*) denotes the local coordinate within the bounding box, and $x^*, y^* = 0, 1, \dots, 6$ have the same value range.

$$m(x^*, y^*) = \sqrt{\frac{\min(x^*, 6 - x^*) \min(y^*, 6 - y^*)}{\max(x^*, 6 - x^*) \max(y^*, 6 - y^*)}} \quad (6)$$

Note that m is between 0 and 1, and a higher value means (x^*, y^*) is near the center of the bounding box. Then, we multiply m onto the RoI feature f and kernel k . Since m has only one channel, it is replicated on channel dimension c to match f . For k , m needs replication on both channel c and hidden dimension d . k also needs expansions on spatial dimension before the multiplication with m , enlarging its size to 7×7 . After the modulation by m , dynamic kernel k is adapted into k_e , giving each element in f a unique kernel.

Our centerness-based PE is also shown in Fig. 2 (c).

The centerness mask m is static with the same size of RoI feature f , which means it has the same value for different f_i and k_i , and it always gives the largest value at the center position. In practice, we try following two strategies to enhance it. First, we make it learnable as model parameters, with its initialization in Eq. (6). Second, we adjust the largest value position by predicting the center coordinate from proposal feature q . However, we find the static m strategy gives better results than the other two. Details are provided in the ablation study in Sec. 4.3.

4. Experiments

4.1. Implementation Details

Datasets. We conduct extensive experiments on MS-COCO 2017 detection dataset [24] with 80 categories in to-

Method	Backbone	TTA	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
AdaMixer [11]	ResNeXt-101-DCN		49.8	69.3	54.3	30.0	52.1	64.3
Dynamic DETR [9]	ResNeXt-101-DCN		49.3	68.4	53.6	30.3	51.6	62.5
Sparse R-CNN	Swin-T		47.0	66.8	51.2	28.3	49.0	60.6
Sparse R-CNN	Swin-B		52.2	72.6	57.2	32.3	54.9	67.2
RecursiveDet	Swin-T		49.1	68.5	53.9	30.5	51.2	61.9
RecursiveDet	Swin-B		53.1	73.9	58.5	33.7	55.8	67.9
RecursiveDet	Swin-B	✓	55.1	75.0	61.8	37.7	57.3	68.5

Table 3. Results on COCO 2017 test-dev set. "TTA" indicates test-time augmentations, following [43].

tal. There are about 118k images in the train2017 set and 5k in val2017. We report the standard MS COCO AP as the main evaluation metric.

Training settings. We adopt the same data augmentation as Sparse R-CNN, including random horizontal flipping, random crop and scale jitter of resizing the input images such that the shortest side is at least 480 and at most 800 pixels while the longest is at most 1333. The training loss is the same as matching cost with loss weight $\lambda_{cls} = 2$, $\lambda_{L_1} = 5$ and $\lambda_{giou} = 2$. The learning rate is divided by 10 at 210K and 250K iterations if the training schedule is 270k iterations, or at 350K and 420K if the training schedule is 450K. The default number of decoding stages is 6.

4.2. Main Results

The proposed RecursiveDet is built on three region-based detectors, including Sparse R-CNN, AdaMixer and DiffusionDet. We provide the performance comparisons of RecursiveDet with well-established methods on COCO 2017 validation set in Tab. 2. Note that, the sampling points in 3D feature space are out of the proposal boxes in AdaMixer, so the centerness-based PE is not implemented into it. It can be seen that the three region-based detectors exceed Cascade R-CNN and the other DETR-series detector, but require more model parameters. Owing to the recursive structure, our method saves a massive amount of parameters. With ResNet-50 backbone, RecursiveDet built on Sparse R-CNN outperforms Sparse R-CNN by 1.5 AP, while parameters are reduced from 106M to 55M. The models built on DiffusionDet and AdaMixer beat their baseline 1.6 AP (47.1 vs. 45.5) and 0.9 AP (47.0 vs. 47.9). RecursiveDet also behaves well when backbone scales up. With ResNet-101, RecursiveDet achieves 47.1, 46.9 and 48.9 based on Sparse R-CNN, DiffusionDet and Adamixer. More surprisingly, RecursiveDet(AdaMixer) and RecursiveDet(Sparse R-CNN) reach 52.1 and 53.1 with the help of Swin-S and Swin-B, respectively. Although RecursiveDet(DiffusionDet) doesn't exceed its baseline, it reduces more than 30% of the model size. Note that, all our models save a large number of parameters. The results of our method on COCO test-dev set are listed in

Tab. 3. RecursiveDet reaches 53.1 AP on Swin-B without bells and whistles. When test-time augmentation is utilized, it achieves 55.1 AP. We also implement our method into Deformable DETR, and make an improvement of 0.8 AP.

4.3. Ablation Study

In this section, we use ResNet-50 to perform ablation studies based on Sparse R-CNN architecture as default. The number of proposal features is 100.

Recursive Structure. Since Sparse R-CNN has multiple cascade decoder stages to refine the predictions progressively, the parameters of it are redundant. We propose to share them in different stages. As shown in Tab. 4, simply sharing the parameters and making the structure recursive reduces the model size from 106M to 40.1M, with a slight performance drop of 0.4 AP. To take full advantage of the decoder's capabilities, we reuse the dynamic layer and out layer to increase the depth of model without adding new parameters. It gains 43.6 AP, surpassing the original Sparse R-CNN 0.8 AP. As the first stage is difficult to locate objects, we keep the initial stage independent with the other recursive ones. Note that the total number of stages remains 6. The performance reaches 43.9 AP, gains 0.3 AP from model 'RecSt + RecDy' with a small cost and negligible latency.

RecSt	RecSt [†]	RecDy	AP	Params(M)	Flops(G)	L(ms)	FPS
			42.8	106	134	49	20
✓			42.4	40	134	48	21
✓		✓	43.6	40	142	50	20
	✓	✓	43.9	53	140	51	20

Table 4. Ablation about recursive structure. 'RecSt' indicates that the parameters of different stages are shared, 'RecDy' means dynamic layer and out layer are reused. '†' implies the first stage's parameters are independent of the remaining stages.

Centerness form in PE. As described in Sec. 3.4, the local coordinate PE of centerness makes model explore location information more precisely. In addition to the static centerness computed within 7×7 region, it can be trained as learnable parameters of the model. The result in Tab. 5

shows a performance drop of about 0.3 AP. As different objects have various shapes and focus on different points, we predict a group of offsets to distinguish where the object center is. Let the value of centerness at this position be the maximum, the further away, the smaller. However, the result is unsatisfactory with 44.5 AP, which is 0.7 less than the static one.

centerness	AP	AP ₅₀	AP ₇₅
static	45.2	63.9	49.4
learnable	44.9	63.5	49.0
adjust	44.5	63.3	48.4

Table 5. Variations of centerness-based PE. 'static' means the centerness is the same for different proposal boxes within the 7×7 region. 'learnable' makes the static centerness be learnable parameters. 'adjust' indicates a set of offsets is generated from ROI feature to adjust the position of the maximum value.

Influences of different modules in RecursiveDet. In this part, we analyze the effectiveness of different components in our model. As shown in Tab. 6, all modules proposed in this paper substantially contribute to the final results. The recursive structure improves the Sparse R-CNN from 42.8 AP to 43.9 AP. PE has not been introduced into region-based detectors before. We encode bounding boxes into embedding to participate in self attention and dynamic conv. It achieves 0.3 AP performance gain. When different stages are shared, it is challenging for the decoder input to discriminate which stage is. Bounding box encoding plays its role by providing geometry information for decoder to distinguish the stage, and it gains 44.4 AP, exceeding the model of recursive structure by 0.5 AP. And the time cost mainly comes from it. Since the bounding box only has the global coordinates and shape sizes, we further employ centerness within the proposal box to introduce a local prior. With the help of centerness, the result finally reaches 45.2 AP.

Recursive	bb-PE	cb-PE	AP	Params(M)	Flops(G)	L(ms)	FPS
			42.8	106	134	49	20
✓			43.9	53	140	51	20
	✓		43.1	110	137	69	15
	✓	✓	43.4	110	138	72	14
✓	✓		44.4	55	145	73	14
✓	✓	✓	45.2	55	149	76	13

Table 6. Ablations about different components (Recursive structure, bounding-box PE(bb-PE) and centerness-based PE(Cb-PE)) in the proposed RecursiveDet.

Visualization. In Fig. 3 we present the convergence speed of Sparse R-CNN, AdaMixer, DiffusionDet and RecursiveDet. It shows that our models are more efficient than their baselines. Since our method shares the decoder parameters, making the cascade detector recursive, we compare the effectiveness of the number of decoder stages at in-

ference between Sparse R-CNN and RecursiveDet in Fig. 4. It can be seen that our method outperforms Sparse R-CNN with any number of stages.




Figure 3. Convergence curves of Sparse R-CNN, AdaMixer, DiffusionDet and their counterpart in RecursiveDet. All models are trained with ResNet-50. Number of proposal feature is 300 for Sparse R-CNN and AdaMixer, and 500 for DiffusionDet.




Figure 4. Effect of the number of stages in cascade and recursive structures. Both models are trained with six decoder stages.

5. Conclusion

This paper investigates the region-based object detector. We propose a RecursiveDet, which increases the detection performance and reduces the model size. There are two types of recursion loops in our detector. We first share the decoder parameters and make it recursive among different stages. We also reuse the dynamic layer and out layer in decoder, and make a short in-stage recursion loop to increase the depth of model. To enhance the adaptability of the decoder, we design bounding box and centerness-based positional encoding, and further utilize them in self attention and dynamic conv. The RecursiveDet is implemented under three typical region-based detectors, including Sparse R-CNN, AdaMixer and DiffusionDet. It achieves consistent improvements on all of them with a lower cost on model parameters.

References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020.
- [3] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [4] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022.
- [5] Zhe Chen, Jing Zhang, and Dacheng Tao. Recurrent glimpse-based decoder for detection with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5260–5269, 2022.
- [6] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [7] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. *Advances in neural information processing systems*, 29, 2016.
- [8] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017.
- [9] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2988–2997, 2021.
- [10] Peng Gao, Minghang Zheng, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fast convergence of detr with spatially modulated co-attention. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3621–3630, 2021.
- [11] Ziteng Gao, Limin Wang, Bing Han, and Sheng Guo. Adamixer: A fast-converging query-based object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5364–5373, 2022.
- [12] Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 303–312, 2021.
- [13] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [14] Ross Girshick, Ilija Radosavovic, Georgia Gkioxari, Piotr Dollár, and Kaiming He. Detectron. <https://github.com/facebookresearch/detectron>, 2018.
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [16] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [17] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1637–1645, 2016.
- [18] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 355–371. Springer, 2020.
- [19] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018.
- [20] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022.
- [21] Ming Liang and Xiaolin Hu. Recurrent convolutional neural network for object recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3367–3375, 2015.
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [25] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022.
- [26] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016.
- [27] Yudong Liu, Yongtao Wang, Siwei Wang, TingTing Liang, Qijie Zhao, Zhi Tang, and Haibin Ling. Cbnet: A novel composite backbone network architecture for object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11653–11660, 2020.
- [28] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang.

- Conditional detr for fast training convergence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3651–3660, 2021.
- [29] Siyuan Qiao, Liang-Chieh Chen, and Alan Yuille. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10213–10224, 2021.
- [30] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [31] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [33] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019.
- [34] Zhiqiang Shen, Zechun Liu, and Eric Xing. Sliced recursive transformer. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 727–744. Springer, 2022.
- [35] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021.
- [36] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3147–3155, 2017.
- [37] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019.
- [38] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in neural information processing systems*, 34:24261–24272, 2021.
- [39] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 2567–2575, 2022.
- [40] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9657–9666, 2019.
- [41] Zhuyu Yao, Jiangbo Ai, Boxun Li, and Chi Zhang. Efficient detr: improving end-to-end object detector with dense prior. *arXiv preprint arXiv:2104.01318*, 2021.
- [42] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel Ni, and Harry Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In *International Conference on Learning Representations*, 2022.
- [43] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. 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*, pages 9759–9768, 2020.
- [44] Jing Zhao, Shengjian Wu, Li Sun, and Qingli Li. Iou-enhanced attention for end-to-end task specific object detection. In *Proceedings of the Asian Conference on Computer Vision*, pages 2901–2917, 2022.
- [45] Anlin Zheng, Yuang Zhang, Xiangyu Zhang, Xiaojuan Qi, and Jian Sun. Progressive end-to-end object detection in crowded scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 857–866, 2022.
- [46] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.
- [47] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 840–849, 2019.
- [48] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9308–9316, 2019.
- [49] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.