

Progressive End-to-End Object Detection in Crowded Scenes

Anlin Zheng^{1,3*}, Yuang Zhang^{2*}, Xiangyu Zhang¹, Xiaojuan Qi³, Jian Sun¹

¹MEGVII Technology ²Shanghai Jiao Tong University ³University of Hong Kong

zyayoung@sjtu.edu.cn, xjq@eee.hku.hk, {zhenganlin, zhangxiangyu, sunjian}@megvii.com

Abstract

In this paper, we propose a new query-based detection framework for crowd detection. Previous query-based detectors suffer from two drawbacks: first, multiple predictions will be inferred for a single object, typically in crowded scenes; second, the performance saturates as the depth of the decoding stage increases. Benefiting from the nature of the one-to-one label assignment rule, we propose a progressive predicting method to address the above issues. Specifically, we first select accepted queries prone to generate true positive predictions, then refine the rest noisy queries according to the previously accepted predictions. Experiments show that our method can significantly boost the performance of query-based detectors in crowded scenes. Equipped with our approach, Sparse RCNN achieves 92.0% AP, 41.4% MR^{-2} and 83.2% JI on the challenging CrowdHuman [35] dataset, outperforming the box-based method MIP [8] that specifies in handling crowded scenarios. Moreover, the proposed method, robust to crowdedness, can still obtain consistent improvements on moderately and slightly crowded datasets like CityPersons [47] and COCO [26]. Code will be made publicly available at <https://github.com/megvii-model/Iter-E2EDET>.

1. Introduction

Crowded object detection is a practical yet challenging research field in computer vision. Many research efforts have been made and achieved impressive progress [28, 5, 6, 46, 49, 8, 33, 14, 23] in the last few decades. However, most of them [28, 5, 6, 46, 49, 8, 33, 14] require hand-craft components, e.g. anchor settings and post-processing, resulted in sub-optimal performance in handling scenes.

Recently, Carion et al. [3] proposed a fully end-to-end object detection framework DETR, which introduces learnable queries to represent objects and achieves competitive performance without any post-processing. It can be categorized as a *query-based* approach to differentiate

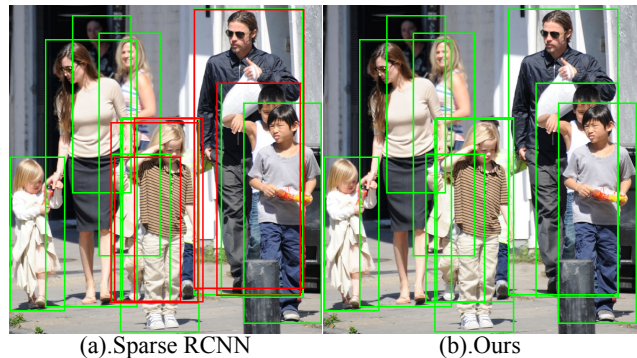


Figure 1: **1a.** Sparse RCNN [37] introduces *false positives* in crowded scenes. **1b.** Our approach can remove those *false positives* and ensure each object can be detected only once. Green boxes indicate *true positives* while red ones represent *false positives*.

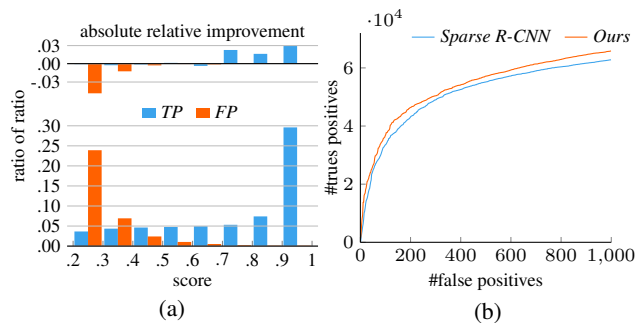


Figure 2: **2a.** The bottom histogram describes the prediction distribution of Sparse RCNN [37] under different confidence scores, while the top one reflects the absolute improvements achieved by our approach compared with Sparse RCNN [37]. **2b.** The FP-TP curve when computing Average Precision (AP).

from the *box-based* [25, 24, 48] and *point-based* [39, 44] methods. Following DETR [3], Sparse RCNN [37] ensures object queries interact with local feature of Region of Interest (RoI), while deformable DETR [52] proposes attention modules that only attend to a small set of key sampling points. They further improve the detection accuracy and mitigate several issues occurred in DETR: slow conver-

* Equal contribution.

gence and high computational overhead.

The above success inspires us to study *query-based* object detection methods in crowded scenes, aiming at designing a more sophisticated end-to-end detection framework. Although these *query-based* approaches [18, 52] can obtain significant results on the slightly crowded datasets like COCO [26], our initial studies show they suffer from several unresolved challenges in crowded scenes: (1). the *query-based* detector tends to infer multiple predictions for a single object, with *false positives* introduced. Figure. 1a shows a common failure case; (2). The performance of a *query-based* detector becomes saturated or even worse as the depth of decoding stage increases, which is depicted in the Appendix.

Our motivations. Further investigations on the *query-based* method, Sparse RCNN [37], yield the following intriguing findings in crowd scenes. As described in Figure. 2a, a large percentage of target objects can be accurately predicted by those predictions with high confidence scores (e.g. higher than a threshold of 0.7), while containing few false positives. These predictions are more likely to be *true positives* that can be taken as *accepted predictions*. While the rest, where a considerable number of *true positives* and *false positives* exist, can be regarded as *noisy predictions*. Naturally, if an object is detected by one accepted prediction, there is no need for noisy predictions to detect it again. Hence, *Why not strengthen the discrimination of those noisy predictions given the context of the accepted predictions?* To this end, the noisy queries can ‘perceive’ whether their targets have been detected or not. If so, their confidence scores will be reduced and then filtered out.

Our contributions Motivated by this, we propose a progressive prediction method equipped with a prediction selector, relation information extractor, query updater, and label assignment to improve the performance of query-based object detectors in handling crowded scenes.

First, we develop a prediction selector to select queries associated with high confidence scores as *accepted queries*, leaving the rest as *noisy queries*. Then, to let the *noisy queries* ‘perceive’ whether their targets have been detected or not, we design a relation extractor for relation modeling between *noisy queries* and their accepted neighbors. Further, a query updater is developed by performing a new local self-attention attending to spatially-related neighbors only. Finally, a new one-to-one label assignment rule is introduced to assign samples among the accepted and refined noisy queries step by step. With the proposed method, the above problems can be well addressed: (1). Each object can be detected only once, which greatly decreases the number of false positives while increasing the number of true positives, as described in Figure. 1b; (2). As depicted in Figure. 2b, the performance is consistently improved compared with its counterparts [37, 52] that have the same depth of

decoding stage.

Our method is generic and can be incorporated into multiple architectures [37, 52], and delivers significant performance improvements of query-based detectors. Equipped with our approach, Sparse RCNN [37] with *ResNet-50* [16] backbone obtains **92.0%** AP, **41.4%** MR⁻² and **83.2%** JI on the challenging dataset *CrowdHuman* [35], outperforming the box-based method MIP [8]. Besides, deformable DETR [52], equipped with our approach, also achieves **92.1%** AP and **84.0%** JI. Moreover, our approach works reasonably well for less crowded scenes, e.g. the Sparse RCNN with our approach can still obtain **1.0%** MR⁻² and **1.1%** AP gains on moderately and slightly crowded datasets *Citypersons* [47] and *COCO* [26], respectively.

2. Related works

End-to-end object detection. RelationNet [20] is one of the pioneering works trying to predict results directly, achieving promising performance compared to their counterparts on several famous benchmarks. DETR [3] introduces learnable queries to represent objects and perform single prediction for each instance directly. Subsequently, deformable-DETR [52] limits the attention field of each query to a local area around the reference points to accelerate the convergence and improve detection performance. Meanwhile, Sparse R-CNN [37] utilizes a fixed set of learnable queries to formulate objects instead of a number of proxy representation, e.g. anchors. Analogous to deformable DETR, RoIAlign [15] is applied to limit the attention field in a local region. Adaptive Clustering Transformer [51] proposes to improve the attention distribution in DETR’s encoder by LSH approximate clustering for convergence acceleration. UP-DETR [9] designs a new self-supervised method to improve the convergence speed of DETR, while TSP[38] analyzes the main factors contributing to slow convergence in DETR. SMCA [13] explores a better information interaction mechanism to further accelerate convergence and improve the performance of DETR.

Object detection in crowded scenes. Research community has poured much interest in exploring occlusion problems on pedestrian detection. Specific methods have been proposed to mitigate this problem, including detecting by parts [28, 5, 6, 46, 49] and improving hand-crafted rules in training target design. Recently, CNN-based methods have dominated the crowded object detection and achieved considerable gains. Several works propose new loss functions to address problems of crowded detection [45, 49]. Besides, the effectiveness of NMS is based on the assumption that multiple instances rarely occur at the same location in an image, which is not true in crowded scenes. But designing duplicate removal for crowded scenes is non-trivial. Soft-NMS [1] and Softer-NMS [17] replace hard removal of nearby proposals with score decay. Several works propose to use a neural network to simulate the function of NMS for

duplicates removal [18, 32]. Others explore NMS-aware training, including NMS with adaptive threshold [19, 27], feature embedding [34] and multiple prediction with set suppression [8, 21], to tackle problem of object detection in crowded scenes.

Recently, PEDR [23] proposes several techniques to improve the performance of *query-based* detectors in coping with crowded detection, which is orthogonal to ours. Their techniques are also applicable to our work.

Relation modeling for object detection. As discussed in [20], early works [10, 12, 41, 42, 30] use object relations as a post-processing step. The detected objects are re-scored by considering object relationships. For example, co-occurrence, which indicates how likely two object classes can exist in the same image, is used by DPM [11] to refine object scores. The subsequent approaches [31, 7] try more complex relation models, by taking additional positions and size into account. These methods achieve moderate success in the pre-deep learning era but do not prove the effectiveness in CNNs. Several recent works perform spatial reasoning [22, 36, 4, 18] to model object relations. Among them, GossipNet [18] and RelationNet [20] are the representative methods. Both share the same spirit of modeling relations among boxes. However, the network of GossipNet [18] is complex (depth>80) and its computation cost is demanding. Although it allows end-to-end learning in principle, no experimental evidence approves. RelationNet [20] utilized the self-attention for feature interaction and obtained a promising improvement in general object detection. Nevertheless, it doesn't show a promising performance in dealing with crowd scenes [35].

Recent works related to ours are PS-RCNN [14] and Iter-Det [33]. They proposed to detect objects according to the previous predicted ones. They need to mask the feature [14] or produce a history map [33] to memorize the previous detections, introducing noise while limiting performance improvement [14] or incur heavy computation [33]. Even so, both of them need a post-processing method to remove duplicates in every iteration.

Recent query-based object detectors [37, 52, 51, 9, 38, 13] utilized learnable queries to represent objects, and take advantage of the self/cross-attention to model the relations among queries, detecting objects in an end-to-end manner. Our work inherits the methodology and boosts their performance in heavily, moderately, and slightly crowded scenes.

3. Methodology

In this section, we first revisit the query-based object detector, e.g. Sparse RCNN [37] briefly. Next, we illustrate our approach primarily deployed on Sparse RCNN explicitly. Finally, the main differences of detector design will be discussed as follows.

3.1. Query Based Object Detector

Our approach can be deployed on most *query-based* object detectors [3, 37, 52]. To illustrate the proposed method, we choose Sparse RCNN [37] as our default instantiation. Figure. 4a depicts its object detection pipeline, which can also be formulated as:

$$\begin{aligned} x_{t-1} &\leftarrow \mathcal{P}^{box}(x^{FPN}, b_{t-1}), \\ q_{t-1}^* &\leftarrow \text{MSA}_{t-1}(q_{t-1}), \\ q_t &\leftarrow \text{DynConv}_{t-1}(q_{t-1}^*, x_{t-1}), \\ b_t &\leftarrow \mathcal{B}_{t-1}(q_t), \end{aligned} \quad (1)$$

where $q \in R^{N \times d}$ denotes the learnable object query. N and d denote the number and dimension of query q , respectively. At stage t , an RoIAlign [15] \mathcal{P}^{box} extracts RoI features from FPN features x^{FPN} , under the guidance of bounding box b_{t-1} predicted by the previous stage. Meanwhile, a multi-head self-attention module MSA_{t-1} is applied to the input query q_{t-1} to get the transformed query q_{t-1}^* . Then, a dynamic convolution module DynConv_{t-1} takes both x_{t-1} and q_{t-1}^* as inputs and performs dynamic convolution to generates q_t for the next stage. Simultaneously, q_t is fed into the box prediction branch \mathcal{B}_{t-1} for current bounding box prediction b_t , which is the input of the next stage t .

3.2. Our Method

As illustrated in Figure. 3, the proposed progressive predicting method consists of several components: prediction selector, relation information extractor, query updater, and label assignment, which will be introduced in detail next.

Prediction selector. For the findings described in Sec. 1, a prediction selector is developed to select those queries prone to generating predictions with high confidence scores as accepted queries, while leaving the rest as noisy ones that need to be further refined. This procedure can be formulated in Equ.(2).

$$\begin{aligned} \mathcal{D}_{t-1}^h &\leftarrow \{b_i | s_i \geq s \wedge b_i \in \mathcal{D}_{t-1}\}, \\ \mathcal{D}_{t-1}^l &\leftarrow \mathcal{D}_{t-1} - \mathcal{D}_{t-1}^h. \end{aligned} \quad (2)$$

where t is the stage number. \mathcal{D}_{t-1} denote the whole predictions produced by the whole queries in the previous $t - 1$ stage. \mathcal{D}_{t-1}^h and \mathcal{D}_{t-1}^l indicate the accepted predictions and noisy predictions generated from the accepted and noisy queries, respectively. b_i and s_i denote the predicted box and its confidence score, respectively. s is the confidence score threshold.

Relation information extractor. As mentioned in Sec. 1, a large percentage of target objects can be accurately predicted by the accepted queries. Therefore, if an object is detected by one accepted prediction, there is no need for

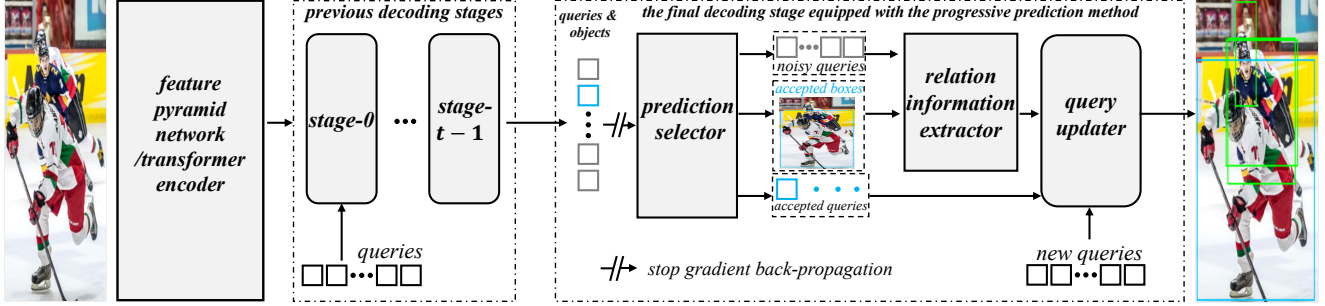
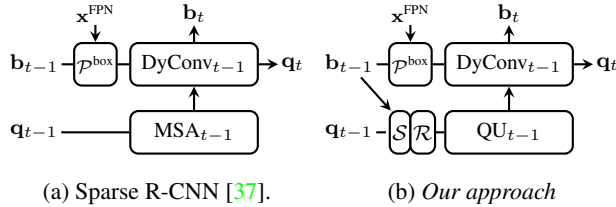


Figure 3: The diagram of the proposed progressive end-to-end object detection framework. First, the *prediction selector* select queries associated with high confidence scores as *accepted queries*, leaving the rest as *noisy queries*. Then *Relation information extractor* models the relations between *noisy queries* and their neighbors from accepted predictions. Next, the queries are fed into the *queries updater* to be further refined by performing a new local self-attention.



(a) Sparse R-CNN [37].

(b) Our approach

Figure 4: The diagram of decoding stage. \mathcal{P} – RoIAlignPool [15], DynConv – Dynamic Convolution, MSA – Multi-head Self-Attention, \mathcal{S} – Prediction Selector, \mathcal{R} – Relation Information Extractor, QU – Query Updater.

noisy predictions to detect it again. In order to equip these noisy queries with the capability of perceiving whether their targets have been detected or not, we develop a relation information extractor to model the relation between the noisy predictions and their accepted neighbors.

The detailed design of the relation information extractor is illustrated in Figure 5, with the procedure formulated in Equ.(3) as well. For each noisy prediction b_i , we first find their accepted neighbors $\mathcal{N}(b_i)$ in \mathcal{D}_{t-1}^h , constructing the spatially-related pairs $(b_i, \mathcal{N}(b_i))$. Then, the encoded pairs together with the *intersection-over-union* (*IoU*) between them are fed to a compact network to obtain the geometry relation features $\mathcal{H}(b_i)$. Since the number of accepted neighbors corresponding to each noisy prediction is uncertain. An aggregation function is employed to reduce $\mathcal{H}(b_i)$ to the same feature dimension, while maintaining the permutation-invariance property. In our approach, we use *max* pooling by default. Besides, the pooled geometry features, fused with the transformed query features, are further activated by a non-linear function.

$$\begin{aligned} \mathcal{N}(b_i) &\leftarrow \{b_j | \mathcal{O}(b_i, b_j) \geq \theta\}, b_i \in \mathcal{D}_{t-1}^l, b_j \in \mathcal{D}_{t-1}^h, \\ \mathcal{H}(b_i) &\leftarrow \mathcal{U}(\mathcal{E}(b_i, \mathcal{N}(b_i))), \quad b_i \in \mathcal{D}_{t-1}^l, \\ \mathcal{R}(b_i) &\leftarrow \mathcal{T}(\text{MaxPool}(\mathcal{H}(b_i)) + \mathcal{F}(q_i)). \end{aligned} \quad (3)$$

where $\mathcal{N}(\cdot)$ represents a function that finds neighbors for a box b_i based on the *IoU* $\mathcal{O}(\cdot, \cdot)$ with a threshold θ . Here, we

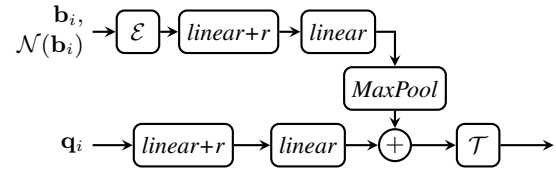


Figure 5: Relation information extractor \mathcal{R} . \mathcal{E} – sine and cosine spatial positional encoding function [20, 43], *linear* – *fc* layer, *r* – ReLU, \mathcal{T} – *fc* layer.

use it to find the accepted neighbors in \mathcal{D}_{t-1}^h for the noisy predictions in \mathcal{D}_{t-1}^l . $\mathcal{E}(\cdot, \cdot)$ refers to the sine and cosine spatial positional encoding function which is the same as that [20, 43]. Also, $\mathcal{U}(\cdot, \cdot)$ denotes a function used to generate geometry relation features $\mathcal{H}(b_i)$ from the encoded inputs. The noisy query q_i corresponds to noisy prediction b_i in \mathcal{D}_{t-1}^l , transformed by the function $\mathcal{F}(\cdot)$. The pooled geometry features and transformed query features $\mathcal{F}(q_i)$ are fused through element-wise summation, followed by a function \mathcal{T} to produce the desired relation features $\mathcal{R}(b_i)$.

As depicted in Figure 5, $\mathcal{U}(\cdot, \cdot)$ consists of two consecutive *fc* layers with *ReLU* [29] activation to increase the non-linearity. Note that $\mathcal{F}(\cdot)$ and $\mathcal{U}(\cdot)$ share the same architecture, but are weight-independent. Here, the gradients of q_i are stopped from back-propagating to the previous stages.

Query updater. To further refine the features of noisy queries, a query updater is developed, which is formulated in Equ.(4). Since the data distribution of \mathcal{D}_{t-1}^l and \mathcal{D}_{t-1}^h is different from that of \mathcal{D}_{t-1} , a new set of learnable queries is first introduced to complement the relation features through element-wise summation. Then the set of complemented noisy queries is taken as the input query q_{t-1} to perform a new local self-attention LMSA_{t-1} and the subsequent dynamic convolution given in Equ.(1).

$$\begin{aligned} q_{t-1} &\leftarrow \{\tilde{q}_i | \tilde{q}_i = \mathcal{R}(b_i) + e_i\}, b_i \in \mathcal{D}_{t-1}^l, e_i \in E \\ q_{t-1}^* &\leftarrow \text{LMSA}_{t-1}(q_{t-1}). \end{aligned} \quad (4)$$

Algorithm 1 Label Assignment for \mathcal{D}_t^l .

Input: $\mathcal{D}_t^l, \mathcal{D}_t^h, \mathcal{G}$;

- 1: \mathcal{D}_t^l : results of \mathcal{D}_{t-1}^l in Equ.(2) from stage t ;
- 2: \mathcal{D}_t^h : results of \mathcal{D}_{t-1}^h in Equ.(2) from stage t ;
- 3: \mathcal{G} : target boxes.

Output: The matched predictions \mathcal{M}_D^l and corresponding targets \mathcal{M}_G^l after assignment.

- 4: Compute matching costs \mathcal{C}_t^h between \mathcal{D}_t^h and \mathcal{G} ;
 - 5: $\mathcal{M}_G^h, \mathcal{M}_D^h = \text{HungarianMatch}(\mathcal{D}_t^h, \mathcal{G}, \mathcal{C}_t^h)$;
 - 6: $\mathcal{G}_t^l = \mathcal{G} - \mathcal{M}_G^h$;
 - 7: Compute matching costs \mathcal{C}_t^l between \mathcal{D}_t^l and \mathcal{G}_t^l ;
 - 8: $\mathcal{M}_G^l, \mathcal{M}_D^l = \text{HungarianMatch}(\mathcal{D}_t^l, \mathcal{G}_t^l, \mathcal{C}_t^l)$;
 - 9: **return** $\mathcal{M}_G^l, \mathcal{M}_D^l$;
-

Since object detection mainly focuses on the local region in an image. We design a new local self-attention module LMSA_t to update the noisy query q_{t-1} . It ensures each query only interacts with local neighbors instead of the whole queries over the full image. The local self-attention first finds those neighbors of each query based on the boxes’ IoUs whose values are greater than 0. Then it performs the ‘qkv’ mechanism in the same way as MSA. To this end, we perform self-attention locally instead of globally.

Different from the neighbor finding rule in [23], we adopt the function $\mathcal{N}(\cdot)$ to select neighbors from \mathcal{D}_{t-1} that are spatially related to q_{t-1} in terms of IoU . Note that, the new local self-attention LMSA_{t-1} is used to replace the MSA_{t-1} in Equ.(1) for feature interaction.

Label assignment Since accepted queries tend to generate true positive predictions, while the noisy ones involve a considerable number of true positives and false positives. Towards end-to-end object detection, we introduce a new one-to-one label assignment rule to assign samples step by step. Specifically, we first match the accepted predictions \mathcal{D}_{t-1}^h with the ground truth set of objects. Then remove those targets that have been matched, and mainly consider the bipartite matching between noisy predictions \mathcal{D}_{t-1}^l and the remaining ground truth set of objects. This matching process is described in Algorithm 1¹, where the matching cost computation is slightly different from the original version [37]. A spatial prior is adopted to compute the matching cost \mathcal{C} , that is, the center of bounding box b_t needs to fall in the corresponding target box. Except for it, the formulation of the matching cost function is identical to the original work.

3.3. Difference of Detector Design

Generally, our approach can be deployed on most query-based object detectors [3, 37, 52]. To illustrate the proposed

¹The HungarianMatch operation in Algorithm 1 is a combinatorial optimization approach that solves the assignment problem, which is commonly used in [3, 37, 52, 44] for one-to-one label assignment.

method, we choose Sparse RCNN [37] as our default instantiation. It consists t ($t = 6$ by default) decoding stages, each of which performs prediction according to Equ. (1). As described in Figure. 3, we keep the first $t - 1$ decoding stages unchanged and only equip the last stage with the proposed method. Therefore, main differences lie in the last decoding stage, which will be described in the following.

Architecture of the last stage. As depicted in Figure. 4b, the last decoding stage t first employs a prediction selector \mathcal{S} to split queries into accepted queries and noisy queries according to the confidence scores of their associated predictions. Then they are input to the relation information extractor \mathcal{R} to extract the relation feature between the noisy predictions and their accepted neighbors. Finally, queries are fed into the query updater QU to be further refined for recognition and localization.

Box prediction Like [37], a box regression branch is used for box prediction in the first $t - 1$ stages. Differently, for the box prediction in the last stage, we directly use the identity mapping results from the previous $t - 1$ stage both in the training and testing phase. This is because, at the latter layers, the predicted bounding boxes are less likely to fluctuate, which is observed in [23]. Meanwhile, the recognition branch remains the same as that of previous stages.

Training loss We adopt the set prediction loss adopted in [37, 52] for training. For the samples to train stage t , we remove those samples from accepted predictions \mathcal{D}_{t-1}^h . to mitigate the class-imbalance issue, we follow the negative sample filtering mechanism [50] to early reject those well-classified negative queries whose confidence scores are lower than 0.05.

4. Experiments

In this section, we evaluate our approach on heavily, moderately, and slightly crowded datasets [35, 47, 26] to demonstrate the generality of the proposed method in diverse scenarios.

Datasets. We adopt three datasets – *CrowdHuman* [35], *CityPersons* [47] and *COCO* [26] – for comprehensive evaluations on heavily, moderately and slightly crowded situations, respectively. Table 1 lists the “instance density” of each dataset. Since our proposed approach mainly aims to improve crowded detections, we perform most of the comparisons and ablation experiments on *CrowdHuman*. The evaluation experiments on *Citypersons* [47] and *COCO* [26] are also conducted to suggest the proposed method is robust to crowdedness.

Evaluation metrics. Following [8], we mainly take three criteria: AP, MR^{-2} and JI as evaluation metrics. Generally, a larger AP, larger JI and smaller MR^{-2} indicates a better performance.

Implementation details. Unless otherwise specified, we take Sparse RCNN [37] as our default instantiation, using

Dataset	# objects/img	# overlaps/img
CrowdHuman [35]	22.64	2.40
CityPersons [47]	6.47	0.32
COCO* [26]	9.34	0.015

Table 1: *Instance density* of each dataset. The threshold for overlap statistics is $\text{IoU} > 0.5$. *Averaged by the number of classes.

standard *ResNet-50* [16] pre-trained on ImageNet as backbone. We train our model with the Adam optimizer with a momentum of 0.9 and weight decay of 0.0001. Models are trained for 50,000 iterations. The initial learning rate is 0.00005 and reduced by a factor of 0.1 at iteration 37,500. The last stage joints the optimization after 5,000 iterations of training. $\lambda_{cls} = 2, \lambda_{L1} = 5, \lambda_{giou} = 2$. The default number of proposal boxes, proposal features, and stages are set to 500, 500, and 6, respectively. Additionally, The dimension of intermediate features in relationship extractor \mathcal{R} is 256. The gradients are detached at proposal boxes from the second stage to stabilize training. Besides, those negative samples, whose *intersection-over-area* (IoA) between any ignore region is higher than a threshold of 0.7, are not involved in training. Further, the hyper-parameters s and θ are 0.7 and 0.4 by default in different query-based detectors [37, 52].

4.1. Experiments on CrowdHuman

CrowdHuman [35] contains 15,000, 4,370 and 5,000 images for training, validation and test, respectively. For a fair comparison, we re-implement most of the involved models [24, 21, 25, 27, 48, 8, 40, 44, 3, 52, 37, 23, 18, 20]. Results are evaluated on the validation set, using the full-body annotations in the dataset.

Main results. We compare with mainstream object detectors, including *box-based*: *one-stage* [25, 48], *two-stages* [8, 24, 21, 27], and *point-based* [40, 44] as well as *query-based* [3, 52, 37, 23].

As shown in Table. 2, our approach outperforms these well-established detectors, achieving significant performance improvements over the *box-based*, *point-based*, and *query-based* counterparts, illustrating the effectiveness of our approach in handling crowded scenes. Specifically, our method achieves 1.8% AP and 0.9% JI gains over the state-of-the-art *box-based* approach MIP [8], which specializes in coping with crowded scenes.

The *query-based* method Sparse RCNN [37], equipped with the proposed method and 500 queries, can achieve 92.0% AP, 41.4% MR^{-2} and 83.2% JI on the challenging CrowdHuman dataset [35], which is 1.3%, 3.3% and 1.8% better than its counterpart – original Sparse RCNN [37]. When increasing the number of queries to 750, our approach can still obtain a better performance of 92.5% AP

and 83.3% JI. This is because more queries can cover more patterns of objects in the image, such as scale, size, position, and other characteristics. Additionally, equipped with our approach, deformable DETR [52]² can also obtain 2.2% MR^{-2} improvements over the original deformable DETR [52]. Moreover, It also achieves 1.4% AP and 1.6% JI gains over the *box-based* method MIP [8], demonstrating the effectiveness of our approach.

Method	#Queries	AP	MR^{-2}	JI
<i>box-based</i>				
RetinaNet [25]	-	85.3	55.1	73.7
ATSS [48]	-	87.0	51.1	75.9
ATSS [48]+MIP [8]	-	88.7	51.6	77.0
FPN [24]+NMS	-	85.8	42.9	79.8
FPN [24]+soft NMS	-	88.2	42.9	79.8
FPN+MIP [8]	-	90.7	41.4	82.4
FPN [†] +NMS	-	84.9	46.3	-
Adaptive NMS [†] [27]	-	84.7	47.7	-
PBN [†] [21]	-	89.3	43.4	-
<i>point-based</i>				
FCOS [40]	-	86.8	54.0	75.7
FCOS [40]+MIP [8]	-	87.3	51.2	77.3
POTO [44]	-	89.1	47.8	79.3
<i>query-based</i>				
DETR [3]	100	75.9	73.2	74.4
PEDR [23]	1000	91.6	43.7	83.3
D-DETR [52]	1000	91.5	43.7	83.1
S-RCNN [37]	500	90.7	44.7	81.4
S-RCNN [37]	750	91.3	44.8	81.3
S-RCNN+Ours	500	92.0	41.4	83.2
S-RCNN+Ours	750	92.5	41.6	83.3
D-DETR+Ours	1000	92.1	41.5	84.0

Table 2: Comparisons of different methods on *CrowdHuman* validation set, +MIP represents multiple instance prediction with set NMS as post-processing. † indicates the approach is implemented by PBM [21]. S-RCNN – Sparse RCNN [37]. D-DETR – deformable DETR [52].

Ablation study of different modules. To explore the effectiveness of the proposed modules in Sec. 3.2, we conduct extensive ablation study of the relation information extractor \mathcal{R} , local self-attention module LMSA and the newly initialized embedding E . All experiments are conducted on Sparse RCNN [37] with 500 queries, *ResNet-50* [16] backbone and evaluated on *CrowdHuman* dataset. Table. 3a shows that the relation information extractor \mathcal{R} can obtain an improvement of 0.8% AP, 1.7% MR^{-2} and 1.6% JI. It indicates its effectiveness in reducing false positives and recalling false negatives. Moreover, when equipped

²The detail implementation of deformable DETR with the proposed schema is illustrated in the Appendix.

with the new local self-attention LMSA, the performance on three evaluation metrics is further boosted, since the local self-attention can reduce duplicates effectively. Further, the newly initialized embeddings, aiming to approximate the new data distribution of noise predictions, can slightly improve MR^{-2} .

Ablation study of hyper-parameter s . To analyze the effect of the confidence score threshold s , we first formulate the relation between detection boxes and target boxes in an image as a bipartite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. It consists of a set $\mathcal{V} = \mathcal{D} \cup \mathcal{G}$ and nodes \mathcal{E} . \mathcal{D} represents a set of predicted boxes whose scores are higher than the pre-defined score threshold, while \mathcal{G} denotes the target boxes. An edge in \mathcal{E} is defined as overlapping when the *IoU* value, between a box in \mathcal{D} and the other one from \mathcal{G} , is higher than 0.5 by default³. Hence, the matching results can be acquired after applying the Hungarian Algorithm. As shown in Figure. 2, as the confidence score increases, the number of *true positives* shows a clean upward trend while the number of *false positives* decreases rapidly. Also, Figure. 7a depicts the performance our method can achieve under different values of s , where the performance increases slightly as s increases. Thus, if not specific, we set s to 0.7 by default.

Ablation study of hyper-parameter θ . Here, we analyze the effect of the hyper-parameter *intersection-over-union* threshold θ . As discussed in Sec. 1, making sure a box candidate can ‘perceive’ its neighbors helps a noisy query decide to decrease its confidence score or not, which is also the prerequisite for our method to work effectively. Different settings of *intersection-over-union* (*IoU*) threshold θ may affect the performance of the whole detector. We perform experiments on the *CrowdHuman* dataset [35] with s frozen as 0.7 while changing the value of θ linearly. From Figure. 7b, we found our approach is robust to the change of *IoU* threshold. This success may attribute to the good approximating feature of the newly designed components.

Comparison with previous relation modeling works. To differentiate the previous works and ours, we evaluate several representative *relation modeling* methods: RelationNet [20], GossipNet [18], IterDet [33]. RelationNet [20] utilized self-attention modules to model the relations among different predictions. Meanwhile, GossipNet [18] uses several hand-designed relation blocks to explore the relationships among the predicted boxes, while IterDet [33] iteratively infers predictions based on a historical map produced from the previous iteration. We re-implement its re-scoring version for RelationNet [20]. For GossipNet⁴ and IterDet⁵, we use their open-source implementations for evaluation. All models use FPN [24] with *ResNet-50* [16] as backbone, following the same training setting in [18, 20, 24].

³Here, we follow the procedure to compute evaluation metric JI.

⁴GossipNet:<https://github.com/hosang/gossipnet>

⁵IterDet:<https://github.com/saic-vul/iterdet>

As shown in Table. 3b, our approach shows better performance when compared with previous relation modeling works. Surprisingly, both RelationNet [20] and GossipNet [18] suffer from a significant drop in AP and MR^{-2} . It could attribute to the sub-optimal label assignment rule. Since both of them choose the prediction with the highest confidence score around one target as the correct box and take the rest as negatives. The predicted coordinates are not involved in computing loss, which might lead to the performance degradation in crowded scenes.

Analysis on false positives. To understand the factors contributing to the performance improvement, we conduct an error analysis on our method. We adopt the recently proposed TIDE [2] to compare our approach with the counterpart Sparse RCNN [37]. We analyzed the composite error at Recall=0.9 for all methods. As illustrated in Figure. 6, our method performs better at removing duplication, providing more accurate localization, and reducing mistaken recognition. Since part of queries can perceive whether their targets are detected or not through the relation information extractor. Also, the local self-attention module ensures queries only interact with their neighbors rather than the whole. To this end, the duplicates could be eliminated efficiently. Besides, with identity mapping plugged in the last regression branch for box prediction, the number of training samples in the previous decoding stage increases, making the optimization much easier. Additionally, benefiting from the new learnable embeddings for data distribution approximation, the representation ability of object queries are further enhanced.

4.2. Experiments on Citypersons

CityPersons [47] is one of the widely used benchmarks for pedestrian detection. It contains 5, 000 images (2, 975 for training, 500 for validation, and 1, 525 for testing, respectively). Each image has a size of 1024×2048 . To improve the overall performance, we proposed to pre-train all models on the *CrowdHuman* dataset and fine-tune them on *CityPersons* (*reasonable*) training subset, then tested on the (*reasonable*) validation subset. For those *box-based* methods, we train and evaluate them with the image resolution enlarged by $1.3 \times$ compared to the original one for better accuracy. The *query-based* approaches are trained and evaluated at the original image size with 500 queries. The other settings remain the same as those of Sparse RCNN [37] and deformable DETR [52].

4.3. Experiments on COCO.

According to Table 1, the crowdness of *COCO* [26] is very low, which is beyond our design purpose. Nevertheless, we still conduct an experiment on this dataset to verify: 1) whether our method generalizes well to multi-class detection; 2) whether our approach can still handle slightly crowded scenarios, especially with *isolated* instances.

\mathcal{R}	LMSA	E	AP	MR^{-2}	J1
			90.7	44.7	81.4
✓			91.5	43.0	83.0
✓	✓		92.0	42.0	83.5
✓	✓	✓	92.0	41.4	83.2

(a) Ablations of different modules.

Method	#Queries	AP	MR^{-2}	J1
GossipNet [18]	-	80.4	49.4	81.6
RelationNet [20]	-	81.6	48.2	74.6
IterDet [33]	-	88.0	47.5	78.0
D-DETR+Ours	500	91.2	42.6	84.0
S-RCNN+Ours	500	92.0	41.4	83.2

(b) Comparisons of different relation modeling approaches.

Method	#Queries	MR^{-2}	AP
FPN+NMS		9.8	94.7
FPN+Soft-NMS [1]	-	9.9	94.9
MIP [8]		8.8	95.8
D-DETR [52]	500	9.4	96.6
S-RCNN [37]	500	10.0	96.8
D-DETR+Ours	500	7.8	96.7
S-RCNN+Ours	500	7.8	97.6

(c) Performance comparisons on *CityPersons*.

Table 3: **3a.** Ablation study of different modules proposed in our approach, taking Sparse RCNN [37] with 500 queries as our default instantiation. **3b.** Comparisons of different relation modeling approaches. All the experiments are conducted on *CrowdHuman* [35] dataset. **3c** Performance comparisons of different methods on *CityPersons* [47]. Both *box-based* [8, 24] and *query-based* approaches [37, 52] are evaluated.

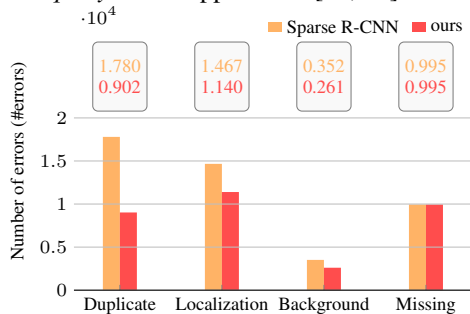


Figure 6: Error analysis on Sparse RCNN [37] and our approach, with *ResNet-50* [16] as backbone. The bar plots show different error types that contribute to the false positives.

Following the common practice of Sparse RCNN [37] with 300 queries, we use a subset of 5000 images in the original validation set (named *minival*) for validation while using the remaining images in the training and validation set for training. Except for the proposed modules and label assignment rule in the last stage, other settings remain the same as the original methods [52, 37]. Table 4 shows the performance comparisons with deformable DETR [52] and Sparse RCNN [37]. Moderate improvements are obtained, e.g. 0.9% AP higher than the deformable DETR [52] and 1.1% AP higher than the Sparse RCNN [37]. The experimental results reflect the effectiveness of our progressive predicting approach in slightly crowded scenarios, proving the proposed method can also solve the performance saturation problem of *query-based* detectors.

5. Conclusion

In this paper, we propose a progressive prediction method to boost the performance of *query-based* object detectors in handling crowded scenes. Equipped with our approach, two representatives *query-based* methods, Sparse RCNN [37] and deformable DETR [52] achieve consistent improvements over the heavily, moderately, as well as slightly crowded datasets [35, 47, 26], which suggests

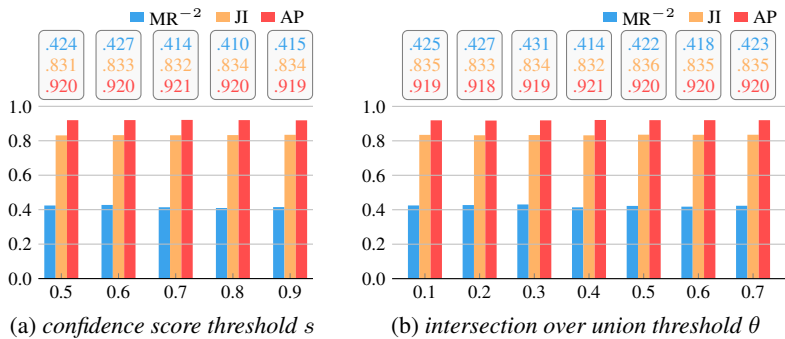


Figure 7: Performance of the proposed method with different configurations of hyper-parameter s and θ on *CrowdHuman* [35] dataset.

Method	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
S-RCNN [37]	45.0	64.2	49.1	27.6	47.5	59.1
D-DETR [52]	45.8	64.5	49.4	28.2	49.0	61.7
S-RCNN+Ours	46.1	65.3	50.6	29.2	48.7	59.9
D-DETR+Ours	46.7	65.3	50.3	28.6	49.8	61.7

Table 4: Performance comparisons of different methods on *COCO* 2017 [26] *minival* set.

our approach is robust to crowdedness. Since Sparse RCNN [37] and deformable DETR [52] require large computing resources, making it difficult for our method to be deployed on devices with limited computing capacity. How to develop a computation-efficacy end-to-end detector is still under exploration. Besides, we found the decision boundary for the noisy queries is unclear. We believe that the performance can be further improved if a better feature engineering method or loss function is adopted. However, it is beyond the purpose of this work.

References

- [1] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S. Davis. Soft-nms – improving object detection with one line of code. *arXiv preprint arXiv:1704.04503*, 2017.
- [2] Daniel Bolya, Sean Foley, James Hays, and Judy Hoffman. Tide: A general toolbox for identifying object detection errors. In *ECCV*, 2020.
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229, 2020.
- [4] Xinlei Chen and Abhinav Gupta. Spatial memory for context reasoning in object detection. *CoRR*, abs/1704.04224, 2017.
- [5] Cheng Chi, Shifeng Zhang, Junliang Xing, Zhen Lei, Stan Li, and Xudong Zou. Pedhunter: Occlusion robust pedestrian detector in crowded scenes. In *AAAI 2020 : The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [6] Cheng Chi, Shifeng Zhang, Junliang Xing, Zhen Lei, Stan Li, and Xudong Zou. Relational learning for joint head and human detection. In *AAAI 2020 : The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- [7] Myung Jin Choi, Antonio Torralba, and Alan S. Willsky. A tree-based context model for object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(2):240–252, 2012.
- [8] Xuangeng Chu, Anlin Zheng, Xiangyu Zhang, and Jian Sun. Detection in crowded scenes: One proposal, multiple predictions. pages 12214–12223, 2020.
- [9] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1601–1610, June 2021.
- [10] Santosh K. Divvala, Derek Hoiem, James H. Hays, Alexei A. Efros, and Martial Hebert. An empirical study of context in object detection. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1271–1278, 2009.
- [11] Pedro F. Felzenszwalb, Ross B. Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010.
- [12] Carolina Galleguillos, Andrew Rabinovich, and Serge Belongie. Object categorization using co-occurrence, location and appearance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [13] Peng Gao, Minghang Zheng, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fast convergence of DETR with spatially modulated co-attention. *CoRR*, abs/2101.07448, 2021.
- [14] Zheng Ge, Zequn Jie, Xin Huang, Rong Xu, and Osamu Yoshie. PS-RCNN: detecting secondary human instances in a crowd via primary object suppression. *CoRR*, abs/2003.07080, 2020.
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. *computer vision and pattern recognition*, 2017.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [17] Yihui He, Xiangyu Zhang, Marios Savvides, and Kris Kitani. Softer-nms: Rethinking bounding box regression for accurate object detection. 2018.
- [18] J. Hosang, R. Benenson, and B. Schiele. Learning non-maximum suppression. In *CVPR*, 2017.
- [19] Jan Hendrik Hosang, Rodrigo Benenson, and Bernt Schiele. A convnet for non-maximum suppression. In *38th German Conference on Pattern Recognition*, pages 192–204, 2016.
- [20] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018.
- [21] Xin Huang, Zheng Ge, Zequn Jie, and Osamu Yoshie. Nms by representative region: Towards crowded pedestrian detection by proposal pairing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10750–10759, 2020.
- [22] Jianan Li, Yunchao Wei, Xiaodan Liang, Jian Dong, Tingfa Xu, Jiashi Feng, and Shuicheng Yan. Attentive contexts for object detection. *IEEE Transactions on Multimedia*, 19(5):944–954, 2017.
- [23] Matthieu Lin, Chuming Li, Xingyuan Bu, Ming Sun, Chen Lin, Junjie Yan, Wanli Ouyang, and Zhidong Deng. Detr for crowd pedestrian detection. *arXiv preprint arXiv:2012.06785*, 2020.
- [24] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, 2017.
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):318–327, 2020.
- [26] Tsung-Yi Lin, Michael Maire, Serge J. 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, 2014.
- [27] Songtao Liu, Di Huang, and Yunhong Wang. Adaptive nms: Refining pedestrian detection in a crowd. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6459–6468, 2019.
- [28] Ruiqi Lu and Huimin Ma. Semantic head enhanced pedestrian detection in a crowd. *arXiv preprint arXiv:1911.11985*, 2019.
- [29] Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Citeseer, 2013.
- [30] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014.

- [31] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014.
- [32] Lu Qi, Shu Liu, Jianping Shi, and Jiaya Jia. Sequential context encoding for duplicate removal. In *Advances in Neural Information Processing Systems*, volume 31, pages 2049–2058, 2018.
- [33] Danila Rukhovich, Konstantin Sofiuk, Danil Galeev, Olga Barinova, and Anton Konushin. Iterdet: Iterative scheme for object detection in crowded environments. *CoRR*, abs/2005.05708, 2020.
- [34] Niels Ole Salscheider. FeatureNMS: Non-maximum suppression by learning feature embeddings. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7848–7854, 2021.
- [35] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018.
- [36] Russell Stewart, Mykhaylo Andriluka, and Andrew Y. Ng. End-to-end people detection in crowded scenes. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2325–2333, 2016.
- [37] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, and Ping Luo. Sparse r-cnn: End-to-end object detection with learnable proposals. *arXiv preprint arXiv:2011.12450*, 2020.
- [38] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris Kitani. Rethinking transformer-based set prediction for object detection. *CoRR*, abs/2011.10881, 2020.
- [39] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9626–9635, 2019.
- [40] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: A simple and strong anchor-free object detector. 2021.
- [41] Torralba, Murphy, Freeman, and Rubin. Context-based vision system for place and object recognition. In *Proceedings Ninth IEEE International Conference on Computer Vision*, pages 273–280 vol.1, 2003.
- [42] Zhuowen Tu. Auto-context and its application to high-level vision tasks. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [44] Jianfeng Wang, Lin Song, Zeming Li, Hongbin Sun, Jian Sun, and Nanning Zheng. End-to-end object detection with fully convolutional network. *arXiv preprint arXiv:2012.03544*, 2020.
- [45] Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, and Chunhua Shen. Repulsion loss: Detecting pedestrians in a crowd. *arXiv preprint arXiv:1711.07752*, 2017.
- [46] Kevin Zhang, Feng Xiong, Peize Sun, Li Hu, Boxun Li, and Gang Yu. Double anchor r-cnn for human detection in a crowd. *arXiv preprint arXiv:1909.09998*, 2019.
- [47] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4457–4465, 2017.
- [48] 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. *arXiv preprint arXiv:1912.02424*, 2019.
- [49] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. Occlusion-aware r-cnn: Detecting pedestrians in a crowd. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 637–653, 2018.
- [50] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. Single-shot refinement neural network for object detection. In *CVPR*, 2018.
- [51] Minghang Zheng, Peng Gao, Xiaogang Wang, Hongsheng Li, and Hao Dong. End-to-end object detection with adaptive clustering transformer. *CoRR*, abs/2011.09315, 2020.
- [52] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR 2021: The Ninth International Conference on Learning Representations*, 2021.