Bucketed Ranking-based Losses for Efficient Training of Object Detectors

Feyza Yavuz¹, Baris Can Cam¹, Adnan Harun Dogan¹, Kemal Oksuz², Emre Akbas^{1,3,*}, and Sinan Kalkan^{1,3,*}

¹ Department of Computer Engineering, METU, Ankara, Türkiye
 ² Five AI Ltd., United Kingdom
 ³ METU ROMER Robotics Center, Ankara, Türkiye
 * Equal senior authorship

Abstract. Ranking-based loss functions, such as Average Precision Loss and Rank&Sort Loss, outperform widely used score-based losses in object detection. These loss functions better align with the evaluation criteria, have fewer hyperparameters, and offer robustness against the imbalance between positive and negative classes. However, they require pairwise comparisons among P positive and N negative predictions, introducing a time complexity of $\mathcal{O}(PN)$, which is prohibitive since N is often large (e.g., 10^8 in ATSS). Despite their advantages, the widespread adoption of ranking-based losses has been hindered by their high time and space complexities. In this paper, we focus on improving the efficiency of ranking-based loss functions. To this end, we propose Bucketed Rankingbased Losses which group negative predictions into B buckets $(B \ll N)$ in order to reduce the number of pairwise comparisons so that time complexity can be reduced. Our method enhances the time complexity. reducing it to $\mathcal{O}(\max(N \log(N), P^2))$. To validate our method and show its generality, we conducted experiments on 2 different tasks, 3 different datasets, 7 different detectors. We show that Bucketed Ranking-based (BR) Losses yield the same accuracy with the unbucketed versions and provide $2 \times$ faster training on average. We also train, for the first time, transformer-based object detectors using ranking-based losses, thanks to the efficiency of our BR. When we train CoDETR, a state-of-the-art transformer-based object detector, using our BR Loss, we consistently outperform its original results over several different backbones. Code will be released.

Keywords: Object Detection · Ranking Losses · Efficient Ranking Losses

1 Introduction

Training an object detector incurs many challenges unique to object detection, including (i) assigning dynamically a large number of object hypotheses with the objects [2, 12, 13, 25, 29, 33, 34], (ii) sampling among these hypotheses to ensure that the background class does not dominate training [19, 27, 29, 32] and



Fig. 1: (a) Existing ranking-based losses (i.e., AP Loss [5,6]) incur significant overhead owing to pairwise comparisons between positives and negatives. (b) We propose bucketing negatives to decrease the number of comparisons, and hence the complexity. Under certain assumptions, our bucketing approach provides equal gradients with conventional ranking-based losses such as AP Loss in (a). (c) Accuracy and efficiency comparison. Our BRS formulation facilitates faster (between $1.9 \times$ and $6.0 \times$) training of visual detectors with similar AP.

(iii) minimizing a multi-task objective function [17,23,24,29]. While the choices on assignment and sampling generally vary depending on the trained object detector, it is very common to combine Cross Entropy or Focal Loss [17] with a regression loss [30,36] as the multi-task loss function. Recently proposed rankingbased loss functions [5,6,23,24] offer an alternative approach for addressing these challenges by formulating the training objective based on the rank of positive examples over negative examples.

Benefits of ranking-based losses. First, they are inherently robust to imbalance [23] and hence, do not require any sampling mechanism under very challenging scenarios [24], e.g., even when the background-foreground ratio is 10K for LVIS [9]. Second, they are shown to generalize over different detectors with diverse architectures – with the exception of transformer-based ones [14, 37, 38], since ranking-based losses further slow down the training of transformer-based detectors, which we address in this paper. Such losses also offer significant performance gain over their score-based counterparts, and having less hyperparameters, they are much easier to tune [23, 24].

The drawback of ranking-based losses. Compared to score-based losses, ranking-based losses are more inefficient as the ranking operation inherently requires each pair of object hypotheses to be compared against each other (Fig. 1(a)), inducing a quadratic complexity on the large number of object hypotheses (e.g., 10^8 for ATSS [34] on COCO [18]). As a result, vectorized implementations for parallel GPU processing are infeasible as such large matrices (e.g., with ~ 10^{16} entries for ATSS) do not fit into GPU memories. This has driven researchers towards alternative ways of computing these matrices, which, in the end, saves from the storage complexity but results in more inefficient algorithms [5,6].

Our approach. The main cause for the aforementioned complexity is the background class, i.e., negatives, forming up to 99.9% of the hypotheses [26]. Considering this, we introduce a novel bucketing approach (Fig. 1(b)) on these negatives to address the main shortcoming of the ranking-based loss functions. Specifically, we first sort all examples based on their scores and bucket the negatives located between successive positives. The negatives within the same bucket are treated as a single example and represented by their average score. We then incorporate our bucketing approach into AP Loss [5,6] and RS Loss [24], and introduce novel mechanisms for calculating the gradients correctly.

The significance of our approach. Theoretically, we prove that our bucketing approach provides the same gradients as the original ranking-based loss. Practically, we show we can reduce loss computation time by up to $\sim 40 \times$ and training time of a detector by up to $6 \times$ (Fig. 1(c)), addressing the main drawback of ranking-based loss functions. Hence, we show that the gap between the score-based and ranking-based loss functions in terms of training time vanishes, enabling us to train strong transformer-based object detectors such as Co-DETR using our loss functions. As ranking-based loss functions now surpass score-based ones or perform on par in every aspect such as accuracy, tuning simplicity, robustness to imbalance and training time, we believe that our work will make their prevalence more dominant in the upcoming object detectors.

Contributions. Our contributions can be summarized as follows:

- We propose a novel bucketing approach to improve the efficiency of computationally-expensive ranking-based losses to train object detectors. Theoretically, our approach yields the same gradients with the original ranking-based losses while decreasing their time complexity. Practically, we enable up to $6 \times$ faster training using our bucketing approach with no accuracy loss.
- For the first time, we incorporate ranking-based loss functions to transformerbased detectors. Specifically, we construct BRS-DETR by replacing the training objective of the state-of-the-art transfomer-based object detector Co-DETR [38] by our Bucketed RS (BRS) Loss.
- Our comprehensive experiments on detection and instance segmentation on 3 different challenging datasets, 5 backbones and 7 detectors show the effectiveness and generalizability of our approach. Our BRS-DETR reaches 50.1 AP on COCO val set with only 12 epochs and ResNet-50 backbone, outperforming all existing detectors using the same backbone and 300 queries. BRS-DETR also outperforms CoDETR with Swin-T and Swin-L backbones, reaching 57.2 AP on the COCO dataset.

2 Related Work

Although score-based loss functions, such as cross-entropy and focal loss, are widely used to train both CNN-based [2, 17, 35] and transformer-based object detectors [14, 37, 38], ranking-based losses for visual detection offer advantages. The primary advantage is their robustness to class imbalance [23]. They yield on par or better performance compared to score-based losses, without the need

for extensive tuning of class weights or multistage sampling techniques. Another advantage they offer is the ease of balancing losses among different tasks. In AP Loss [5,6], aLRP Loss [23] and RS Loss [24], the authors use the ratio of classification and regression losses to balance these tasks. As a result, rankingbased losses have much fewer hyperparameters during training, compared to score-based ones.

Other Ranking-based Losses. The literature has explored other rankingbased loss functions with different goals and assumptions. For example, DR Loss [28] constrains a margin to be satisfied between positives and negatives, not taking the recall into account. Smooth AP [1] introduces a smooth, differentiable approximation of AP, with the assumption that $|\mathcal{N}|$ is not large. Several studies focused on improving the efficiency of ranking-based losses for use in Support Vector Machines (SVMs) [21, 22]. However, these solutions have been limited to training linear SVMs and hence, problems that can be solved using linear functions. In this paper, we focus on improving ranking-based loss functions [5, 6, 24] that can be used for training complex deep networks such as object detectors.

Comparative Summary. Ranking-based training objective for visual detectors proved their superiority over score-based loss functions in visual detection tasks [5, 6, 23, 24]. However, due to the requirement of pairwise ranking their time complexity is higher than the score-based alternatives. Such a disadvantage of ranking-based loss functions prevents their applicability to larger visual detectors (i.e., transformer-based object detectors).

3 Background on Ranking-based Losses

Ranking-based losses for object detection rely on pairwise comparisons between the scores of different detections to determine the rank of a detection among positives and negatives. Denoting the score of the i^{th} detection by s_i , we can compare the scores of two detections i and j with a simple difference transform: $x_{ij} = s_j - s_i$. By counting the number of instances with $x_{ij} > 0$, we can determine i^{th} detection's rank among positives (\mathcal{P}) and all detections (positives and negatives: $\mathcal{P} \cup \mathcal{N}$) as follows:

$$\operatorname{rank}^{+}(i) = \sum_{j \in \mathcal{P}} \bar{H}(x_{ij}), \text{ and } \operatorname{rank}(i) = \sum_{j \in \mathcal{P} \cup \mathcal{N}} \bar{H}(x_{ij}),$$
(1)

where $\overline{H}(x)$ is one if x > 0 and zero otherwise. Since $d\overline{H}(x)/dx$ is either infinite (at x = 0) or zero (for $x \neq 0$), a smoothed version is used (δ : a hyper-parameter):

$$H(x) = \begin{cases} 0, & x < -\delta \\ \frac{x}{2\delta} + 0.5, & -\delta \le x \le \delta \\ 1, & \delta < x \end{cases}$$
(2)

3.1 Revisiting Average Precision (AP) Loss

Given the definitions in Eq. 1, AP Loss [5,6] can be defined as:

$$\mathcal{L}_{AP} = 1 - \mathrm{AP} = 1 - \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \operatorname{precision}(i) = 1 - \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \frac{\operatorname{rank}^+(i)}{\operatorname{rank}(i)}.$$
 (3)

Chen et al. [5,6] simplified Eq. 3 by rewriting it in terms of pairwise comparisons between positive and negative detections:

$$\mathcal{L}_{AP} = \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} \frac{H(x_{ij})}{\operatorname{rank}(i)} = \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} \sum_{j \in \mathcal{N}} L_{ij}^{AP},$$
(4)

where $L_{ij}^{AP} = H(x_{ij})/\operatorname{rank}(i)$ is called the primary term. Note that L_{ij}^{AP} is nondifferentiable since the step function $(H(\cdot))$ applied on x_{ij} is non-differentiable, which will be discussed next.

3.2 Revisiting Identity Update

Oksuz et al. [24] showed that various ranking-based losses (including AP Loss in Eqs. 3 & 4) can be written in a general form as:

$$\mathcal{L} = \frac{1}{Z} \sum_{i \in \mathcal{P} \cup \mathcal{N}} (\ell(i) - \ell^*(i)), \tag{5}$$

where $\ell(i)$ is the ranking-based error (e.g., precision error) computed on the *i*th detection, $\ell^*(i)$ is the target ranking-based error (the lowest error possible) and Z is the normalization constant. The loss in Eq. 5 can be computed and optimized as follows [24]:

1. Computation of the Loss. First, each pair of logits $(s_i \text{ and } s_j)$ are compared by calculating their difference transforms as $x_{ij} = s_j - s_i$. With the step function $H(\cdot)$ (Eq. 2), the number of detections higher than s_i (and therefore its precision error, $\ell(i)$) can be easily calculated (see Sect. 3.1). Given x_{ij} , the loss can be rewritten and calculated in terms of primary terms L_{ij} as $\mathcal{L} = \frac{1}{Z} \sum_{i \in \mathcal{P} \cup \mathcal{N}} \sum_{j \in \mathcal{P} \cup \mathcal{N}} L_{ij}$,

by taking:

$$L_{ij} = (\ell(i) - \ell^*(i)) \, p(j|i), \tag{6}$$

where p(j|i) is a probability mass function (pmf) that distributes the error computed on the *i*th example over the *j*th example in order to determine the pairwise primary term L_{ij} . p(j|i) is commonly taken as a uniform distribution [23, 24]. **2. Optimization of the Loss.** The gradient of the primary term wrt. to the difference transform $(\partial L_{ij}/\partial x_{ij})$ is non-differentiable. Denoting this term by Δx_{ij} , we have [5]:

$$\frac{\partial \mathcal{L}}{\partial s_i} = \sum_{j,k} \frac{\partial \mathcal{L}}{\partial L_{jk}} \Delta x_{jk} \frac{\partial x_{jk}}{\partial s_i} = \frac{1}{Z} \Big(\sum_j \Delta x_{ji} - \sum_j \Delta x_{ij} \Big).$$
(7)

Therefore, optimizing a ranking-based loss function reduces to determining Δx_{ij} . Inspired by Chen et al. [5], Oksuz et al. [24] employ Perceptron Learning [31] and show that Δx_{ij} in Eq. 7 is simply the primary term itself: $\Delta x_{ij} = -(L_{ij}^* - L_{ij}) =$ $-(0-L_{ij}) = L_{ij}$, hence the name Identity Update. Plugging this into Eq. 7 yields (see Oksuz et al. [24] for the steps of the derivation):

$$\frac{\partial \mathcal{L}}{\partial s_i} = \frac{1}{Z} \Big(\sum_j L_{ji} - \sum_j L_{ij} \Big).$$
(8)

Therefore, both computing and optimizing the loss reduces determining the primary term L_{ij} .

3.3 Revisiting AP Loss and RS Loss with Identity Update

AP Loss [5, 6]. For defining AP Loss with identity update, the errors in Eq. 5 can be derived as $\ell_{\rm R}(i) = \frac{N_{\rm FP}(i)}{\operatorname{rank}(i)}$, $\ell_{\rm R}^*(i) = 0$ and $Z = |\mathcal{P}|$. Furthermore, defining $p_R(j|i)$ as a uniform pmf, that is $p_R(j|i) = \frac{H(x_{ij})}{N_{\rm FP}(i)}$, we can obtain the primary terms of AP Loss using Identity Update, completing the derivation for computation and optimization (Eq. 8):

$$L_{ij}^{AP} = \begin{cases} (\ell_{\mathrm{R}}(i) - \ell_{\mathrm{R}}^{*}(i)) p_{R}(j|i), & \text{for } i \in \mathcal{P}, j \in \mathcal{N} \\ 0, & \text{otherwise,} \end{cases}$$
(9)

RS Loss [24]. RS Loss includes an additional sorting objective $(\ell_S(i))$ to promote better-localized positives to be ranked higher than other positives:

$$\mathcal{L}_{\rm RS} = \frac{1}{|\mathcal{P}|} \sum_{i \in \mathcal{P}} [\underbrace{(\ell_{\rm R}(i) - \ell_{\rm R}^*(i))}_{\rm ranking} + \underbrace{(\ell_{\rm S}(i) - \ell_{\rm S}^*(i))}_{\rm sorting}].$$
(10)

The primary terms of RS Loss can be obtained in a similar manner to AP Loss (see the Supp.Mat. for derivations), yielding:

$$L_{ij}^{RS} = \begin{cases} (\ell_{\mathrm{R}}(i) - \ell_{\mathrm{R}}^{*}(i)) p_{R}(j|i), & \text{for } i \in \mathcal{P}, j \in \mathcal{N} \\ (\ell_{\mathrm{S}}(i) - \ell_{\mathrm{S}}^{*}(i)) p_{S}(j|i), & \text{for } i \in \mathcal{P}, j \in \mathcal{P}, \\ 0, & \text{otherwise}, \end{cases}$$
(11)

where $p_S(j|i)$ is the sorting pmf. Note that RS Loss has the same primary terms with AP Loss for $i \in \mathcal{P}, j \in \mathcal{N}$. The primary terms for $i \in \mathcal{P}, j \in \mathcal{P}$ directly targets promoting the better-localized positives.

3.4 Complexity of Ranking-based Loss Functions

These loss functions originally have the space complexity of $\mathcal{O}((|\mathcal{P}|+|\mathcal{N}|)^2)$ [5,6] as they need to compare each pair. This quadratic space complexity makes it infeasible compute these loss functions using modern GPUs as the number of logits in object detection is very large. To alleviate that, Chen et al. [5] introduced two tricks at the expense of making the computation inefficient: (1) Loop

on Positives: By implementing a loop over the positive examples, Chen et al. managed to reduce the time complexity to $\mathcal{O}(|\mathcal{P}|(|\mathcal{P}| + |\mathcal{N}|)) \approx \mathcal{O}(|\mathcal{P}||\mathcal{N}|)$ since $\mathcal{N} \gg \mathcal{P}$ and the space complexity to $\mathcal{O}(|\mathcal{P}| + |\mathcal{N}|)$. (2) Discard Trivial Negatives: AP Loss considers negative examples with lower rank than all positives as trivial and disregards them.

4 Bucketed Ranking-based Loss Functions for Efficient Optimisation of Object Detectors

Training object detectors with existing ranking losses has several advantages as we presented before, however, they suffer from significantly large training time. We now address this by our Bucketing Approach, and propose our Bucketed AP and RS Losses.

4.1 Bucketing Negatives for Efficient Optimisation

Our main intuition is to bucket the sequential negative examples considering that they will have very similar or equal gradients. Fig. 1(a) demonstrates this phenomenon in which n_1, n_2 both are assigned equal gradients as well as each negative among n_3, n_4, n_5, n_6 in the standard AP Loss. Formally, we sort all given logits $s_i \in \mathcal{P} \cup \mathcal{N}$ using a conventional sorting methods. Let us denote this sorted permutation of the logits by $\hat{s}_1, \hat{s}_2, \dots, \hat{s}_{|\mathcal{P} \cup \mathcal{N}|}$, i.e., $\hat{s}_1 > \hat{s}_2 > \dots > \hat{s}_{|\mathcal{S}|}$. Given these sorted logits and denoting the *i*th positive logit in the ordering by \hat{s}_i^+ , the buckets of negatives $B_1, \dots, B_{|\mathcal{P}+1|}$ can be obtained by:

$$B_1 > \hat{s}_1^+ > B_2 > \hat{s}_2^+ > \dots > \hat{s}_{|\mathcal{P}|}^+ > B_{|\mathcal{P}+1|}.$$
(12)

We will denote the size of the *i*th negative bucket by b_i . Having created the buckets, we now determine a single logit value for each bucket, which we call as the prototype logit as it is necessary while assigning the ranking error. We denote the prototype logit for the *i*th bucket by s_i^b . Note that if $\delta = 0$ in Eq. 2, then any logit satisfying $\hat{s}_{i-1}^+ > s_i^b > \hat{s}_i^+$ can be a prototype logit as the ranking stays the same. However, in the case of $\delta > 0$, the boundaries between the logits are smoothed. That's why, practically, we find it effective to use the mean logit of a bucket as its prototype logit.

Note that this bucketing approach reduces the number of logits (positive and prototype negative) to a maximum of $2|\mathcal{P}| + 1$. As a result, the pairwise errors can now fit into the memory. Consequently, the loop in the red box of Alg. S1 (in Suppl. Mat.) is no longer necessary, which gives rise to efficient ranking-based loss functions which we discuss in the following.

4.2 Bucketed Ranking-based Loss Functions

Here, we introduce Bucketed versions of AP and RS Losses. Please refer to Alg. S2 (in Suppl. Mat.) for the algorithm.

Bucketed AP Loss Definition. To define Bucketed version of AP Loss, we need to define current $(\ell_{\rm R}^b(i))$ and target ranking errors $(\ell_{\rm R}^{b,*}(i))$ for the *i*th positive. Similar to previous works [5, 6, 24], we use a target ranking error of 0, i.e., $\ell_{\rm R}^{b,*}(i) = 0$. Defining the current ranking error, similar to Eq. 3, requires rank(*i*) and $N_{\rm FP}(i)$. These two quantities can easily be defined using the bucket size b_i and the prototype logit. Different from the conventional AP Loss, we obtain the pairwise relation $H(x_{ij}^b)$ using the prototype logit s_i^b . Particularly, with the ordering in Eq. 12, rank(*i*) = $\sum_{j=1}^{i} H(x_{ij}) + H(x_{ij}^b)b_j$ and $N_{\rm FP}(i) = \sum_{j=1}^{i} H(x_{ij}^b)b_j$ where $x_{ij}^b = s_j^b - s_i$. Then, the resulting ranking error computed on *i*th positive is:

$$\ell_{\rm R}^b(i) = \frac{N_{\rm FP}(i)}{\operatorname{rank}(i)} = \frac{\sum_{j=1}^i H(x_{ij}^b) b_j}{\sum_{j=1}^i H(x_{ij}) + H(x_{ij}^b) b_j}, \quad i \in \mathcal{P}.$$
 (13)

Optimisation. Here, we need to define the primary terms, as the product between the error and the pmf following Identity Update [24]. Unlike the previous work, the weights of each prototype negative are not equal as a bucket includes varying number of negatives. Therefore, we use a weighted pmf while distributing the ranking error over negatives. Formally, if j is the jth prototype negative, then we define the pmf as $p(j^b|i) = b_j/N_{\rm FP}(i)$. The resulting primary term between the *i*th positive and j^b th negative is then:

$$L_{ij}^{b} = \frac{\sum_{j=1}^{i} H(x_{ij}^{b}) b_{j}}{\sum_{j=1}^{i} H(x_{ij}) + H(x_{ij}^{b}) b_{j}} \times \frac{b_{j}}{N_{\rm FP}(i)}, \quad i \in \mathcal{P}, j \in \tilde{\mathcal{N}}$$
(14)

where $\tilde{\mathcal{N}}$ is the set of prototype negatives. However, we need the primary terms for the *i*th negative, not for the prototype ones. Given L_{ij}^b , the gradients of the prototype negatives and actual negatives can easily be obtained following Identity Update. However, we still need to find the gradients of the actual negatives. To do so, we simply normalise the prototype gradient by its bucket size b_j and distribute the gradients to actual negatives, completing our method. In Theorem 1, we establish that our Bucketed AP Loss provides exactly the same gradients with AP Loss when $\delta = 0$.

Bucketed RS Loss Given that we obtained Bucketed AP Loss, converting RS Loss into a bucketed form is more straightforward. This is because (i) the primary term and its gradients of RS Loss are equal to those of AP Loss if $i \in \mathcal{P}$, $j \in \mathcal{N}$ and we simply use L_{ij}^b in such cases; and (ii) if $i \in \mathcal{P}$, $j \in \mathcal{P}$, then an additional term is included in RS Loss to estimate the pairwise relations between positives (Eq. 11), which is not affected by bucketing.

4.3 Theoretical Discussion

The first theorem ensures that the gradients provided by our loss functions are identical with their conventional counterparts under certain circumstances. The second states that our algorithm is theoretically faster than the conventional AP and RS Losses.

9

Theorem 1. Bucketed AP Loss and Bucketed RS Loss provide exactly the same gradients with AP and RS Losses respectively when $\delta = 0$.

Theorem 2. Bucketed RS and Bucketed AP Losses have $\mathcal{O}(\max((|\mathcal{P}|+|\mathcal{N}|)\log(|\mathcal{P}|+|\mathcal{N}|), |\mathcal{P}|^2))$ time complexity.

As a result, we observe the same performance empirically in less amount of time. Proofs are provided in Supp. Mat.

5 Bucketed Rank-Sort (BRS) DETR

While transformer-based detectors [8, 14, 16, 37, 38] have been providing the best performance in several object detection benchmarks, optimizing such detectors using ranking-based losses based on performance measures has not been investigated. Here, we address this gap by incorporating our BRS Loss into Co-DETR [38] as the current SOTA detector on the common COCO benchmark [18].

Co-DETR increases the number and variation of the positive examples in the transformer head by leveraging one-to-many assignment strategies in anchorbased detectors (such as ATSS [34] and Faster R-CNN [29]) as auxiliary heads. As the number of proposals is large in such detectors, they can easily provide more and diverse examples, resulting in better performance of the transformer head. Co-DETR is originally supervised with the following loss function:

$$\sum_{l=1}^{L} \left(\tilde{\mathcal{L}}_{l}^{Dec} + \lambda_{1} \sum_{k=1}^{K} \tilde{\mathcal{L}}_{k,l}^{Aux} + \lambda_{2} \sum_{k=1}^{K} \mathcal{L}_{k,l}^{Aux} \right),$$
(15)

where λ_1 and λ_2 weigh each loss component. The first two components, $\tilde{\mathcal{L}}_l^{Dec}$ and $\tilde{\mathcal{L}}_{k,l}^{Aux}$, are the losses of the l^{th} decoder layer in the following form:

$$\lambda_{cls} \mathcal{L}_{cls} + \lambda_{bbox} \mathcal{L}_{bbox} + \lambda_{IoU} \mathcal{L}_{IoU}, \tag{16}$$

and \mathcal{L}_{cls} , \mathcal{L}_{bbox} and \mathcal{L}_{IoU} are Focal Loss, L1 Loss and GIoU Loss respectively [38]. $\tilde{\mathcal{L}}_{l}^{Dec}$ and $\tilde{\mathcal{L}}_{k,l}^{Aux}$ differ in their input queries. That is, while $\tilde{\mathcal{L}}_{l}^{Dec}$ follows standard DETR-based models [14, 37] with one-to-one matching of the queries, $\tilde{\mathcal{L}}_{k,l}^{Aux}$ is computed based on positive anchors of k^{th} auxiliary head with one-to-many assignment. And $\mathcal{L}_{k,l}^{Aux}$ is the conventional loss of the k^{th} auxiliary head, e.g., the weighted sum of classification, localisation and centerness losses for ATSS [34].

In order to align the training and evaluation objectives better, we replace $\tilde{\mathcal{L}}_{l}^{Dec}$ and $\tilde{\mathcal{L}}_{k,l}^{Aux}$ in Eq. 16 by:

$$\mathcal{L}_{BRS} + \lambda_{bbox} \mathcal{L}_{bbox} + \lambda_{IoU} \mathcal{L}_{IoU}, \tag{17}$$

and set λ_{bbox} and λ_{IoU} dynamically during training to $\mathcal{L}_{BRS}/\mathcal{L}_{bbox}$ and $\mathcal{L}_{BRS}/\mathcal{L}_{IoU}$ to simplify the hyperparameter tuning following Oksuz et al. [24]. Similarly, we replace the loss functions of the auxiliary ATSS [34] and Faster R-CNN [29] heads ($\mathcal{L}_{k,l}^{Aux}$ in Eq. 15) by our BRS Loss. This modification also aligns with the

objective of the auxiliary heads by the performance measure, hence results in more accurate auxiliary heads [24]. Furthermore, as using BRS Loss with Faster R-CNN does not require sampling thanks to its robustness to imbalance, no limitation is imposed on the number of positives in Faster R-CNN. Hence, the main aim of Co-DETR, that is to introduce more positive examples to the transformer head, is corroborated. As a result, our BRS DETR enables significantly efficient training compared to RS Loss also by improving the performance of Co-DETR.

6 Experiments

We analyze our contributions in three primary sections. First, we evaluate the effectiveness of our bucketing approach by contrasting it with RS Loss and AP Loss across different CNN-based visual detectors, on object detection and instance segmentation tasks. Next, we thoroughly examine how our bucketed RS Loss performs on transformer-based object detectors with different backbones. This is the first time a ranking-based loss is applied to transformer-based object detectors, thanks to the time-efficiency of our bucketing method. Finally, we present comparisons with score-based loss functions and other ranking-based loss functions in the literature. Our experiments clearly show that in terms of training time, accuracy and ease of tuning, our BRS Loss is more preferable over any existing loss function to train object detectors.

Dataset and Performance Measures. Unless otherwise noted, we train all models on COCO *trainval35k* (115k images) and test them on *minival* (5k images). We use COCO-style Average Precision (AP) and also report AP₅₀, AP₇₅ as the APs at IoUs 0.50 and 0.75; and AP_S, AP_M and AP_L to present the accuracy on small, medium and large objects.

6.1 Efficiency and Effectiveness of Bucketed Ranking Losses

To comprehensively analyze the effectiveness of our bucketing method, we design experiments on real-world data as well as synthetic data. For the former, we aim to present that our method significantly decreases the training time of the detection and segmentation methods (up to $\sim 5\times$). As training the entire detector does not isolate the runtime of the loss function, on which our main contribution is, we also design an experiment with synthetic data. This set of experiments show that our bucketing approach, in fact, decreases the loss function runtime by up to 40×, thereby resulting in shorter training time of the detectors.

Experimental Setup. We follow the experimental setup used in RSLoss [24] only by replacing the loss function by our bucketed loss function to ensure comparability and consistency of results. We report our findings on efficiency and accuracy using the models trained on 4 Tesla A100 GPUs with a batch size of 16, i.e., 4 images/GPU. Unless stated otherwise, we train the models for 12 epochs and test them on images with a size of 1333×800 following the common convention [12, 24, 34]. In order to comprehensively demonstrate the efficiency of our approach, we use five different detectors: Faster R-CNN [29], Cascade R-CNN [2],

Mathad	ethod Detector		C	Time		Accuracy	,
Method	Detector		2	$T_{iter}(s)\downarrow$	$AP\uparrow$	$AP_{50}\uparrow$	$AP_{75}\uparrow$
		D 50	RS	0.50	39.4	59.5	43.0
	Easter B CNN	11-50	BRS	0.17 3.0x ↓	39.5	59.5	42.8
Multi Store	raster n-onin	D 101	RS	0.75	47.3	67.4	51.2
Mutti Stage		11-101	BRS	0.38 2.0x ↓	47.7	67.8	51.5
	Cascade R-CNN	R-50	RS	1.28	41.1	58.7	44.1
			BRS	0.24 5.3x ↓	41.1	58.8	44.2
	ATSS	D 50	AP	0.32	38.1	58.2	41.0
			BAP	0.15 2.1x ↓	38.5	58.6	41.3
			RS	0.36	39.8	58.9	42.6
One Stage			BRS	0.15 2.4x ↓	39.8	58.8	42.9
One Stage		11-50	AP	0.45	37.3	54.3	41.2
	DAA		BAP	0.30 1.5x ↓	37.2	56.2	40.2
	IAA		RS	0.57	40.8	58.8	44.6
			BRS	0.30 1.9x ↓	40.9	59.0	44.4

Table 1: Bucketed Losses (BRS, BAP) vs. RS & AP Losses on COCO object detection for object detectors. R-X: ResNet-X. Models with R-101 are trained 36 epochs.

ATSS [34] and PAA [12] for object detection as well as Mask R-CNN [10] for instance segmentation.

Bucketed Ranking Based Losses Significantly Decrease the Training Time with Similar Accuracy. Multi-stage detectors: Among multi-stage detectors, we train the commonly-used Faster R-CNN and Cascade R-CNN with RS Loss and BRS Loss, and present average iteration time as well as AP of the trained models. In order to evaluate our contribution comprehensively, we train Faster R-CNN also with a stronger setting, in which, we use ResNet-101, train it for 36 epochs using multi-scale training similar to [24]. The results are presented in Table 1, in which we can clearly see that BRS Loss consistently reduces the training time of all three detectors as well as preserves (or slightly improve) their performance. Especially for Cascade R-CNN, which is still a very popular and strong object detector along with its variants such as HTC [4], the training time decreases by $5.3 \times$. This is because the loss function is applied to the logits for three times based on the cascaded nature of this detector, and therefore, our contribution can be easily noticed. Moreover, it might seem that the efficiency gain decreases once the size of the backbone increases from R-50 to R-101 in Faster R-CNN. However, this is an expected result as the overall feature extraction time increases due to the larger number of parameters in the backbone, and hence, as we will discuss in the synthetic experiments, this is independent from our bucketing approach, operating on the loss function after the feature extraction.

One-stage detectors: To show that our gains generalize to one-stage detectors and to AP Loss, we train the common ATSS and PAA detectors with our BAP and BRS Losses, and compare them with AP and RS Losses respectively. Table 1 validates our previous claims: Our bucketed losses obtain similar performance in around half of the training time required for AP and RS Losses.

Table 2: Comparison with RS Loss on instance segmentation task on COCO val using Mask R-CNN. BB: Backbone.

DD	C Efficiency		Accuracy				
ЪБ	L	$T_{iter}(s)\downarrow$	$AP\uparrow$	$AP_{50}\uparrow$	$AP_{75}\uparrow$		
R 50	RS	0.56	36.3	57.2	38.8		
11-50	BRS	0.24 2.3x ↓	36.2	57.2	38.8		
R 101	RS	0.59	40.2	61.8	43.5		
10-101	BRS	0.27 2.2x ↓	40.3	62.0	43.8		

Table 3: Comparison on different inst. segm. datasets using Mask R-CNN. AP_{75} is N/A as it is not used for Cityscapes.

Dataset	c Efficiency		Accuracy			
Dataset	~	$T_{iter}(s)\downarrow$	$AP\uparrow$	$AP_{50}\uparrow$	$AP_{75}\uparrow$	
Citysoppos	RS	0.43	43.5	68.1	N/A	
Cityscapes	BRS	0.19 2.3x ↓	43.3	67.7	N/A	
IVIS	RS	0.87	25.6	39.2	27.3	
LVIS	BRS	0.35 2.5x ↓	25.8	39.6	27.4	

Table 4: BRS Loss further improves performance with similar training budget. Wereport mask AP for instance segmentation

Task	Detector	L	$AP\uparrow$	$\mathrm{AP}_{50}\uparrow$	$\mathrm{AP_{75}}\uparrow$
Object Det.	C. R-CNN	RS BRS	$\frac{41.1}{\textbf{42.3}} + 1.2$	58.7 60.2	44.1 45.2
Inst. Seg.	M. R-CNN	RS BRS	36.3 37.3 +1.0	57.2 58.4	38.8 40.2

Instance Segmentation: Given that our loss functions is efficient in object detectors, one simple extension is to see their generalization to instance segmentation methods. To show that, we train Mask R-CNN, a common baseline, with our BRS Loss on three different dataset from various domains: (i) COCO (Table 2), (ii) Cityscapes [7] as an autonomous driving dataset (Table 3) and LVIS [9] a long-tailed dataset with more than 1K classes (Table 3). The results confirm our previous finding that our BRS Loss decreases the training time significantly by up to $\sim 2.5 \times$ compared to RS Loss by preserving its accuracy.

Given Similar Training Budget, Bucketed Ranking Based Loss Functions Significantly Improve the Accuracy. There might be several benefits of reducing the training time of the detector as more GPU time is saved. Here, we show a use-case in which we ask the following question: what would happen if we allocate the similar amount of training time to both bucketed and non-bucketed losses? To answer that, we train Cascade R-CNN, a detector, and Mask R-CNN, an instance segmentation method, using our BRS Loss. Considering our speed-ups on these models $(6.0 \times \text{ and } 2.3 \times)$, we simply increased their training epochs from 12 to 36 and 27 respectively. Note that, for Cascade R-CNN, we still spend significantly less amount of training time. Table 4 shows that models trained with the BRS Loss outperform (with +1.2 and 1.0 AP) their counterparts trained with the RS Loss in both cases.

Experiments on Synthetic Data. Up to now, to measure any speed-up, we considered the entire pipeline and did not decouple the loss computation step. However, we only improve the loss computation step. Hence, an analysis focusing only on this part can reveal our contribution more clearly. Therefore, we design a simple experiment using synthetic data. Particularly, we randomly generate L logits such that m% of these logits are positive (see Supp.Mat. for data generation details). To cover the wide range of existing detectors, we set $L \in \{10K, 100K, 1M\}$ and $m \in \{0.1, 1.0, 2.0, 5.0\}$. Different choices for L and m

 Table 5: Comparison of our BRS-DETR
 with Co-DETR (trained with its original loss function) on different backbones on COCO val set.

В	ackbone	Detector	AP	AP_{50}	AP_{75}	AP_s	AP_{m}	AP_1
R	eeNet50	Co-DETR	49.3	67.2	54.0	32.1	52.6	63.8
	carvetoo	BRS-DETR	50.1	67.4	54.6	31.9	53.9	65.0
	Swin T	Co-DETR	51.7	69.6	56.4	34.4	54.9	66.8
Γ,	Swiii-1	BRS-DETR	52.3	69.5	57.1	34.6	55.7	68.0
	Carries T	Co-DETR	56.9	75.5	62.6	40.1	61.2	73.3
1	Swiii-L	BRS-DETR	57.2	75.0	62.5	39.4	61.7	74.1

Table 6: Comparison of our BRS-DETR with DETR variants (trained with their original loss functions) w ResNet-50 on COCO val set. Qu: Queries. Ep: Epochs.

Detector	Qu	Ep	AP	AP_{50}	AP_{75}	AP_s	AP_{m}	AP_{l}
DETR [14]	100	300	42.0	62.4	44.2	20.5	45.8	61.1
DN-DETR [15]	300	50	48.6	67.4	52.7	31.0	52.0	63.7
DINO [3]	900	12	49.4	66.9	53.8	32.3	52.5	63.9
H-DETR [11]	300	12	48.7	66.4	52.9	31.2	51.5	63.5
Co-DETR [38]	300	12	49.3	67.2	54.0	32.1	52.6	63.8
BRS-DETR	300	12	50.1	67.4	54.6	31.9	53.9	65.0

correspond to the logit distributions of common object detectors. We compute BAP and AP Loss on this synthetic data and compare their runtime to focus on loss computation time. We observed that our BAP Loss has a significant, up to $\sim 40 \times$, speed-up compared to AP Loss. The speed-up becomes more pronounced with increasing number of logits. Supp. Mat. presents detailed results.

6.2 The Effectiveness of Our BRS-DETR

While we incorporate our BRS Loss into Co-DETR, we use all original settings unless otherwise noted. Please refer to Supp.Mat. for the details.

BRS-DETR Outperforms Existing Transformer-based Detectors Consistently. Here, we first employ ResNet-50 backbone as it has been used by many detectors, enabling us to compare our method with many different DETR-based manner in a consistent manner. In this fair comparison setting, our BRS-DETR outperforms all existing DETR variants reaching 50.1 AP as shown in Table 6. For example, our BRS DETR outperforms DN-DETR by 1.5AP with less training epochs and improves Co-DETR as we discuss next.

Training Co-DETR with RS Loss takes 6x less time compared to using **RS** Loss. Finally, we also compare the training efficiencies of BRS Loss and RS Loss on Co-DETR. We observe that our BRS Loss decreases the training time of Co-DETR by $6.0 \times$ (from 4.14s per iteration to 0.69s) in comparison to RS Loss. We note that this is the largest training time gain of our BRS Loss. This is because Co-DETR consists of multiple transformer-based as well as auxiliary heads, requiring multiple loss estimations.

BRS-DETR Improves Co-DETR over Different Backbones Consistently. In order to show the effectiveness of our BRS Loss, we compare our results with Co-DETR in Table 5 using different backbones. We note that we improve the baseline Co-DETR consistently in all settings. For example, our improvement on ResNet-50 is 0.8AP, which is a notable improvement. However, as the backbone gets larger (e.g., Swin-L [20]), our gains decrease, which is expected and commonly observed in the literature, e.g., [12].

Table 7: Comparison with the scorebased losses. CE: Cross-entropy. #H: Number of hyperparameters. Ours has very similar $T_{iter}(s)$ also by being more accurate and simple-to-tune.

Detector	\mathcal{L}	$T_{iter}(s)\downarrow$	$\mathrm{AP}\uparrow$	#H
	CE+L1	0.14	37.6	9
Faster R-CNN	RS	0.50	39.4	3
	BRS (Ours)	0.17	39.5	3
	Focal L.+GIoU	0.14	39.3	5
ATSS	RS	0.36	39.8	1
	BRS (Ours)	0.15	39.8	1

 Table 8: Comparison with other rankingbased losses. Our approach performs better on both accuracy and efficiency.

Detector	L	$T_{iter}(s)\downarrow$	$\mathrm{AP}\uparrow$	#H
Faster R-CNN	aLRP [23] BRS (Ours)	0.28	37.4 39 5	3
ATSS	AP [5]	0.32	38.1	5
	DR [28]	0.20	38.1	5
	$\begin{tabular}{ c c c c c c } \hline Or & \mathcal{L} & $T_{iter}(s) \downarrow AP \uparrow $$\\ \hline $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $$	1 1		
BetinaNet	DR [28]	0.29	37.4	5
	BRS (Ours)	0.23	38.3	1

6.3 Comparison with Score-based and Other Ranking-based Losses

In previous sections, we showed the efficiency of our loss functions compared to AP and RS Losses as their counterparts. Here, we compare our bucketed losses with the score based losses, i.e., Cross-entropy and Focal Loss, and other ranking based losses including DR Loss [28] and aLRP Loss [23]. To do so, in Tables 7 and 8, we report average training time, AP for accuracy and number of hyper-parameters to capture the tuning simplicity of the loss functions. Compared to score-based loss functions in Table 7, our BRS Loss improves performance of both approaches, is significantly simpler-to-tune with a low number of hyperparameters and has similar training iteration time. Hence, our BRS Loss is either superior or on par with existing score-based losses. As for other ranking-based loss function for all three detectors. As an example, compared to ATSS trained with DR Loss, our BRS Loss yields 1.8AP better accuracy with 25% less training time and significantly less number of hyperparameters. Therefore, our bucketed loss functions are now promising alternatives to train object detectors.

7 Conclusion

In this paper, we introduced a novel method to improve the efficiency of rankingbased loss functions by binning negatives into buckets and implementing positivenegative pairwise comparisons between positives and buckets of negatives. We showed that this method improves the computational complexity to a level where pairwise comparisons can be stored as a matrix in memory and running time becomes very close to that of score-based loss functions. Also, for the first time, we integrated a ranking-based loss to Co-DETR, a DETR-based detector, which was possible thanks to our method's lower complexity, and reported improvements on different backbones. Our comprehensive experiments comprising 2 different tasks, 3 different dataset and 7 different detectors – all with positive outcomes – show the general applicability of our approach.

Acknowledgements

We gratefully acknowledge the computational resources provided by TÜBİTAK ULAKBİM High Performance and Grid Computing Center (TRUBA), METU Center for Robotics and Artificial Intelligence (METU-ROMER) and METU Image Processing Laboratory. Dr. Akbas is supported by the "Young Scientist Awards Program (BAGEP)" of Science Academy, Türkiye.

References

- Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-ap: Smoothing the path towards large-scale image retrieval. In: European Conference on Computer Vision. pp. 677–694. Springer (2020)
- Cai, Z., Vasconcelos, N.: Cascade R-CNN: Delving into high quality object detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers (2021)
- Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: Hybrid task cascade for instance segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Chen, K., Li, J., Lin, W., See, J., Wang, J., Duan, L., Chen, Z., He, C., Zou, J.: Towards accurate one-stage object detection with ap-loss. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Chen, K., Lin, W., li, J., See, J., Wang, J., Zou, J.: Ap-loss for accurate one-stage object detection. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) pp. 1–1 (2020)
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Dai, Z., Cai, B., Lin, Y., Chen, J.: Unsupervised pre-training for detection transformers. IEEE Transactions on Pattern Analysis and Machine Intelligence p. 1–11 (2022). https://doi.org/10.1109/tpami.2022.3216514, http://dx.doi.org/ 10.1109/TPAMI.2022.3216514
- Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask R-CNN. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2017)
- 11. Jia, D., Yuan, Y., He, H., Wu, X., Yu, H., Lin, W., Sun, L., Zhang, C., Hu, H.: Detrs with hybrid matching (2023)
- 12. Kim, K., Lee, H.S.: Probabilistic anchor assignment with iou prediction for object detection. In: The European Conference on Computer Vision (ECCV) (2020)
- 13. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: The European Conference on Computer Vision (ECCV) (2018)

- 16 Yavuz et al.
- Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13619–13627 (June 2022)
- 15. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising (2022)
- Li*, L.H., Zhang*, P., Zhang*, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., Chang, K.W., Gao, J.: Grounded language-image pretraining. In: CVPR (2022)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 42(2), 318–327 (2020)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: The European Conference on Computer Vision (ECCV) (2014)
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S.E., Fu, C., Berg, A.C.: SSD: single shot multibox detector. In: The European Conference on Computer Vision (ECCV) (2016)
- 20. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows (2021)
- Mohapatra, P., Jawahar, C., Kumar, M.P.: Efficient optimization for average precision svm. Advances in Neural Information Processing Systems 27 (2014)
- Mohapatra, P., Rolínek, M., Jawahar, C., Kolmogorov, V., Pawan Kumar, M.: Efficient optimization for rank-based loss functions. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Oksuz, K., Cam, B.C., Akbas, E., Kalkan, S.: A ranking-based, balanced loss function unifying classification and localisation in object detection. In: Advances in Neural Information Processing Systems (NeurIPS) (2020)
- Oksuz, K., Cam, B.C., Akbas, E., Kalkan, S.: Rank & sort loss for object detection and instance segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3009–3018 (2021)
- Oksuz, K., Cam, B.C., Kahraman, F., Baltaci, Z.S., Kalkan, S., Akbas, E.: Maskaware iou for anchor assignment in real-time instance segmentation. In: The British Machine Vision Conference (BMVC) (2021)
- Oksuz, K., Cam, B.C., Kalkan, S., Akbas, E.: Imbalance problems in object detection: A review. IEEE transactions on pattern analysis and machine intelligence 43(10), 3388–3415 (2020)
- Pang, J., Chen, K., Shi, J., Feng, H., Ouyang, W., Lin, D.: Libra R-CNN: Towards balanced learning for object detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Qian, Q., Chen, L., Li, H., Jin, R.: Dr loss: Improving object detection by distributional ranking. In: CVPR (2020)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 39(6), 1137–1149 (2017)
- Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 31. Rosenblatt, F.: The perceptron: A probabilistic model for information storage and organization in the brain. Psychological Review pp. 65–386 (1958)

- Shrivastava, A., Gupta, A., Girshick, R.: Training region-based object detectors with online hard example mining. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- 33. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., Luo, P.: SparseR-CNN: End-to-end object detection with learnable proposals. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Zhang, X., Wan, F., Liu, C., Ji, R., Ye, Q.: Freeanchor: Learning to match anchors for visual object detection. In: Advances in Neural Information Processing Systems (NeurIPS) (2019)
- Zheng, Z., Wang, P., Liu, W., Li, J., Rongguang, Y., Dongwei, R.: Distance-iou loss: Faster and better learning for bounding box regression. In: AAAI Conference on Artificial Intelligence (2020)
- 37. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. ArXiv abs/2010.04159 (2020), https://api.semanticscholar.org/CorpusID:222208633
- Zong, Z., Song, G., Liu, Y.: Detrs with collaborative hybrid assignments training (2022)