

# Revisiting AP Loss for Dense Object Detection: Adaptive Ranking Pair Selection

Dongli Xu, Jinhong Deng, Wen Li \*

School of Computer Science and Engineering & Shenzhen Institute for Advanced Study  
University of Electronic Science and Technology of China

{dongliixu, jhdeng1997, liwenbnu}@gmail.com

## Abstract

*Average precision (AP) loss has recently shown promising performance on the dense object detection task. However, a deep understanding of how AP loss affects the detector from a pairwise ranking perspective has not yet been developed. In this work, we revisit the average precision (AP) loss and reveal that the crucial element is that of selecting the ranking pairs between positive and negative samples. Based on this observation, we propose two strategies to improve the AP loss. The first of these is a novel Adaptive Pairwise Error (APE) loss that focusing on ranking pairs in both positive and negative samples. Moreover, we select more accurate ranking pairs by exploiting the normalized ranking scores and localization scores with a clustering algorithm. Experiments conducted on the MS-COCO dataset support our analysis and demonstrate the superiority of our proposed method compared with current classification and ranking loss. The code is available at <https://github.com/Xudangliatiger/APE-Loss>.*

## 1. Introduction

Object detection is one of the fundamental computer vision tasks and aims to predict both the category labels and the bounding-box coordinates of all objects in a given image [3, 19, 21, 27, 30]. It also plays an important role in many down-stream applications such as instance segmentation [1, 9, 31, 33] and face detection [6]. Modern object detectors can be divided into two-stage approaches [3, 27] and one-stage approaches [35, 39]. One-stage detectors adopt dense prediction without a region proposal phase; thus they are also known as dense object detectors. One-stage object detectors are also naturally faster than two-stage detectors and popular in real-world applications such as that on edge devices [35, 39]. Most one-stage object detectors [16, 19, 21, 26, 30, 38, 40] rely on a classification task to discriminate the category of objects, *i.e.*, directly pre-

dicting the category probability for every patch in an image. However, it often suffers from an extreme imbalance problem, as huge numbers of background patches (*i.e.*, negative samples) can overwhelm the sufficient loss gradients for foreground patches (*i.e.*, positive samples). Recently, to alleviate this sensitivity to the ratio of negative and positive samples, Average Precision Loss [4] converts the classification task into a ranking task by explicitly modeling sample relationships that are calculated by comparing all sample pairs.

Despite AP loss performing well at addressing the sample imbalance problem, the essential mechanism of how AP loss affects the detector is still veiled. Accordingly, in this paper, we delve into AP loss from a pairwise ranking perspective [2]; specifically, the crucial part of AP loss is the minimizing of the pairwise error between the positive and negative samples. To find out the essential properties of AP loss, we analyze components of pairwise error separately and reveal that the key factor is the accurate and complete ranking pair selection. Consequently, we improve the AP loss by adjusting the way in which ranking pairs are constructed and selected. Proper ranking pair selection achieves notable performance gains for object detection.

More specifically, we first derive a reformulation of the AP loss and show that it contains three main components of generalized pairwise error: distance function, balance constant and ranking pair selection. We conduct detailed experiments to verify these different components' effects and determine that proper ranking pair selection plays a vital role in producing accurate detection results. We then investigate the barriers to good performance encountered by existing ranking pair selection strategies and identify the following issues: (1) Traditional strategies [13, 19, 30, 38] ignore the ranking pairs between positive samples, which may lead to inaccurate category probability prediction, and harms the Non-Maximum Suppression (NMS) result; (2) The probability of image content in the ranking task and accuracy in the localization task, *i.e.*, ranking score and localization score have different distributions, which brings imbalanced attentions to those two tasks during the pair selection pro-

\* Corresponding author.

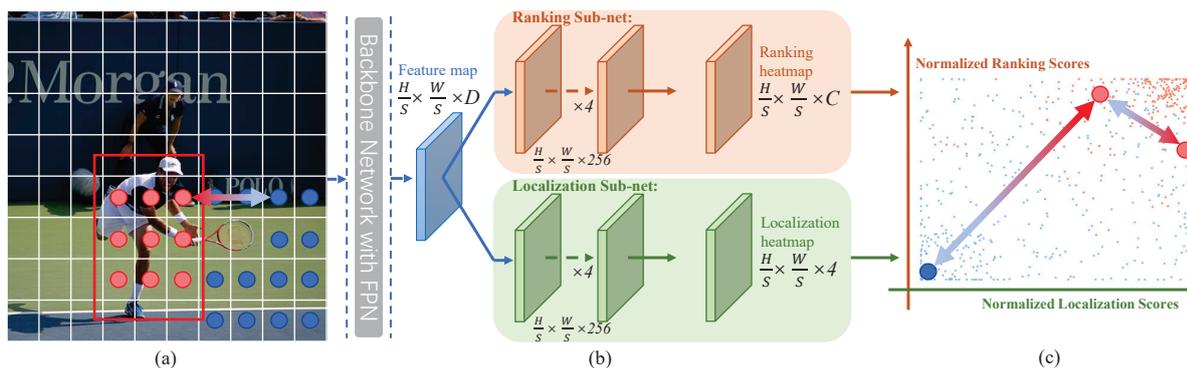


Figure 1. Comparison between our ranking pair selection and a traditional methods. Here, the large red and blue dots are used to highlight positive and negative samples, respectively, while the two-way arrows are used to highlight typical ranking pairs. (a) Current pairwise ranking methods ignore the pair between positive positions and focus on the anchor-based IoU to select ranking pairs. The red and white boxes are the ground truth and preset anchor. (b) FCOS detector (this paper omits the center-ness branch). (c) Our method can make up for the loss of ranking pairs between positive samples and focus on the distribution of prediction scores to adaptively select more accurate ranking pairs.

cessing as in Probabilistic Anchor Assignment (PAA) [13].

To alleviate these barriers, we propose a simple yet effective Adaptive Ranking Pair Selection (ARPS) approach to provide complete and more accurate ranking pairs for calculating pairwise error. To begin with, we build extra ranking pairs from the positive set to calculate Adaptive Pairwise Error (APE), which is differentiable and easy to implement and can exploit the ranking information between positive samples. It is worth noting that APE loss can also be considered a more accurate AP loss formulation. Second, we align the instance-level ranking scores and localization scores via normalization and feed them into a clustering algorithm (*e.g.*, Gaussian Mixture Model in PAA [13]) to create a better split between the positive and the negative set. After that, ranking pairs can be easily obtained from every combination pair of two clusters. Experimental results show that our method can optimize the training procedure of ranking tasks with higher accuracy.

The contributions of this work are three-fold and can be summarized as follows: 1) We conduct thorough experiments to verify each part of AP loss from a pairwise perspective and reveal that inappropriate ranking pair selection is the main obstacle; 2) we propose a ranking pair selection algorithm, *i.e.*, ARPS, to exploit complete and more accurate ranking pairs automatically; 3) our method achieves competitive performance when evaluated against all other existing classification and ranking methods.

## 2. Related Work

**Dense Object Detection:** In object detection frameworks, two-stage detectors commonly utilize cross-entropy loss [3, 27] for the classification purpose, while many dense detectors equip GHM loss [14] and focal loss [19] to alleviate the sample imbalance problem. Moreover, a large number of works have attempted to boost dense object detection,

using methods that include exploiting information from different layers [8, 18], feature alignment [25, 37] and accurate category probability [11, 13, 16, 36, 37].

**Accurate Object Detection:** Recent studies [11, 16, 30, 36] have integrated the quality information of positive samples, *i.e.*, Intersection over Union (IoUs) between predicted boxes and ground-truth boxes, into the classification task or additional IoU (or center-ness) prediction task (*e.g.*, IoUs are employed as the labels for the CE loss on classification task). When detectors are able to predict the qualities of the positive samples, more accurately predicted category probabilities (*e.g.*, IoU-aware classification scores in [37]) will greatly improve the results of the box selection algorithm NMS [11].

**Ranking for Object Detection:** Recently, AP loss [4] has been proposed to avoid the imbalance problem in dense object detection by proposing a novel ranking framework to replace the classification task. aLRP Loss [22] extends AP loss to optimize unified detection metric LRP and also addresses the localization task. DR loss [24] re-distributes positive and negative samples so that only one expectation ranking pair is used in loss; however, it ignores localization qualities, *i.e.*, IoUs of positive samples. Learning to ranking model [29] adds an extra branch in two-stage detectors to solve the ranking of localization quality. Similarly, RankDetNet [20] considers the task as a separated ranking task and tackles it with an IoU-guided ranking loss. Adopting a different approach to these works, we verify each part of AP loss from a pairwise perspective through experimental analysis and determine that inappropriate ranking pair selection is the main obstacle. Hence we propose a ranking pair selection algorithm, ARPS, to automatically produce complete and more accurate ranking pairs and merge the localization qualities into a more accurate AP metric optimization problem.

From the perspective of localization quality ranking, a

related and concurrent work is RS loss [23], which also proposes to leverage the information among positive samples by applying a sorting loss. Although RS loss [23] considers the relation among positive samples, it differs from our proposed method in essential ways. They combines IoUs into error-driven update to tackle the non-differentiable nature of ranking and sorting; by contrast we utilize IoUs information to construct extra ranking pairs from positive samples into a differentiable pairwise error, and can work without error-driven update.

**Training Sample Selection:** Throughout this paper, we will use the terms “ranking pair selection” and “training sample selection” interchangeably, with the understanding that the ranking pairs are formulated as different pairs from two sample sets; *i.e.*, the positive and negative. RetinaNet [19] defines positive samples using a manually adjusted IoU threshold of preset anchor box, while ATSS [38] provides an adaptive IoU threshold. Recent studies [12, 13, 15] adaptively separate positive and negative sets depending on the model’s learning status.

### 3. Preliminary

Based on traditional detector designs, let  $\mathbf{O} \in R^{H \times W \times 3}$  be an input image of width  $W$  and height  $H$ . Given an image  $\mathbf{O}$ , we use a backbone network  $\mathcal{F}_{FPN}$  with Feature Pyramid Net (FPN) [18] to obtain  $n$  feature maps as follows:

$$\mathbf{F}_0, \mathbf{F}_1, \dots, \mathbf{F}_i, \dots, \mathbf{F}_n = \mathcal{F}_{FPN}(\mathbf{O}) \quad (1)$$

where  $\mathbf{F}_i \in R^{\frac{H}{S_i} \times \frac{W}{S_i} \times D_i}$  is the feature map at layer  $i \in \{0, 1, \dots, n\}$ ,  $S_i$  is the output stride and  $D_i$  is the dimensional depth of  $\mathbf{F}_i$ . The ranking sub-net  $\mathcal{F}_{rank}$  and localization sub-net  $\mathcal{F}_{loc}$  are applied on each  $\mathbf{F}_i$ , such that our objective heatmaps are obtained as follows:

$$\hat{\mathbf{P}}_i = \mathcal{F}_{rank}(\mathbf{F}_i) \quad \hat{\mathbf{B}}_i = \mathcal{F}_{loc}(\mathbf{F}_i) \quad (2)$$

where  $\hat{\mathbf{P}}_i \in R^{\frac{H}{S_i} \times \frac{W}{S_i} \times M}$  and  $\hat{\mathbf{B}}_i \in R^{\frac{H}{S_i} \times \frac{W}{S_i} \times 4}$  represent category probability prediction and localization bounding box prediction, respectively; moreover, 4 denotes the encoding of four coordinates for localization. Furthermore,  $M$  is equal to  $C \times A$ , where  $C$  is the number of classes ( $C$  is 80 for the COCO [17] dataset) and  $A$  is the number of anchors.

Specifically, with  $N_{pos}$  positive samples, our training objective function  $L_{rank}$  can be described as the average precision loss  $L_p$  of each positive sample. The  $L_{rank}$  and localization loss  $L_{loc}$  on these two sub-nets can be expressed as follows:

$$L_{rank} = \frac{1}{N_{pos}} \sum_{i^+} \sum_{c^+} \sum_{x^+, y^+} L_p(\hat{\mathbf{P}}_{i^+, x^+, y^+, c^+}, \hat{\mathbf{P}}) \quad (3)$$

$$L_{loc} = \frac{1}{N_{pos}} \sum_{i^+} \sum_{x^+, y^+} L_{GloU}(\hat{\mathbf{B}}_{i^+, x^+, y^+}, \mathbf{B}_{i^+, x^+, y^+}) \quad (4)$$

where  $\hat{\mathbf{P}}_{i^+, x^+, y^+, c^+} \in R$  and  $\hat{\mathbf{P}}$  are the predictions for each sample on a positive location  $\{i^+, x^+, y^+, c^+\}$  and all the samples on every location, respectively. We train  $M$  binary predictors rather than one multi-class predictor, following [19].  $\hat{\mathbf{B}}_{i^+, x^+, y^+} \in R^4$  and  $\mathbf{B}_{i^+, x^+, y^+} \in R^4$  are the prediction and the ground truth label for localization, respectively, while  $\{i^+, x^+, y^+\}$  indicates the position of positive samples. As for localization,  $L_{GloU}$  is the GloU loss [28].

### 4. A Pairwise Ranking Perspective on AP Loss

Note that the predicted probabilities  $\hat{\mathbf{P}}$  (before sigmoid function) of all samples are fed into the loss function  $L_p$ . This loss can accordingly exploit the ranking pairs between the positive and negative samples. We formulate the precision loss as follows:

$$\begin{aligned} L_p(\hat{\mathbf{P}}_u, \hat{\mathbf{P}}) &= 1 - \text{Precision}(\hat{\mathbf{P}}_u, \hat{\mathbf{P}}) = \frac{FP}{TP + FP} \\ &= \frac{\text{rank}^-(u)}{\text{rank}^+(u) + \text{rank}^-(u)} \\ &= \frac{\sum_{v \in \mathcal{N}} H(\hat{\mathbf{P}}_v - \hat{\mathbf{P}}_u)}{\sum_{v \in \mathcal{P}, v \neq u} H(\hat{\mathbf{P}}_v - \hat{\mathbf{P}}_u) + \sum_{v \in \mathcal{N}} H(\hat{\mathbf{P}}_v - \hat{\mathbf{P}}_u)} \end{aligned} \quad (5)$$

where index  $u$  is used to replace  $\{i^+, x^+, y^+, c^+\}$  for notation convenience, while  $\text{rank}^+$  and  $\text{rank}^-$  denote the ranking position of sample  $u$  in the positive sample set  $\mathcal{P}$  and negative sample set  $\mathcal{N}$  respectively. They also indicate the number of True Positives ( $TP$ ) and False Positives ( $FP$ ), respectively. Here, to approximately calculate the ranking position, AP loss adopts the distance function  $\mathbf{H}(\cdot)$  between sample pairs, which is a piece-wise step function and can be written as follows:

$$H(x) = \begin{cases} 0, & x < -\delta \\ x/2\delta, & -\delta \leq x \leq \delta \\ 1, & \delta < x \end{cases} \quad (6)$$

where  $\delta$  is a tuning hyper-parameter used to control the slope of  $\mathbf{H}(\cdot)$  in  $[-\delta, +\delta]$ .

However, the vector field of prediction errors of  $\mathbf{H}(\cdot)$  is not conservative. Thus [4] proposed an Error-Driven Update method for AP loss by manually setting the error  $g_u$  as the gradient of precision loss  $L_p$  w.r.t. a positive prediction  $\hat{\mathbf{P}}_u$ . Here,  $g_u$  is defined as follows:

$$g_u = \frac{-\sum_{v \in \mathcal{N}} H(\hat{\mathbf{P}}_v - \hat{\mathbf{P}}_u)}{\text{rank}^+(u) + \text{rank}^-(u)} \quad (7)$$

where  $g_u$  is used in the same way as  $\frac{\partial L_p(\hat{\mathbf{P}}_u, \hat{\mathbf{P}})}{\partial \hat{\mathbf{P}}_u}$  in back propagation. Similarly,  $g_v$ , the gradients w.r.t. a negative prediction  $\hat{\mathbf{P}}_v$  can be formulated as:  $g_v = H(\hat{\mathbf{P}}_v -$

Table 1. Comparison of distance functions.

distance function	Error-Driven Update	AP	AP <sub>50</sub>	AP <sub>75</sub>
$\mathbf{H}(\cdot)$	✓	37.4	57.5	39.2
$\mathbf{S}(\cdot)$	✓	37.3	57.4	38.9
$CE(\mathbf{S}(\cdot), 0)$		37.3	57.2	39.1

$\hat{P}_u)/(rank^+(u) + rank^-(u))$ , where  $g_v$  is used in the same way as  $\frac{\partial L_p(\hat{P}_u, \hat{\mathbf{P}})}{\partial \hat{P}_v}$  in back propagation. It is worth noting that  $\sum_{v \in \mathcal{N}} g_v = -g_u$ ; thus, pairwise design can alleviate the imbalance of positive and negative samples.

In fact, the denominator term of the backpropagation gradient  $g_u$  shown in Eq. (7) is the same as the one of precision loss in Eq. (5) and can be considered as a balance constant. It is evident that optimizing AP loss is equivalent to minimizing the numerator in Eq. (5) [22], which can also be considered as the sum of the distance error between each pair of prediction scores, *i.e.*, pairwise error. Therefore, to reduce the redundancy of analyzing ranking tasks, we can rewrite the precision loss as pairwise error loss with the unified formalized description following [2]. We here present this unified pairwise error loss  $L_{PE}$  as follows:

$$L_{PE}(\hat{P}_u, \hat{\mathbf{P}}) = -\frac{1}{BC} \sum_{v \in \mathcal{N}} D(\hat{P}_v - \hat{P}_u) \quad (8)$$

where  $BC$  denotes the balance constant, while  $\mathbf{D}(\cdot)$  denotes the distance function adopted as the pairwise error. Finally, we have our ranking loss  $L_{rank}$  for all positive samples in a mini-batch:  $L_{rank} = \sum_u L_{PE}(\hat{P}_u, \hat{\mathbf{P}})/N_{pos}$ .

## 5. What Contributes to Pairwise Error?

As shown in Eq. (8), the reformulation of the pairwise ranking loss consists of three essential parts: (1) the distance function  $\mathbf{D}(\cdot)$ ; (2) the balance constant  $BC$ ; (3) the pair selection of  $\hat{P}_v$  and  $\hat{P}_u$ . To find out the essential properties of different designs in ranking-based losses, we analyze each part separately and verify the effect of different strategies for pairwise ranking methods [2, 22, 29]:

**Distance Function:** [2, 29] adopt sigmoid  $\mathbf{S}(\cdot)$  as distance function, which can be described as follows:  $S(\hat{P}_v - \hat{P}_u) = 1/(1 + \exp(\lambda(\hat{P}_u - \hat{P}_v)))$ , where  $\lambda$  is a tuning hyper-parameter that controls the slope of the sigmoid function near the origin. By contrast, AP loss uses a non-differentiable function  $\mathbf{H}(\cdot)$ , described in Eq. (6) and employs Error-Driven Update for backpropagation. However, as shown in the first and second rows of Table 1, when  $\mathbf{H}(\cdot)$  in AP loss is replaced with  $\mathbf{S}(\cdot)$ , the influence on performance is minor. Additionally, we also use a cross entropy loss  $CE(\cdot, 0)$  to replace Error-Driven Update (see appendix for more details). The results are shown in Table 2, where it can be seen that the usage of Error-Driven Update provides the same performance as CE loss (*i.e.*, 37.3). These results illustrate that the influence on the performance of different

Table 2. Comparison of balance constants.

Balance Constant	Threshold	AP	AP <sub>50</sub>	AP <sub>75</sub>
$rank^+(i) + rank^-(u)$		37.3	57.4	38.9
$N_{neg}$	✓	37.3	56.7	39.4

Table 3. Varying for sampling strategy on  $L_{PE}$ .

Sampling Strategy	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
IoU threshold	37.3	57.4	38.9	19.4	41.7	51.8
ATSS	39.2	59.4	41.8	21.4	43.2	53.7

distance functions employed in recent pairwise loss [2, 4, 29] are minor.

**Balance Constant:** Consequently, we need a balance constant to normalize the sum of all pairwise errors. AP loss uses  $rank^+(i) + rank^-(u)$  in Eq. (7) as its balance constant. While [2] sets the balance constant as the number of valid negative samples  $N_{neg}$ , which can keep the loss function simple. Note that, only if  $N_{neg}$  is small enough, we can replace balance constant with  $N_{neg}$ . Hence, we use a threshold to define which sample is valid (for more details, please refer to the appendix). We conduct empirical studies with these two balance constants. As shown in Table 1, the two different balance constants also yield similar performances (37.3 *v.s.* 37.3). Although the results suggest that these two balance constant choices are robust, the balance constant design of AP loss is more adaptive and does not require a valid sample threshold hyper-parameter.

**Ranking Pair Selection:** As for the selection of ranking pairs, *i.e.*,  $(\hat{P}_v, \hat{P}_u)$ , ALRP loss [22] adopts ATSS [38] while AP loss adopts an IoU threshold, following [19]. When we use the ATSS [38] to select ranking pairs, the performance of pairwise loss is significantly promoted; specifically, it is boosted from 37.3 to 39.2. This result may be caused by some important ranking pairs being ignored or many of the selected ranking pairs being inaccurate. Thus, inappropriate ranking pair selection is a significant barrier to higher accuracy.

**Analysis of Inappropriate Ranking Pair Selection:** Current state-of-the-art pair selection strategies [4, 22, 24] only depend on IoUs between preset anchor boxes and ground truths; thus, there is a lack of exploiting image content. As can be seen from Fig. 2 (a), the background and foreground may have the same priority, since the distance to the centre of this ground truth box is the same for both of them, leading to an inferior ranking pair selection (*e.g.*, there will be a ranking pair between the two grass positions in Fig. 2 (a)). Recent studies on adaptive sampling selection [12, 13, 15] integrate both the predicted category probabilities and localization results into positive sample selection, since the category predictions can represent the image content, while accurate localization results can also achieve better detection accuracy. To address this barrier, we employ PAA [13] which utilize both classification loss and localization loss

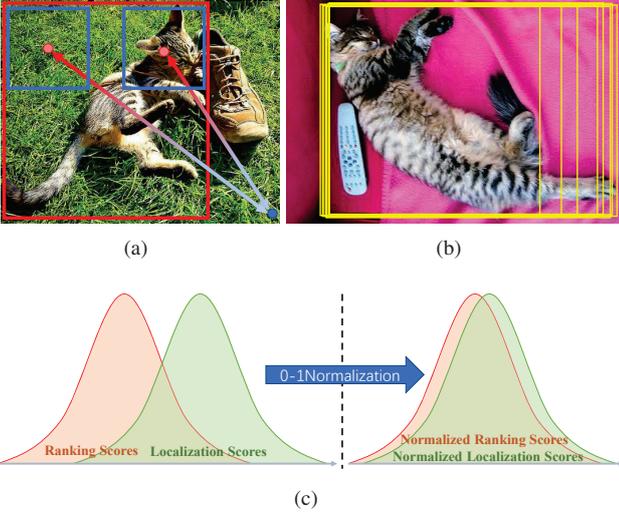


Figure 2. (a) Different positions in the ground truth box may have the same IoU between their preset anchor box and ground truth. The blue boxes are used to denote anchor boxes, while two-way arrows are used to highlight two ranking pairs. (b) Dense object detection can predict many candidate boxes from positive positions. The box that perfectly predicts this cat is clearly performing better than those that ignore the cat’s foot. (c) **Left:** The distribution of original ranking scores and localization scores. Although they are in the range of  $[0, 1]$ , they have different distributions. **Right:** The distribution of 0-1 normalized ranking scores and localization scores.

for adaptive sampling selection. However, two major problems remain with the existing sampling algorithms when incorporated into the pairwise ranking method:

(1) existing ranking pair selection strategies naturally ignore all important ranking pairs between positive samples (*i.e.*,  $v$  in Eq. (8) is limited in the negative sample set  $\mathcal{N}$ ). As shown in Fig. 2 (b), the detection model may predict many candidate boxes for a specific positive position. This reveals the existence of localization accuracy disparities between positive samples. It is accordingly necessary to exploit the localization information for ranking loss. If ranking pairs between the positive samples are ignored, the detector is more likely to select unreliable boxes, which has a severely detrimental impact on the final accuracy.

(2) Ranking scores (*i.e.*, the predicted ranking probabilities  $S(\hat{\mathbf{P}})$ ) used in the sampling processing have different distributions with localization scores (*i.e.*, IoUs), leading to the attentions of the two tasks being uncoordinated. We illustrate this phenomenon in Fig. 2 (c). The localization scores typically have values close to 1, while the ranking scores generally have relatively lower values. This is because the training stops before the ranking scores of positive samples goes to 1. In particular, when the  $\hat{P}_u$  of positive sample is relatively higher than  $\hat{P}_v$ , the distance function  $D(\cdot)$  will drop to 0. To this end, PAA [13] will focus more

---

### Algorithm 1: Adaptive Ranking Pair Selection

---

**Input:**  
 $\mathcal{P}$  is a set of positive positions  
 $\mathcal{N}$  is a set of negative positions  
 $\hat{\mathbf{I}}$  are the IoUs between boxes  $\hat{\mathbf{B}}$  predicted at the corresponding position and ground truth  
**Output:**  
 $\mathcal{A}$  is a set of adaptive negative position sets

```

1  $\mathcal{A} \leftarrow \emptyset$ ;
2 forall  $u \in \mathcal{P}$  do
3   build an empty negative sample set for each positive
   sample position  $u$ :  $\mathcal{A}_u \leftarrow \emptyset$ ;
4   forall  $v \in \mathcal{P}, v \neq u$  do
5     if  $\hat{I}_v < \hat{I}_u$  then
6       If its IoU is smaller than that of  $p$ , this
       sample should be considered an adaptive
       False Positive (aFP) sample in original  $\mathcal{P}$  set
        $\mathcal{A}_u = \mathcal{A}_u \cup v$ ;
7     end
8   end
9    $\mathcal{A}_u = \mathcal{A}_u \cup \mathcal{N}$ ;
10   $\mathcal{A} = \mathcal{A} \cup \{\mathcal{A}_u\}$ ;
11 end
12 return  $\mathcal{A}$ ;
```

---

on image content than localization accuracy, because classification loss will obtain a higher value than localization loss at the end of the training.

## 6. Adaptive Ranking Pair Selection

We propose ARPS to calculate Adaptive Pairwise Error between positive samples. Our method focuses on dynamically selecting negative samples  $\mathcal{A}_u$  in positive set  $\mathcal{P}$  according to their localization qualities. Notably, each positive sample  $u$  is assigned with a different set  $\mathcal{A}_u$  to replace  $\mathcal{N}$  in Eq. (8). Here, APE loss  $L_{APE}$  can be written as follows:

$$L_{APE}(\hat{\mathbf{P}}_u, \hat{\mathbf{P}}) = -\frac{1}{BC} \sum_{v \in \mathcal{A}_u} D(\hat{P}_v - \hat{P}_u) \quad (9)$$

where the balance constant  $BC$  is set as  $rank^+(u) + rank^-(u)$ , while the distance function is set as cross entropy with sigmoid  $CE(\mathbf{S}(\cdot), 0)$ . In a departure from previous works [11, 16, 30, 36], adaptive pairwise error integrates localization quality-related information into extra ranking pairs between positive samples rather than regressing a certain value (e.g., center-ness in [30] and IoU in [16]). Our method is adaptive and does not bring any hyperparameters, which eases training overloads for different applications. With the dynamic expansion of the negative sample set, our proposed APE loss considers the ranking pairs in the positive samples and further boosts the detection accuracy.

**Description:** Alg. 1 describes how the proposed method works for an existing negative sample set and dynamically expands it. For each positive position  $u \in \mathcal{P}$  with IoU  $\hat{I}_u$ , we add the samples  $v$  with lower IoU  $\hat{I}_v$  than  $\hat{I}_u$  into a dynamically expanded negative sample set  $\mathcal{A}_u$  as shown in Alg. 1. After we store all additional negative samples from  $\mathcal{P}$ , we merge  $\mathcal{A}_u$  with  $\mathcal{N}$  in lines 9 to 10. After we employ the ARPS to calculate adaptive pairwise error in Eq. (9), it would be interesting to explore the obstacle part in the original AP metric. Hence, we re-write APE loss in AP formula and analyse which parts of AP loss are modified by ARPS.

$$\begin{aligned}
 L_{\text{APE}}(\hat{P}_u, \hat{P}) &= -\frac{1}{BC} \sum_{v \in \mathcal{A}_u} D(\hat{P}_v - \hat{P}_u) \\
 &= \frac{\sum_{v \in \mathcal{N}} D(\hat{P}_v - \hat{P}_u) + \sum_{v \in \mathcal{P}, I_v < I_u} D(\hat{P}_v - \hat{P}_u)}{\sum_{v \in \mathcal{P}, v \neq u} D(\hat{P}_v - \hat{P}_u) + \sum_{v \in \mathcal{N}} D(\hat{P}_v - \hat{P}_u)} \quad (10) \\
 &== \frac{FP + aFP}{(TP - aFP) + (FP + aFP)}
 \end{aligned}$$

where  $BC$  is also set as  $rank^+(u) + rank^-(u)$ ; moreover, a sample should be considered as an adaptive False Positive (aFP) sample if its IoU is smaller than that of  $p$ .

In conclusion, the question of how to define TPs and FPs is an important one for dense object detection, as not all positive samples are the best detection results. ARPS accordingly identifies some adaptive False Positive (aFP) samples in the original TP set, and can thus be considered to provide a more accurate AP metric.

Additionally, to reduce the inaccurate ranking pairs in ARPS that are selected by the preset anchor’s IoU, we focus on image contents by investigating ranking scores and localization scores. Considering there are massive numbers of candidate boxes before the NMS algorithm, positive samples around objects can receive comparatively higher localization scores while the IoUs of negative samples are normally close to 0. This means that a large localization score gap exists between positive and negative samples. Instead, it can be inferred from Fig. 3 (b) that the ranking scores of the positive and negative samples only have a small gap between them (e.g., the value of the distance function will be close to 0 when  $\hat{P}_u - \hat{P}_v$  goes to 1, which means that after sigmoid function, their gap  $S(\hat{P}_u) - S(\hat{P}_v)$  is far smaller than 1).

We here employ PAA [13] to exploit the information of these pair gaps. Specifically, we first use 0-1 normalization for each instance to reduce the scale differences between the ranking score and localization score. We can then observe from Fig. 3 (c) that samples naturally gather close to two ends (i.e., (0, 0) and (1, 1)), and the blank area between the blue and red shades is an ideal separation that we expected. However, there are still many outliers of the two clusters, i.e., the small blue and red dots in the gap. If the ranking

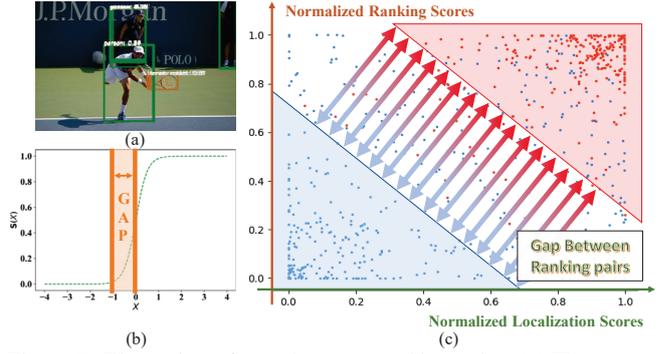


Figure 3. Illustration of gaps between ranking pairs. (a) The massive candidate box before NMS can have a high IoU with the ground truth, while all remaining negative samples have 0 IoU. (b)  $S(\cdot)$  and its design, which lead to a small gap between ranking scores. (c) Ideal gap between ranking pairs in the view of normalized ranking and localization scores.

pairs are formulated as every different potential combination of two samples from the two clusters (the bottom-left and top-right region in Fig. 3 (c)), the pairwise error will in turn enlarge this gap and control the outliers as the training progresses. Therefore, we convert the ranking pair selection problem into a two-cluster discrimination problem on the two normalized scores. Following PAA [13], Gaussian Mixture Model (GMM) is used to construct positive and negative clusters.

## 7. Experiments

### 7.1. Dataset and Implementation Details

Experiments are conducted on the MS-COCO 2017 dataset [17]. We use a split set `train-2017` (115k images) for training and validation set `val-2017` (5k images) for the ablation study and compare with other losses. `test-dev` split is used to test advanced additions. The COCO-style average precision (AP) is employed as the performance metric. The implementation is based on Pytorch 1.4, CUDA 10 and mmdetection [5]. Unless otherwise specified, all hyper-parameters are set to default, as in RetinaNet [19] and FCOS [30]. The ranking score threshold of NMS is 0.15 and IoU threshold of NMS is 0.6. The batch size is 16 and 4 GPUs are used for training (we use 4  $\times$  GTX 2080Ti-11G for training on ResNet-50 [10] and 4  $\times$  RTX Titan-24G for training with advanced methods, i.e., advanced backbone and multi-scale training). SGD is used as an optimizer and warmed up for the first 500 iterations.

**RetinaNet 512:** Here, we follow the baseline, i.e., AP loss, and employ SSD-style [21] data augmentations (including Photo Metric Distortion and Min IoU Random Crop [5]) for comparable evaluations. Moreover, to improve this baseline, we employ GIoU [28] loss to obtain more accurate localization results. Here, the max pair number  $Q$  (see appendix) here is 1, 000, 000. The reso-

Table 4. Overall contribution of the proposed methods over baseline.  $\diamond$  denotes the usage of a higher learning rate. \* denotes replacing the clustering input of PAA [13] with normalized ranking and localization scores.

Ranking Loss	Sampling Strategy	AP	AP <sub>50</sub>	AP <sub>75</sub>
AP Loss $\diamond$	IoU threshold	37.3	57.4	38.9
Adaptive Pairwise Error	IoU threshold	38.3	55.8	41.1
Adaptive Pairwise Error	ATSS	39.9	58.3	42.6
Adaptive Pairwise Error	PAA*	41.1	59.6	43.7

Table 5. Contribution of APE loss in association with localization scores.

Method	PCC	SCC	KCC
Pairwise Error	0.3742	0.3790	0.2538
Adaptive Pairwise Error	0.4940	0.5034	0.3449

lution of the input images is  $512 \times 512$  and the initial learning rate is 0.004. For 48-epochs training, we divide learning rate by 10 at 32 epochs and again at 44 epochs.

**FCOS 800:** We here omit the center-ness branch and make the pipeline more elegant. To ensure focus on more important samples, we also use category prediction to weight the GIoU loss following [16], and use IoU to weight APE loss following [32]. The weight of ranking loss is set to 1.0. The max pair number  $Q$  here is 100,000 (see appendix). The resolution of the input images is  $800 \times 1333$  with the same ratio of raw images. The initial learning rate is 0.01. For 24-epochs training, we divide the learning rate by 10 at 18 and again at 22 epochs.

## 7.2. Ablation Study

We conduct ablation experiments on ResNet-50-FPN [10, 18] to evaluate the effect of each component of our proposed method. To begin with, we implement our approach on RetinaNet to evaluate individual components relative to the baseline method, *i.e.*, AP loss [4]. However, RetinaNet is an obsolete detector with lower performance than the currently most popular method, *i.e.*, FCOS. Therefore, we also test our method on the state-of-the-art detector FCOS (ATSS) to illustrate the superiority of APE loss over recent classification losses.

**Effect of Adaptive Pairwise Error:** We conduct experiments with our implementation of Pairwise Error on RetinaNet, where  $\lambda$  is set to 8 (see appendix for the varying of  $\lambda$ ). As shown in Table 4, by replacing AP loss with APE loss, we can improve the AP and AP<sub>75</sub> by 1.0 and 2.2, respectively. This demonstrates that our proposed APE loss improves detection accuracy. Moreover, to better understand the reasons why APE loss is effective, we adopt IoUs and ranking scores of detection results with higher localization scores (*i.e.*, IoUs larger than 0.5) as two lists of data and calculate their Pearson Correlation Coefficient (PCC), Spearman Correlation Coefficient (SCC) and Kendall rank Correlation Coefficient (KCC). The results are presented

Table 6. Comparison of different clustering input scores of PAA. Here all scores are 0-1 normalized for each instance.

GMM Clustering Input	0-1 Normalization	AP	AP <sub>50</sub>	AP <sub>75</sub>
Focal Loss + GIoU Loss		39.4	58.1	41.4
Ranking + Localization Scores		40.5	57.6	43.5
Ranking + Localization Scores	✓	41.1	59.6	43.7

Table 7. Comparison of different detectors with our method.

Detector	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
RetinaNet(512)	41.1	59.6	43.7	22.9	45.2	57.3
FCOS(800)	41.5	59.2	44.7	23.6	46.1	54.8

Table 8. Training time comparison of different training conditions.

Ranking Loss	Resolution	Learning Rate	Epoch
AP Loss	512	0.001	100
AP Loss	512	0.004	48
Adaptive Pairwise Error	512	0.004	48
Adaptive Pairwise Error	800	0.01	24

in Table 5. From the table, it can be seen that APE loss promotes PCC by 0.12, SCC by 0.12 and KCC by 0.09. As quantitative measurements of correlation, the three coefficients are promoted significantly, indicating that APE loss can associate the category predictions with localization qualities and thereby boost NMS accuracy.

**Effect of Normalized Scores Input of PAA:** To study the necessity of two normalized scores for PAA, we test the original PAA [13] that adopts Focal loss [19] and GIoU loss [28] for clustering. As the results in Table 6 show, when we replace normalized ranking and classification scores with Focal loss and GIoU loss, the performance drops from 41.1 to 39.4. We contend that it is because the scales of ranking and localization gap in these losses are not well aligned; localization loss would be smaller, and its information could be easily ignored. However, PAA works well in classification methods owing to the fact that these two scores are trained to be equal [16, 37] and could have similar distribution scales. We additionally test different clustering input features; normalized ranking scores and localization scores independently achieve 38.9 and 33.9 AP, illustrating the association of two scores brings 2.2 AP performance gain. Moreover, as shown in Table 6, when normalization processing on the two scores is omitted, the performance also drops to 40.5, illustrating the normalization brings 0.6 AP performance gain.

**Effect of Higher Resolution and Learning Rate:** To facilitate comparative evaluation of our method with state-of-the-art losses, we incorporate it into the most popular detector with the higher input resolution, *i.e.*, FCOS(800). As the results in Table 7 show, training on FCOS(800) can improve our method by 0.4 AP. Note that we here discard SSD-style augmentation. We argue that this promotion occurs due to the higher input resolution, since the main performance gap between RetinaNet and FCOS is addressed by ATSS [38]. We also record the comparison of training

Table 9. Here we omit center-ness branch [30] or IoU branch [13] in our method. Data Aug. denotes the SSD-style [21] data augmentation. † denotes multi-scale training, ‡ denotes multi-scale testing.

Loss Function	Method	Data Aug.	Epoch	Backbone	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
<i>Classification Loss (on val-2017):</i>										
Focal Loss [19] (on test-dev)	Retina(500)	✗	12	ResNet-50-FPN	32.5	50.9	34.8	13.9	35.8	46.7
GHM Loss [14]	Retina(800)	✗	12	ResNet-50-FPN	35.8	55.5	38.1	19.6	39.6	46.7
Focal Loss [38]	ATSS(800)	✗	12	ResNet-50-FPN	39.2	56.5	40.7	20.6	42.1	49.1
Quality Focal Loss [16]	ATSS(800)	✗	12	ResNet-50-FPN	39.9	58.5	43.0	22.4	43.9	52.7
Varifocal Loss [37]	ATSS(800)	✗	12	ResNet-50-FPN	40.2	58.2	44.0	23.9	44.1	52.2
Focal Loss [13]	PAA(800)	✗	18	ResNet-50-FPN	41.1	59.4	44.3	23.5	45.4	54.3
<i>Ranking Loss (on val-2017):</i>										
Average Precision Loss [4, 22]	Retina(512)	✓	100	ResNet-50-FPN	35.4	58.1	37.0	19.2	39.7	49.2
Distributional Ranking Loss [24]	Retina(800)	✗	12	ResNet-50-FPN	37.4	56.0	40.0	20.8	41.2	50.5
aLRP Loss [22]	ATSS(512)	✓	100	ResNet-50-FPN	41.0	61.2	43.3	22.9	44.7	55.8
RankDetNet Losses [20]	ATSS(800)	✗	24	ResNet-50-FPN	40.7	58.4	44.0	23.8	44.8	52.8
Rank & Sort Loss [23]	PAA(800)	✗	12	ResNet-50-FPN	41.0	59.1	44.5	23.9	44.9	53.5
<i>Ours (on val-2017):</i>										
Adaptive Pairwise Error	Retina(512)	✓	48	ResNet-50-FPN	38.3	55.8	41.1	21.0	42.5	53.1
Adaptive Pairwise Error	ATSS(512)	✓	48	ResNet-50-FPN	39.9	58.3	42.6	21.9	44.1	55.0
Adaptive Pairwise Error	ATSS(800)	✗	24	ResNet-50-FPN	40.8	58.4	44.0	23.0	44.5	54.5
Adaptive Pairwise Error	PAA*(512)	✓	48	ResNet-50-FPN	41.1	59.6	43.7	22.9	45.2	57.3
Adaptive Pairwise Error	PAA*(512)	✓	96	ResNet-50-FPN	42.9	61.3	46.1	25.0	47.1	59.1
Adaptive Pairwise Error	PAA*(800)	✗	24	ResNet-50-FPN	41.5	59.2	44.7	23.6	46.1	54.8
<i>Losses with Adv. Tricks (on test-dev):</i>										
Average Precision Loss <sup>‡</sup> [4]	Retina(512)	✓	100	ResNet-101-FPN	42.1	63.5	46.4	25.6	45.0	53.9
Distributional Ranking Loss <sup>†‡</sup> [24]	Retina(800)	✗	18	ResNeXt-101-FPN	44.7	63.8	48.7	28.2	47.4	56.2
Focal Loss <sup>†</sup> [38]	ATSS(800)	✗	24	ResNeXt-101-FPN-DCN	47.7	66.5	51.9	29.7	50.8	59.4
Generalize Focal Loss <sup>†</sup> [16]	ATSS(800)	✗	24	ResNeXt-101-FPN-DCN	48.2	67.4	52.6	29.2	51.7	60.2
<i>Ours with Adv. Tricks (on test-dev):</i>										
Adaptive Pairwise Error <sup>†</sup>	PAA*(800)	✗	24	ResNeXt-101-FPN-DCN	49.0	67.1	53.0	29.9	52.9	63.4
Adaptive Pairwise Error <sup>†</sup>	PAA*(800)	✗	36	Res2Net-101-FPN-DCN	<b>50.0</b>	<b>67.7</b>	<b>53.9</b>	<b>31.8</b>	<b>55.3</b>	<b>66.7</b>

epochs as shown in Table 8. It can be observed that our method only needs a quarter of the AP loss training epochs, *i.e.*, 24 rather than 100. However, due to the slow converge speed of baseline AP loss, we still train APE loss with more epochs than recent losses [13, 16, 19, 37]. If a higher learning rate (*i.e.*, 0.04) is used, we can also train AP loss within 48 epochs, the same as ours. Thus, we argue that the promotion of training efficiency occurs largely because of the higher learning rate and higher image resolution than the baseline.

### 7.3. Comparison with State-of-the-Art Losses

Finally, we compare APE loss with state-of-the-art loss functions on val-2017, as shown in Table 9. Here, every method only uses single-model and single-scale testing, omits 1.5 stage regression (as in [25, 37]) and utilizes ResNet-50-FPN [10, 18]. APE loss achieves 41.5 on the FCOS [30] detector without center-ness branch. We also verify our method on ATSS, PAA and RetinaNet. The performances of APE loss are superior to that of recent state-of-the-art losses. Furthermore, we test our method with advanced additions on test-dev, *i.e.*, multi-scale training and stronger backbones (ResNext [34], Res2Net [7]) with Deformable ConvNets v2 [41]. As shown in Table 9, those additions boost our methods to 50.0 AP.

## 8. Conclusion

In this paper, we revisit the AP loss from a pairwise ranking perspective for dense object detection. In the process, we reveal an essential fact that proper ranking pair selection plays an important role in producing accurate detection results compared with the distance function and balance constant. Therefore, we propose a novel strategy, Adaptive Ranking Pair Selection (ARPS), by providing more complete and accurate ranking pairs. We first exploit the localization information into extra rank pairs with the Adaptive Pairwise Error, which can also be considered as a more accurate form of AP loss. We then use normalized ranking scores and localization scores to split the positive and negative samples. The proposed method is very simple and achieves performance comparable to existing classification and ranking methods.

## Acknowledgement

This work is supported by the Major Project for New Generation of AI under Grant No. 2018AAA0100400, the National Natural Science Foundation of China (Grant No. 62176047) and Beijing Natural Science Foundation (Z190023).

## References

- [1] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 9157–9166, 2019. [1](#)
- [2] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proc. Int. Conf. Machine Learning (ICML)*, pages 89–96, 2005. [1, 4](#)
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 6154–6162, 2018. [1, 2](#)
- [4] Kean Chen, Jianguo Li, Weiyao Lin, John See, Ji Wang, Lingyu Duan, Zhibo Chen, Changwei He, and Junni Zou. Towards accurate one-stage object detection with ap-loss. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5119–5127, 2019. [1, 2, 3, 4, 7, 8](#)
- [5] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. [6](#)
- [6] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 5203–5212, 2020. [1](#)
- [7] Shanghua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip HS Torr. Res2net: A new multi-scale backbone architecture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2):652–662, 2021. [8](#)
- [8] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 7036–7045, 2019. [2](#)
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 2961–2969, 2017. [1](#)
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 770–778, 2016. [6, 7, 8](#)
- [11] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yunying Jiang. Acquisition of localization confidence for accurate object detection. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 784–799, 2018. [2, 5](#)
- [12] Wei Ke, Tianliang Zhang, Zeyi Huang, Qixiang Ye, Jianzhuang Liu, and Dong Huang. Multiple anchor learning for visual object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 10206–10215, 2020. [3, 4](#)
- [13] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 355–371, 2020. [1, 2, 3, 4, 5, 6, 7, 8](#)
- [14] Buyu Li, Yu Liu, and Xiaogang Wang. Gradient harmonized single-stage detector. In *Proc. AAAI Conf. Artificial Intelligence (AAAI)*, pages 8577–8584, 2019. [2, 8](#)
- [15] Hengduo Li, Zuxuan Wu, Chen Zhu, Caiming Xiong, Richard Socher, and Larry S Davis. Learning from noisy anchors for one-stage object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 10588–10597, 2020. [3, 4](#)
- [16] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 21002–21012, 2020. [1, 2, 5, 7, 8](#)
- [17] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Microsoft coco: Common objects in context. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 2117–2125, 2014. [3, 6](#)
- [18] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 2117–2125, 2017. [2, 3, 7, 8](#)
- [19] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Focal loss for dense object detection. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 2117–2125, 2017. [1, 2, 3, 4, 6, 7, 8](#)
- [20] Ji Liu, Dong Li, Rongzhang Zheng, Lu Tian, and Yi Shan. Rankdetnet: Delving into ranking constraints for object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 264–273, 2021. [2, 8](#)
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 21–37, 2016. [1, 6, 8](#)
- [22] Kemal Oksuz, Baris Can Cam, Emre Akbas, and Sinan Kalkan. A ranking-based, balanced loss function unifying classification and localisation in object detection. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020. [2, 4, 8](#)
- [23] Kemal Oksuz, Baris Can Cam, Emre Akbas, and Sinan Kalkan. Rank & sort loss for object detection and instance segmentation. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 3009–3018, 2021. [3, 8](#)
- [24] Qi Qian, Lei Chen, Hao Li, and Rong Jin. Dr loss: Improving object detection by distributional ranking. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 12164–12172, 2020. [2, 4, 8](#)
- [25] Han Qiu, Yuchen Ma, Zeming Li, Songtao Liu, and Jian Sun. Borderdet: Border feature for dense object detection. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 549–564, 2020. [2, 8](#)
- [26] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 779–788, 2016. [1](#)
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 91–99, 2015. [1, 2](#)

- [28] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 658–666, 2019. 3, 6, 7
- [29] Zhiyu Tan, Xuecheng Nie, Qi Qian, Nan Li, and Hao Li. Learning to rank proposals for object detection. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 8273–8281, 2019. 2, 4
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, pages 9627–9636, 2019. 1, 2, 5, 6, 8
- [31] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pages 649–665, 2020. 1
- [32] Shengkai Wu, Jinrong Yang, Xinggang Wang, and Xiaoping Li. Iou-balanced loss functions for single-stage object detection. *arXiv preprint arXiv:1908.05641*, 2019. 7
- [33] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 12193–12202, 2020. 1
- [34] Saining Xie, Ross Girshick, Piotr Dollar, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 1492–1500, 2017. 8
- [35] Yunyang Xiong, Hanxiao Liu, Suyog Gupta, Berkin Akin, Gabriel Bender, Yongzhe Wang, Pieter-Jan Kindermans, Mingxing Tan, Vikas Singh, and Bo Chen. Mobiledets: Searching for object detection architectures for mobile accelerators. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 3825–3834, 2021. 1
- [36] Dongli Xu, Jian Guan, Pengming Feng, and Wenwu Wang. Association loss for visual object detection. *IEEE Signal Processing Letters*, 27:1435–1439, 2020. 2, 5
- [37] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sunderhauf. Varifocalnet: An iou-aware dense object detector. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 8514–8523, 2021. 2, 7, 8
- [38] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 9759–9768, 2020. 1, 3, 4, 7, 8
- [39] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 6848–6856, 2018. 1
- [40] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 840–849, 2019. 1
- [41] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 9308–9316, 2019. 8