

DR Loss: Improving Object Detection by Distributional Ranking

Qi Qian¹ Lei Chen² Hao Li² Rong Jin¹

Alibaba Group

¹Bellevue, WA, 98004, USA

²Hangzhou, China

{qi.qian, fanjiang.cl, lihao.lh, jinrong.jr}@alibaba-inc.com

Abstract

Most of object detection algorithms can be categorized into two classes: two-stage detectors and one-stage detectors. Recently, many efforts have been devoted to one-stage detectors for the simple yet effective architecture. Different from two-stage detectors, one-stage detectors aim to identify foreground objects from all candidates in a single stage. This architecture is efficient but can suffer from the imbalance issue with respect to two aspects: the inter-class imbalance between the number of candidates from foreground and background classes and the intra-class imbalance in the hardness of background candidates, where only a few candidates are hard to be identified. In this work, we propose a novel distributional ranking (DR) loss to handle the challenge. For each image, we convert the classification problem to a ranking problem, which considers pairs of candidates within the image, to address the inter-class imbalance problem. Then, we push the distributions of confidence scores for foreground and background towards the decision boundary. After that, we optimize the rank of the expectations of derived distributions in lieu of original pairs. Our method not only mitigates the intra-class imbalance issue in background candidates but also improves the efficiency for the ranking algorithm. By merely replacing the focal loss in RetinaNet with the developed DR loss and applying ResNet-101 as the backbone, mAP of the single-scale test on COCO can be improved from 39.1% to 41.7% without bells and whistles, which demonstrates the effectiveness of the proposed loss function.

1. Introduction

The performance of object detection has been improved dramatically with the development of deep neural networks in the past few years. Most of detection algorithms fall into two categories: two-stage detectors [4, 11, 12, 15] and one-stage detectors [3, 14, 16, 18, 21, 25, 29]. For the two-stage schema, the procedure of the algorithms can be divided




Figure 1. Illustration of the proposed distributional ranking loss. First, we push the distributions of confidence scores towards the decision boundary by re-weighting examples. Then, we try to rank the expectation of the derived distribution of foreground above that of background by a large margin.

into two parts. In the first stage, a region proposal method filters most of background candidate bounding boxes and keeps only a small set of candidates. In the following stage, these candidates are classified as specific foreground classes or background and the bounding boxes will be further refined by minimizing a regression loss. Two-stage detectors demonstrate the superior performance on real-world data sets while the efficiency can be an issue in practice, especially for the devices with limited computing resources, e.g., smartphones, cameras, etc.

Therefore, one-stage detectors are developed for the efficient detection. Different from two-stage detectors, one-stage methods consist of a single phase and have to identify foreground objects from all candidates directly. The procedure of one-stage detectors is straightforward and efficient. However, one-stage detectors can suffer from the imbalance problem that can reside in the following two aspects. First, the numbers of candidates between classes are imbalanced. Without a region proposal phase, the number of background candidates can easily overwhelm that of foreground ones. Second, the hardness of identification for background candidates is imbalanced. Most of them can be easily identified from foreground objects while only a few of them are hard to be classified.

To mitigate the imbalance problem, SSD [18] adopts hard negative mining in training, which is a popular strategy [24, 27] to keep a small set of background candidates

with the highest loss. By eliminating simple background candidates, the strategy balances the number of candidates between classes and the hardness of background simultaneously. However, certain information from background can be lost, and thus the detection performance can degrade as illustrated in [16]. RetinaNet [16] proposes to keep all background candidates but assign different weights for the loss functions of candidates. The weighted cross entropy loss is referred as focal loss. It makes the algorithm focus on the hard candidates while reserving the information from all candidates. This strategy improves the performance of one-stage detectors significantly. Despite the success of focal loss, it re-weights classification losses in a heuristic way and can be insufficient to address the imbalance problem. Moreover, focal loss is designed for a single *candidate* and is image-independent while object detection aims to identify objects in a single *image*. Focal loss lacks the exploration for each image as a whole and the inconsistency can make the performance suboptimal.

In this work, we propose an image-dependent ranking loss to handle the imbalance challenge. First, to mitigate the effect of the inter-class imbalance problem, we convert the classification problem to a ranking problem, which considers ranks of pairs. Since each pair consists of a foreground candidate and a background candidate, it is well balanced. Moreover, considering the intra-class imbalance in hardness of background candidates, we design the distributional ranking (DR) loss to rank the distribution of confidence scores for foreground above that for background candidates. As illustrated in Fig. 1, we first push the original distributions towards the decision boundary with appropriate constraints. After obtaining the drifted distributions, we can rank the expectations of distributions in lieu of original examples to identify foreground from background, which improves the efficiency by reducing the number of pairs from $\mathcal{O}(n^2)$ to $\mathcal{O}(1)$ in ranking, where n is the number of candidates in an image. Compared with focal loss, DR loss is image-dependent and can explore the information within each image sufficiently.

We conduct experiments on COCO [17] to demonstrate the proposed DR loss. Since the focal loss is designed as the classification loss in RetinaNet, we adopt RetinaNet as the base detector for a fair comparison. Specifically, we merely replace the focal loss with the DR loss while keeping other components unchanged. With ResNet-101 [12] as the backbone, minimizing our loss function can boost the mAP of RetinaNet from 39.1% to 41.7%, which confirms the effectiveness of the proposed loss.

The rest of this paper is organized as follows. Section 2 reviews the related work in object detection. Section 3 describes the details of the proposed DR loss. Section 4 compares our method to others on COCO detection task. Finally, Section 5 concludes this work.

2. Related Work

Detection is a fundamental task in computer vision. In conventional methods, hand crafted features, e.g., HOG [5] and SIFT [19], are used for detection either with a sliding-window strategy which holds a dense set of candidates, e.g., DPM [6] or with a region proposal method which keeps a sparse set of candidates, e.g., Selective Search [26]. Recently, deep neural networks have shown the dominating performance in classification tasks [13], and the features obtained from neural networks are leveraged for detection tasks.

R-CNN [9] equips the region proposal stage and works as a two-stage algorithm. It first obtains a sparse set of regions by selective search. In the next stage, a deep convolutional neural network is applied to extract features for each region. Finally, regions are classified with a conventional classifier, e.g., SVM. R-CNN improves the performance of detection by a large margin but the procedure is too slow for real-world applications. Hence, many variants are developed to accelerate it [8, 22]. To further improve the accuracy, Mask-RCNN [11] adds a branch for object mask prediction to boost the performance with the additional information from multi-task learning. Besides the two-stage structure, Cascade R-CNN [2] develops a multi-stage strategy to promote the quality of detectors after the region proposal stage in a cascade fashion.

One-stage detectors are developed for efficiency [3, 18, 20, 23, 29]. Since there is no region proposal phase to sample background candidates, one-stage detectors can suffer from the imbalance issue from both the inter-class imbalance between foreground and background candidates and intra-class imbalance in the background candidates. To address the challenge, SSD [18] adopts hard negative mining, which only keeps a small set of hard background candidates for training. Recently, focal loss [16] is proposed to handle the problem in RetinaNet. Unlike SSD, it keeps all background candidates but re-weights them such that the hard examples are assigned with a large weight. Focal loss improves the performance of one-stage detection explicitly, but the imbalance problem in detection is still not sufficiently explored. Besides those anchor-based algorithms, anchor-free one-stage detectors [25, 29] have been developed, where focal loss is also applied for classification. The work closest to ours is the AP-loss in [3], where a ranking loss is designed to optimize the average precision. However, the loss focuses on the original pairs and is non-differentiable. A specific algorithm has to be developed to minimize the AP-loss. In this work, we develop the DR loss that ranks the expectations of distributions in lieu of original pairs. DR loss is differentiable and can be optimized with stochastic gradient descent (SGD) in the standard training pipeline. Therefore, our loss can work in a plug and play manner, which is important for real-world applications.

3. DR Loss

Given a set of candidate bounding boxes from an image, a detector has to identify the foreground objects from background ones with a classification model. Let θ denote a classifier and it can be learned by optimizing the problem

$$\min_{\theta} \sum_i^N \sum_{j,k} \ell(p_{i,j,k}) \quad (1)$$

where N is the number of total images. In this work, we employ sigmoid function to predict the probability for each candidate. $p_{i,j,k}$ is the prediction from θ and indicates the estimated probability that the j -th candidate in the i -th image is from the k -th class. $\ell(\cdot)$ is the loss function. In most of detectors, the classifier is learned by minimizing the cross entropy loss or its variants.

The objective in Eqn. 1 is prevalent but can suffer from the inter-class imbalance problem. The problem can be demonstrated by rewriting the original problem as

$$\min_{\theta} \sum_i^N \left(\sum_{j_+}^{n_+} \ell(p_{i,j_+}) + \sum_{j_-}^{n_-} \ell(p_{i,j_-}) \right) \quad (2)$$

where j_+ and j_- denote the positive (i.e., foreground) and negative (i.e., background) examples (e.g., anchors), respectively. n_+ and n_- are the corresponding number of examples. When $n_- \gg n_+$, the accumulated loss from the latter term will dominate. This issue is from the fact that the losses for positive and negative examples are separated and the contribution from positive examples will be overwhelmed by negative ones. One heuristic to handle the problem is emphasizing positive examples, which can change the weights for the corresponding losses. In this work, we aim to address the problem in a fundamental way. For brevity, we will omit the index of image (i.e., i) from the next subsection.

3.1. Ranking

To mitigate the challenge from the imbalance between classes, we consider to optimize the rank between positive and negative examples. Given a pair of positive and negative examples, an ideal ranking model can rank the positive example above the negative one with a large margin

$$\forall j_+, j_- \quad p_{j_+} - p_{j_-} \geq \gamma$$

where γ is a non-negative constant. Compared with the objective in Eqn. 1, the ranking model optimizes the relationship between individual positive and negative examples, which is well balanced.

The objective of ranking for a single image can be written as

$$\min_{\theta} \sum_{j_+}^{n_+} \sum_{j_-}^{n_-} \ell(p_{j_-} - p_{j_+} + \gamma) \quad (3)$$

where the hinge loss is applied as the loss function

$$\ell_{\text{hinge}}(z) = [z]_+ = \begin{cases} z & z > 0 \\ 0 & \text{o.w.} \end{cases}$$

The objective can be interpreted by the equivalent form

$$\begin{aligned} & \frac{1}{n_+ n_-} \sum_{j_+}^{n_+} \sum_{j_-}^{n_-} \ell(p_{j_-} - p_{j_+} + \gamma) \\ & = E_{j_+, j_-} [\ell(p_{j_-} - p_{j_+} + \gamma)] \end{aligned} \quad (4)$$

It demonstrates that the objective measures the expectation of the ranking loss on a randomly sampled pair.

The ranking loss addresses the inter-class imbalance issue by comparing the rank of each positive example to negative examples. However, it ignores a phenomenon in object detection, where the hardness of negative examples is also imbalanced. Besides, the ranking loss introduces a new challenge, that is, the vast number of pairs. We tackle them in the following subsection.

3.2. Distributional Ranking

As indicated in Eqn. 4, the ranking loss in Eqn. 3 punishes a mis-ranking for a uniformly sampled pair. In detection, most of negative examples can be easily ranked well, that is, a randomly sampled pair will not incur the ranking loss with high probability. Therefore, we consider to optimize the ranking boundary to avoid the trivial solution

$$\min_{\theta} \ell(\max_{j_-} \{p_{j_-}\} - \min_{j_+} \{p_{j_+}\} + \gamma) \quad (5)$$

If we can rank the positive example with the lowest score above the negative one with the highest confidence, the whole set of examples in an image are perfectly ranked. The pair in Eqn. 5 is referred as the worst-case scenario, which will incur the largest loss among all pairs. Compared with the conventional ranking loss, optimizing the loss from the worst-case scenario is much more efficient, which reduces the number of pairs from $n_+ n_-$ to 1. Moreover, it clearly eliminates the inter-class imbalance issue since only a single pair of positive and negative examples is required for each image. However, this formulation is very sensitive to the selected pair, which can result in the degraded detection model.

To improve the robustness, we first introduce the distribution of confidence scores for the positive and negative examples and obtain the expectation as

$$P_+ = \sum_{j_+}^{n_+} q_{j_+} p_{j_+}; \quad P_- = \sum_{j_-}^{n_-} q_{j_-} p_{j_-}$$

where $\mathbf{q}_+ \in \Delta$ and $\mathbf{q}_- \in \Delta$ denote the distributions over positive and negative examples, respectively. P_+ and

P_- represent the expected scores under the corresponding distribution. Δ is the simplex as $\Delta = \{q : \sum_j q_j = 1, \forall j, q_j \geq 0\}$. When \mathbf{q}_+ and \mathbf{q}_- are the uniform distribution, P_+ and P_- demonstrate the expectation from the original distribution.

With these definitions, the distribution corresponding to the worst-case scenario can be derived as

$$P_+ = \min_{\mathbf{q}_+ \in \Delta} \sum_{j_+}^{n_+} q_{j_+} p_{j_+}; \quad P_- = \max_{\mathbf{q}_- \in \Delta} \sum_{j_-}^{n_-} q_{j_-} p_{j_-}$$

We can rewrite the problem in Eqn. 5 in the equivalent form

$$\min_{\theta} \ell(P_- - P_+ + \gamma)$$

which can be considered as ranking the distributions between positive and negative examples in the worst-case scenario.

By investigating the new formulation, it is obvious that optimizing the worst-case scenario is not robust due to the fact that the domain of the generated distribution is unconstrained. Consequently, it will concentrate on a single example while ignoring the influence of the original distribution that contains massive information. Hence, we improve the robustness of the loss by regularizing the freedom of the derived distribution

$$\begin{aligned} P_- &= \max_{\mathbf{q}_- \in \Delta, \Omega(\mathbf{q}_- || \mathbf{o}_-) \leq \epsilon_-} \sum_{j_-}^{n_-} q_{j_-} p_{j_-} \\ -P_+ &= \max_{\mathbf{q}_+ \in \Delta, \Omega(\mathbf{q}_+ || \mathbf{o}_+) \leq \epsilon_+} \sum_{j_+}^{n_+} q_{j_+} (-p_{j_+}) \end{aligned}$$

where \mathbf{o}_+ , \mathbf{o}_- denote the original distributions for positive and negative examples, respectively. $\Omega(\cdot)$ is a regularizer for the diversity of the distribution to prevent the distribution from the trivial one-hot solution. It measures the similarity between the generated distribution and the original distribution, and some popular similarity function can be applied, e.g., L_p distance, Rényi entropy, Shannon entropy, etc. ϵ_- and ϵ_+ are constants to control the freedom of derived distributions.

To obtain the constrained distribution, we consider the subproblem

$$\begin{aligned} &\max_{\mathbf{q}_- \in \Delta} \sum_{j_-} q_{j_-} p_{j_-} \\ \text{s.t.} \quad &\Omega(\mathbf{q}_- || \mathbf{o}_-) \leq \epsilon_- \end{aligned}$$

According to the dual theory [1], given ϵ_- , we can find the parameter λ_- to obtain the optimal \mathbf{q}_- by solving the problem

$$\max_{\mathbf{q}_- \in \Delta} \sum_{j_-} q_{j_-} p_{j_-} - \lambda_- \Omega(\mathbf{q}_- || \mathbf{o}_-)$$

We observe that the former term is linear in \mathbf{q}_- . Hence, if $\Omega(\cdot)$ is convex in \mathbf{q}_- , the problem can be solved efficiently by first order algorithms [1]. In this work, we adopt KL-divergence as the regularizer and have the closed-form solution as follows

Proposition 1. *For the problem*

$$\max_{\mathbf{q}_- \in \Delta} \sum_{j_-} q_{j_-} p_{j_-} - \lambda_- \text{KL}(\mathbf{q}_- || \mathbf{o}_-)$$

we have the closed-form solution as

$$q_{j_-} = \frac{1}{Z_-} o_{j_-} \exp\left(\frac{p_{j_-}}{\lambda_-}\right); \quad Z_- = \sum_{j_-} o_{j_-} \exp\left(\frac{p_{j_-}}{\lambda_-}\right)$$

Proof. It can be proved directly from K.K.T. condition [1]. \square

For the distribution over positive examples, we have the similar result as

Proposition 2. *For the problem*

$$\max_{\mathbf{q}_+ \in \Delta} \sum_{j_+} q_{j_+} (-p_{j_+}) - \lambda_+ \text{KL}(\mathbf{q}_+ || \mathbf{o}_+)$$

we have the closed-form solution as

$$q_{j_+} = \frac{1}{Z_+} o_{j_+} \exp\left(\frac{-p_{j_+}}{\lambda_+}\right); \quad Z_+ = \sum_{j_+} o_{j_+} \exp\left(\frac{-p_{j_+}}{\lambda_+}\right)$$

Remark 1 These Propositions show that the harder the example, the larger the weight of the example. Besides, the weight is image-dependent and will be affected by other examples in the same image.

The original distributions (i.e., \mathbf{o}_- and \mathbf{o}_+) can also influence the derived distributions by weighting each candidate. Therefore, the prior knowledge about the problem can be encoded into the original distributions, which makes generating new distributions more flexible. Here we take \mathbf{o}_- as an example to illustrate different distributions.

- **Uniform distribution:** It means that $\forall j, o_{j_-} = 1/n_-$. With the constant value, the closed-form solution can be simplified as $q_{j_-} = \frac{1}{Z_-} \exp\left(\frac{p_{j_-}}{\lambda_-}\right)$; $Z_- = \sum_{j_-} \exp\left(\frac{p_{j_-}}{\lambda_-}\right)$
- **Hard negative mining:** In this scenario, we assume $\forall j, o_{j_-} \in \{0, 1/\hat{n}_-\}$, where \hat{n}_- denotes the number of non-zero elements in \mathbf{o}_- . According to Proposition 1, only candidates selected by \mathbf{o}_- will be accumulated to derive the new distribution. Therefore, our formulation can incorporate with hard negative mining by setting the weights in \mathbf{o}_- appropriately.

To keep the loss function simple, we adopt the uniform distribution in this work. Fig. 2 illustrates the changing of the distribution with the proposed strategy. The derived distribution approaches the distribution corresponding to the worst-case scenario when decreasing λ . Note that the original distributions in Fig. 2 (a) and Fig. 2 (b) have the same mean but different variance. For the distribution with the small variance as in Fig. 2 (a), we can observe that the regularizer λ should be small to change the distribution effectively. When the distribution has the large variance, Fig. 2 (b) shows that a large λ is sufficient to change the shape of the distribution dramatically. Considering that the distributions of positive and negative examples have different variances, Fig. 2 implies that different weights for the regularizers should be assigned.




Figure 2. Illustration of the drifting in the distribution. We randomly sample $1e7$ points from a Gaussian distribution with different variances to mimic scores of anchors. We change the weights of examples according to the proposed strategy as in Proposition 1 and then plot the curves of different probability density functions (PDF) when varying λ .

With the closed-form solutions of distributions, the expectation of distributions can be computed as

$$\hat{P}_- = \sum_{j^-} q_{j^-} p_{j^-} = \sum_{j^-} \frac{1}{Z_-} \exp\left(\frac{p_{j^-}}{\lambda_-}\right) p_{j^-} \quad (6)$$

$$\hat{P}_+ = \sum_{j^+} q_{j^+} p_{j^+} = \sum_{j^+} \frac{1}{Z_+} \exp\left(\frac{-p_{j^+}}{\lambda_+}\right) p_{j^+}$$

Finally, smoothness is crucial for the convergence of non-convex optimization [7]. So we apply the smooth approximation instead of the original hinge loss as the loss function for pairs. The popular substitutes to the hinge loss include quadratic loss and logistic loss

$$\ell_{\text{quad}}(z) = \begin{cases} z & z \geq \rho \\ \frac{(z+\rho)^2}{4\rho} & -\rho < z < \rho \\ 0 & z \leq -\rho \end{cases} \quad (7)$$

$$\ell_{\text{logistic}}(z) = \frac{1}{L} \log(1 + \exp(Lz)) \quad (8)$$

where ρ and L control the approximation error of the function. The larger the L is, the closer to the hinge loss the

approximation is. ρ works in an opposite direction. Fig. 3 compares the hinge loss to its smooth variants. Explicitly, these functions share the similar shape and we adopt the logistic loss in this work.




Figure 3. Illustration of the hinge loss and its smooth variants.

Incorporating all of these components, our distributional ranking loss can be defined as

$$\min_{\theta} \mathcal{L}_{\text{DR}}(\theta) = \sum_i^N \ell_{\text{logistic}}(\hat{P}_{i,-} - \hat{P}_{i,+} + \gamma) \quad (9)$$

where $\hat{P}_{i,-}$ and $\hat{P}_{i,+}$ are given in Eqn. 6 and $\ell_{\text{logistic}}(\cdot)$ is in Eqn. 8. If there is no positive examples in an image, we will let $\hat{P}_{i,+} = 1$. Compared with the conventional ranking loss, we rank the expectations of two distributions. It shrinks the number of pairs to 1 that leads to the efficient optimization.

The gradient of the objective in Eqn. 9 is easy to compute. The detailed calculation of the gradient can be found in the supplementary.

If optimizing the DR loss by the standard SGD with mini-batch as $\theta_{t+1} = \theta_t - \eta \frac{1}{m} \sum_{s=1}^m \nabla \ell_t^s$, we can show that it can converge as in the following theorem. The norm of the gradient is applied to measure the convergence, which is a standard criterion for non-convex optimization [7]. The detailed proof is cast to the supplementary.

Theorem 1. Let θ_t denote the model obtained from the t -th iteration with SGD optimizer and the size of mini-batch is m . If we assume the objective \mathcal{L} in Eqn. 9 is μ -smoothness and the variance of the gradient is bounded as $\forall s, \|\nabla \ell_t^s - \nabla \mathcal{L}_t\|_F \leq \delta$, when setting the learning rate as $\eta = \frac{\sqrt{2m\mathcal{L}(\theta_0)}}{\delta\sqrt{\mu T}}$ and $\eta \leq \frac{1}{\mu}$, we have

$$\frac{1}{T} \sum_t \|\nabla \mathcal{L}(\theta_t)\|_F^2 \leq \frac{2\delta\sqrt{2\mu}}{\sqrt{mT\mathcal{L}(\theta_0)}}$$

Remark 2 Theorem 1 implies that the learning rate depends on the mini-batch size and the number of iterations as $\eta = \mathcal{O}(\sqrt{m/T})$ and the convergence rate is $\mathcal{O}(1/\sqrt{mT})$, where mT/N is the number of training epochs.

We can obtain a scaling strategy for the learning rate. Let η_0, m_0 and T_0 denote a default configuration for training. If we change the mini-batch size as $m' = m_0/\alpha$, where

α is a non-negative constant, and keep the same number of epochs for training (i.e., $T' = \alpha T_0$), the convergence rate remains the same. However, the learning rate becomes $\eta' = \mathcal{O}(\sqrt{m'/T'}) = \eta_0/\alpha$. It shows that to obtain the same performance with a different mini-batch size, we have to rescale the learning rate with a corresponding factor, which is consistent with the observation in [10]. Besides, the learning rate should be no larger than $1/\mu$, which means that the scaling strategy is inapplicable when the mini-batch size is too large.

3.3. Recover Classification from Ranking

Detection is to identify foreground objects from background. Therefore, the results from ranking have to be converted to classification. A straightforward way is to set a threshold for all ranking scores. However, the range of ranking scores from different images can vary due to the image-dependent mechanism, and should be calibrated for classification. We investigate the bound for the ranking scores of positive and negative examples as follows.

Theorem 2. *When optimizing the ranking problem as*

$$\forall j_+, j_-, \quad p_{j_+} - p_{j_-} \geq \gamma$$

it implies

$$\forall j_+, \quad p_{j_+} > \gamma; \quad \forall j_-, \quad p_{j_-} \leq 1 - \gamma$$

Therefore, we can recover the standard classification criterion by setting a large margin.

Corollary 1. *If setting the margin of ranking as $\gamma = 0.5$, we can recover the classification criterion for ranking scores*

$$\forall j_+, \quad p_{j_+} > 0.5; \quad \forall j_-, \quad p_{j_-} \leq 0.5$$

With these appropriate settings, our final objective for detection can be summarized as

$$\min \sum_i^N \tau \ell_{\text{DR}}^i + \ell_{\text{Reg}}^i$$

where ℓ_{Reg} is the original regression loss in RetinaNet and we keep it unchanged. τ is the parameter for balancing the weights between classification and regression. We fix it as $\tau = 4$ in the experiments.

4. Experiments

4.1. Implementation Details

We evaluate the proposed DR loss on COCO data set [17], which contains about 118k images for training, 5k images for validation, and 40k images for test. To focus on the comparison of loss functions, we employ the structure of

RetinaNet [16] as the backbone and only substitute the corresponding focal loss. For a fair comparison, we implement our algorithm in a public codebase¹. Besides, we train the model with the same configuration as RetinaNet. Specifically, the model is learned with SGD on 8 GPUs and the mini-batch size is set as 16 where each GPU can hold 2 images at each iteration. Most of experiments are trained with 90k iterations that is denoted as “1×”. The initial learning rate is 0.01 and is decayed by a factor of 10 after 60k iterations and then 80k iterations. For anchor density, we apply the same setting as in [16], where each location has 3 scales and 3 aspect ratios. The standard COCO evaluation criterion is used to compare the performance of different methods.

4.2. Parameters in DR Loss

From the definition in Eqn. 9, DR loss has three parameters λ_+ , λ_- and L to be tuned. λ_+ and λ_- regularize the distribution of scores for positive and negative examples, respectively. L controls the smoothness of the loss function. The margin γ is fixed as 0.5 according to Corollary 1. Compared with the focal loss [16], DR loss has one more parameter. However, RetinaNet lacks optimizing the relationship between positive and negative distributions, and it has an additional parameter to initialize the output probability of the classifier (i.e., 0.01) to fit the distribution of background. In contrast, we initialize the probability of the sigmoid function at 0.5, which is the default threshold for binary classification scenario without any prior knowledge. It verifies that the proposed DR loss can handle the imbalance problem better. Consequently, DR loss roughly has the same number of parameters as that in focal loss.

We will have the ablation study on these parameters to illustrate the influence in the next subsections. Note that RetinaNet applies Feature Pyramid Network (FPN) [15] to obtain multiple scale features. To compute DR loss in one image, we collect anchors from multiple pyramid levels and obtain a single distribution for positive and negative anchors, respectively.

4.3. Effect of Parameters

We conduct ablation experiments to evaluate the effect of multiple parameters on the *minival* set. All experiments in this subsection are implemented with a single image scale of 800 for training and test. ResNet-50 [12] is applied as the backbone for comparison. Only horizontal flipping is adopted as the data augmentation in this subsection.

Effect of λ_+ and λ_- : First, we evaluate the effect of λ_+ and λ_- in Eqn. 6. These parameters constrain the freedom of the derived distributions. As illustrated in Fig. 2,

¹<https://github.com/facebookresearch/maskrcnn-benchmark>

variances of distributions will have the impact on selecting appropriate weights for regularizers. We investigate the variance of positive and negative anchors, and observe that the standard deviation of positive anchors is about 10 times larger than that of negative ones. So we roughly set $\lambda_+ = 1$ and $\lambda_- = 0.1$ and fine-tune them as $\lambda_+ = 1/\log(h_+)$ and $\lambda_- = 0.1/\log(h_-)$. It is easy to show that this strategy is equivalent to fixing λ_+ and λ_- as 1 and 0.1, and changing the base in the definition of the KL-divergence as $\text{KL}(\mathbf{q}||\mathbf{o}) = \sum_j q_j \log_h \frac{q_j}{o_j}$.

We vary h_+ and h_- and summarize the results in Table 1. First, we observe that the default setting with $\lambda_+ = 1$ and $\lambda_- = 0.1$ can outperform focal loss by 1%, which demonstrates the effectiveness of the proposed DR loss. Second, the performance of our loss is quite stable in a reasonable range. Finally, the distribution of positive anchors is more sensitive to a small λ_+ , which is consistent with the illustration in Fig. 2. We keep the best settings in the following experiments.

h_+	h_-	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
e	e	37.1	56.1	39.4	19.7	40.9	50.1
e	3.5	37.4	56.0	40.0	20.8	41.2	50.5
e	5.5	37.2	55.7	40.0	19.6	41.2	50.4
e	20	36.6	54.7	39.4	19.6	40.5	50.3
5.5	3.5	36.7	55.2	39.4	19.9	40.4	50.0
20	3.5	35.6	54.2	38.1	19.2	39.7	48.3

Table 1. Comparison of λ_+ and λ_- as in Eqn. 6. Note that we tune the parameters in the form of $\lambda_+ = 1/\log(h_+)$ and $\lambda_- = 0.1/\log(h_-)$. We adopt $1 \times$ iterations and ResNet-50 as the backbone in training. Performance on the *minival* is reported for the ablation study.

Effect of Smoothness: L controls the smoothness of the loss function in Eqn. 8. We compare the model with different L 's in Table 2. We also include the results for the quadratic loss function in Eqn. 7 with different ρ 's for comparison. The original hinge loss is denoted as ‘‘Hinge’’. First, all smooth loss functions outperform hinge loss. It confirms that smoothness is important for non-convex optimization. Second, the smooth variants of hinge loss surpass focal loss with a significant margin. It is because that DR loss leverages the information from an image rather than an anchor, which can handle the imbalance issue better. Since quadratic loss and logistic loss have the similar performance, we adopt the logistic loss with $L = 6$ in the rest experiments.

Effect of Pairing Strategy: In DR loss, we propose to rank a single pair consisting of expectations from positive and negative distributions. To evaluate the pairing strategy, we compare it to different strategies for ranking. Specifically, we denote optimizing all pairs in Eqn. 3 as ‘‘All’’,

Loss	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
Focal	36.1	55.0	38.7	19.5	39.5	49.0
Hinge	35.8	54.0	38.3	19.3	39.5	47.6
$\rho = 0.2$	36.9	55.5	39.5	21.1	40.7	49.6
$\rho = 0.5$	37.2	56.0	39.8	21.1	41.1	50.4
$L = 4$	37.2	55.9	39.9	20.3	41.0	50.3
$L = 6$	37.4	56.0	40.0	20.8	41.2	50.5
$L = 8$	37.1	55.7	39.7	19.5	41.2	50.5
$L = 10$	36.8	55.4	39.4	20.0	40.7	50.0

Table 2. Comparison of different loss functions. ρ and L are from Eqn. 7 and Eqn. 8, respectively.

Pair	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
All	12.9	23.0	12.6	8.8	16.7	15.0
NegOnly	37.0	55.5	39.5	19.8	40.7	50.5
DR	37.4	56.0	40.0	20.8	41.2	50.5

Table 3. Comparison of different pairing strategies.

which is corresponding to the standard ranking problem. We also include a variant of DR loss as ‘‘NegOnly’’ that pushes distributions for negative anchors only. The objective of NegOnly on a single image can be written as

$$\min_{\theta} \frac{1}{n_+} \sum_{j_+} \ell_{\text{logistic}}(\hat{P}_- - p_{j_+} + \gamma)$$

The result of optimizing the worst-case scenario in Eqn. 5 is not included since training with that fails to obtain the meaningful result.

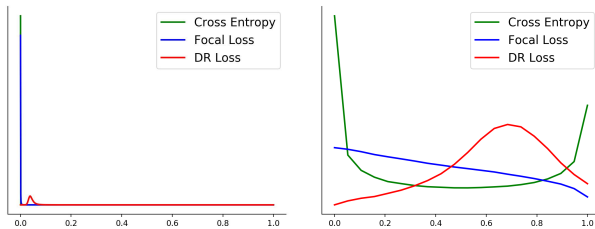
The comparison is summarized in Table 3. As illustrated in Section 3.1, the conventional ranking algorithm suffers from the intra-class imbalance in the hardness of negative anchors, which results in the poor performance for detection. By mitigating this issue with the proposed strategy, NegOnly can outperform focal loss. It confirms that handling the imbalance in the negative anchors is important and the proposed strategy can serve the purpose well. Finally, we observe that a tailored distribution for positive anchors can further improve the performance as in DR loss.

Effect of DR Loss: To illustrate the effectiveness of DR loss, we collect the confidence scores of anchors from all images in *minival* and compare the empirical probability density in Fig. 4. We include the results from cross entropy loss and focal loss in the comparison.

First, we observe that most of examples have an extremely low confidence after minimizing cross entropy loss. It is because the number of negative examples overwhelms that of positive ones and it will classify most of examples to be negative to obtain a small loss as demonstrated in Eqn. 2. Second, focal loss is better than cross entropy loss by improving the distribution of positive anchors. However, the expectation of the foreground distribution is still close to

Methods	Backbone	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
<i>two-stage detectors</i>							
Faster R-CNN+++ [12]	ResNet-101-C4	34.9	55.7	37.4	15.6	38.7	50.9
Faster R-CNN w FPN [15]	ResNet-101-FPN	36.2	59.1	39.0	18.2	39.0	48.2
Deformable R-FCN [4]	Aligned-Inception-ResNet	37.5	58.0	40.8	19.4	40.1	52.5
Mask R-CNN [11]	ResNet-101-FPN	38.2	60.3	41.7	20.1	41.1	50.2
<i>one-stage detectors</i>							
YOLOv2 [21]	DarkNet-19	21.6	44.0	19.2	5.0	22.4	35.5
SSD513 [18]	ResNet-101-SSD	31.2	50.4	33.3	10.2	34.5	49.8
AP-Loss [3]	ResNet-101-FPN	37.4	58.6	40.5	17.3	40.8	51.9
RetinaNet [16]	ResNet-101-FPN	39.1	59.1	42.3	21.8	42.7	50.2
RetinaNet [16]	ResNeXt-32x8d-101-FPN	40.8	61.1	44.1	24.1	44.2	51.2
CornerNet [14]	Hourglass-104	40.5	56.5	43.1	19.4	42.7	53.9
FSAF [29]	ResNet-101-FPN	40.9	61.5	44.0	24.0	44.2	51.3
FCOS [25]	ResNet-101-FPN	41.5	60.7	45.0	24.4	44.8	51.6
FCOS [25]	ResNeXt-32x8d-101-FPN	42.7	62.2	46.1	26.0	45.6	52.6
Dr. Retina	ResNet-101-FPN	41.7	60.9	44.8	23.5	44.9	53.1
Dr. Retina	ResNeXt-32x8d-101-FPN	43.1	62.8	46.4	25.6	46.2	54.0
Dr. Retina (<i>multi-scale test</i>)	ResNet-101-FPN	43.4	62.1	47.0	26.7	46.1	55.0
Dr. Retina (<i>multi-scale test</i>)	ResNeXt-32x8d-101-FPN	44.7	63.8	48.7	28.2	47.4	56.2

Table 4. Comparison with the state-of-the-art methods on COCO *test-dev* set.



(a) Negative Anchors Distribution (b) Positive Anchors Distribution
Figure 4. Illustration of empirical PDF of distributions that are computed from images in the *minival*.

that of background, and it interprets the fact that focal loss has to initialize the probability of the classifier to be small (i.e., 0.01). Compared to cross entropy and focal loss, DR loss improves the foreground distribution significantly. By optimizing our ranking loss with a large margin, the expectation of the positive anchors is larger than 0.5 while that of background is smaller than 0.1. It confirms that DR loss can address the imbalance between classes well. Besides, the hardness of negative anchors with DR loss is more balanced than that with cross entropy or focal loss. It verifies that with the image-dependent mechanism, DR loss can handle the intra-class imbalance in background examples and focus on the hard negative examples appropriately. More analysis can be found in the supplementary.

4.4. Comparison with State-of-the-Art

We denote RetinaNet with DR loss as “Dr. Retina” and compare it to the state-of-the-art detectors on COCO *test-dev* set. We follow the setting in [16] to increase the number

of training iterations to $2\times$, which contains $180k$ iterations, and applies scale jitter in $[640, 800]$ as the additional data augmentation for training. Note that we still use a single image scale and a single crop for test as above. Table 4 summarizes the comparison for Dr. Retina. With ResNet-101 as the backbone, we can observe that Dr. Retina improves mAP from 39.1% to 41.7%. It illustrates that DR loss can explore the imbalance issue in detection more sufficiently than focal loss. Equipped with ResNeXt-32x8d-101 [28] and $1.5\times$ iterations (i.e., $135k$ iterations), the performance of Dr. Retina can achieve 43.1% as a one-stage detector on COCO detection task without bells and whistles. Note that we only replace focal loss with DR loss to obtain the significant gain, which implies that DR loss can be a good substitute of focal loss. Finally, the multi-scale test with scales from $\{400, 500, 600, 700, 800, 900, 1000, 1100, 1200\}$ can further improve the performance as expected.

5. Conclusion

In this work, we introduce the distributional ranking loss to address the imbalance challenge in one-stage object detection. We first convert the original classification problem to a ranking problem, which balances the positive and negative classes. After that, we propose to push the original distributions to the decision boundary and rank the expectations of derived distributions in lieu of original examples to focus on the hard examples, which balances the hardness of background examples. Experiments on COCO verify the effectiveness of the proposed loss function.

References

- [1] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: delving into high quality object detection. In *CVPR*, pages 6154–6162, 2018.
- [3] 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 *CVPR*, pages 5119–5127, 2019.
- [4] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, pages 764–773, 2017.
- [5] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, pages 886–893, 2005.
- [6] Pedro F. Felzenszwalb, David A. McAllester, and Deva Ramanan. A discriminatively trained, multi-scale, deformable part model. In *CVPR*, 2008.
- [7] Saeed Ghadimi and Guanghui Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [8] Ross B. Girshick. Fast R-CNN. In *ICCV*, pages 1440–1448, 2015.
- [9] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014.
- [10] Priya Goyal, Piotr Dollár, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch SGD: training imagenet in 1 hour. *CoRR*, abs/1706.02677, 2017.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, pages 2980–2988, 2017.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [13] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1106–1114, 2012.
- [14] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, pages 765–781, 2018.
- [15] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, 2017.
- [16] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2999–3007, 2017.
- [17] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, pages 740–755, 2014.
- [18] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV*, pages 21–37, 2016.
- [19] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [20] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [21] Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. In *CVPR*, pages 6517–6525, 2017.
- [22] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015.
- [23] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *ICLR*, 2014.
- [24] Abhinav Shrivastava, Abhinav Gupta, and Ross B. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, pages 761–769, 2016.
- [25] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [26] Jasper R. R. Uijlings, Koen E. A. van de Sande, Theo Gevers, and Arnold W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, 104(2):154–171, 2013.
- [27] Paul A. Viola and Michael J. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, pages 511–518, 2001.
- [28] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 5987–5995, 2017.
- [29] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *CVPR*, pages 840–849, 2019.