Is Heuristic Sampling Necessary in Training Deep Object Detectors?

Joya Chen[®], *Member, IEEE*, Dong Liu[®], *Senior Member, IEEE*, Tong Xu[®], *Member, IEEE*, Shiwei Wu, Yifei Cheng[®], and Enhong Chen[®], *Senior Member, IEEE*

Abstract—To train accurate deep object detectors under the extreme foreground-background imbalance, heuristic sampling methods are always necessary, which either re-sample a subset of all training samples (hard sampling methods, e.g. biased sampling, OHEM), or use all training samples but re-weight them discriminatively (soft sampling methods, e.g. Focal Loss, GHM). In this paper, we challenge the necessity of such hard/soft sampling methods for training accurate deep object detectors. While previous studies have shown that training detectors without heuristic sampling methods would significantly degrade accuracy, we reveal that this degradation comes from an unreasonable classification gradient magnitude caused by the imbalance, rather than a lack of re-sampling/re-weighting. Motivated by our discovery, we propose a simple yet effective Sampling-Free mechanism to achieve a reasonable classification gradient magnitude by initialization and loss scaling. Unlike heuristic sampling methods with multiple hyperparameters, our Sampling-Free mechanism is fully data diagnostic, without laborious hyperparameters searching. We verify the effectiveness of our method in training anchor-based and anchor-free object detectors, where our method always achieves higher detection accuracy than heuristic sampling methods on COCO and PASCAL VOC datasets. Our Sampling-Free mechanism provides a new perspective to address the foreground-background imbalance. Our code is released at https://github.com/ChenJoya/sampling-free.

Index Terms—Object detection, foreground-background imbalance, heuristic sampling, sampling-free.

Manuscript received December 4, 2020; revised May 9, 2021 and June 29, 2021; accepted July 20, 2021. Date of publication August 31, 2021; date of current version October 12, 2021. This work was supported in part by the National Key Research and Development Program of China under Grant 2018YFB1402600, in part by the National Natural Science Foundation of China under Contract 61727809 and Contract 62072423, and in part by the Graphics Processing Unit (GPU) Cluster built by Anhui Province Key Laboratory of Big Data Analysis and Application, University of Science and Technology of China. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Zhenzhong Chen. (*Corresponding author: Enhong Chen.*)

Joya Chen, Tong Xu, and Enhong Chen are with Anhui Province Key Laboratory of Big Data Analysis and Application, University of Science and Technology of China (USTC), Hefei 230052, China, and also with the School of Computer Science and Technology, University of Science and Technology of China (USTC), Hefei 230052, China (e-mail: chenjoya@mail.ustc.edu.cn; tongxu@ustc.edu.cn; cheneh@ustc.edu.cn).

Dong Liu is with the Department of Electronic Engineering and Information Science, University of Science and Technology of China (USTC), Hefei 230052, China (e-mail: dongeliu@ustc.edu.cn).

Shiwei Wu and Yifei Cheng are with Anhui Province Key Laboratory of Big Data Analysis and Application, University of Science and Technology of China (USTC), Hefei 230052, China, and also with the School of Data Science, University of Science and Technology of China (USTC), Hefei 230052, China (e-mail: dwustc@mail.ustc.edu.cn; chengyif@mail.ustc.edu.cn).

Digital Object Identifier 10.1109/TIP.2021.3106802

I. INTRODUCTION

WITH the development of deep learning [1], [2], recent years have witnessed remarkable advancement in object detection [3]. Among them, representative successes include two-stage R-CNN detectors [4]–[15]: their first stage uses a region proposal network (RPN [4]) to generate some candidates from dense, predefined bounding-boxes (*i.e.* anchors), then the second stage uses a region-of-interest subnetwork (RoI-subnet) for object classification and localization. To pursue higher efficiency, one-stage approaches [16]–[23] directly recognize objects from dense anchors rather than generating candidate proposals. Both two-stage and one-stage detectors adopt the anchoring scheme, where massive anchors ($\sim 10^5$) are uniformly sampled over an image.

Nevertheless, when training these anchor-based detectors, only a few anchors ($\sim 10^2$) that highly overlap with objects will be assigned to foreground samples, which always results in an extreme imbalance between foreground and background (*i.e. fg-bg* imbalance) within the anchors. In previous studies [17], [24], such imbalance may impede the training from convergence, as well as limit the detection accuracy. More recently, anchor-free object detectors [11], [14], [25]–[33] have gained much attention due to the replacement of anchors by points (*e.g.* corner/center points), but they still suffer from the similar imbalance within the points.

To address the *fg-bg* imbalance, several heuristic methods have been proposed to train deep object detectors in recent years. These methods can be divided into two categories. The first category re-samples a subset of training samples, *e.g.* biased sampling [4], online hard example mining [34] (OHEM), IoU-balanced sampling [12]. The second category re-weights training samples discriminatively, *i.e.* assigns different weights to different training samples, like Focal Loss [17], gradient harmonizing mechanism [24] (GHM), prime sample attention mechanism [35] (PISA). According to [36], these two categories can be named "hard sampling methods" and "soft sampling methods," respectively. We also use the term "heuristic sampling methods" to refer to them in the following.

Although deep object detectors are always equipped with heuristic sampling methods, it is still very difficult to design a suitable hard/soft sampling strategy. Each heuristic sampling methods have a different re-sampling/re-weighting method as it is unknown which sample and what weighting value is better. For example, in GHM [24], the authors hold the opinion that the optimal distribution of gradient is hard to define

1941-0042 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.



Fig. 1. This figure illustrates the differences between heuristic sampling methods and the Sampling-Free mechanism in the treatment of training samples. Here we use the bounding-boxes to denote training samples (*e.g.* anchors) in object detection. (a) Hard sampling (*e.g.* biased sampling [4], OHEM [34], IoU-balanced sampling [12]) re-samples a subset of training samples; (b) Soft sampling (*e.g.* Focal Loss [17], GHM [24], PISA [35]) uses all training samples but focuses on some of them by re-weighting. For instance, thicker boxes in (b) denote training samples with higher weights. (c) Sampling-Free equally uses all training samples.

and requires further research. Moreover, heuristic sampling methods always introduce multiple hyperparameters, which requires laborious searching.

Can we discard heuristic sampling methods when training deep object detectors? In the past, it was demonstrated [17], [24] that the detector without heuristic sampling methods will suffer from the extreme fg-bg imbalance, which would trail the detector with heuristic sampling methods about 20% detection accuracy. Some methods [23], [33], [37]–[40] adaptively define foreground/background labels to anchors/points, but they still rely on heuristic sampling methods to address the fg-bg imbalance. Other ranking-based methods [41]–[43] try to avoid the fg-bg imbalance by transforming the classification task into the ranking task, but they select pairs of N samples to train, thus have much greater computational cost than heuristic sampling methods ($O(N^2)$ vs. O(N)). It seems impossible to make a cost-free replacement of heuristic sampling methods when training a deep object detector.

In this paper, we discover that a reasonable classification gradient magnitude is the key to address the fg-bgimbalance, rather than hard/soft sampling. Motivated by this, we propose a simple yet effective *Sampling*-*Free* mechanism that adaptively controls the classification gradient magnitude by initialization and loss scaling techniques, which enables discarding heuristic sampling methods but achieves better accuracy. Specifically, at the start of the training, the optimal bias initialization is used to reduce the excessive classification gradient magnitude caused by fg-bg imbalance. During the training process, we leverage the bounding-box regression loss to adjust the classification loss, to achieve an adaptive adjustment for the classification gradient magnitude. As shown in Fig. 1, unlike heuristic sampling methods, our method treats all training samples equally, without any hyperparameters introduced.

Experimental results on COCO [44] and PASCAL VOC [45] datasets have demonstrated that our method is effective for both anchor-based and anchor-free object detectors, which always achieves higher detection accuracy than heuristic sampling methods. By replacing Focal Loss [17] with Sampling-Free in the adaptive label assignment strategy [39], we obtain a state-of-the-art 49.6 AP on COCO test-dev, without bells and whistles. Sampling-Free is also generalized for the instance segmentation task, which helps Mask R-CNN to obtain better segmentation accuracy. Moreover, no hyperparameter is introduced in our method. Our Sampling-Free mechanism provides a new perspective to address the fg-bg imbalance.

Our main contributions are as follows:

• For the first time, we discover what prevents detectors without heuristic sampling methods from achieving good accuracy — the unreasonable classification gradient magnitude under the fg-bg imbalance, rather than the lack of re-sampling/re-weighting on training samples.

• We propose a novel Sampling-Free mechanism that enables training deep object detectors without heuristic sampling methods. It adaptively controls the classification gradient magnitude by initialization and loss scaling, which is easy to implement and introduces no hyperparameters.

• Collaborating with Sampling-Free mechanism, it is feasible to train deep object detectors without any hard/soft sampling methods and achieve better results on COCO and PASCAL VOC benchmarks.

II. RELATED WORK

Classical object detectors [46], [47] usually rely on handcrafted feature extractor, which is hard to design. With the development of deep learning [1], [2], deep object detectors quickly come to dominate object detection. In this section, we introduce the development of deep object detectors, then introduce the concept, cause, and solution of the fg-bg imbalance. Finally, we discuss the relations and differences between our work and previous works.

A. Deep Object Detection

Among deep object detectors, the anchor-based approach is the most popular approach, which tiles massive default bounding-boxes (*i.e.* anchors) on an image to cover objects. There are mainly two types of anchor-based approaches:

1) Two-Stage Anchor-Based: It is popularized by Faster R-CNN [4], which firstly generates some candidates from massive anchors by region proposal network (RPN [4]), then determines the accurate location and object category by a subnetwork (RoI-subnet [48], [49]). A large number of Faster R-CNN variants [5]–[15] appear over the years, yielding a large improvement in detection accuracy.

2) One-Stage Anchor-Based: It is popularized by SSD [50], which performs much faster than two-stage detectors due to the elimination of the RPN, but usually achieves lower accuracy than two-stage detectors. A series of advances [16]–[23], [40] in recent years promote one-stage anchor-based detector to be more accurate.

In recent years, a large number of anchor-free approaches are proposed, which detect objects by points or regions rather than anchors. Most anchor-free object detectors follow the one-stage detection pipeline but eliminate the usage and hyperparameters of "anchor boxes", which shows better simplicity. Although the differences among anchor-free detectors are much smaller than the differences between one-stage and twostage anchor-based detectors, there are some subtle differences in the definition of training samples among different anchorfree detectors. Specifically, some of the anchor-free approaches detect objects by generating bounding-boxes according to predefined or self-learned keypoints. The early attempt of predefined points is CornerNet [25], which uses the top-left corner and bottom-right corner to represent objects. After that, researchers use various pre-defined points to represent objects, such as extreme points [26], grid points [11], center points [27], [28], and self-learned points [29], [30]. Others [14], [31]–[33] try to learn the position of the object center, and then regress the distances from the center to the four sides of the object bounding-box for detection. The most popular center-based anchor-free detector is FCOS [31], which regards all the locations around the center of objects as foreground examples. GA-RPN [14] successfully designs an anchor-free RPN in the two-stage pipeline.

B. Foreground-Background Imbalance Problem

Training a deep object detector involves two tasks: classification and localization. For classification, the number of background examples is much larger than foreground examples, which is known as the fg-bg imbalance [17]. We introduce it in the different label assignment strategy:

1) fg-bg Imbalance in Fixed Label Assignment: In the fixed label assignment [4], [25], [31], [50], there is a pre-defined rule to assign a training sample to a fg/bg example. For instance, the anchor-based approach usually considers anchors that have large intersection-over-union (IoU) with ground-truths as foreground examples (e.g. IoU>0.5). The anchor-free approach usually regards points around the center as foreground examples. However, the total number of anchors/points is always huge, which may be $10^3 \sim 10^4$ times more than foreground anchors/points. Therefore, during training, the fg-bg imbalance inevitably occurs in the classification task.

2) fg-bg Imbalance in Adaptive Label Assignment: Recently, several adaptive label assignment methods [23], [33], [37]–[40] are proposed to overcome the limitations of fixed label assignment. However, they still suffer from extreme fg-bgimbalance. For example, FreeAnchor [23] claims that it faces an even more serious sample imbalance than RetinaNet [17]. These methods still rely on Focal Loss [17] to address the fg-bg imbalance in the classification task.

C. Solutions for Foreground-Background Imbalance

As we can see, the fg-bg imbalance always exists in training deep object detectors, which impedes deep object detectors from achieving higher accuracy as reported in [17], [24]. In previous works, there are three solutions:

1) Sampling Methods: It is the most common solution for fg-bg imbalance, which has two groups [36] — hard sampling and soft sampling. The hard sampling method re-samples a set of training samples by some strategies. For example, biased sampling [4] randomly samples 256 examples with 1:1 fg-to-bg ratio during training RPN. OHEM [34] and IoU-balanced sampling [12] selects training samples according to loss and IoU values, respectively. The objectness [16], [19], [20] modules, generative methods [51], [52] can also be regarded as hard sampling methods. Soft sampling re-weights training samples discriminatively by some strategies. Focal Loss [17] dynamically assigns higher weight to the hard training samples (*i.e.* with high loss value). Similar to Focal Loss, GHM [24] suppresses gradients originating from easy and very hard training samples (i.e. with low loss value). PISA [35] re-weights training samples according to the IoU between training samples and ground-truths.

2) Classification to Ranking: As the fg-bg imbalance always exists in classification task in deep object detectors, AP Loss [41] and DR Loss [42] propose to convert the classification task into ranking task. These methods train a pair of samples rather than independent sample. Specifically, the predicted score of one training sample is transformed into the difference between the predicted scores of two training samples. These methods are also quite in line with the detection evaluation metric (average precision, AP).

To date, almost all deep object detectors are equipped with sampling methods during training. The ranking task, however, trains pairs of N examples thus has $\mathcal{O}(N^2)$ computational cost, which is much higher than $\mathcal{O}(N)$ cost of the classification

TABLE I Comparison of Different Classification Loss Functions in Object Detection. This Table Comes From "Gradient Harmonized Single-Stage Detector" [24]

| method | AP | AP ₅₀ | AP ₇₅ | AP_S | AP_M | AP_L |
|-----------------|------|------------------|------------------|--------|--------|--------|
| CE loss | 28.6 | 43.3 | 30.7 | 11.4 | 30.7 | 40.7 |
| OHEM [34] | 31.1 | 47.2 | 33.2 | - | - | - |
| Focal Loss [17] | 35.6 | 55.6 | 38.2 | 19.1 | 39.2 | 46.3 |
| GHM-C [24] | 35.8 | 55.5 | 38.1 | 19.6 | 39.6 | 46.7 |

task. Although sampling methods are popular, the re-sampling/ re-weighting strategy is hard to design, and both sampling methods and ranking-based methods require laborious hyperparameter tuning. Our work overcomes these shortcomings, which discards sampling methods in the classification task without any hyperparameters introduced.

III. INVESTIGATION

As shown in Table I, it was believed [17], [24] that under the fg-bg imbalance, training a deep object detector without heuristic sampling methods will lead to a nearly 20% decrease in the detection accuracy. In this section, we will investigate why this is the case. It will be revealed that the decrease should be attributed to the unreasonable gradient magnitude, rather than the lack of re-sampling/re-weighting on training samples.

In the following, we first mathematically introduce fg-bg imbalance and sampling methods, then theoretically analyze their impact on the training process. Finally, we experimentally demonstrate that the gradient magnitude is the key factor affecting the detector accuracy.

A. Concepts

1) Foreground-Background Imbalance: In general, deep object detectors tend to generate numerous samples to cover as many objects as possible. Although there are various label assignment strategies [4], [17], [23], [38], [39] to define foreground and background samples, the imbalance between foreground and background samples will be inevitably caused due to the rarity of objects and the majority of samples, namely foreground-background (fg-bg) imbalance. In other words, the number of foreground samples N^f is much smaller than that of background samples N^b (*i.e.* $N^f \ll N^b$). Unlike the common class imbalance caused by the biased dataset, the fg-bg imbalance is more likely to be introduced by the "numerous" sample generation strategy. Thus, for a deep object detector, the fg-bg imbalance is similarly distributed during training and inference, as the detector always shares the same sample generation strategy in those two phases.

2) *Heuristic Sampling Methods:* Whether hard or soft sampling, the essence is to re-sample or re-weight training samples in the loss computation. If we regard hard sampling as soft sampling with weights 0 or 1, both of them be summarized as

$$L = s \sum_{i}^{N} w_{i} l_{i}, \tag{1}$$

where L denotes the overall training loss for a batch, and s is a scaling term. w_i and l_i are the weight and the loss of

i-th sample in a batch, respectively. In general, deep object detectors uses cross-entropy (CE) loss as l_i in the classification task, *i.e.*

$$L = -s \sum_{i}^{N} w_{i} [y_{i} \log(p_{i}) + (1 - y_{i}) \log(1 - p_{i})], \quad (2)$$

where y_i is the ground-truth label for *i*-th sample: $y_i = 1$ if it is foreground, otherwise $y_i = 0$. p_i is the confidence score, ranging from 0 to 1. Since the weights for foreground and background samples usually have different forms, Equation 2 can be further rewritten as

$$L = -s \sum_{i}^{N} y_{i} w_{i}^{f} \log(p_{i}) + (1 - y_{i}) w_{i}^{b} \log(1 - p_{i}), \quad (3)$$

where the notations f and b denote foreground and background, respectively. Equation 3 can also describe the training loss without sampling methods. In that case, we have $w_i^f = w_i^b = 1$ for all training samples.

B. Analysis

As most deep object detectors are trained with mini-batch stochastic gradient descent (mini-batch SGD), we discuss here the effect of the fg-bg imbalance and heuristic sampling methods on mini-batch SGD training. For each iteration, the learnable parameters Θ of the detector will be updated in the direction of the gradient, *i.e.*

$$\Theta^{t+1} = \Theta^t - \eta \frac{\partial L}{\partial \Theta^t},\tag{4}$$

where Θ^t denotes the parameters in *t*-th step, and η is the learning rate. According to Equation 3, the gradient can be further expressed as

$$\frac{\partial L}{\partial \Theta^{t}} = -s \sum_{i}^{N} y_{i} \frac{\partial \left[w_{i}^{f} \log(p_{i})\right]}{\partial \Theta^{t}} + (1 - y_{i}) \frac{\partial \left[w_{i}^{b} \log(1 - p_{i})\right]}{\partial \Theta^{t}}.$$
 (5)

As the exact form of the parameters Θ^{t} is unknown, a quantitative analysis for mini-batch SGD training seems impossible, especially after multiple training iterations. Therefore, in the following, we turn to analyze the case at the start of the training. We denote the learning parameters at the start of the training as Θ_{s} . At this point, Θ_{s} cannot distinguish foreground samples from background samples. In other words, the detector outputs similar confidence scores $p_{i} \approx p$ for all samples. Then, we have

$$\frac{\partial L}{\partial \Theta^s} \approx -s \sum_{i}^{N} y_i \frac{\partial [w_i^f \log(p)]}{\partial \Theta^s} + (1 - y_i) \frac{\partial [w_i^b \log(1 - p)]}{\partial \Theta^s} \approx -s \sum_{i}^{N} y_i [\frac{\partial w_i^f}{\partial \Theta^s} \log(p) + \frac{w_i^f}{p} \frac{\partial p}{\partial \Theta^s}] + (1 - y_i) [\frac{\partial w_i^b}{\partial \Theta^s} \log(1 - p) - \frac{w_i^b}{1 - p} \frac{\partial p}{\partial \Theta^s}]. \quad (6)$$

If the weights w_i^f and w_i^b are constants, Equation 6 would be simple as $\frac{\partial w_i^f}{\partial \Theta^s} = \frac{\partial w_i^b}{\partial \Theta^s} = 0$. However, in most soft sampling methods [17], [24], [35], the weight of a training sample is usually dynamic that may depend on confidence scores, IoU scores, and training iterations. Therefore, we would like to discuss the training cases of constant weights and dynamic weights, respectively.

1) Constant Weights: In this case, we have $w_i^f = w^f$, $w_i^b = w^b$ for any *i*, where w^f and w^b are constants. Thus, we also have $\frac{\partial w_i^f}{\partial \Theta^s} = \frac{\partial w_i^b}{\partial \Theta^s} = 0$, and Equation 6 can be derived as

$$\frac{\partial L}{\partial \Theta^s} \approx -s \sum_{i}^{N} y_i \frac{w^f}{p} \frac{\partial p}{\partial \Theta^s} - (1 - y_i) \frac{w^b}{1 - p} \frac{\partial p}{\partial \Theta^s}$$
$$= -s (N^f \frac{w^f}{p} - N^b \frac{w^b}{1 - p}) \frac{\partial p}{\partial \Theta^s}, \tag{7}$$

where N^f and N^b are the number of foreground and background training samples in a training iteration, respectively. Following [53], we use $||\cdot||$ to denote the L2-norm of a gradient vector, which represent its magnitude. From Equation 7, the gradient magnitude is

$$||\frac{\partial L}{\partial \Theta^s}|| \approx s |N^f \frac{w^f}{p} - N^b \frac{w^b}{1-p}| \cdot ||\frac{\partial p}{\partial \Theta^s}||.$$
(8)

If Θ_s is not biased, the initial estimate for both foreground and background samples are p = 0.5, then we have

$$||\frac{\partial L}{\partial \Theta_s}(p=0.5)|| \approx 2s|N^f w^f - N^b w^b| \cdot ||\frac{\partial p}{\partial \Theta^s}||.$$
(9)

When heuristic sampling methods are not used, *i.e.* $w_i^f = w_i^b = 1$, the gradient magnitude of the fg-bgimbalance case $(N_f \ll N_b)$ will be much larger than that of the balanced case $(N_f \approx N_b)$. Thus, if we train a detector without heuristic sampling methods, the fg-bg imbalance will result in a much larger gradient magnitude at the start of the training. If the scaling term s is not set properly, it may lead to training divergence. Compared with the gradient magnitude with heuristic sampling methods, the weighting terms w^f and w^b can alleviate the imbalance between N_f and N_b , thus leading to better stability at the start of the training.

2) Dynamic Weights: In this case, we take the well-known Focal Loss [17] as an example, which proposes a unified representation of the weighting term as $\alpha_t (1 - p_t)^{\gamma}$. When we apply it separately for foreground and background samples, we have $w_i^f = \alpha (1 - p_i)^{\gamma}$ and $w_i^b = (1 - \alpha) p_i^{\gamma}$. Here α and γ are the hyperparameters in Focal Loss for adaptively re-weighting training samples. As $p_i \approx p$ at the start of the training, we have $w_i^f \approx \alpha (1 - p)^{\gamma}$ and $w_i^b \approx (1 - \alpha) p^{\gamma}$. Then, Equation 6 can be derived as

$$\frac{\partial FL}{\partial \Theta^s} \approx -s \sum_{i}^{N} y_i \frac{\partial p}{\partial \Theta^s} \left[\frac{\partial w_i^f}{\partial p} \log(p) + \frac{w_i^f}{p} \right] \\ + (1 - y_i) \frac{\partial p}{\partial \Theta^s} \left[\frac{\partial w_i^b}{\partial p} \log(1 - p) - \frac{w_i^b}{1 - p} \right]$$

$$\approx -s \frac{\partial p}{\partial \Theta^{s}} \sum_{i}^{N} y_{i} \alpha (1-p)^{\gamma-1} [-\gamma \log(p) + \frac{1-p}{p}] \\ + (1-y_{i})(1-\alpha) p^{\gamma-1} [\gamma \log(1-p) - \frac{p}{1-p}] \\ \approx -s \frac{\partial p}{\partial \Theta^{s}} \{ N^{f} \alpha (1-p)^{\gamma-1} [-\gamma \log(p) + \frac{1-p}{p}] \\ + N^{b} (1-\alpha) p^{\gamma-1} [\gamma \log(1-p) - \frac{p}{1-p}] \}, \quad (10)$$

where *FL* denotes Focal Loss. As reported in [17], the best setting of Focal Loss is $\alpha = 0.25$, $\gamma = 2$ on COCO dataset [44], and Focal Loss uses a biased initialization to ensure $p \approx 0.01$. With these values, the gradient magnitude can be computed as

$$||\frac{\partial FL}{\partial \Theta^s}|| \approx 2s|10N^f - 10^{-4}N^b| \cdot ||\frac{\partial p}{\partial \Theta^s}||.$$
(11)

As we can see, the RHS (right hand side) of Equation 11 is equal to the RHS of Equation 9 by setting $w^f = 10$, $w^b = 10^{-4}$. Coincidentally, for COCO [44] dataset, a training anchor will learn 80 binary classifies for 80 object classes. In our observation, the *fg-to-bg* ratio of training anchor is 1 : 10³, thus $\frac{N^b}{N^f} \approx 8 \times 10^4$, which is close to $\frac{w^f}{w^b} = 10^5$. As we can see, Focal Loss also tries to alleviate the imbalance between N_f and N_b , to obtain a reasonable gradient magnitude.

However, this does not mean that it is impossible to train the detector without heuristic sampling methods. In fact, we can reduce *s* to lower the excessive gradient magnitude. But this in turn creates the dilemma of too small gradient magnitude. Specifically, when we set a small s_m as the scaling factor to train a detector without heuristic sampling methods, as $N^f \ll N^b$, *p* will rapidly approach $p \approx \frac{N^f}{N^f + N^b}$ to achieve a minimal loss value, and the gradient magnitude

$$||\frac{\partial L}{\partial \Theta_s}(w^f = w^b = 1, s = s_m)||$$

$$\approx s_m |\frac{N^f}{p} - \frac{N^b}{1-p}| \cdot ||\frac{\partial p}{\partial \Theta^s}|| \quad (12)$$

will be greatly decreased. At this point, s_m becomes unreasonable and we should set a large scaling factor. It is not surprising why most effective heuristic sampling methods are dynamic, like Focal Loss [17] and GHM [24].

To sum up, it is essential to control the gradient magnitude in the classification task when training a detector without heuristic sampling methods. As illustrated in [43], [54], the gradient magnitude will have a significant impact on the performance of multi-task learning. Object detection usually involves two or more tasks, the unreasonable gradient magnitude on the classification task affects not only itself but also other tasks. But in fact, we have various ways to control the gradient magnitude, and the heuristic sampling method is not the only choice.

Furthermore, the fg-bg imbalance, as we illustrated in Sec. III-A, also has the similar distribution in training and inference. If we use heuristic sampling methods during training, then it is equivalent to breaking the consistency of this distribution. In other words, Θ^e is obtained from the weighted

imbalance distribution, which may not perform well in the vanilla imbalance distribution. Next, we will experimentally explore how to train deep object detectors without heuristic sampling methods.

C. Experimental Verification

To investigate the accuracy gap between the detector with and without sampling methods, we experimentally investigate the difference between well-known Focal Loss [17] and CE loss. Focal Loss is widely used to address the fg-bg imbalance in the one-stage anchor-based and anchor-free object detectors [14], [17], [22], [23], [25]–[33], [40]. In previous studies [17], [24], Focal Loss helps RetinaNet [17] to yield $4\sim7$ higher AP on COCO [44] than CE loss. For simplicity, we denote the RetinaNet with Focal Loss and CE loss as *RetinaNet-FL* and *RetinaNet-CE*, respectively.

Two differences exist between *RetinaNet-FL* and *RetinaNet-CE*, one of which is the classification loss. Specifically, *RetinaNet-FL* and *RetinaNet-CE* use Focal Loss (L^{FL}) and CE loss (L^{CE}) in the classification task, respectively. Following the notations in Sec. III-A and Sec. III-B, we have

$$L^{FL} = \frac{1}{N^f} \sum_{i=1}^{N} w_i l_i, \quad L^{CE} = \frac{1}{N^f} \sum_{i=1}^{N} l_i, \quad (13)$$

where w_i is the weighting term, and $w_i = \alpha (1 - p_i)^{\gamma}$ if *i*-th training sample is foreground, otherwise $w_i = (1 - \alpha)p_i^{\gamma}$.

Another difference between *RetinaNet-FL* and *RetinaNet-CE* is the initialization. Focal Loss uses a biased initialization that initializes the final classification convolutional layer with the bias $b = -\log \frac{1-\pi}{\pi}$. Then, at the start of the training, $p_i = \pi$ is tenable for every training anchor as p_i is computed by sigmoid activation. A heuristic value $\pi = 10^{-2}$ is used in Focal Loss to avoid network diverging. Unfortunately, this will result in network diverging during training *RetinaNet-CE*. We will also explore how to adjust π in the following.

1) Classification Loss: Our first discovery is that RetinaNet-CE has poor stability on the classification loss, which reflects the unreasonable classification gradient magnitude. This phenomenon corresponds to our discovery in Sec. III-B. Training RetinaNet-CE with default $\pi = 10^{-2}$ causes a large classification loss, with the network diverging in a few iterations. See Fig. 2, we observe that RetinaNet-FL can be stably trained, but training RetinaNet-CE needs to carefully tune π and s to avoid the large classification loss. Only with specific settings, we can obtain several converged models, as shown in Fig. 2.

2) Detection Accuracy: Our second discovery is that the classification loss scale will greatly influence the accuracy of the *RetinaNet-CE*. As shown in Table IIa, if a model can be stably trained, then its accuracy will be similar when s is fixed, but tuning s leads to great changes in AP. Table IIb shows that *RetinaNet-CE* with $\pi = 10^{-5}$, $s = 10^{-1}$ has already achieved the comparable accuracy of *RetinaNet-FL* (36.3 AP vs. 36.4 AP). This is an inspiring result, as previous works [17], [24] reported there is a 4~7 AP gap between CE loss and Focal Loss. Our investigation successfully narrows

TABLE IIDETECTION ACCURACY OF CE LOSS ON COCO minival, WITHDIFFERENT s AND π . "N/A" REFERS TOTHE NETWORK DIVERGING

(a) Accuracy of *RetinaNet-CE* in different π, s

| AP | s = 1 | $s = 10^{-1}$ | $s = 10^{-2}$ | $s = 10^{-3}$ |
|-----------------|-------|---------------|---------------|---------------|
| $\pi = 10^{-2}$ | n/a | n/a | n/a | 1.6 |
| $\pi = 10^{-3}$ | n/a | n/a | n/a | 1.6 |
| $\pi = 10^{-4}$ | n/a | n/a | 26.8 | 1.6 |
| $\pi = 10^{-5}$ | n/a | 36.3 | 26.8 | 1.6 |

(b) Accuracy comparison between RetinaNet-FL and RetinaNet-CE

| Model | AP | AP ₅₀ | AP ₇₅ |
|-----------------------------------------------|------|------------------|------------------|
| RetinaNet-FL | 36.4 | 55.0 | 39.0 |
| RetinaNet-CE ($\pi = 10^{-5}, s = 10^{-1}$) | 36.3 | 54.8 | 38.7 |

this gap. However, it is still unknown why adjusting π and *s* can help to train *RetinaNet-CE*. Next, we will perform an analysis of this.

3) Analysis: At the start of the training, the bias initialization ensures $p_i \approx \pi$ for each training sample. According to Eq. 2 and Eq. 13, we can estimate L^{FL} as

$$L^{FL} \approx -\alpha (1-\pi)^{s} \log(\pi) - (1-\alpha)\pi^{s} \frac{N^{b}}{N^{f}} \log(1-\pi), \quad (14)$$

and estimate L^{CE} as

$$sL^{CE} \approx -s\log(\pi) - s\frac{N^b}{N^f}\log(1-\pi).$$
(15)

For COCO [44] dataset, a training anchor will learn 80 binary classifies for 80 object classes. In our observation, the *fg-to-bg* ratio of training anchor is 1 : 10³, thus $\frac{N^b}{N^f} \approx 10^3 \times 80$. With $\alpha = 0.25$, s = 2, $\pi = 10^{-2}$ in training *RetinaNet-FL*, and $s = 10^{-1}$, $\pi = 10^{-5}$ in training *RetinaNet-CE*, we can estimate L^{FL} and L^{CE} quantitatively as

$$L^{FL} \approx 1.19$$
, and $L^{CE} \approx 1.23$, (16)

where the two values are very close. Since the loss scale can reflect the gradient magnitude, we believe that the similar gradient magnitude is the reason for *RetinaNet-CE* with $s = 10^{-1}, \pi = 10^{-5}$ achieves AP comparable to that of *RetinaNet-FL*. Moreover, this setting also facilitates multi-task learning. See Fig. 3a, *RetinaNet-FL* achieves a quite good balance between the localization loss L^{LOC} and L^{FL} , whereas Fig. 3b shows that only *RetinaNet-CE* with $\pi = 10^{-5}$, $w = 10^{-1}$ can obtain this balance. Other settings, however, will lead the training to be dominated by the localization task, which is harmful to training an accurate object detector. In conclusion, our investigation reveals that the classification gradient magnitude is the central cause of the accuracy gap, and we can adjust it by initialization and loss scaling.

IV. METHOD

Our investigation reveals that by tuning s and π , the detector without heuristic sampling methods can achieve a similar detection accuracy to that with heuristic sampling methods. However, tuning them is always laborious. In this



Fig. 2. Loss curves of Focal Loss L^{FL} and cross-entropy loss L^{CE} . "X" means the network diverging. The detector is RetinaNet with ResNet-50-FPN [6], [55] backbone, trained on COCO train2017 [44] with 1× learning schedule [56] (12 epochs), implemented on maskrcnn-benchmark [57]. We only show the first 1k iterations for better visualization.

section, we propose a novel Sampling-Free mechanism, which addresses the *fg-bg* imbalance by adaptively setting *s* and π , thus adaptively controlling the classification gradient magnitude.

4) Discarding Heuristic Sampling Methods: As sampling methods are always a default part in training deep object detectors, the first step of our Sampling-Free mechanism is discarding heuristic sampling methods during training. For one-stage anchor-based object detectors [16]–[23], [40], soft sampling methods (*e.g.* Focal Loss [17], GHM [24], PISA [35]) is widely used for re-weighting training samples in the classification task. In our method, we use the standard CE loss to train the classification task, which treats all training samples equally.

For two-stage anchor-based approaches [4]–[15], hard sampling methods (*e.g.* biased sampling [4], OHEM [34]) are widely used for re-sampling training samples. In our method, we train all training samples in RPN and RoI-subnet. For example, a common implementation of biased sampling [4]

in training Faster R-CNN is: (1) RPN randomly selects 256 anchors with a biased 1:1 *fg-to-bg* ratio, (2) RoI-subnet randomly selects 512 proposals with a biased 1:3 *fg-to-bg* ratio. In our method, we train all examples in both RPN and RoI-subnet. That is to say, we train RPN with all foreground/background anchors ($\sim 10^5$ per-image) and train RoI-subnet with all foreground/background proposals ($\sim 10^3$ per-image). We use CE loss as the classification loss both in RPN and RoI-subnet.

For anchor-free object detectors [11], [14], [25]–[33], they regard "points" as the training samples rather than "anchors" in anchor-based object detectors, and most of them use soft sampling methods to address the fg-bg imbalance. We follow the principle of Sampling-Free in anchor-based object detectors that equally use all training samples during training, *i.e.* we use the standard CE loss in the classification task.

5) Optimal Bias Initialization: Sec. III-C has shown that adjusting π can help to avoid network diverging. However, it is difficult to determine π . We propose optimal bias initialization



Fig. 3. Loss curves of localization (L^{LOC}) and classification $(L^{FL} \text{ or } L^{CE})$, which are from the same models in Fig. 2. Note that the curves of L^{LOC} with different s in (b) are very similar, thus we only show the curve of L^{LOC} with $s = 10^{-1}$ in (b).

to compute π from data statistics rather than tuning it. Our idea is to initialize the bias of the last convolutional layer to obtain a minimal classification loss value. The derivative of L^{CE} is

$$\frac{\partial L^{CE}}{\partial \pi} = -\frac{1}{\pi} + (\frac{N}{N^f} - 1)\frac{1}{1 - \pi}.$$
 (17)

When $\pi = \frac{N^f}{N}$, $\frac{\partial L^{CE}}{\partial \pi} = 0$, and L^{CE} will attain the minimal value. As the predicted score is predicted with the sigmoid activation, we can obtain the optimal initial bias as

$$b = -\log\frac{1-\pi}{\pi} = -\log(\frac{N}{N^f} - 1).$$
 (18)

Here $\frac{N}{Nf}$ can be computed by pre-defined anchors, thus the computation is efficient as it does not require network forwarding. In our observation, $\frac{N}{Nf} \approx 10^5$, which corresponds to $\pi = 10^{-5}$ that performs best in our experiments. It is worth noting that the accuracy of the model is robust to our initialization strategy, as the model can "utilize" the imbalanced distribution to obtain a lower loss. We initialize the model to ensure the stability of the classification loss.

6) Guided Loss Scaling: Usually, the overall loss function to train a deep object detector is composed of a localization loss term L^{LOC} and a classification loss term L^{CLS} . Let L^t denote the overall loss in the *t*-th training step. When we use CE loss as the classification loss, we have

$$L^{t} = (L^{LOC})^{t} + (L^{CLS})^{t} = (L^{LOC})^{t} + s^{t} (L^{CE})^{t}, \quad (19)$$

where s^t is used to scale the $(L^{CE})^t$ as the CE loss scale is unreasonable under the fg-bg imbalance. As mentioned in III-C, it is essential to control the classification loss scale to be close to the localization loss scale. A straightforward way is to adjust s^t . However, it results in a new hyperparameter. Our key idea is to adjust s^t dynamically during training. That says, instead of using a constant s^t , we define a guided term

$$g^{t} = \frac{(L^{LOC})^{t}}{(L^{CE})^{t}},$$
 (20)

and let $s^t = g^t$, which suggests using the localization loss scale of the current mini-batch as the target of the rescaled CE loss scale. Thus, this technique is termed "guided loss scaling". It is worth noting that g^t is only used for scaling the classification loss, *i.e.* its gradient is ignored in the backpropagation. Therefore, the overall gradient is

$$\frac{\partial L^{t}}{\partial \Theta^{t}} = \frac{\partial (L^{LOC})^{t}}{\partial \Theta^{t}} + g^{t} \frac{\partial (L^{CE})^{t}}{\partial \Theta^{t}}, \qquad (21)$$

which ignores the gradient calculation of s^t .

Our guided loss scaling can be interpreted threefold. First, according to Sec. III-C (especially Fig. 3), it appears a good choice to let the localization loss scale and the classification loss scale be similar, where the classification loss is either Focal Loss or CE loss. Second, it is convenient to use the localization loss as guidance, because the localization loss is already there for object detection. Third, the classification loss without sampling methods (*i.e.* CE loss) is greatly influenced by the fg-bg imbalance, but localization loss is little influenced as it is computed merely for foreground anchors. Thus, the localization loss is helpful to control the unreasonable classification loss due to the fg-bg imbalance.

However, in our experiments, we find that the detector may not achieve the best detection accuracy when the classification loss is simply equal to the localization loss. Fortunately, the well-known uncertainty weighting [58] proposes a simple method to weigh two losses from the perspective of Bayesian uncertainty. When we apply the method to our case, the overall loss would be

$$L^{t} = \frac{1}{(\sigma_{1}^{t})^{2}} (L^{LOC})^{t} + \frac{1}{(\sigma_{2}^{t})^{2}} (L^{CE})^{t} + 2\log(\sigma_{1}^{t}\sigma_{2}^{t}), \quad (22)$$

where σ_1^t and σ_2^t are learnable parameters, and they are initialized as $\sigma_1^0 = \sigma_2^0 = 1$. $2\log(\sigma_1^t \sigma_2^t)$ is the normalization term to avoid the degradation of $\frac{1}{(\sigma_1^t)^2} \to 0$ and $\frac{1}{(\sigma_2^t)^2} \to 0$. But if we train the detector in this way, the training would be quickly failed, as the classification loss would be much larger

(b) Alletten studies of Generaline Energies ECOC [21]

TABLE III

Ablation Studies of Our Sampling-Free Mechanism on COCO minival. "FL → CE" Denotes "Focal Loss → CE Loss", and "Init" Denotes "Initialization". From the Following Three Subtables, Better AP Can Be Achieved With Sampling-Free, but the Hybrid of the Guided Loss Scaling and Heuristic Sampling Methods Cannot Improve the Detection Accuracy

| (a) Ablation | studies | 01 5a | mpning | g-riee | III Ketillahet | [1/] | (b) Ablatic | m studi | es or . | sampn | ng-rie | e in reos [s | ,1] |
|---------------------|---------|-------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|------------------------|-------------|---------------------|---------|---------|-------|--------|------------------------|-------------|
| Components | | R | etinaNe | et (Resl | Net-50-FPN, $1 \times$ | :) | Components | | | FCOS | (ResNe | t-50-FPN, $1 \times$) | |
| $FL \rightarrow CE$ | X | ~ | < | 1 | X | 1 | $FL \rightarrow CE$ | X | 1 | 1 | 1 | X | 1 |
| Optimal Bias Init | X | X | Image: A start of the start of | X | 1 | 1 | Optimal Bias Init | X | X | 1 | X | 1 | ✓ |
| Guided Loss Scaling | × | X | X | 1 | 1 | / | Guided Loss Scaling | X | X | X | 1 | 1 | > |
| AP | 36.4 | n/a | n/a | n/a | 36.5 (+0.1) | 37.0 (+0.6) | AP | 37.1 | n/a | n/a | n/a | 37.1 (+0.0) | 37.6 (+0.5) |
| AP_{50} | 55.0 | n/a | n/a | n/a | 55.5 (+0.5) | 56.5 (+1.5) | AP_{50} | 56.0 | n/a | n/a | n/a | 56.2 (+0.2) | 57.4 (+1.4) |
| AP_{75} | 39.0 | n/a | n/a | n/a | 38.8 (-0.2) | 39.2 (+0.2) | AP ₇₅ | 39.8 | n/a | n/a | n/a | 39.7 (-0.1) | 40.3 (+0.5) |
| AP_S | 19.9 | n/a | n/a | n/a | 20.1 (+0.2) | 20.3 (+0.4) | AP_S | 21.3 | n/a | n/a | n/a | 21.0 (-0.3) | 21.9 (+0.6) |
| AP_M | 40.3 | n/a | n/a | n/a | 40.1 (-0.2) | 40.5 (+0.2) | AP_M | 41.0 | n/a | n/a | n/a | 41.3 (+0.3) | 41.2 (+0.2) |
| AP_L | 48.9 | n/a | n/a | n/a | 48.1 (-0.8) | 49.5 (+0.6) | AP_L | 47.8 | n/a | n/a | n/a | 47.9 (+0.1) | 48.5 (+0.7) |

| (c) Adiation studies of Sampling-Free in Faster R-CIVI |
|--------------------------------------------------------|
|--------------------------------------------------------|

| Sampling Free Mechanism | Faster R-CNN (ResNet-50-FPN [6], [55], $1 \times$) | | | | | | | | | | | | | |
|---------------------------------------------------|-----------------------------------------------------|-----|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|-------|
| Sampling-Free Weenamsin | RPN | RoI | RPN | RoI | RPN | RoI | RPN | RoI | RPN | RoI | RPN | RoI | RPN | RoI |
| Biased Sampling \rightarrow Non-sampling | X | X | X | X | X | X | X | X | 1 | X | X | 1 | 1 | 1 |
| Optimal Bias Initialization + Guided Loss Scaling | X | × | 1 | X | × | 1 | 1 | 1 | 1 | × | × | 1 | 1 | 1 |
| AP | 36 | .8 | 36.5 | (-0.3) | 36.4 | (-0.4) | 36.8 (| +0.0) | 37.5 (| (+0.7) | 38.1 (| (+1.3) | 38.4 (| +1.6) |
| AP ₅₀ | 58 | .4 | 58.2 | (-0.2) | 57.9 (| (-0.5) | 58.7 (| +0.3) | 59.0 (| +0.6) | 59.6 (| (+1.2) | 59.9 (| +1.5) |
| AP ₇₅ | 40 | .0 | 39.6 | (-0.4) | 39.4 | (-0.6) | 40.0 (| +0.0) | 40.4 (| +0.4) | 41.6 (| (+1.6) | 41.7 (| +1.7) |
| AP_S | 20 | .7 | 21.2 (| +0.5) | 21.1 (| +0.4) | 21.1 (| +0.4) | 21.5 (| (+0.8) | 22.2 (| (+1.5) | 22.3 (| +1.6) |
| AP_M | 39 | .7 | 39.4 | (-0.3) | 39.1 | (-0.6) | 40.0 (| +0.3) | 40.7 (| (+1.0) | 41.2 (| (+1.5) | 41.6 (| +1.9) |
| APL | 47 | .9 | 47.6 | (-0.3) | 47.5 | (-0.4) | 47.8 (| (-0.1) | 48.8 (| +0.9) | 50.0 (| (+2.1) | 50.9 (| +3.0) |

 $(\sim 10 \times)$ than the localization loss at the start of the training (see Figure 3). Hence, the guided term g^t is necessary, and the overall loss should be

(a) Alletian studies of Generalian East in DetineNet [17]

$$L^{t} = \frac{1}{(\sigma_{1}^{t})^{2}} (L^{LOC})^{t} + \frac{g^{t}}{(\sigma_{2}^{t})^{2}} (L^{CE})^{t} + 2\log(\sigma_{1}^{t}\sigma_{2}^{t}).$$
(23)

To keep the consistency with the original training loss, we can identify that the localization loss does not require weighting ($\sigma_1^t = 1$). We can also use δ^t to denote $\frac{1}{(\sigma_2^t)^2}$, then the overall loss would be very simple, *i.e.*

$$L^{t} = (L^{LOC})^{t} + g^{t} \delta^{t} (L^{CE})^{t} - \log \delta^{t}.$$

$$(24)$$

We notice that there have been several works [54], [58], [59] for adaptive multi-task loss scaling. Our guided loss scaling is different from them in three points: 1) it is aimed at controlling the classification loss under the fg-bg imbalance, which belongs to the single-task loss weighting rather than the multi-task loss weighting; 2) it is specifically designed for deep object detectors as it requires the localization loss to guide the classification loss; 3) it converts the class imbalance problem to the loss scaling problem, which seems not reported before in the literature, to our best knowledge.

V. EXPERIMENTS

In this section, we will perform extensive experiments to validate our Sampling-Free mechanism. Before that, we first describe the experimental details about datasets and baselines. Then, we perform ablation studies on anchor-based and anchor-free object detectors. Finally, we compare our method with existing heuristic sampling methods, and present experimental results on public datasets.

A. Implementation Details

1) Datasets: We use the well-known COCO [44] and PASCAL VOC [45] datasets to validate our method. For

COCO dataset, following common practices [4], [17], we train models on the train2017 split and perform ablation studies on minival split, and report detection accuracy on test-dev split, where COCO-style average precision (AP) is used as the evaluation metrics. For PASCAL VOC dataset, we also follow common practices [4], [20] that training models on a union set of PASCAL VOC 2007 and 2012 set (07+12 split), and evaluated on PASCAL VOC 2007 test set (07test split), where VOC-style mean average precision (mAP) is used as the evaluation metrics.

2) Baselines: We use three object detectors — RetinaNet [17] (one-stage anchor-based), Faster R-CNN [4] (two-stage anchor-based), FCOS [31] (anchor-free) that implemented on maskrcnn-benchmark [57] to perform experiments, where we follow the public standard training configurations to implement them, which means that we have not made any changes for the hyperparameters of them. Besides, we also use Mask R-CNN [7] to validate Sampling-Free in the instance segmentation task.

B. Ablation Studies

1) Detection Accuracy: As shown in Table III, we perform experiments of Sampling-Free on RetinaNet, FCOS, and Faster R-CNN, respectively. See Table IIIa and Table IIIb, the experimental phenomena of Sampling-Free on RetinaNet and FCOS are similar. Only when we discard heuristic sampling methods (Focal Loss \rightarrow CE Loss) and use optimal bias initialization with guided loss scaling at the same time, we can obtain meaningful detection accuracy improvements. Specifically, Sampling-Free achieves 0.6 AP and 0.5 AP higher than Focal Loss in RetinaNet and FCOS, respectively. This is an impressive improvement as the well-know GHM-C [24] only outperforms Focal Loss 0.2 AP. More importantly, this is the first time that CE Loss has achieved better performance than the soft sampling method in object detection.



Fig. 4. Mask R-CNN [7] (37.8 box AP, 34.2 mask AP on COCO minival) vs. Mask R-CNN with Sampling-Free (39.0 box AP, 34.9 mask AP on COCO minival) in ResNet-50-FPN backbone. The latter exhibits better detection and segmentation results.

TABLE IV

TRAINING SPEED AND MEMORY COST OF SAMPLING-FREE ON FASTER R-CNN, WHICH IS EVALUATED ON A SINGLE Nvidia-Titan-Xp GPU WITH BATCH SIZE 1

| Detector | Method | Speed | Memory |
|-----------------------------|-----------------|-------|--------|
| Faster R-CNN | Biased Sampling | 172ms | 1714MB |
| (ResNet-50-FPN, $1\times$) | Sampling-Free | 184ms | 1669MB |

For Faster R-CNN (See Table IIIc), when we discard heuristic sampling methods without optimal bias initialization and guided loss scaling, the training of the detector will be failed. When we only use the optimal bias initialization and guided loss scaling without discarding heuristic sampling methods, the detector also cannot obtain improvement on detection accuracy. Only when we discard heuristic sampling methods, and use optimal bias initialization with guided loss scaling at the same time, we will observe obvious gains in AP. Sampling-Free improves 0.7 AP and 1.3 AP in RPN and RoI-subnet, respectively. When we use Sampling-Free in both RPN and RoI-subnet, an impressive 1.6 AP improvement can be obtained, with the gains from all AP metrics.

We notice that Sampling-Free exhibits more improvements to Faster R-CNN than RetinaNet and FCOS, which may be due to biased sampling wasting most background examples, but Sampling-Free allows all foreground and background examples to be trained synchronously.

2) Training Speed and Memory Cost: As both Focal Loss and Sampling-Free use all samples to train, the training speed and the memory cost of *RetinaNet-FL* and *RetinaNet-CE* are very similar. However, in Faster R-CNN, Sampling-Free allows more background samples to be trained. As shown in Table IV, we measure the performance of Faster R-CNN with Sampling-Free on a single Nvidia-Titan-Xp GPU with batch size 1. Interestingly, although the training speed becomes slower for Faster R-CNN with Sampling-Free

TABLE V RESULTS OF SAMPLING-FREE ON COCO test-dev

| Detector | Method | AP | AP ₅₀ | AP ₇₅ |
|-------------------------------|-----------------|------|------------------|------------------|
| Faster R-CNN | Biased Sampling | 39.3 | 61.4 | 42.7 |
| (ResNet-101-FPN, $1 \times$) | Sampling-Free | 40.7 | 62.4 | 44.3 |
| RetinaNet | Focal Loss | 38.8 | 58.4 | 41.7 |
| (ResNet-101-FPN, $1 \times$) | Sampling-Free | 39.4 | 58.2 | 42.9 |
| Faster R-CNN | Biased Sampling | 39.6 | 61.3 | 43.0 |
| (ResNet-101-FPN, $2\times$) | Sampling-Free | 41.0 | 62.6 | 44.6 |
| RetinaNet | Focal Loss | 38.9 | 58.4 | 41.7 |
| (ResNet-101-FPN, $2\times$) | Sampling-Free | 39.5 | 58.3 | 43.2 |

TABLE VI RESULTS OF SAMPLING-FREE ON PASCAL VOC 07test SPLIT

| Detector | Method | mAP |
|--------------------------------|-----------------|------|
| RetinaNet | Focal Loss | 79.3 |
| (ResNet-50-FPN, $0.2\times$) | Sampling-Free | 80.1 |
| Faster R-CNN | Biased Sampling | 80.9 |
| (ResNet-50-FPN, $0.2 \times$) | Sampling-Free | 81.5 |

(as more background proposals are trained in RoI-subnet), its memory cost is reduced, which is owing to the operation of biased sampling also requires considerable memory costs.

C. Experimental Results

1) Results on COCO and PASCAL VOC: For COCO dataset, we have demonstrated that the effectiveness of Sampling-Free on ResNet-50-FPN backbone and $1 \times$ learning schedule. We further verify our methods on the larger backbone and the longer learning schedule. As shown in Table V, for ResNet-101-FPN backbone, Sampling-Free still shows impressive detection accuracy improvements, which can improve Faster R-CNN and RetinaNet about 1.5 AP and 0.5 AP, respectively. Even with the $2 \times$ learning schedule, we observe a steady increase in AP as well.

TABLE VII

This Table Illustrates the Comparison Between Heuristic Sampling Methods and Our Sampling-Free Mechanism. ΔAP and $\Delta Hyperparameters$ Denote the Change in Detection Accuracy and Hyperparameters Relative to the Baseline Method (Focal Loss in RetinaNet, Biased Sampling in Faster R-CNN). Sampling-Free Achieves the Best ΔAP Without Any Hyperparameter Introduced

| Solutions | Abbreviation | Δ AP in RetinaNet (R-50-FPN, 1×) | ΔAP in Faster R-CNN (R-50-FPN, 1×) | Δ hyperparameters |
|---------------|----------------------------|----------------------------------------------|----------------------------------------------|--------------------------|
| Hard Sampling | OHEM [34] | n/a | $36.4 \rightarrow 36.6 \ (+0.2 \ \text{AP})$ | $2 \rightarrow 2$ |
| mard Sampling | IoU-balanced sampling [12] | n/a | $36.4 \rightarrow 36.8 \ (+0.4 \ \text{AP})$ | $2 \rightarrow 3$ |
| Soft Sampling | GHM-C [24] | $35.6 \rightarrow 35.8 \ (+0.2 \ \text{AP})$ | n/a | $2 \rightarrow 1$ |
| Soft Sampling | ISR [35] | n/a | $36.4 \rightarrow 37.9 \ (+1.5 \ \text{AP})$ | $2 \rightarrow 4$ |
| Non-Sampling | Sampling-Free | $36.4 \rightarrow 37.0 \ (+0.6 \ AP)$ | $36.8 ightarrow 38.4 \ (+1.6 \ AP)$ | 2 ightarrow 0 |



(a) RetinaNet with Sampling-Free exhibits better performance than RetinaNet with Focal Loss [17] (baseline)



(b) Faster R-CNN with Sampling-Free exhibits better performance than Faster R-CNN with biased sampling [4] (baseline)

Fig. 5. Visualized detection results of different object detectors with and without our Sampling-Free mechanism, which show that the detector with Sampling-Free mechanism performs better. The backbone of these detectors is ResNet-50-FPN [6], [55].

For PASCAL VOC dataset, as shown in Table VI, Sampling-Free improves 0.8 mAP and 0.6 mAP for RetinaNet and Faster R-CNN, respectively. These results illustrate the robustness of our Sampling-Free mechanism.

2) Comparison: We compare Sampling-Free with existing heuristic sampling methods to illustrate our advantages. As the implementations of Sampling-Free and other methods may be on different platforms (*e.g.* maskrcnn-benchmark vs. mmdetection), we mainly compare their changes in performance relative to the baseline method (Focal Loss in RetinaNet, Biased Sampling in Faster R-CNN). As presented in Table VII, compared with heuristic sampling methods, Sampling-Free has the following three advantages:

• Sampling-Free achieves the best relative detection accuracy improvement, where 0.6 AP and 1.6 AP improvement are obtained in RetinaNet and Faster R-CNN, respectively.

• Sampling-Free has no hyperparameters to search — In contrast, other heuristic sampling methods have at least one hyperparameter. The ISR [35], which is closest to us in detection accuracy, introduces 4 hyperparameters.

(a) AP of ATSS (ResNeXt-64x4d-101-DCN, $2\times$) on COCO test-dev

| Method | AP | AP_{50} | AP_{75} | AP_S | AP_M | AP_L |
|------------------|------|-----------|-----------|--------|--------|--------|
| ATSS | 47.7 | 66.5 | 51.9 | 29.7 | 50.8 | 59.4 |
| w. Sampling-Free | 48.2 | 66.9 | 52.4 | 30.3 | 51.3 | 59.9 |

(b) AP of PAA (ResNeXt-64x4d-101-DCN, $2\times$) on COCO test-dev

| Method | AP | AP ₅₀ | AP ₇₅ | AP_S | AP_M | AP_L |
|------------------|------|------------------|------------------|--------|--------|--------|
| PAA | 49.0 | 67.8 | 53.3 | 30.2 | 52.8 | 62.2 |
| w. Sampling-Free | 49.6 | 68.3 | 53.8 | 30.6 | 53.7 | 63.1 |

TABLE IX

RESULTS OF SAMPLING-FREE IN THE INSTANCE SEGMENTATION TASK

(a) Box AP on COCO minival (ResNet-50-FPN, $1 \times$)

| Method | AP | AP ₅₀ | AP_{75} | AP_S | AP_M | AP_L |
|------------------|------|------------------|-----------|--------|--------|--------|
| Mask R-CNN | 37.8 | 59.3 | 41.1 | 21.5 | 41.1 | 49.9 |
| w. Sampling-Free | 39.0 | 60.3 | 42.5 | 22.5 | 41.9 | 51.2 |

(b) Mask AP on COCO minival (ResNet-50-FPN, $1\times$)

| Method | AP | AP ₅₀ | AP_{75} | AP_S | AP_M | AP_L |
|------------------|------|------------------|-----------|--------|--------|--------|
| Mask R-CNN | 34.2 | 55.9 | 36.3 | 15.6 | 36.8 | 50.6 |
| w. Sampling-Free | 34.9 | 56.8 | 37.1 | 16.2 | 37.3 | 51.2 |

• Sampling-Free has better versatility — it is applicable to both one-stage and two-stage deep object detectors.

3) Results on Adaptive Label Assignment Strategies: To date, the state-of-the-art detection accuracy is achieved by adaptive label assignment methods [38], [39], where the definition of foreground/background training sample is very different from that in deep object detectors. To validate the effectiveness of Sampling-Free in them, we replace Focal Loss with Sampling-Free in ATSS [38] and PAA [39]. As presented in Table VIII, we successfully verify that the state-of-the-art models of ATSS and PAA can be further improved with Sampling-Free.

4) Instance Segmentation: Although Sampling-Free can help detectors achieve better detection accuracy, it is unknown whether detection results produced with Sampling-Free mechanism can facilitate other tasks in practice. Therefore, we introduce Sampling-Free into Mask R-CNN to observe whether it can achieve higher accuracy in instance segmentation. Note that in Mask R-CNN, the heuristic sampling method is not used in its segmentation branch, thus we only apply our Sampling-Free mechanism on the classification branch. As shown in Table IX, Mask R-CNN with Sampling-Free achieves 1.2 box AP and 0.7 mask AP gains. We visualize the detection and segmentation results in Fig. 5, which suggests that the detection results produced with Sampling-Free can also improve the instance segmentation task.

5) Visualization: More visualization results are in Fig. 5.

VI. CONCLUSION

In this paper, we explored whether heuristic sampling methods are necessary for training accurate deep object detectors under the fg-bg imbalance. Our investigation revealed that the key to training without heuristic sampling methods under the fg-bg imbalance is to control the classification gradient magnitude. Inspired by this, we proposed a novel Sampling-Free mechanism to control the classification gradient magnitude from initialization and loss scaling, without new hyperparameters introduced. Extensive experiments demonstrated the effectiveness of Sampling-Free in various object detectors, which also yields considerable gains in the instance segmentation task and the state-of-the-art label assignment strategies. Our Sampling-Free mechanism provides a new perspective to address the fg-bg imbalance.

Although Sampling-Free can support training with crossentropy loss, it is not designed for the detection metrics, which may limit further performance increases. Specifically, the average precision metric expects samples to have an IoU-related confidence score. There has been some work on this [60], [61], but they still use a variant of Focal Loss to train the classification task. Thus, a future study direction is to use only cross-entropy loss to model this unified confidence score. On a larger scale, one can try to use cross-entropy loss with metric-specific design to achieve better performance.

REFERENCES

- A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 2, pp. 84–90, Jun. 2012.
- [2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [3] L. Liu et al., "Deep learning for generic object detection: A survey," Int. J. Comput. Vis., vol. 128, pp. 1405–1573, Feb. 2019.
- [4] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, Jun. 2017.
- [5] J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object detection via region-based fully convolutional networks," in *Proc. NeurIPS*, 2016, pp. 379–387.
- [6] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 936–944.
- [7] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask R-CNN," in Proc. Int. Conf. Comput. Vis. (ICCV), 2017, pp. 2980–2988.
- [8] H. Hu, J. Gu, Z. Zhang, J. Dai, and Y. Wei, "Relation networks for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3588–3597.
- [9] Z. Cai and N. Vasconcelos, "Cascade R-CNN: Delving into high quality object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6154–6162.
- [10] B. Jiang, R. Luo, J. Mao, T. Xiao, and Y. Jiang, "Acquisition of localization confidence for accurate object detection," in *Proc. ECCV*, 2018, pp. 816–832.
- [11] X. Lu, B. Li, Y. Yue, Q. Li, and J. Yan, "Grid R-CNN," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2019, pp. 7363–7372.
- [12] J. Pang, K. Chen, J. Shi, H. Feng, W. Ouyang, and D. Lin, "Libra R-CNN: Towards balanced learning for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 821–830.
- [13] K. Chen et al., "Hybrid task cascade for instance segmentation," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2019, pp. 4974–4983.
- [14] J. Wang, K. Chen, S. Yang, C. C. Loy, and D. Lin, "Region proposal by guided anchoring," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2965–2974.
- [15] Y. Li, Y. Chen, N. Wang, and Z.-X. Zhang, "Scale-aware trident networks for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6054–6063.

- [16] T. Kong, F. Sun, A. Yao, H. Liu, M. Lu, and Y. Chen, "RON: Reverse connection with objectness prior networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5244–5252.
- [17] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2999–3007.
- [18] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Jul. 2017, pp. 6517–6525.
- [19] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," 2018, arXiv:1804.02767. [Online]. Available: http:// arxiv.org/abs/1804.02767
- [20] S. Zhang, L. Wen, X. Bian, Z. Lei, and S. Z. Li, "Single-shot refinement neural network for object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4203–4212.
- [21] S. Liu, D. Huang, and Y. Wang, "Receptive field block net for accurate and fast object detection," in *Proc. ECCV*, 2018, pp. 404–419.
- [22] C. Zhu, Y. He, and M. Savvides, "Feature selective anchor-free module for single-shot object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 840–849.
- [23] X. Zhang, F. Wan, C. Liu, R. Ji, and Q. Ye, "FreeAnchor: Learning to match anchors for visual object detection," in *Proc. NeurIPS*, 2019, pp. 147–155.
- [24] B. Li, Y. Liu, and X. Wang, "Gradient harmonized single-stage detector," in *Proc. AAAI*, 2019, pp. 8577–8584.
- [25] H. Law and J. Deng, "CornerNet: Detecting objects as paired keypoints," in Proc. Eur. Conf. Comput. Vis., 2018, pp. 765–781.
- [26] X. Zhou, J. Zhuo, and P. Krähenbühl, "Bottom-up object detection by grouping extreme and center points," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 850–859.
- [27] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," 2019, arXiv:1904.07850. [Online]. Available: http://arxiv.org/abs/1904.07850
- [28] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "CenterNet: Keypoint triplets for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6569–6578.
- [29] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin, "RepPoints: Point set representation for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9657–9666.
- [30] Z. Yang *et al.*, "Dense reppoints: Representing visual objects with dense point sets," in *Proc. ECCV*, 2020, pp. 227–244.
- [31] Z. Tian, C. Shen, H. Chen, and T. He, "FCOS: Fully convolutional one-stage object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.* (*ICCV*), Oct. 2019, pp. 9627–9636.
- [32] T. Kong, F. Sun, H. Liu, Y. Jiang, L. Li, and J. Shi, "FoveaBox: Beyound anchor-based object detection," *IEEE Trans. Image Process.*, vol. 29, pp. 7389–7398, 2020.
- [33] C. Zhu, F. Chen, Z. Shen, and M. Savvides, "Soft anchor-point object detection," in *Proc. ECCV*, 2020, pp. 91–107.
- [34] A. Shrivastava, A. Gupta, and R. Girshick, "Training region-based object detectors with online hard example mining," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 761–769.
- [35] Y. Cao, K. Chen, C. C. Loy, and D. Lin, "Prime sample attention in object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11583–11591.
- [36] K. Oksuz, B. C. Cam, S. Kalkan, and E. Akbas, "Imbalance problems in object detection: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Mar. 19, 2020, doi: 10.1109/TPAMI.2020.2981890.
- [37] H. Li, Z. Wu, C. Zhu, C. Xiong, R. Socher, and L. S. Davis, "Learning from noisy anchors for one-stage object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10585–10594.
- [38] S. Zhang, C. Chi, Y. Yao, Z. Lei, and S. 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), Jun. 2020, pp. 9756–9765.
- [39] K. Kim and H. S. Lee, "Probabilistic anchor assignment with iou prediction for object detection," in *Proc. ECCV*, 2020.
- [40] W. Ke, T. Zhang, Z. Huang, Q. Ye, J. Liu, and D. Huang, "Multiple anchor learning for visual object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10206–10215.
- [41] K. Chen et al., "Towards accurate one-stage object detection with APloss," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2019, pp. 5119–5127.
- [42] Q. Qian, L. Chen, H. Li, and R. Jin, "DR loss: Improving object detection by distributional ranking," in *Proc. CVPR*, 2020, pp. 12161–12169.

- [43] K. Oksuz, B. C. Cam, E. Akbas, and S. Kalkan, "A ranking-based, balanced loss function unifying classification and localisation in object detection," in *Proc. NeurIPS*, 2020.
- [44] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 740–755.
- [45] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, Sep. 2009.
- [46] P. F. Felzenszwalb, R. B. Girshick, and D. McAllester, "Cascade object detection with deformable part models," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2010, pp. 2241–2248.
- [47] P. Viola and M. J. Jones, "Robust real-time face detection," Int. J. Comput. Vis., vol. 57, no. 2, pp. 137–154, 2004.
- [48] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 580–587.
- [49] R. Girshick, "Fast R-CNN," in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Dec. 2015, pp. 1440–1448.
- [50] W. Liu et al., "SSD: Single shot multibox detector," in Proc. ECCV, 2016, pp. 21–37.
- [51] K. Oksuz, B. C. Cam, E. Akbas, and S. Kalkan, "Generating positive bounding boxes for balanced training of object detectors," in *Proc.* WACV, 2020, pp. 883–892.
- [52] X. Wang, A. Shrivastava, and A. Gupta, "A-Fast-RCNN: Hard positive generation via adversary for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3039–3048.
- [53] Y. E. Nesterov, "Introductory lectures on convex optimization—A basic course," in *Applied Optimization*, vol. 87. Springer, 2004. [Online]. Available: https://dblp.org/rec/books/sp/Nesterov04.bib, doi: 10.1007/978-1-4419-8853-9.
- [54] Z. Chen, V. Badrinarayanan, C. Lee, and A. Rabinovich, "GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *Proc. ICML*, 2018, pp. 793–802.
- [55] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.* (CVPR), Jun. 2016, pp. 770–778.
- [56] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, and K. He. (2018). *Detectron*. [Online]. Available: https://github.com/ facebookresearch/detectron
- [57] M. Francisco and G. Ross. (2018). MaskrCNN-Benchmark. [Online]. Available: https://github.com/facebookresearch/maskrcnn-benchmark
- [58] R. Cipolla, Y. Gal, and A. Kendall, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7482–7491.
- [59] O. Sener and V. Koltun, "Multi-task learning as multi-objective optimization," in *Proc. NeurIPS*, 2018, pp. 525–536.
- [60] H. Zhang, Y. Wang, F. Dayoub, and N. Sünderhauf, "Varifocal-Net: An IoU-aware dense object detector," in *Proc. CVPR*, 2021, pp. 8514–8523.
- [61] X. Li et al., "Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection," in Proc. NeurIPS, 2020.



Joya Chen (Member, IEEE) received the bachelor's degree from the School of Automotive Engineering, Wuhan University of Technology (WUT), in June 2018. To chase his AI dream, he took the National Postgraduate Entrance Examination and received the 1st place in the School of Computer Science and Technology, University of Science and Technology of China (USTC), in March 2018. He also received the master's degree from USTC in June 2021. His research interests focus on image/video/3D understanding.



Dong Liu (Senior Member, IEEE) received the B.S. and Ph.D. degrees in electrical engineering from the University of Science and Technology of China (USTC), Hefei, China, in 2004 and 2009, respectively. He was a Research Staff Member with Nokia Research Center, Beijing, China, from 2009 to 2012. He joined USTC in 2012 and became a Professor in 2020. He has authored or coauthored more than 100 papers in international journals and conferences. He has 19 granted patents. His research interests include image and video processing, coding, analy-

sis, and data mining. He is a Senior Member of CCF and CSIG, and an Elected Member of MSA-TC of the IEEE CAS Society. He received the 2009 IEEE TCSVT Best Paper Award and the VCIP 2016 Best 10% Paper Award. He serves or had served as the Chair for the IEEE Future Video Coding Study Group (FVC-SG), the Publicity Co-Chair for ICME 2021, the Registration Co-Chair for ICME 2019, and the Symposium Co-Chair for WCSP 2014.



Tong Xu (Member, IEEE) received the Ph.D. degree from USTC in 2016. He is currently working as an Associate Professor with Anhui Province Key Laboratory of Big Data Analysis and Application, University of Science and Technology of China (USTC), Hefei, China. He has authored more than 70 top-tier journal articles and conference papers in the fields of social network and management computing analysis, including IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, IEEE TRANSACTIONS ON MOBILE COMPUTING, IEEE

TRANSACTIONS ON MULTIMEDIA, *ACM TOMM*, KDD, SIGIR, WWW, AAAI, and IJCAI. He was a recipient of the Best Paper Award of KSEM 2020. Besides, he has served on numerous conferences, such as the Track Chair for IEEE ICKG 2020, the Session Chair for CCKS 2018 and SMP 2020, a Co-Organizer for workshops held in KDD 2018-2021 and SDM 2020-2021, and the Program Committee Member for several international conferences, like KDD, AAAI, WWW, and IJCAI.



Yifei Cheng received the B.S. degree from the University of Science and Technology of China (USTC), Hefei, China, in 2015, where he is currently pursuing the Ph.D. degree. He is also working with the Key Laboratory of Big Data Analysis and Application. His major research interests include optimization and federated learning.



Enhong Chen (Senior Member, IEEE) received the B.Sc. degree from Anhui University in 1989, the master's degree from Hefei University of Technology in 1992, and the Ph.D. degree in computer science from USTC in 1996. He is currently a Professor, the Executive Dean of the School of Data Science, and the Vice Dean of the School of Computer Science, University of Science and Technology of China (USTC). He is also the Vice Director of the National Engineer Laboratory for Speech and Language Information Processing, the Director of

Anhui Province Key Laboratory of Big Data Analysis and Application, and the Chairman of Anhui Province Big Data Industry Alliance. His current research interests are data mining and machine learning, especially social network analysis and recommender systems. He has published more than 200 papers on many journals and conferences, including international journals such as IEEE TRANSACTIONS, ACM Transactions, and important data mining conferences, such as KDD, ICDM, and NIPS. He is a CCF Fellow. He is a member of the Decision Advisory Committee of Shanghai Since June 2018. His research was supported by the National Natural Science Foundation of China and the National High Technology Research and Development Program 863 of China. He won the Best Application Paper Award on KDD2008 and the Best Research Paper Award on ICDM2011. He is a Winner of the National Science Fund for Distinguished Young Scholars in 2013 and the Scientific and Technological Innovation Leading Talent of Ten Thousand Talent Program in 2017.



Shiwei Wu received the B.E. degree from Xiamen University (XMU), Xiamen, China, in 2018. He is currently pursuing the Ph.D. degree with USTC, Hefei, China. He is also working with the Key Laboratory of Big Data Analysis and Application. His major research interests include video understanding and multimedia.