

# Localization Distillation for Object Detection

Zhaohui Zheng, Rongguang Ye, Qibin Hou , *Member, IEEE*, Dongwei Ren , *Member, IEEE*, Ping Wang , Wangmeng Zuo , *Senior Member, IEEE*, and Ming-Ming Cheng , *Senior Member, IEEE*

**Abstract**—Previous knowledge distillation (KD) methods for object detection mostly focus on feature imitation instead of mimicking the prediction logits due to its inefficiency in distilling the localization information. In this paper, we investigate whether logit mimicking always lags behind feature imitation. Towards this goal, we first present a novel localization distillation (LD) method which can efficiently transfer the localization knowledge from the teacher to the student. Second, we introduce the concept of valuable localization region that can aid to selectively distill the classification and localization knowledge for a certain region. Combining these two new components, for the first time, we show that logit mimicking can outperform feature imitation and the absence of localization distillation is a critical reason for why logit mimicking under-performs for years. The thorough studies exhibit the great potential of logit mimicking that can significantly alleviate the localization ambiguity, learn robust feature representation, and ease the training difficulty in the early stage. We also provide the theoretical connection between the proposed LD and the classification KD, that they share the equivalent optimization effect. Our distillation scheme is simple as well as effective and can be easily applied to both dense horizontal object detectors and rotated object detectors. Extensive experiments on the MS COCO, PASCAL VOC, and DOTA benchmarks demonstrate that our method can achieve considerable AP improvement without any sacrifice on the inference speed. Our source code and pretrained models are publicly available at <https://github.com/HikariTJU/LD>.

**Index Terms**—Knowledge distillation, localization distillation, object detection, rotated object detection.

## I. INTRODUCTION

AS A model compression technology, knowledge distillation (KD) [24], [58] has been an efficient technique in learning compact models to mitigate the computational burden.

Manuscript received 12 April 2022; revised 4 December 2022; accepted 14 February 2023. Date of publication 24 February 2023; date of current version 30 June 2023. This work was supported in part by the NSFC under Grants 62176130 and 62272311, in part by CAST through Young Elite Scientist Sponsorship Program under Grant YESS20210377, in part by the Fundamental Research Funds for the Central Universities, Nankai University, under Grant 63223049, and in part by the the Natural Science Foundation of Heilongjiang Province under Grant YQ2022F004. Recommended for acceptance by A. Kovashka. (Corresponding author: Qibin Hou.)

Zhaohui Zheng, Qibin Hou, and Ming-Ming Cheng are with the TMCC, CS, Nankai University, Tianjin 300071, China (e-mail: zh\_zheng@mail.nankai.edu.cn; andrewhoux@gmail.com; cmm@nankai.edu.cn).

Rongguang Ye and Ping Wang are with the School of Mathematics, Tianjin University, Tianjin 300072, China (e-mail: ementon@tju.edu.cn; wang\_ping@tju.edu.cn).

Dongwei Ren and Wangmeng Zuo are with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin 150001, China (e-mail: rendongwei@hit.edu.cn; cswmzuo@gmail.com).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2023.3248583>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2023.3248583

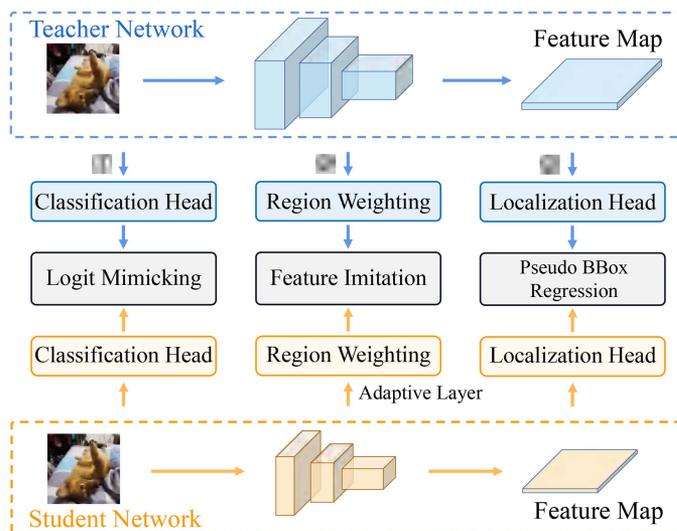


Fig. 1. Existing KD pipelines for object detection. ① Logit Mimicking: classification KD in [24]. ② Feature Imitation: recent popular methods distill intermediate features based on various distillation regions, which usually need adaptive layers to align the size of the student's feature map. ③ Pseudo BBox Regression: treating teachers' predicted bounding boxes as additional regression targets [6], [60].

It has been widely validated to be useful for boosting the performance of small-sized student networks by transferring the generalized knowledge captured by large-sized teacher networks [24], [29], [31], [58], [63], [78]. Speaking of KD in object detection, there are mainly three popular KD pipelines as shown in Fig. 1. Logit mimicking [24], also known as classification KD, is originally designed for image classification, where the KD process operates on the logits of the teacher-student pair. Feature imitation, motivated by the pioneer work FitNet [58], aims to enforce the consistency of the feature representations between the teacher-student pair. The last one, namely the pseudo bounding box regression, uses the predicted bounding boxes from the teacher as an addition supervision to the bounding box prediction branch of the student.

Among these methods, the original logit mimicking technique [24] for classification is often inefficient as it only transfers the classification knowledge while neglects the importance of localization knowledge distillation. Therefore, existent KD methods for object detection mostly focus on feature imitation, and demonstrate that distilling the feature representations is more advantageous than distilling the logits [30], [66], [81]. We summarize three crucial reasons for this phenomenon: First of all, the effectiveness of logit mimicking partially relies on the number of classes which may vary in different application



Fig. 2. Bottom edge for “elephant” and right edge for “surfboard” are ambiguous to locate.

scenarios [66]. Second, the logit mimicking can only be applied to the classification head, which cannot distill the localization information. At last, in the framework of multi-task learning, feature imitation can transfer the hybrid knowledge of classification and localization which can benefit the downstream classification and localization tasks.

In this work, we examine the aforementioned common belief in object detection KD, and challenge whether feature imitation always stays ahead of logit mimicking? For this purpose, we firstly present a simple yet effective localization distillation (LD) method which is inspired by an interesting observation that the bounding box distributions generated by the teacher [38], [51] can serve as a strong supervision to the student detector. The bounding box distribution [38], [51] is originally designed to model the real distributions of bounding boxes, an efficient way to solve the localization ambiguity as shown in Fig. 2. With the discretized probability distribution representations, the localizer can reflect the localization ambiguity by the flatness and sharpness of the distribution, which is not held in the conventional Dirac delta representation of bounding boxes [42], [53], [56], [62]. This allows our LD to efficiently transfer richer localization knowledge from the teacher to the student than using pseudo bounding box regression (right part in Fig. 1).

Combining the proposed LD and the classification KD yields a unified KD method based on a pure logit mimicking framework for both the classification branch and the localization branch. As logit mimicking enables us to separately distill the classification knowledge and the localization one, we found that these two sub-tasks favor different distillation regions. Motivated by this, we introduce the concept of valuable localization region (VLR) and propose to conduct distillation in a selective region distillation manner. We will show the advantage of using VLR in our distillation framework in the experiment section.

Furthermore, we comprehensively discuss the technical details of LD and elaborate on the behavior of logit mimicking and feature imitation. Intriguingly, we observe that logit mimicking can outperform feature imitation for the first time, which indicates that the absence of localization distillation is actually the key reason why logit mimicking under-performs in object detection for years. Another observation is that we find the reason why logit mimicking works is not because of the consistency of the feature representations between the teacher-student pair. Conversely, the student learns significantly different feature representations from the teacher’s in terms of

the  $l_n$  distance and linear correlation. We also observe that if the student is trained with feature imitation, it tends to produce a sharp AP score landscape in the feature subspace, and aggravates the training difficulty in the early training stage.

The above observations reflect the great potentials of logit mimicking over feature imitation:

- 1) being able to separately transfer different types of knowledge,
- 2) learning more robust feature representations, and
- 3) easing the training difficulty.

Our method is simple and can be easily equipped with in both horizontal and rotated object detectors to improve their performance without introducing any inference overhead. Extensive experiments on MS COCO show that without bells and whistles, we can lift the AP score of the strong baseline GFocal [38] with ResNet-50-FPN backbone from 40.1 to 42.1, and  $AP_{75}$  from 43.1 to 45.6. Our best model using ResNeXt-101-32x4d-DCN backbone can achieve a single-scale test of 50.5 AP, which surpasses all existing detectors under the same backbone, neck, and test settings. PyTorch [49] and Jittor [26] version of the source code and pretrained models are publicly available at <https://github.com/HikariTJU/LD>.

The main contributions of this paper are four-fold:

- 1) We present a novel localization distillation method that greatly improve the distillation efficiency of logit mimicking in object detection.
- 2) We provide exploratory experiments and analysis for the behavior of logit mimicking and feature imitation. To our best knowledge, this is the first work revealing the great potential of logit mimicking over feature imitation.
- 3) We present a selective region distillation based on the newly introduced valuable localization region to better distill the student detector.
- 4) We extend our LD to a rotated version so that it can be applied to arbitrary-oriented object detection.

This paper is a substantial extension of its previous conference version [87]. In particular, (a) We provide theoretical connection for the proposed LD and the classification KD that they share the equivalent optimization effects; (b) We conduct more detailed and insightful analysis for logit mimicking and feature imitation, including the different characteristics of the learned feature representations and logits, and the training difficulty of feature imitation; (c) We extend the original LD to a more generic version, namely rotated LD, which can distill arbitrary-oriented object detectors.

## II. RELATED WORK

### A. Knowledge Distillation

Knowledge distillation [2], [24], [46], [48], [59], [78], as a hot research topic, has been deeply studied recently. The fundamental idea is to use a well-performed large-sized teacher network to transfer the captured knowledge to the small-sized student network. Logit mimicking, *a.k.a.* classification KD, was first introduced by Hinton et al. [24], where the logit outputs of the student classifier are supervised by those of the teacher classifier. Later, FitNet [58] extends the teacher-student learning framework by mimicking the intermediate-level hints from the

hidden layers of the teacher model. Knowledge distillation was first applied to object detection in [6], where the hint learning, classification KD, and pseudo bounding box regression were simultaneously used for multi-class object detection. However, an object detector requires not only precise classification ability, but also strong localization ability. The absence of localization knowledge distillation limits the performance of the conventional KD method.

To tackle the above issue, many feature imitation methods have been developed, most of which focus on where to distill and loss function weighting. Among these, Li et al. [36] proposed to mimic the features within the region proposal for Faster R-CNN. Wang et al. [66] imitated the fine-grained features on close anchor box locations. Recently, Dai et al. [10] introduced the General Instance Selection Module to mimic deep features within the discriminative patches between teacher-student pairs. DeFeat [19] leverages different loss weights when conducting feature imitation on the object regions and the background region. There are also various feature imitation methods from the perspective of weighted imitation loss, including Gaussian mask weighted [60], feature richness weighted [88], and prediction-guided imitation loss [34]. Unlike the aforementioned methods, our work introduces localization distillation and demonstrate that logit mimicking can outperform feature imitation for KD in object detection.

### B. Object Localization

Object localization is a fundamental issue in object detection [18], [33], [44], [64], [65], [73], [74], [75], [80], [93]. Bounding box regression is the most popular way so far for localization in object detection [13], [21], [42], [53], [56], where the Dirac delta distribution representation has been used for years. R-CNN series [4], [47], [56], [79] adopt multiple regression stages to refine the detection results, while YOLO series [3], [53], [54], [55], SSD series [14], [42], [89], and FCOS series [38], [62] adopt one-stage regression. In [57], [77], [85], [86], IoU-based loss functions are proposed to improve the localization quality of bounding boxes. Recently, bounding box representation has evolved from Dirac delta distribution [42], [53], [56] to Gaussian distribution [9], [23], and further to probability distribution [38], [51]. The probability distribution of bounding boxes is more comprehensive for describing the uncertainty of bounding boxes, and has been validated to be the most advanced bounding box representation so far.

### C. Localization Quality Estimation

As the name suggests, Localization Quality Estimation (LQE) predicts a score that measures the localization quality of the bounding boxes predicted by the detector. LQE is usually used to cooperate with the classification task during training [37], *i.e.*, enhancing the consistency between classification and localization. It can also be applied in joint decision-making during post-processing [28], [53], [62], which considers both the classification score and LQE when performing NMS. Early research can be dated to YOLOv1 [53], where the predicted object confidence is used to penalize the classification score.

Then, box/mask IoU [27], [28] and box/polar centerness [62], [69] are proposed to model the uncertainty of detections in object detection and instance segmentation, respectively. For bounding box representation, Softer-NMS [23] and Gaussian YOLOv3 [9] predict variances for each edge of the bounding boxes. LQE is a preliminary approach to model localization ambiguity.

### D. Arbitrary-Oriented Object Detection

Driven by the success of object detection, rotated object detection has become a hot topic in computer vision recently [90]. The mainstream rotated object detectors, such as RRPN [45], generate rotated proposals based on Faster R-CNN [56], while Rotated-RetinaNet [40] directly predicts an additional rotated angle based on RetinaNet. To address the boundary discontinuity and square-like problems, SCRDet [74] and RSDet [50] propose IoU-smooth L1 loss and modulated loss respectively for attaining smoother boundary loss, and CSL [72] proposes to use angle classification instead of angle regression.

Different from the horizontal bounding box regression which can easily leverage the IoU-based losses (*e.g.*, GIoU [57], DIoU [85], and CIoU [86]) to enhance localization ability, the Skew IoU loss for rotated bounding box regression is quite difficult to implement due to the complexity of the backward propagation in existing deep learning libraries [1], [26], [49]. PIoU loss [8] approximates the Skew IoU by accumulating the pixels of the intersection and union of two rotated bounding boxes. GWD [73] and KLD [75] model the rotated bounding boxes via the 2D Gaussian Distribution representation and propose to use the Gaussian Wasserstein distance and KL divergence to simulate the Skew IoU loss, respectively. More recently, based on the 2D Gaussian distribution representation of rotated bounding boxes, Yang et al. [76] proposed the KFIOU loss by exploiting the Kalman filter formulation to mimic the Skew IoU in the trend level. To sum up, the rotated regression-based detectors are still dominating this task owing to their simplicity and strong performance.

## III. APPROACH

To begin with, we revisit the knowledge distillation background, including logit mimicking and feature imitation. Next, we describe our simple yet effective localization distillation (LD) and explain how to apply LD for rotated object detection. Then, we analyze the property of the proposed LD loss, especially the theoretical connection to the classification KD. In addition, we also introduces the concept of valuable localization region for better distilling the localization knowledge in our framework. Finally, we describe the selective region distillation based on the newly introduced valuable localization region and give the optimization objective.

### A. Preliminaries

In the KD pipeline of object detection, the input image is fed into two object detectors, *i.e.*, the student detector and the frozen teacher detector. The distillation process forces the outputs of the

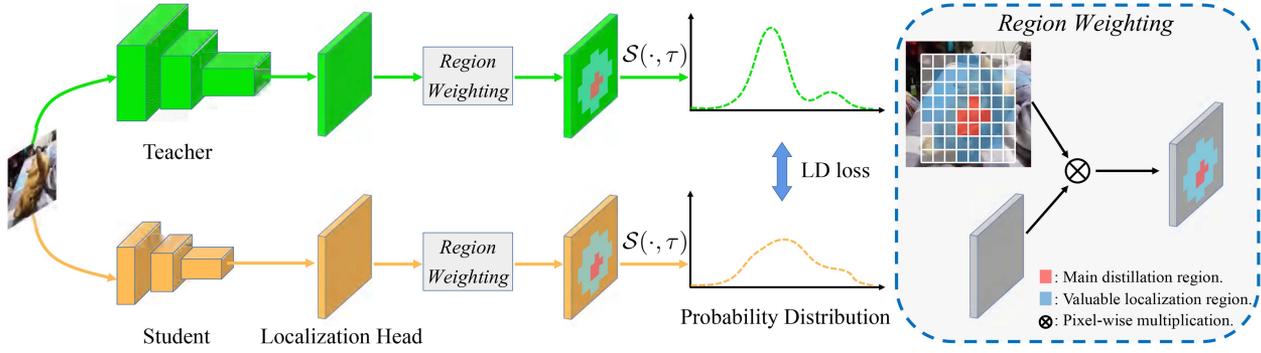


Fig. 3. Illustration of localization distillation (LD) for an edge  $e \in \mathcal{B}$ . Only the localization branch is visualized here.  $\mathcal{S}(\cdot, \tau)$  is the generalized SoftMax function with temperature  $\tau$ . For a given detector, we first switch the bounding box representation to probability distribution. Then, we determine where to distill via region weighting on the main distillation region and the valuable localization region. Finally, we calculate the LD loss between two probability distributions predicted by the teacher and the student.

student to mimic those of the teacher. There are two mainstream paradigms of KD methods in object detection.

*Logit mimicking:* The logit mimicking (LM) is first developed for image classification [24], in which the student model can be improved by mimicking the soft output of the teacher classifier. Let  $z_S, z_T \in \mathbb{R}^{W \times H \times C}$  be the logits predicted by the student and the teacher, respectively.  $W$  and  $H$  represent the output size of the logit maps.  $C$  denotes the number of classes. These logits are then transformed into probability distributions  $p_\tau$  and  $q_\tau$  by using the generalized SoftMax function. We can train the network by minimizing the loss:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{KD}} \quad (1)$$

$$= \mathcal{H}(p, g) + \lambda \mathcal{H}(p_\tau, q_\tau), \quad (2)$$

where  $p$  is the predicted probability vectors,  $g = \{0, 1\}^n$  is the one-hot ground-truth label,  $\mathcal{H}$  is the cross-entropy loss, and  $\lambda$  balances the two loss terms. For object detection, the distillation can be carried out on some pre-defined distillation region  $\mathcal{R}$ .

*Feature imitation.* Recently, it has been found that feature imitation (FI), which aims to transfer knowledge by imitating the deep features between teacher-student pairs, works better than the classification KD [58], [66]. Mathematically, the feature imitation procedure can be formulated as:

$$\mathcal{L}_{\text{FI}} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \|\tilde{M}_S(r) - M_T(r)\|_2, \quad (3)$$

where  $\mathcal{R}$  is the imitation region, and  $|\cdot|$  is the cardinality of the region. Note that an adaptive layer is needed to transform the size of student's feature map  $M_S$  to be the same as the teacher's  $M_T$ , so that  $\tilde{M}_S, M_T \in \mathbb{R}^{W \times H \times D}$ .

*Bounding box representation:* For a given bounding box  $\mathcal{B}$ , the conventional representations have two forms, i.e.,  $\{\delta_x, \delta_y, \delta_w, \delta_h\}$  (encoding the coordinate mappings of the central point, the width and the height from the anchor box to the ground-truth box) [42], [53], [56] and  $\{t, b, l, r\}$  (the distances from the sampled point to the top, bottom, left, and right edges) [62]. These two forms actually follow the Dirac delta distribution that only focuses on the ground-truth locations but cannot model

the ambiguity of bounding boxes as shown in Fig. 2. This is also clearly demonstrated in some previous works [23], [38].

## B. Localization Distillation

In this subsection, we present localization distillation (LD), a new way to enhance the distillation efficiency for object detection. Our LD is evolved from the view of probability distribution representation of bounding boxes (anchor free [38] and anchor-based [51]), which is originally designed for generic object detection and carries abundant localization information. The working principle of our LD can be seen in Fig. 3. The procedure is the same to both anchor-based and anchor-free detectors.

Given an object detector, we follow [38], [51] to convert the bounding box representation from a quaternary representation to a probability distribution. Let  $e \in \mathcal{B}$  be one of the regression variables of bounding box, whose regression range is  $[e_{\min}, e_{\max}]$ . The bounding box distribution quantizes the continuous regression range into a uniform discretized variable  $e = [e_0, e_1, \dots, e_n] \in \mathbb{R}^{n+1}$  with  $n$  sub-intervals, where  $e_0 = e_{\min}$  and  $e_n = e_{\max}$ . The localization head predicts  $n + 1$  logits  $z = \{z_0, z_1, \dots, z_n\}$ , corresponding to the endpoints of the subintervals  $\{e_0, e_1, \dots, e_n\}$ . Each edge of the given bounding box can be represented as a probability distribution by using the SoftMax function. For the number of the subinterval  $n$ , we follow the settings of GFocal [38], and a recommended choice of  $n$  is  $8 \sim 16$ . Different from [38], [51], we transform  $z_S$  and  $z_T$  into the probability distributions  $p_\tau$  and  $q_\tau$  using the generalized SoftMax function  $\mathcal{S}(\cdot, \tau)$ . Note that when  $\tau = 1$ , it is equivalent to the original SoftMax function. When  $\tau \rightarrow 0$ , it tends to be a Dirac delta distribution. When  $\tau \rightarrow \infty$ , it will be a uniform distribution. Empirically,  $\tau > 1$  is set to soften the distribution, making the bounding box distribution carry more information. The localization distillation for measuring the similarity between the two probability vectors  $p_\tau, q_\tau \in \mathbb{R}^n$  for one of the bounding box representation  $e$  is attained by:

$$\mathcal{L}_{\text{LD}}^e = \mathcal{H}(p_\tau, q_\tau) \quad (4)$$

$$= \mathcal{H}(\mathcal{S}(z_S, \tau), \mathcal{S}(z_T, \tau)). \quad (5)$$

Then, LD for all the four edges of some bounding box  $\mathcal{B}$  can be formulated as:

$$\mathcal{L}_{LD}(\mathcal{B}_S, \mathcal{B}_T) = \sum_{e \in \mathcal{B}} \mathcal{L}_{LD}^e, \quad (6)$$

where  $\mathcal{B}_S, \mathcal{B}_T$  are the predicted bounding boxes of the student and the teacher, respectively.

### C. Rotated LD

Our LD can also be flexibly used to distill rotated bounding box detectors. Parametric regression is the most popular manner in the classical dense regression-based rotated object detection [71], [73], [74], [75].  $\mathcal{B} = \{\delta_x, \delta_y, \delta_w, \delta_h, \delta_\theta\}$  is commonly used to represent a rotated bounding box, where  $\delta_\theta$  denotes the encoded rotated angle. To conduct rotated localization distillation, we firstly generate the lower and upper bounds of the regression range  $[e_{\min}, e_{\max}]$ , where  $e \in \mathcal{B}$ .

Note that the rotated angle prediction  $\delta_\theta$  usually has a different regression range from  $\delta_x, \delta_y, \delta_w, \delta_h$ . Thus, different lower and upper bounds of regression ranges are set for them. In practice,  $[e_{\min}, e_{\max}] \subset [-5, 5]$  will be an acceptable choice. Then, we convert the rotated bounding box to rotated bounding box distributions, as Section III-B describes. Finally, the LD loss is calculated according to (6) for the rotated bounding box distributions.

### D. Property of LD

We can see that our LD holds the formulation of the standard logit mimicking. The question one may ask is: Does LD also inherit the property of the classification KD, especially for the optimization process? Different from the classification task where a unique integer is treated as the ground-truth label, the ground-truth label of the localization task is a float point number  $e^*$ , whose value, for instance, is ranged in an interval  $[e_i, e_{i+1}]$ . In the following, we show an important property of LD, demonstrating that it can inherit the optimization effects held by the classification KD.

*Proposition 1:* Let  $\mathbf{s}$  be the student's predicted probability vector, and  $u_1, u_2$  are two constants with  $u_1 + u_2 = 1$ . Then, we have:

- 1) If  $\mathbf{p}, \mathbf{q}$  are two classification probability vectors, LD effect on the linear combination  $\mathbf{l} = u_1\mathbf{p} + u_2\mathbf{q}$  is equal to the linear combination of KD effects on  $\mathbf{p}, \mathbf{q}$ ;
- 2) If  $\mathbf{l}$  is a localization probability vector, LD effect on  $\mathbf{l}$  is equal to two KD effects on its decomposition  $\mathbf{p}$  and  $\mathbf{q}$ .

The above two share the same expression,

$$\partial LD_i^l = u_1 \partial KD_i^p + u_2 \partial KD_i^q, \quad (7)$$

where  $\partial KD_i^p$  denotes the derivatives of the KD loss of two probabilities  $\mathbf{s}, \mathbf{p}$  w.r.t. a given logit  $z_i$ , and  $\partial LD_i^p$  likewise for the LD loss.

The proof can be found in the supplementary materials. Proposition 1 provides the theoretical connection between LD and the classification KD. It shows that the optimization effects of LD on a float point number localization problem is functionally equivalent to two KD effects on the integer position classification

---

### Algorithm 1: Valuable Localization Region.

---

**Require:** A set of anchor boxes  $\mathcal{B}^a = \{\mathcal{B}_i^a\}$  and a set of ground truth boxes  $\mathcal{B}^{gt} = \{\mathcal{B}_j^{gt}\}$ ,  $1 \leq i \leq I$ ,  $1 \leq j \leq J$ . Positive threshold  $\alpha_{pos}$  of label assignment.

**Ensure:**  $\mathbf{V} = \{v_{ij}\}_{I \times J}$ ,  $v_{ij} \in \{0, 1\}$  encodes final location of VLR, where 1 denotes VLR and 0 indicates ignore.

1: Compute DIoU matrix  $\mathbf{X} = \{x_{ij}\}_{I \times J}$  with

$$x_{ij} = DIoU(\mathcal{B}_i^a, \mathcal{B}_j^{gt}).$$

2:  $\alpha_{vl} = \gamma \alpha_{pos}$ .

3: Select locations with  $\mathbf{V} = \{\alpha_{vl} \leq \mathbf{X} \leq \alpha_{pos}\}$ .

4: **return**  $\mathbf{V}$

---

problems. Therefore, as a direct corollary of [61], LD holds the gradient rescaling to the distribution focal loss (DFL) [38] w.r.t. the relative prediction confidence at two near positions. For the details, we refer to the supplementary materials.

### E. Valuable Localization Region

Previous works mostly force the deep features of the student to mimic those of the teacher by minimizing the  $l_2$  loss. However, a straightforward question arises: Should we use the whole imitation regions without discrimination to distill the hybrid knowledge? According to our observation, the answer is no. In this subsection, we describe the valuable localization region (VLR) to further improve the distillation efficiency, which we believe will be a promising way to train better student detectors.

Specifically, the distillation region is divided into two parts, the main distillation region and the valuable localization region. The main distillation region is intuitively determined by label assignment, *i.e.*, the positive locations of the detection head. The valuable localization region can be obtained by Algorithm 1. First, we calculate the DIoU [85] matrix  $\mathbf{X}$  between all the anchor boxes  $\mathcal{B}^a$  and the ground-truth boxes  $\mathcal{B}^{gt}$ . Then, we set the lower bound of DIoU to be  $\alpha_{vl} = \gamma \alpha_{pos}$ , where  $\alpha_{pos}$  is the positive IoU threshold of label assignment. The VLR can be defined as  $\mathbf{V} = \{\alpha_{vl} \leq \mathbf{X} \leq \alpha_{pos}\}$ . Our method has only one hyperparameter  $\gamma \leq 1$ , which controls the range of the VLRs. When  $\gamma = 0$ , all the locations whose DIoUs between the preset anchor boxes and the GT boxes satisfy  $0 \leq x_{ij} \leq \alpha_{pos}$  will be determined as VLRs. When  $\gamma \rightarrow 1$ , the VLR will gradually shrink to empty. Here we use DIoU [85] since it gives higher priority to the locations close to the center of the object.

Similar to label assignment, our method assigns attributes to each location across multi-level FPN. In this way, some of locations outside the GT boxes will also be considered. So, we can actually view the VLR as an outward extension of the main distillation region. Note that for anchor-free detectors, like FCOS, we can use the preset anchors on feature maps and do not change its regression form, so that the localization learning maintains to be the anchor-free type. While for anchor-based detectors which usually set multiple anchors per location, we unfold the anchor boxes to calculate the DIoU matrix, and then assign their attributes.

### F. Selective Region Distillation

Given the above descriptions, the total loss of logit mimicking for training the student  $\mathcal{S}$  can be represented as:

$$\begin{aligned} \mathcal{L} = & \lambda_0 \mathcal{L}_{\text{cls}}(\mathcal{C}_S, \mathcal{C}^{gt}) + \lambda_1 \mathcal{L}_{\text{reg}}(\mathcal{B}_S, \mathcal{B}^{gt}) + \lambda_2 \mathcal{L}_{\text{DFL}}(\mathcal{B}_S, \mathcal{B}^{gt}) \\ & + \lambda_3 \mathbb{I}_{\text{Main}} \mathcal{L}_{\text{LD}}(\mathcal{B}_S, \mathcal{B}_T) + \lambda_4 \mathbb{I}_{\text{VL}} \mathcal{L}_{\text{LD}}(\mathcal{B}_S, \mathcal{B}_T) \\ & + \lambda_5 \mathbb{I}_{\text{Main}} \mathcal{L}_{\text{KD}}(\mathcal{C}_S, \mathcal{C}_T) + \lambda_6 \mathbb{I}_{\text{VL}} \mathcal{L}_{\text{KD}}(\mathcal{C}_S, \mathcal{C}_T), \end{aligned} \quad (8)$$

where the first three terms are exactly the same to the classification and bounding box regression branches for any regression-based detector, *i.e.*,  $\mathcal{L}_{\text{cls}}$  is the classification loss,  $\mathcal{L}_{\text{reg}}$  is the bounding box regression loss and  $\mathcal{L}_{\text{DFL}}$  is the distribution focal loss [38].  $\mathbb{I}_{\text{Main}}$  and  $\mathbb{I}_{\text{VL}}$  are the distillation masks for the main distillation region and the valuable localization region, respectively.  $\mathcal{L}_{\text{KD}}$  is the KD loss [24],  $\mathcal{C}_S$  as well as  $\mathcal{C}_T$  denote the classification head output logits of the student and the teacher, respectively, and  $\mathcal{C}^{gt}$  is the ground-truth class label.

All the distillation losses will be weighted by the same weight factors according to their types. More clearly, the weight factor of the LD loss follows that of the bbox regression term and the weight factor of the KD loss follows that of the classification term. Also, it is worth mentioning that the DFL loss term can be disabled since LD loss has sufficient guidance ability. In addition, we can enable or disable the four types of distillation losses so as to distill the student in different regions selectively.

## IV. EXPERIMENT

In this section, we conduct comprehensive ablation studies and analysis to demonstrate the superiority of the proposed LD and distillation scheme on the challenging large-scale MS COCO [41] benchmark, PASCAL VOC [12], and aerial image DOTA dataset [68].

### A. Experiment Setup

**MS COCO:** The train2017 (118 K images) is utilized for training and val2017 (5 K images) is used for validation. We also obtain the evaluation results on MS COCO test-dev 2019 (20 K images) by submitting to the COCO server. The experiments are conducted under the mmDetection [7] framework. Unless otherwise stated, we use ResNet [22] with FPN [39] as our backbone and neck networks, and the FCOS-style [62] anchor-free head for classification and localization. The training schedule for ablation experiments is set to single-scale  $1 \times$  mode (12 epochs). For other training and testing hyper-parameters, we follow exactly the GFocal [38] protocol, including QFL loss for classification and GIoU loss for bbox regression, *etc.* We use the standard COCO-style measurement, *i.e.*, average precision (AP), for evaluation. All the baseline models are retrained by adopting the same settings so as to fairly compare them with our LD.

**PASCAL VOC:** We also provide experimental results on another popular object detection benchmark, *i.e.*, PASCAL VOC [12]. We use the VOC 07+12 training protocol, *i.e.*, the union of VOC 2007 trainval set and VOC 2012 trainval set (16551 images) for training, and VOC 2007 test set (4952

images) for evaluation. The initial learning rate is 0.01 and the total training epochs are set to 4. The learning rate decreases by a factor of 10 after the 3 rd epoch. For comprehensively evaluating the localization performance, the average precision (AP) along with 5 mAP across different IoU thresholds are reported, *i.e.*, AP<sub>50</sub>, AP<sub>60</sub>, AP<sub>70</sub>, AP<sub>80</sub> and AP<sub>90</sub>.

**DOTA.** As for the evaluation of rotated LD, we report the detection results on the classic aerial image dataset DOTA [68]. We follow the standard mmRotate [90] training and testing protocol. The train set and validation set consist of 1403 images and 468 images, respectively, which are randomly selected in literature. These huge images are cropped into smaller subimages with shape  $600 \times 600$ , which is in line with the cropping protocol in official implementation. In practice, we obtain about 15,700 training and 5,300 validation patches. Unless otherwise stated, all the hyper-parameters follow the default settings of mmRotate for a fair comparison. We report results in terms of AP and 5 mAPs under different IoU thresholds, which is consistent with PASCAL VOC. Due to the memory limitation, the teachers are ResNet-34 FPN with  $2 \times$  training schedule (24 epochs), and the students are ResNet-18 FPN with  $1 \times$  training schedule (12 epochs).

### B. Ablation Study

**Temperature  $\tau$  in LD:** Our LD introduces a hyper-parameter, *i.e.*, the temperature  $\tau$ . Table I(a) reports the results of LD with various temperatures, where the teacher model is ResNet-101 with AP 44.7 and the student model is ResNet-50. Here, only the main distillation region is adopted. Compared to the first row in Table I(a), different temperatures consistently lead to better results. In this paper, we simply set the temperature in LD as  $\tau = 10$ , which is fixed in all the other experiments.

**LD vs . Pseudo BBox Regression:** The teacher bounded regression (TBR) loss [6] is a preliminary attempt to enhance the student on the localization head, *i.e.*, the pseudo bbox regression in Fig. 1, which is represented as:

$$\mathcal{L}_{\text{TBR}} = \lambda \mathcal{L}_{\text{reg}}(\mathcal{B}^s, \mathcal{B}^{gt}), \text{ if } \ell_2(\mathcal{B}^s, \mathcal{B}^{gt}) + \varepsilon > \ell_2(\mathcal{B}^t, \mathcal{B}^{gt}), \quad (9)$$

where  $\mathcal{B}^s$  and  $\mathcal{B}^t$  denote the predicted boxes of student and teacher respectively,  $\mathcal{B}^{gt}$  denotes the ground truth boxes,  $\varepsilon$  is a predefined margin, and  $\mathcal{L}_{\text{reg}}$  represents the GIoU loss [57]. Here, only the main distillation region is adopted. From Table I(b), we can see that the TBR loss does yield performance gains (+0.4 AP and +0.7 AP<sub>75</sub>) when using proper threshold  $\varepsilon = 0.1$  in (9). However, it uses the coarse bbox representation, which does not contain any localization uncertainty information of the detector, leading to sub-optimal results. On the contrary, our LD directly produces 41.1 AP and 44.9 AP<sub>75</sub>, since it utilizes the probability distribution of bounding boxes which contains rich localization knowledge.

**Various  $\gamma$  in VLR:** The newly introduced VLR has the parameter  $\gamma$  which controls the range of VLR. As shown in Table I(c), AP is stable when  $\gamma$  ranges from 0 to 0.5. The variation in AP in this range is around 0.1. As  $\gamma$  increases, the VLR gradually shrinks to empty. The performance also gradually drops to 41.1,

TABLE I  
ABLATIONS. WE SHOW ABLATION EXPERIMENTS FOR LD AND VLR ON MS COCO val2017

$\tau$	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	$\varepsilon$	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>	$\gamma$	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
–	40.1	58.2	43.1	23.3	44.4	52.5	–	40.1	58.2	43.1	23.3	44.4	52.5	–	40.1	58.2	43.1	23.3	44.4	52.5
1	40.3	58.2	43.4	22.4	44.0	52.4	0.1	40.5	58.3	43.8	23.0	44.2	52.7	1	41.1	58.7	44.9	23.8	44.9	53.6
5	40.9	58.2	44.3	23.2	45.0	53.2	0.2	40.2	58.2	43.6	23.1	44.0	53.0	0.75	41.2	58.8	44.9	23.6	45.4	53.5
10	<b>41.1</b>	<b>58.7</b>	<b>44.9</b>	<b>23.8</b>	44.9	<b>53.6</b>	0.3	40.1	58.4	43.1	23.6	43.9	52.5	0.5	41.7	59.4	45.3	24.2	45.6	54.2
15	40.7	58.5	44.2	23.5	44.3	53.3	0.4	40.3	58.4	43.4	22.8	44.0	52.6	0.25	<b>41.8</b>	<b>59.5</b>	<b>45.4</b>	24.2	45.8	<b>54.9</b>
20	40.5	58.3	43.7	<b>23.8</b>	44.1	53.5	LD	<b>41.1</b>	<b>58.7</b>	<b>44.9</b>	<b>23.8</b>	<b>44.9</b>	<b>53.6</b>	0	41.7	<b>59.5</b>	<b>45.4</b>	<b>24.5</b>	<b>45.9</b>	54.0

(a) **Temperature  $\tau$  in LD:** The generalized Softmax function with large  $\tau$  brings considerable gains. We set  $\tau = 10$  by default. The teacher is ResNet-101 and the student is ResNet-50.

(b) **LD vs. Pseudo BBox Regression [6]:** The localization knowledge can be more efficiently transferred by our LD compared to the pseudo bbox regression. The teacher is ResNet-101 and the student is ResNet-50.

(c) **Role of  $\gamma$  in VLR:** Conducting LD on valuable localization region has a positive effect on performance. We set  $\gamma = 0.25$  by default. The teacher is ResNet-101 and the student is ResNet-50.

TABLE II  
EVALUATION OF **SELECTIVE REGION DISTILLATION** FOR KD AND OUR LD. THE TEACHER-STUDENT PAIR IS RESNET-101→RESNET-50 FOR COCO, AND RESNET-101→RESNET-18 FOR VOC 07+12

LD		KD		MS COCO val2017			VOC 07+12		
Main	VLR	Main	VLR	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>
				40.1	58.2	43.1	51.8	75.8	56.3
✓				41.1	58.7	44.9	53.0	75.9	57.6
✓	✓			41.8	59.5	45.4	53.4	76.3	<b>58.3</b>
✓	✓	✓		<b>42.1</b>	<b>60.3</b>	<b>45.6</b>	53.1	76.8	57.6
✓	✓	✓	✓	42.0	60.0	45.4	<b>53.7</b>	<b>77.3</b>	58.2

*i.e.*, conducting LD on the main distillation region only. The sensitivity analysis experiments on the parameter  $\gamma$  indicate that conducting LD on the VLR has a positive effect on performance. In the rest experiments, we set  $\gamma$  to 0.25 for simplicity.

**Selective Region Distillation:** There are several interesting observations regarding the roles of KD and LD and their preferred regions. We report the relevant ablation study results in Table II, where “Main” means that the logit mimicking is conducted on the main distillation region, *i.e.*, the positive locations of label assignment, and “VLR” denotes the valuable localization region. For MS COCO, it can be seen that conducting “Main LD,” “VLR LD,” and “Main KD” all benefits the student’s performance. This indicates that the main distillation regions contain the valuable knowledge for both classification and localization and the classification KD benefits less compared to LD. Then, we impose the classification KD on a larger range, *i.e.*, the VLR. However, we observe that further incorporating “VLR KD” yields no improvement (the last two rows of Table II). This is the main reason why we adopt the proposed selective region distillation for COCO.

Next, we check the roles of KD and LD on PASCAL VOC. Table II shows that it is beneficial to transfer the localization knowledge to both the main distillation region and the VLR. However, due to the different knowledge distribution patterns, it shows a similar degradation of the classification KD. Comparing the 3rd row and the 4th row of Table II, “Main KD” leads to a performance drop, while “VLR KD” produces a positive effect to the student. This indicates that the selective region distillation can take the advantages of both KD and LD on their respective favorable regions.

TABLE III  
QUANTITATIVE RESULTS OF LD FOR LIGHTWEIGHT DETECTORS. THE TEACHER IS RESNET-101. THE RESULTS ARE REPORTED ON MS COCO val2017

Student	LD	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
ResNet-18		35.8	53.1	38.2	18.9	38.9	47.9
	✓	37.5	54.7	40.4	20.2	41.2	49.4
ResNet-34		38.9	56.6	42.2	21.5	42.8	51.4
	✓	41.0	58.6	44.6	23.2	45.0	54.2
ResNet-50		40.1	58.2	43.1	23.3	44.4	52.5
	✓	42.1	60.3	45.6	24.5	46.2	54.8

TABLE IV  
QUANTITATIVE RESULTS OF LD ON VARIOUS POPULAR DENSE OBJECT DETECTORS. THE TEACHER IS RESNET-101 AND THE STUDENT IS RESNET-50. THE RESULTS ARE REPORTED ON MS COCO val2017

Student	LD	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
RetinaNet [40]		36.9	54.3	39.8	21.2	40.8	48.4
	✓	39.0	56.4	42.4	23.1	43.2	51.1
FCOS [62]		38.6	57.2	41.5	22.4	42.2	49.8
	✓	40.6	58.4	44.1	24.3	44.1	52.3
ATSS [84]		39.2	57.3	42.4	22.7	43.1	51.5
	✓	41.6	59.3	45.3	25.2	45.2	53.3

**LD for Lightweight Detectors:** Table III reports the results of our distillation scheme (“Main LD + VLR LD + Main KD” on COCO), where a series of lightweight students are distilled, including ResNet-18, ResNet-34, and ResNet-50. For all given students, our LD can stably improve the detection performance without any bells and whistles. From these results, we can see that our LD improves the students ResNet-18, ResNet-34, ResNet-50 by +1.7, +2.1, +2.0 in AP, and +2.2, +2.4, +2.5 in AP<sub>75</sub>, respectively.

**Application to Other Dense Object Detectors:** Our LD can be flexibly applied to other dense object detectors, including either anchor-based or anchor-free types. We employ LD with the divide-and-conquer distillation scheme to several recently popular detectors, such as RetinaNet [40] (anchor-based), FCOS [62] (anchor-free) and ATSS [83] (anchor-based). According to the results in Table IV, we can see that our LD can consistently improve the baselines by around 2 AP scores.

**Arbitrary-Oriented Object Detectors:** As a direct extension of our LD, the rotated bounding box requires an additional

TABLE V  
QUANTITATIVE RESULTS OF ROTATED LD ON THE POPULAR  
ARBITRARY-ORIENTED OBJECT DETECTORS. THE TEACHER IS RESNET-34 AND  
THE STUDENT IS RESNET-18. THE RESULTS ARE REPORTED ON THE  
VALIDATION SET OF DOTA-V1.0

Student	AP	AP <sub>50</sub>	AP <sub>60</sub>	AP <sub>70</sub>	AP <sub>80</sub>	AP <sub>90</sub>
R-RetinaNet [40]	33.7	58.0	54.5	42.3	22.9	4.7
LD (ours)	39.1	63.8	61.1	48.8	28.7	8.8
GWD [73]	37.1	63.1	60.1	46.7	24.7	6.2
LD (ours)	40.2	66.4	63.6	50.3	28.2	8.5

probability distribution, i.e., the rotated angle distribution. We make the necessary and minimum modification to two arbitrary-oriented object detectors, 1) the foundation of dense regression-based rotated detector—Rotated-RetinaNet [40] and 2) the recently popular 2D Gaussian distribution modeling detector—GWD [73]. We follow the mmRotate [90] training and testing protocols. We use ResNet-34 as the teacher and ResNet-18 as the student for GPU memory saving. The results are reported on the validation set of DOTA-v1.0 [68].

The results have been shown in Table V, which demonstrates that our LD can also be successfully applied to rotated object detectors and attain considerable improvement in aerial image detection. Particularly, we obtain impressive improvements for the mAP under more rigorous IoU thresholds, e.g., AP<sub>70</sub>, AP<sub>80</sub>, AP<sub>90</sub>. This shows the excellent compatibility of our LD, which can be applied to not only horizontal bounding boxes but also the rotated ones. In addition, it is worth mentioning that our LD does not rely on the representations of bounding boxes and the optimization way of modeling (IoU-based loss for horizontal bounding box prediction [57], [85] and 2D Gaussian modeling for rotated bounding box prediction [73]).

### C. Logit Mimicking Vs. Feature Imitation.

Thus far, we have validated the effectiveness of our LD and the selective region distillation in distilling different types of object detectors. The proposed LD along with the classification KD provides a unified logit mimicking framework. It naturally raises several interesting questions:

- In terms of detection performance, how does logit mimicking perform compared to feature imitation? Does feature imitation stay ahead of logit mimicking?
- What are the characteristics of these two different distillation techniques? Are the deep feature representations and logits learned different?

In this subsection, we shall provide answers to the above questions.

*Quantitative Comparison on Numerical Results:* We first compare our proposed LD with several state-of-the-art feature imitation methods. We adopt the selective region distillation, i.e., performing KD and LD on the main distillation region, and performing LD on the VLR. Since modern detectors are usually equipped with FPN [39], following previous works [10], [19], [66], we re-implement their methods and impose all the feature imitations on multi-level FPN for a fair comparison.

TABLE VI  
LOGIT MIMICKING Vs. FEATURE IMITATION. “OURS” MEANS WE USE THE  
SELECTIVE REGION DISTILLATION, I.E., “MAIN LD + VLR LD + MAIN KD”.  
“\*” DENOTES WE REMOVE THE “MAIN KD”. THE TEACHER IS RESNET-101  
AND THE STUDENT IS RESNET-50 [22]. THE RESULTS ARE REPORTED ON MS  
COCO VAL2017

Method	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
Baseline (GFocal [38])	40.1	58.2	43.1	23.3	44.4	52.5
FitNets [58]	40.7	58.6	44.0	23.7	44.4	53.2
Inside GT Box	40.7	58.6	44.2	23.1	44.5	53.5
Main Region	41.1	58.7	44.4	24.1	44.6	53.6
Fine-Grained [66]	41.1	58.8	44.8	23.3	45.4	53.1
DeFeat [19]	40.8	58.6	44.2	24.3	44.6	53.7
GI Imitation [10]	41.5	59.6	45.2	24.3	45.7	53.6
Ours	42.1	60.3	45.6	24.5	46.2	54.8
Ours + FitNets	42.1	59.9	45.7	25.0	46.3	54.4
Ours + Inside GT Box	42.2	60.0	45.9	24.3	46.3	55.0
Ours + Main Region	42.1	60.0	45.7	24.6	46.3	54.7
Ours + Fine-Grained	42.4	60.3	45.9	24.7	46.5	55.4
Ours* + Fine-Grained	42.1	59.7	45.6	24.8	46.1	54.8
Ours + DeFeat	42.2	60.0	45.8	24.7	46.1	54.4
Ours + GI Imitation	42.4	60.3	46.2	25.0	46.6	54.5

Here, “FitNets” [58] distills the whole feature maps. “DeFeat” [19] means the loss weights of feature imitation outside the GT boxes are larger than those inside the GT boxes. “Fine-Grained” [66] distills the deep features on the close anchor box locations. “GI Imitation” [10] selects the distillation regions according to the discriminative predictions of the student and the teacher. “Inside GT Box” means we select the ground-truth boxes with the same stride on the FPN layers as the feature imitation regions. “Main Region” means we imitate the features within the main distillation region.

From Table VI, we can see that distillation within the whole feature maps attains +0.6 AP gains. By setting a larger loss weight for the locations outside the GT boxes (DeFeat [19]), the performance is slightly better than that using the same loss weight for all locations. Fine-Grained [66] focusing on the locations near GT boxes, produces 41.1 AP, which is comparable to the results of feature imitation using the Main Region. GI imitation [10] searches the discriminative patches for feature imitation and gains 41.5 AP. Due to the large gap in predictions between student and teacher, the imitation regions may appear anywhere.

Despite the notable improvements of these feature imitation methods, they do not explicitly consider the knowledge distribution patterns. On the contrary, our method can transfer the knowledge via a selective region distillation, which directly produces 42.1 AP. It is worth noting that our method operates on logits instead of deep features, indicating that our LD is a critical component for logit mimicking to outperform the feature imitation. Moreover, our method is orthogonal to the aforementioned feature imitation methods. Table VI shows that with these feature imitation methods, our performance can be further improved. Particularly, with GI imitation, we improve the strong GFocal baseline by +2.3 AP and +3.1 AP<sub>75</sub>.

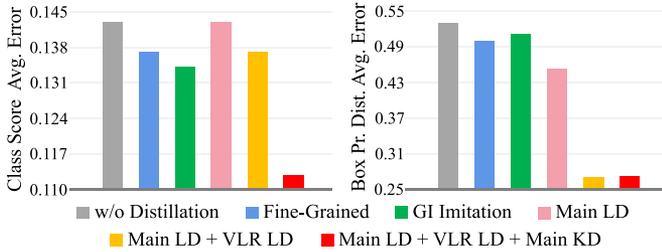


Fig. 4. Visual comparisons of SOTA feature imitation and our LD. We show the average L1 error of classification scores and box probability distributions between teacher and student at the P4, P5, P6 and P7 FPN levels. The teacher is ResNet-101 and the student is ResNet-50. The results are evaluated on MS COCO val2017.

*Teacher-Student Error Comparison:* We first check the average teacher-student errors of the classification scores and the box probability distributions, as shown in Fig 4. One can see that the Fine-Grained feature imitation [66] and GI imitation [10] reduce the two errors as expected, since the classification knowledge and localization knowledge are mixed on feature maps. Our “Main LD” and “Main LD + VLR LD” have comparable or larger classification score average errors than Fine-Grained [66] and GI imitation [10] but lower box probability distribution average errors. This indicates that these two settings with only LD can significantly reduce the box probability distribution distance between the teacher and the student but they cannot reduce this error for the classification head. If we impose the classification KD on the main distillation region, yielding “Main LD + VLR LD + Main KD,” both the classification score average error and the box probability distribution average error can be reduced.

We also visualize the L1 error summation of the localization head logits between the student and the teacher for each location at the P5 and P6 FPN levels. As shown in Fig. 5, comparing to “Without Distillation,” we can see that the GI imitation [10] does decrease the localization discrepancy between the teacher and the student. Notice that we particularly choose a model (“Main LD + VLR LD”) with slightly better AP performance than GI imitation for visualization. Our method can clearly reduce this error and alleviate the localization ambiguity.

In Fig. 6, we plot the average errors between the student and the teacher in terms of deep feature, class logit and bbox logit, respectively. It can be seen that these three types of errors show an almost consistent trend as the test resolution changes. Interestingly, we find that even though the logit mimicking can shrink the errors of both the bbox logits and the classification ones, it learns complete different feature representations from the teacher’s. From the left side of Fig. 6, our method enlarges the distance between the student’s feature representations and those of the teacher. Moreover, Table VII shows that the logit mimicking produces a nearly zero Pearson correlation coefficient for the feature representations between the teacher-student pair. This indicates that if the student is only trained with logit mimicking, it produces a far different and nonlinearly correlated feature representation to teacher’s. Be that as it may, we can still attain well-performed logits for good generalization. The last column of Table VII and Fig. 6 show that the logit mimicking

TABLE VII  
THE AVERAGE PEARSON CORRELATION COEFFICIENT BETWEEN THE TEACHER-STUDENT PAIR. ‘GI’: GI IMITATION. ‘OURS’: OUR LOGIT MIMICKING SCHEME WITH THE SELECTIVE REGION DISTILLATION. THE RESULTS ARE EVALUATED ON MS COCO val2017

	w/o distillation	GI	Ours	Ours + GI
deep features	-0.0042	0.8175	-0.0031	0.8373
bbox logits	0.9222	0.9326	0.9733	0.9745

is able to approach the teacher’s logits not only in distance but also in linear correlation.

*AP Landscape:* Distilling an object detector from either the feature level or the logit level is a high-dimensional non-convex optimization problem, which is easy in practice but hard in theory. To better understand the behavior of logit mimicking and feature imitation, we present a new visualization method, termed AP landscape, which is especially designed for object detection to observe the AP changes caused by minute perturbations in the learnt feature representations. A canonical approach was taken in [35], who studied the loss surface visualization by linearly interpolating the parameters of two networks.

In our visualization, we are particularly curious about the empirical characterization of the feature representations and how they affect the final performance. Considering two feature representations  $M_f, M_l$  which are learnt by the detectors trained with feature imitation and logit mimicking, respectively, we visualize the AP landscapes within the 2D projected space  $M_f \oplus M_l$ . We use two scalar parameters  $\alpha$  and  $\beta$  to obtain a new feature representation by using the weighted sum  $M(\alpha, \beta) = \alpha M_f + \beta M_l$ . Note that when  $\alpha = 0$  and  $\beta = 1$ , it represents that the feature representations are predicted by the logit mimicking method and inversely the feature imitation when  $\alpha = 1$  and  $\beta = 0$ . Then, we feed  $M(\alpha, \beta)$  to the downstream heads and plot the final AP score. Due to the computational burden, we set  $\alpha, \beta \in [-0.5, 1.5]$  to visualize the 2D AP landscapes.

From Fig. 7, we see that logit mimicking learns robust feature representations, i.e., the red pentagram at (0,1), which is surrounded by a flat and well-performed region of AP score. Second, we observe that the GI imitation produces a much sharper AP landscape than logit mimicking. We attribute the landscape sharpness of the GI imitation to the hard  $l_2$  loss supervision. In this case, it is hard for the student to imitate the high-level and advanced feature representations from the teacher, which corresponds to a heavy detector with a longer training schedule and higher accuracy. On the contrary, the logit mimicking gives the feature representations much more liberty to learn, leading to a better generalization. As shown in Fig. 8, logit mimicking can also reduce the optimization difficulty in the early training stage, while feature imitation converges slower and has a poor generalization in the early training stage.

*Summary:* Based on the above results and observations, we can draw the following conclusions:

- Logit mimicking can outperform feature imitation in object detection when the localization knowledge distillation is explicitly distilled.

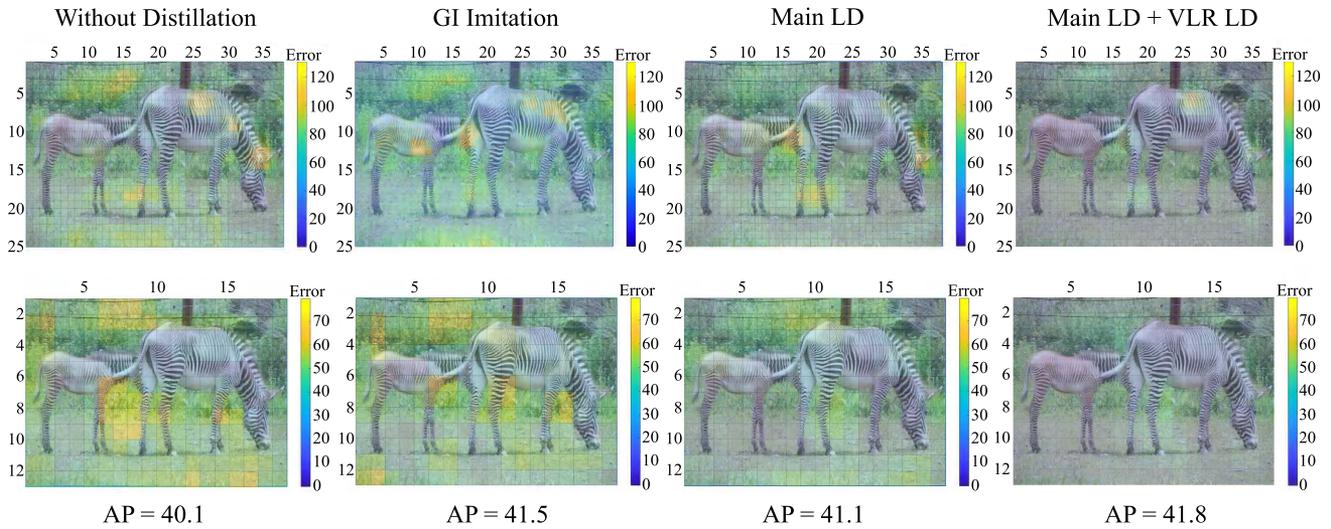


Fig. 5. Visual comparisons between the state-of-the-art feature imitation and our LD. We show the per-location L1 error summation of the localization heads logits between the teacher and the student as the P5 (first row) and P6 (second row) FPN levels. The teacher is ResNet-101 and the student is ResNet-50. We can see that compared to the GI imitation [10], our method (“Main LD + VLR LD”) can significantly reduce the errors for almost all the locations. Darker is better. Best viewed in color.

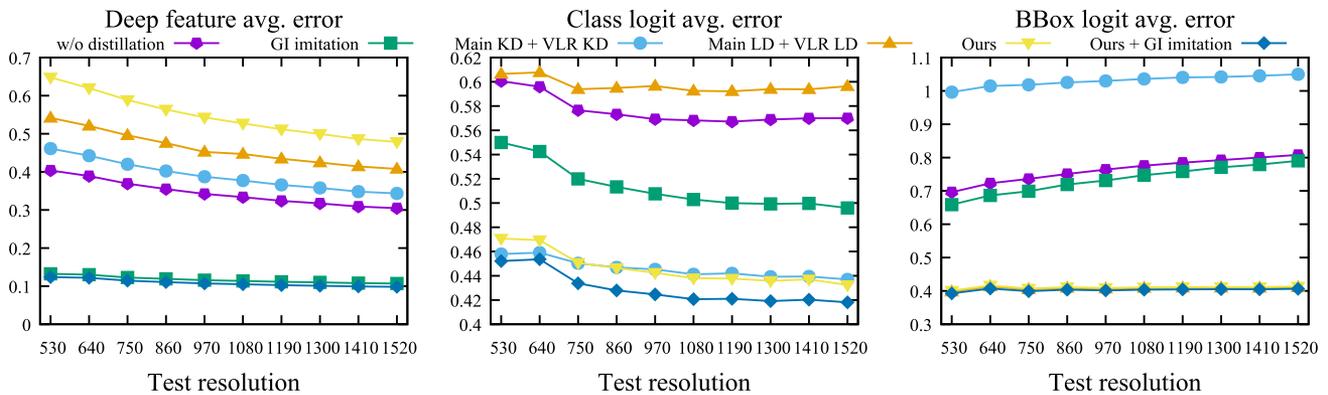


Fig. 6. Average teacher-student error on (left) deep feature representation, (middle) class logits, and (right) bbox logits. “Ours” denotes “Main LD + VLR LD + Main KD”. The curves are evaluated on MS COCO val2017.

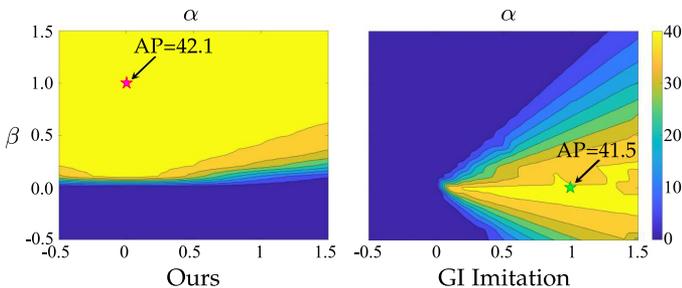


Fig. 7. The 2D contour plots of AP landscapes in feature subspace. The AP landscapes are evaluated on MS COCO val2017.

training convergence. Logit mimicking with the selective region distillation can significantly increase the consistency of the logits between the teacher-student pair, keep the learning liberty of features, and thereby speed up training process and benefit the KD performance more. This indicates that the consistency of feature representations between the teacher-student pair is not the crucial factor of improving the KD performance.

- Feature imitation can increase the consistency of the feature representations between the teacher-student pair, but come some drawbacks such as less feature robustness and slow

#### D. Comparison With the State-of-The-Arts

We compare our LD with the state-of-the-art dense object detectors by using our LD to further boost GFocalV2 [37]. For COCO val2017, since most previous works use ResNet-50-FPN backbone with the single-scale  $1\times$  training schedule (12 epochs) for validation, we also report the results under this setting for a fair comparison. For COCO test-dev 2019, following a previous

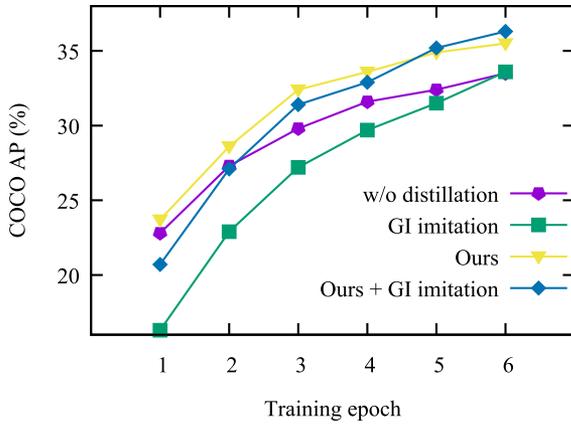


Fig. 8. The average precision (AP) during the early training stage. The feature imitation significantly slows down the convergence and gets a sub-optimal generalization. Logit mimicking (Ours) can reduce the training difficulty in the early training stage.

TABLE VIII  
COMPARISON WITH STATE-OF-THE-ART METHODS ON COCO VAL2017 AND TEST-DEV2019. TS: TRAINING SCHEDULE. '1×': SINGLE-SCALE TRAINING 12 EPOCHS. '2×': MULTI-SCALE TRAINING 24 EPOCHS

Method	TS	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
<b>ResNet-50 backbone on val2017</b>							
RetinaNet [40]	1×	36.9	54.3	39.8	21.2	40.8	48.4
FCOS [62]	1×	38.6	57.2	41.5	22.4	42.2	49.8
SAPD [93]	1×	38.8	58.7	41.3	22.5	42.6	50.8
ATSS [84]	1×	39.2	57.3	42.4	22.7	43.1	51.5
BorderDet [52]	1×	41.4	59.4	44.5	23.6	45.1	54.6
AutoAssign [92]	1×	40.5	59.8	43.9	23.1	44.7	52.9
PAA [32]	1×	40.4	58.4	43.9	22.9	44.3	54.0
OTA [17]	1×	40.7	58.4	44.3	23.2	45.0	53.6
GFocal [38]	1×	40.1	58.2	43.1	23.3	44.4	52.5
GFocalV2 [37]	1×	41.1	58.8	44.9	23.5	44.9	53.3
LD (ours)	1×	<b>42.7</b>	<b>60.2</b>	<b>46.7</b>	<b>25.0</b>	<b>46.4</b>	<b>55.1</b>
<b>ResNet-101 backbone on test-dev 2019</b>							
RetinaNet [40]	2×	39.1	59.1	42.3	21.8	42.7	50.2
FCOS [62]	2×	41.5	60.7	45.0	24.4	44.8	51.6
SAPD [93]	2×	43.5	63.6	46.5	24.9	46.8	54.6
ATSS [84]	2×	43.6	62.1	47.4	26.1	47.0	53.6
BorderDet [52]	2×	45.4	64.1	48.8	26.7	48.3	56.5
AutoAssign [92]	2×	44.5	64.3	48.4	25.9	47.4	55.0
PAA [32]	2×	44.8	63.3	48.7	26.5	48.8	56.3
OTA [17]	2×	45.3	63.5	49.3	26.9	48.8	56.1
GFocal [38]	2×	45.0	63.7	48.9	27.2	48.8	54.5
GFocalV2 [37]	2×	46.0	64.1	50.2	27.6	49.6	56.5
LD (ours)	2×	<b>47.1</b>	<b>65.0</b>	<b>51.4</b>	<b>28.3</b>	<b>50.9</b>	<b>58.5</b>
<b>ResNeXt-101-32x4d-DCN backbone on test-dev 2019</b>							
SAPD [93]	2×	46.6	66.6	50.0	27.3	49.7	60.7
GFocal [38]	2×	48.2	67.4	52.6	29.2	51.7	60.2
GFocalV2 [37]	2×	49.0	67.6	53.4	29.8	52.3	61.8
LD (ours)	2×	<b>50.5</b>	<b>69.0</b>	<b>55.3</b>	<b>30.9</b>	<b>54.4</b>	<b>63.4</b>

work [37], the LD models with the  $1333 \times [480 : 960]$  multi-scale  $2 \times$  training schedule (24 epochs) are included. The training is carried on a machine node with 8 GPUs with a batch size of 2 per GPU and initial learning rate 0.01 for a fair comparison. During inference, single-scale testing ( $[1333 \times 800]$  resolution)

is adopted. For different students ResNet-50, ResNet-101 and ResNeXt-101-32x4d-DCN [70], [94], we also choose different networks ResNet-101, ResNet-101-DCN and Res2Net-101-DCN [16] as their teachers, respectively.

Table VIII reports the quantitative results. It can be seen that our LD improves the AP score of the SOTA GFocalV2 by +1.6 and the AP<sub>75</sub> score by +1.8 when using the ResNet-50-FPN backbone. When using the ResNet-101-FPN and ResNeXt-101-32x4d-DCN with multi-scale  $2 \times$  training, we achieve the highest AP scores, 47.1 and 50.5, which outperform all existing dense object detectors under the same backbone, neck and test settings. More importantly, our LD does not introduce any additional network parameters or computational overhead and hence can guarantee exactly the same inference speed as GFocalV2.

## V. CONCLUSION

In this article, we propose a flexible localization distillation for dense object detection and a selective region distillation based on a new valuable localization region. We show that 1) logit mimicking can be better than feature imitation; and 2) the selective region distillation for transferring the classification and localization knowledge is important when distilling object detectors. We hope our method could provide new research intuitions for the object detection community to develop better distillation strategies. In the future, the applications of LD to sparse object detectors (DETR [5] series), the heterogeneous detector pairs, and other relevant fields, e.g., instance segmentation, object tracking and 3D object detection, warrant further research. Besides, since our LD shares the equivalent optimization effect to classification KD, some improved KD methods may also bring gain to LD, e.g., Relational KD [48], Self-KD [15], [82], Teacher Assistant KD [46], and Decoupled KD [84], etc. Cross architecture distillation using recent state-of-the-art classification models [11], [20], [25], [43], [67] as teachers is also an interesting direction to explore.

## REFERENCES

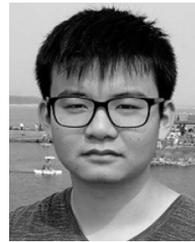
- [1] M. Abadi et al., "TensorFlow: A system for large-scale machine learning," in *Proc. 12th USENIX Symp. Operating Syst. Des. Implementation*, 2016, pp. 265–283.
- [2] J.-H. Bae, D. Yeo, J. Yim, N.-S. Kim, C.-S. Pyo, and J. Kim, "Densely distilled flow-based knowledge transfer in teacher-student framework for image classification," *IEEE Trans. Image Process.*, vol. 29, pp. 5698–5710, 2020.
- [3] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," 2020, *arXiv:2004.10934*.
- [4] Z. Cai and N. Vasconcelos, "Cascade R-CNN: Delving into high quality object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 6154–6162.
- [5] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 213–229.
- [6] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, "Learning efficient object detection models with knowledge distillation," *Adv. Neural Inform. Process. Syst.*, vol. 30, 2017.
- [7] K. Chen et al., "MMDetection: Open MMLab detection toolbox and benchmark," 2019, *arXiv:1906.07155*.
- [8] Z. Chen et al., "PIoU loss: Towards accurate oriented object detection in complex environments," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 195–211.

- [9] J. Choi, D. Chun, H. Kim, and H.-J. Lee, "Gaussian YOLOv3: An accurate and fast object detector using localization uncertainty for autonomous driving," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 502–511.
- [10] X. Dai et al., "General instance distillation for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 7842–7851.
- [11] Z. Dai, H. Liu, Q. Le, and M. Tan, "Coatnet: Marrying convolution and attention for all data sizes," *Adv. Neural Inform. Process. Syst.*, vol. 34, pp. 3965–3977, 2021.
- [12] 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, 2010.
- [13] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1627–1645, Sep. 2010.
- [14] C.-Y. Fu, W. Liu, A. Ranga, A. Tyagi, and A. C. Berg, "DSSD: Deconvolutional single shot detector," 2017, *arXiv:1701.06659*.
- [15] T. Furlanello, Z. Lipton, M. Tschannen, L. Itti, and A. Anandkumar, "Born again neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1607–1616.
- [16] S.-H. Gao et al., "Res2net: A new multi-scale backbone architecture," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 2, pp. 652–662, Feb. 2021.
- [17] Z. Ge, S. Liu, Z. Li, O. Yoshie, and J. Sun, "OTA: Optimal transport assignment for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 303–312.
- [18] S. Gidaris and N. Komodakis, "Locnet: Improving localization accuracy for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 789–798.
- [19] J. Guo et al., "Distilling object detectors via decoupled features," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 2154–2164.
- [20] M.-H. Guo et al., "PCT: Point cloud transformer," *Comput. Vis. Media*, vol. 7, no. 2, pp. 187–199, 2021.
- [21] L. Han, P. Tao, and R. R. Martin, "Livestock detection in aerial images using a fully convolutional network," *Comput. Vis. Media*, vol. 5, no. 2, pp. 221–228, 2019.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [23] Y. He, C. Zhu, J. Wang, M. Savvides, and X. Zhang, "Bounding box regression with uncertainty for accurate object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 2888–2897.
- [24] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015, *arXiv:1503.02531*.
- [25] Q. Hou, C.-Z. Lu, M.-M. Cheng, and J. Feng, "Conv2former: A simple transformer-style convnet for visual recognition," 2022, *arXiv:2211.11943*.
- [26] S.-M. Hu, D. Liang, G.-Y. Yang, G.-W. Yang, and W.-Y. Zhou, "Jittor: A novel deep learning framework with meta-operators and unified graph execution," *Sci. China Inf. Sci.*, vol. 63, no. 222103, pp. 1–21, 2020.
- [27] Z. Huang, L. Huang, Y. Gong, C. Huang, and X. Wang, "Mask scoring R-CNN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 6409–6418.
- [28] B. Jiang, R. Luo, J. Mao, T. Xiao, and Y. Jiang, "Acquisition of localization confidence for accurate object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 784–799.
- [29] X. Jin et al., "Knowledge distillation via route constrained optimization," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 1345–1354.
- [30] Z. Kang, P. Zhang, X. Zhang, J. Sun, and N. Zheng, "Instance-conditional knowledge distillation for object detection," *Adv. Neural Inform. Process. Syst.*, vol. 34, pp. 16468–16480, 2021.
- [31] J. Kim, S. U. Park, and N. Kwak, "Paraphrasing complex network: Network compression via factor transfer," *Adv. Neural Inform. Process. Syst.*, vol. 31, pp. 2765–2774, 2018.
- [32] K. Kim and H. S. Lee, "Probabilistic anchor assignment with IoU prediction for object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 355–371.
- [33] T. Kong, F. Sun, C. Tan, H. Liu, and W. Huang, "Deep feature pyramid reconfiguration for object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 169–185.
- [34] G. Li, X. Li, Y. Wang, S. Zhang, Y. Wu, and D. Liang, "Knowledge distillation for object detection via rank mimicking and prediction-guided feature imitation," in *Proc. AAAI Conf. Artif. Intell.*, 2022, vol. 36, pp. 1306–1313.
- [35] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the loss landscape of neural nets," *Adv. Neural Inform. Process. Syst.*, vol. 31, 2018.
- [36] Q. Li, S. Jin, and J. Yan, "Mimicking very efficient network for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6356–6364.
- [37] X. Li, W. Wang, X. Hu, J. Li, J. Tang, and J. Yang, "Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 11632–11641.
- [38] X. Liet et al., "Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection," *Adv. Neural Inform. Process. Syst.*, vol. 33, 2020, pp. 21002–21012.
- [39] 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.*, 2017, pp. 2117–2125.
- [40] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 2980–2988.
- [41] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [42] W. Luet et al., "SSD: Single shot multibox detector," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 21–37.
- [43] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, and S. Xie, "A convnet for the 2020s," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 11976–11986.
- [44] X. Lu, B. Li, Y. Yue, Q. Li, and J. Yan, "Grid R-CNN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 7363–7372.
- [45] J. Ma et al., "Arbitrary-oriented scene text detection via rotation proposals," *IEEE Trans. Multimedia*, vol. 20, no. 11, pp. 3111–3122, Nov. 2018.
- [46] S. I. Mirzadeh, M. Farajtabar, A. Li, N. Levine, A. Matsukawa, and H. Ghasemzadeh, "Improved knowledge distillation via teacher assistant," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, pp. 5191–5198.
- [47] 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 Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 821–830.
- [48] W. Park, D. Kim, Y. Lu, and M. Cho, "Relational knowledge distillation," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 3967–3976.
- [49] A. Paszke et al., "Automatic differentiation in pytorch," *Adv. Neural Inform. Process. Syst.*, 2017.
- [50] W. Qian, X. Yang, S. Peng, Y. Guo, and J. Yan, "Learning modulated loss for rotated object detection," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, pp. 2458–2466.
- [51] H. Qiu, H. Li, Q. Wu, and H. Shi, "Offset bin classification network for accurate object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 13188–13197.
- [52] H. Qiu, Y. Ma, Z. Li, S. Liu, and J. Sun, "Borderdet: Border feature for dense object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 549–564.
- [53] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 779–788.
- [54] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 7263–7271.
- [55] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," 2018, *arXiv:1804.02767*.
- [56] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *Adv. Neural Inform. Process. Syst.*, vol. 28, 2015.
- [57] H. Rezaeifighi, N. Tsoi, J. Y. Gwak, A. Sadeghian, I. Reid, and S. Savarese, "Generalized intersection over union: A metric and a loss for bounding box regression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 658–666.
- [58] A. Romero et al., "Fitnets: Hints for thin deep net," in *Proc. Int. Conf. Learn. Represent.*, 2015.
- [59] W. Son, J. Na, J. Choi, and W. Hwang, "Densely guided knowledge distillation using multiple teacher assistants," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 9395–9404.
- [60] R. Sun, F. Tang, X. Zhang, H. Xiong, and Q. Tian, "Distilling object detectors with task adaptive regularization," 2020, *arXiv:2006.13108*.
- [61] J. Tanget et al., "Understanding and improving knowledge distillation," 2020, *arXiv:2002.03532*.
- [62] Z. Tian, C. Shen, H. Chen, and T. He, "FCOS: Fully convolutional one-stage object detection," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 9627–9636.
- [63] G.-H. Wang, Y. Ge, and J. Wu, "Distilling knowledge by mimicking features," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 11, pp. 8183–8195, Nov. 2021.

- [64] J. Wang, K. Chen, S. Yang, C. C. Loy, and D. Lin, "Region proposal by guided anchoring," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 2965–2974.
- [65] J. Wanget al., "Side-aware boundary localization for more precise object detection," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 403–419.
- [66] T. Wang, L. Yuan, X. Zhang, and J. Feng, "Distilling object detectors with fine-grained feature imitation," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4933–4942.
- [67] Y.-H. Wu, Y. Liu, X. Zhan, and M.-M. Cheng, "P2T: Pyramid pooling transformer for scene understanding," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published, doi: [10.1109/TPAMI.2022.3202765](https://doi.org/10.1109/TPAMI.2022.3202765).
- [68] G.-S. Xiaet al., "Dota: A large-scale dataset for object detection in aerial images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 3974–3983.
- [69] E. Xieet al., "Polarmask: Single shot instance segmentation with polar representation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 12193–12202.
- [70] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1492–1500.
- [71] X. Yang, Q. Liu, J. Yan, A. Li, Z. Zhang, and G. Yu, "R3det: Refined single-stage detector with feature refinement for rotating object," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, pp. 3163–3171.
- [72] X. Yang and J. Yan, "Arbitrary-oriented object detection with circular smooth label," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 677–694.
- [73] X. Yang, J. Yan, Q. Ming, W. Wang, X. Zhang, and Q. Tian, "Rethinking rotated object detection with Gaussian Wasserstein distance loss," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 11830–11841.
- [74] X. Yanget al., "Scrdet: Towards more robust detection for small, cluttered and rotated objects," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 8232–8241.
- [75] X. Yanget al., "Learning high-precision bounding box for rotated object detection via Kullback-Leibler divergence," *Adv. Neural Inform. Process. Syst.*, vol. 34, pp. 18381–18394, 2021.
- [76] X. Yanget al., "The KFIoU loss for rotated object detection," in *Proc. Int. Conf. Learn. Represent.*, 2023.
- [77] J. Yu, Y. Jiang, Z. Wang, Z. Cao, and T. Huang, "UnitBox: An advanced object detection network," in *Proc. ACM Int. Conf. Multimedia*, 2016, pp. 516–520.
- [78] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [79] H. Zhang, H. Chang, B. Ma, N. Wang, and X. Chen, "Dynamic R-CNN: Towards high quality object detection via dynamic training," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 260–275.
- [80] H. Zhang, Y. Wang, F. Dayoub, and N. Sünderhauf, "Varifocalnet: An IoU-aware dense object detector," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8514–8523.
- [81] L. Zhang and K. Ma, "Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors," in *Proc. Int. Conf. Learn. Represent.*, 2020.
- [82] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, and K. Ma, "Be your own teacher: Improve the performance of convolutional neural networks via self distillation," in *Proc. Int. Conf. Comput. Vis.*, 2019, pp. 3713–3722.
- [83] S. Zhang, C. Chi, Y. Yao, Z. Lei, and Stan Z. Li, "Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 9759–9768.
- [84] B. Zhao, Q. Cui, R. Song, Y. Qiu, and J. Liang, "Decoupled knowledge distillation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 11953–11962.
- [85] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, and D. Ren, "Distance-IoU loss: Faster and better learning for bounding box regression," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, pp. 12993–13000.
- [86] Z. Zhenget al., "Enhancing geometric factors in model learning and inference for object detection and instance segmentation," *IEEE Trans. Cybern.*, vol. 52, no. 8, pp. 8574–8586, Aug. 2022.
- [87] Z. Zhenget al., "Localization distillation for dense object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 9407–9416.
- [88] D. Zhixing et al., "Distilling object detectors with feature richness," *Adv. Neural Inform. Process. Syst.*, vol. 34, pp. 5213–5224, 2021.
- [89] P. Zhou, B. Ni, C. Geng, J. Hu, and Y. Xu, "Scale-transferrable object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 528–537.
- [90] Y. Zhouet al., "Mmrotate: A rotated object detection benchmark using pytorch," in *Proc. ACM Int. Conf. Multimedia*, 2022, pp. 7331–7334.
- [91] B. Zhu et al., "Autoassign: Differentiable label assignment for dense object detection," 2020, *arXiv:2007.03496*.
- [92] C. Zhu, F. Chen, Z. Shen, and M. Savvides, "Soft anchor-point object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 91–107.
- [93] C. Zhu, Y. He, and M. Savvides, "Feature selective anchor-free module for single-shot object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 840–849.
- [94] X. Zhu, H. Hu, S. Lin, and J. Dai, "Deformable convnets v2: More deformable, better results," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 9308–9316.



**Zhaohui Zheng** received the MS degree in computational mathematics from Tianjin University, Tianjin, China, in 2021. He is currently working toward the Ph.D. degree with the School of Computer Science, Nankai University, Tianjin. His research interests include object detection, instance segmentation, and knowledge distillation.



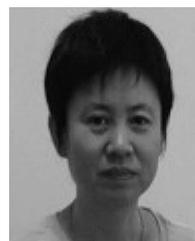
**Rongguang Ye** received the BS and MS degrees from the School of Mathematics, Tianjin University, Tianjin, China, in 2019 and 2022, respectively. He is currently with Intel Asia-Pacific Research and Development Ltd., as an AI framework engineer. His research interests include object detection and computer vision.



**Qibin Hou** (Member, IEEE) received the PhD degree from the School of Computer Science, Nankai University, Tianjin, China. Then, he was with the National University of Singapore, Singapore, as a research fellow. He is currently an associate professor with the School of Computer Science, Nankai University. He has authored or coauthored more than 30 papers on top conferences/journals, including T-PAMI, CVPR, ICCV, and NeurIPS. His research interests include deep learning and computer vision.



**Dongwei Ren** (Member, IEEE) received the two Ph.D. degrees in computer application technology from the Harbin Institute of Technology, Harbin, China, and The Hong Kong Polytechnic University, Hong Kong, in 2017 and 2018, respectively. From 2018 to 2021, he was an assistant professor with the College of Intelligence and Computing, Tianjin University. He is currently an associate professor with the School of Computer Science and Technology, Harbin Institute of Technology. His research interests include computer vision and deep learning.



**Ping Wang** received the B.S., M.S., and Ph.D. degrees in computer science from Tianjin University, Tianjin, China, in 1988, 1991, and 1998, respectively. She is currently a professor with the School of Mathematics, Tianjin University. Her research interests include image processing and machine learning.



**Wangmeng Zuo** (Senior Member, IEEE) received the Ph.D. degree from the Harbin Institute of Technology, Harbin, China, in 2007. He is currently a professor with the School of Computer Science and Technology, Harbin Institute of Technology. He has authored or coauthored more than 100 papers in top tier journals and conferences. His publications have been cited more than 30,000 times in literature. His research interests include image enhancement and restoration, image and face editing, object detection, visual tracking, and image classification. He is on the editorial boards of *IEEE Transactions on Pattern Analysis and Machine Intelligence* and *IEEE Transactions on Image Processing*.



**Ming-Ming Cheng** (Senior Member, IEEE) received the PhD degree from Tsinghua University in 2012. Then, he did two years research fellow, with Prof. Philip Torr in Oxford. He is currently a professor with Nankai University, leading the Media Computing Lab. His research interests include computer graphics, computer vision, and image processing. He was the recipient of research awards including National Science Fund for Distinguished Young Scholars and ACM China Rising Star Award. He is on the editorial boards of *IEEE Transactions on Pattern Analysis and Machine Intelligence* and *IEEE Transactions on Image Processing*.