This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Localization Distillation for Dense Object Detection

Zhaohui Zheng^{1*}, Rongguang Ye^{2*}, Ping Wang², Dongwei Ren³, Wangmeng Zuo³, Qibin Hou^{1†}, Ming-Ming Cheng¹ ¹TMCC, CS, Nankai University ²School of Mathematics, Tianjin University ³School of Computer Science and Technology, Harbin Institute of Technology

Abstract

Knowledge distillation (KD) has witnessed its powerful capability in learning compact models in object detection. Previous KD methods for object detection mostly focus on imitating deep features within the imitation regions instead of mimicking classification logit due to its inefficiency in distilling localization information and trivial improvement. In this paper, by reformulating the knowledge distillation process on localization, we present a novel localization distillation (LD) method which can efficiently transfer the localization knowledge from the teacher to the student. Moreover, we also heuristically introduce the concept of valuable localization region that can aid to selectively distill the semantic 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 localization knowledge distillation is more important and efficient than semantic knowledge for distilling object detectors. Our distillation scheme is simple as well as effective and can be easily applied to different dense object detectors. Experiments show that our LD can boost the AP score of GFocal-ResNet-50 with a single-scale $1 \times$ training schedule from 40.1 to 42.1 on the COCO benchmark without any sacrifice on the inference speed. Our source code and pretrained models are publicly available at https://github.com/HikariTJU/LD.

1. Introduction

Localization is a fundamental issue in object detection [15, 24, 33, 49, 50, 55-57, 61, 68]. Bounding box regression is the most popular manner so far for localization in object detection [10, 32, 39, 42], where Dirac delta distribution representation is intuitive and popular for years. However, localization ambiguity where objects cannot be confidently located by their edges is still a common issue. For exam-



Figure 1. Bottom edge for "*elephant*" and right edge for "*surf-board*" are ambiguous.

ple, as shown in Fig. 1, the bottom edge for "elephant" and the right edge for "surfboard" are ambiguous to locate. This issue is even worse for lightweight detectors. One way to alleviate this problem is the knowledge distillation (KD), which, as a model compression technology, has been widely validated to be useful for boosting the performance of the small-sized student network by transferring the generalized knowledge captured by the large-sized teacher network.

Speaking of KD in object detection, previous works [22, 52, 62] have pointed out that the original logit mimicking technique [19] for classification is inefficient as it only transfers the semantic knowledge (*i.e.*, classification), while neglects the importance of localization knowledge distillation. Therefore, existent KD methods for object detection mostly focus on enforcing the consistency of the deep features between the teacher-student pair, and exploit various imitation regions for distillation [5, 8, 16, 25, 52]. Fig. 2 exhibits three popular KD pipelines for object detection. However, as the semantic knowledge and the localization one are mixed on the feature maps, it is hard to tell whether it is beneficial to the performance to transfer the hybrid knowledge for each location and which regions are conducive to the transfer of a certain type of knowledge.

Motivated by the aforementioned questions, in this paper, instead of simply distilling the hybrid knowledge on the feature maps, we propose a novel divide-and-conquer distillation strategy that transfers the semantic and localization knowledge separately. For semantic knowledge, we use the

^{*}Equal contribution.

[†]Corresponding author.



Figure 2. Existing KD pipelines for object detection. ① Logit Mimicking: classification KD in [19]. 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.

original classification KD [19]. For localization knowledge, we reformulate the knowledge transfer process on localization and present a simple yet effective localization distillation (LD) method by switching the bounding box to probability distribution [28, 37]. This is quite different from previous works [5, 47] that treat the teacher's outputs as additional regression targets (i.e., the Pseudo BBox Regression in Fig. 2). Benefiting from the probability distribution representation, our LD can efficiently transfer rich localization knowledge learnt by the teacher to the student. In addition, based on the proposed divide-and-conquer distillation strategy, we further introduce valuable localization region (VLR) to help efficiently judge which regions are conducive to classification or localization learning. Through a series of experiments, we, for the first time, show that the original logit mimicking can be better than feature imitation and localization knowledge distillation is more important and more efficient than semantic knowledge. We believe that separately distilling the semantic and localization knowledge based on their respective favorable regions could be a promising way to train better object detectors.

Our method is simple and can be easily equipped with in any dense 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 [28] with ResNet-50-FPN backbone from 40.1 to 42.1, and AP₇₅ from 43.1 to 45.6. Our best model using ResNeXt-101-32x4d-DCN backbone can achieve a singlescale test of 50.5 AP, which surpasses all existing detectors under the same backbone, neck, and test settings.

2. Related Work

In this section, we give a brief review on the related works, including bounding box regression, localization quality estimation, and knowledge distillation.

2.1. Bounding Box Regression

Bounding box regression is the most popular method for localization in object detection. R-CNN series [3,35,42,60] adopt multiple regression stages to refine the detection results, while [2, 32, 39–41, 48] adopt one-stage regression. In [43,58,64,65], IoU-based loss functions are proposed to improve the localization quality of bounding box. Recently, bounding box representation has evolved from Dirac delta distribution [32, 39,42] to Gaussian distribution [7, 18], and further to probability distribution [28, 37]. The probability distribution of bounding box is more comprehensive for describing the uncertainty of bounding box, and is validated to be the most advanced bounding box representation so far.

2.2. Localization Quality Estimation

As the name suggests, Localization Quality Estimation (LQE) predicts a score that measures the localization quality of the bounding box predicted by the detector. LQE is usually used to cooperate with classification task during training [27], *i.e.*, enhancing the consistency between classification and localization. It can also be applied in joint decision-making during post-processing [21, 39, 48], *i.e.*, considering both classification score and LQE when performing NMS. Early research can be dated to YOLOv1 [39], where the predicted object confidence is used to penalize the classification score. Then, box/mask IoU [20,21] and box/polar center-ness [48, 53] are proposed to model the uncertainty of detections for object detection and instance segmentation, respectively. From the perspective of bounding box representation, Softer-NMS [18] and Gaussian YOLOv3 [7] predict variance for each edge of bounding box. LQE is a preliminary approach to model localization ambiguity.

2.3. Knowledge Distillation

Knowledge distillation [1, 19, 34, 36, 45, 59] aims to learn compact and efficient student models guided by excellent teacher networks. FitNets [44] proposes to mimic the intermediate-level hints from the hidden layers of the teacher model. Knowledge distillation was first applied to object detection in [5], where the hint learning and KD are both used for multi-class object detection. Then, Li et al. [25] proposed to mimic the feature within the region proposal for Faster R-CNN. Wang et al. [52] mimicked fine-grained features on close anchor box locations.



Figure 3. Illustration of localization distillation (LD) for an edge $e \in \mathcal{B} = \{t, b, l, r\}$. Only the localization branch is visualized here. $\mathcal{S}(\cdot, \tau)$ is the generalized SoftMax function with temperature τ . 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.

Recently, Dai et al. [8] introduced the General Instance Selection Module to mimic deep features within the discriminative patches between teacher-student pairs. DeFeat [16] leverages different loss weights when conducting feature imitation on the object regions and the background region. Different from the aforementioned method based on feature imitation, our work introduces localization distillation and proposes to separately transfer the classification and localization knowledge based on the valuable localization region to make the distillation more efficient.

3. Proposed Method

In this section, we introduce the proposed distillation method. Instead of distilling the hybrid knowledge on feature maps, we present a novel divide-and-conquer distillation strategy that separately distills the semantic and localization knowledge based on their respective preferred regions. To transfer the semantic knowledge, we simply adopt the classification KD [19] on the classification head while for localization knowledge, we propose a simple yet effective localization distillation (LD). Both techniques operate on the logits of individual heads rather than deep features. Then, to further improve the distillation efficiency, we introduce valuable localization region (VLR) that can help judge which type of knowledge is conducive to transfer for different regions. In what follows, we first briefly revisit the probability distribution representation of bounding box and then transit to the proposed method.

3.1. Preliminaries

For a given bounding box \mathcal{B} , conventional representations have two forms, *i.e.*, $\{x, y, w, h\}$ (central point coordinates, width and height) [32, 39, 42] and $\{t, b, l, r\}$ (distance from the sampling point to the top, bottom, left and right edges) [48]. 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. 1. This is also clearly demonstrated in previous works [7, 18, 28, 37].

In our method, we use the recent probability distribution representation of bounding box [28, 37] which is more comprehensive for describing the localization uncertainty of bounding box. Let $e \in \mathcal{B}$ be an edge of a bounding box. Its value can be generally represented as

$$\hat{e} = \int_{e_{\min}}^{e_{\max}} x \Pr(x) dx, \quad e \in \mathcal{B},$$
(1)

where x is the regression coordinate ranged in $[e_{\min}, e_{\max}]$, and $\Pr(x)$ is the corresponding probability. The conventional Dirac delta representation is a special case of Eqn. (1), where $\Pr(x) = 1$ when $x = e^{gt}$, otherwise $\Pr(x) = 0$. By quantizing the continuous regression range $[e_{\min}, e_{\max}]$ into the uniform discretized variable $e = [e_1, e_2, \cdots, e_n]^T \in \mathbb{R}^n$ with n subintervals, where $e_1 = e_{\min}$ and $e_n = e_{\max}$, each edge of the given bounding box can be represented as the probability distribution by using the SoftMax function.

3.2. 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 [28] of bounding box which is originally designed for the generic object detection and carries abundant localization information. The ambiguous and clear edges in Fig. 1 will be respectively reflected by the flatness and sharpness of distribution.

The working principle of our LD can be seen in Fig. 3. Given an arbitrary dense object detector, following [28], we first switch the bounding box representation from a quaternary representation to a probability distribution. We choose

 $\mathcal{B} = \{t, b, l, r\}$ as the basic form for bounding box. Unlike the $\{x, y, w, h\}$ form, the physical meaning of each variable in the $\{t, b, l, r\}$ form is consistent, which is convenient for us to restrict the probability distribution of each edge to the same interval range. According to [63], there is no performance difference between the two forms. Thus, when the $\{x, y, w, h\}$ form is given, we will first switch it to the $\{t, b, l, r\}$ form.

Let z be the n logits predicted by the localization head for all possible positions of edge e, denoted by z_T and z_S for the teacher and the student, respectively. Different from [28, 37], we transform z_T and z_S into probability distributions p_T and p_S using the generalized SoftMax function $S(\cdot, \tau) = \text{SoftMax}(\cdot/\tau)$. Note that when $\tau = 1$, it is equivalent to the original SoftMax function. When $\tau \to 0$, it tends to be a Dirac delta distribution. When $\tau \to \infty$, it will degrade to be a uniform distribution. Empirically, $\tau > 1$ is set to soften the distribution, making the probability distribution carry more information.

The localization distillation for measuring the similarity between the two probabilities $p_T, p_S \in \mathbb{R}^n$ is attained by:

$$\mathcal{L}_{\rm LD}^e = \mathcal{L}_{\rm KL}(\boldsymbol{p}_S^{\tau}, \boldsymbol{p}_T^{\tau}) \tag{2}$$

$$= \mathcal{L}_{\mathrm{KL}} \left(\mathcal{S}(\boldsymbol{z}_S, \tau), \mathcal{S}(\boldsymbol{z}_T, \tau) \right), \qquad (3)$$

where \mathcal{L}_{KL} represents the KL-Divergence loss. Then, LD for all the four edges of bounding box \mathcal{B} can be formulated as:

$$\mathcal{L}_{\mathrm{LD}}(\mathcal{B}_S, \mathcal{B}_T) = \sum_{e \in \mathcal{B}} \mathcal{L}^e_{\mathrm{LD}}.$$
 (4)

Discussion. Our LD is the first attempt to adopt logit mimicking to distill localization knowledge for object detection. Though the probability distribution representation for boxes has been proven useful in the generic object detection task [28], no one has explored its performance in localization knowledge distillation. We combine the probability distribution representation for boxes and the KL-Divergence loss and demonstrate that such a simple logit mimicking technique performs well in improving the distillation efficiency of object detectors. This also makes our LD quite different from previous relevant works that, on the contrary, emphasize the importance of feature imitation. In our experiment section, we will show more numerical analysis on the advantages of the proposed LD.

3.3. 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 should be: Should we use the whole imitation regions without discrimination to distill the hybrid knowledge? According to our observation, the answer is no. Previous works [11, 13, 26, 46, 51] Algorithm 1 Valuable Localization Region

- **Require:** A set of anchor boxes $B_l^a = \{\mathcal{B}_{i_l}^a\}$ and a set of ground truth boxes $B^{gt} = \{\mathcal{B}_j^a\}, 1 \leq i_l \leq I_l, 1 \leq j \leq J, I_l = W_l \times H_l$. Positive threshold α_{pos} of label assignment. W_l and H_l are the sizes of *l*-th FPN level.
- **Ensure:** $V_l = \{v_{i_l j}\}_{I_l \times J}, v_{i_l j} \in \{0, 1\}$ encodes final location of VLR, where 1 denotes VLR and 0 indicates ignore.
- 1: Compute DIoU matrix $X_l = \{x_{i_l j}\}_{I_l \times J}$ with $x_{i_l j} = DIoU(\mathcal{B}_{i_l}^a, \mathcal{B}_j^{gt})$.
- 2: $\alpha_{vl} = \gamma \alpha_{pos}$.
- 3: Select locations with $V_l = \{ \alpha_{vl} \leq X_l \leq \alpha_{pos} \}.$
- 4: return V_l

have pointed out that the knowledge distribution patterns are different for classification and localization. Therefore, we, in this subsection, 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, for the *l*-th FPN level, we calculate the DIoU [64] matrix X_l between all the anchor boxes B_1^a and the ground-truth boxes B^{gt} . Then, we set the lower bound of DIoU to be $\alpha_{vl} = \gamma \alpha_{pos}$, where α_{pos} is the positive IoU threshold of label assignment. The VLR can be defined as $V_l = \{\alpha_{vl} \leq X_l \leq \alpha_{pos}\}$. Our method has only one hyperparameter γ , 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 $0 \leq x_{i_1 i_j} \leq \alpha_{pos}$ will be determined as VLRs. When $\gamma \rightarrow 1$, the VLR will gradually shrink to empty. Here we use DIoU [64] 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 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 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.

au	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
-	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
5	40.9	58.2	44.3	23.2	45.0	53.2
10	41.1	58.7	44.9	23.8	44.9	53.6
15	40.7	58.5	44.2	23.5	44.3	53.3
20	40.5	58.3	43.7	23.8	44.1	53.5

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

ε	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
-	40.1	58.2	43.1	23.3	44.4	52.5
0.1	40.5	58.3	43.8	23.0	44.2	52.7
0.2	40.2	58.2	43.6	23.1	44.0	53.0
0.3	40.1	58.4	43.1	23.6	43.9	52.5
0.4	40.3	58.4	43.4	22.8	44.0	52.6
LD	41.1	58.7	44.9	23.8	44.9	53.6

(b) **LD** *vs.* **Pseudo BBox Regression [5]**: The localization knowledge can be more efficiently transferred by our LD. The teacher is ResNet-101 and the student is ResNet-50.

γ	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
-	40.1	58.2	43.1	23.3	44.4	52.5
1	41.1	58.7	44.9	23.8	44.9	53.6
0.75	41.2	58.8	44.9	23.6	45.4	53.5
0.5	41.7	59.4	45.3	24.2	45.6	54.2
0.25	41.8	59.5	45.4	24.2	45.8	54.9
0	41.7	59.5	45.4	24.5	45.9	54.0

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

Table 1. Ablations.	We show ablation	experiments for	LD and	VLR on MS	COCO val20)17.
---------------------	------------------	-----------------	--------	-----------	------------	------

3.4. Overall Distillation Process

The total loss for training the student S can be represented as:

$$\mathcal{L} = \lambda_{0} \mathcal{L}_{cls}(\mathcal{C}_{S}, \mathcal{C}^{gt}) + \lambda_{1} \mathcal{L}_{reg}(\mathcal{B}_{S}, \mathcal{B}^{gt}) + \lambda_{2} \mathcal{L}_{DFL}(\mathcal{B}_{S}, \mathcal{B}^{gt}) + \lambda_{3} \mathbb{I}_{Main} \mathcal{L}_{LD}(\mathcal{B}_{S}, \mathcal{B}_{T}) + \lambda_{4} \mathbb{I}_{VL} \mathcal{L}_{LD}(\mathcal{B}_{S}, \mathcal{B}_{T}) + \lambda_{5} \mathbb{I}_{Main} \mathcal{L}_{KD}(\mathcal{C}_{S}, \mathcal{C}_{T}) + \lambda_{6} \mathbb{I}_{VL} \mathcal{L}_{KD}(\mathcal{C}_{S}, \mathcal{C}_{T}),$$
(5)

where the first three terms are exactly same to the classification and bounding box regression branches for any regression-based detector, *i.e.*, \mathcal{L}_{cls} is the classification loss, \mathcal{L}_{reg} is the bounding box regression loss and \mathcal{L}_{DFL} is the distribution focal loss [28]. \mathbb{I}_{Main} and \mathbb{I}_{VL} are the distillation masks for the main distillation region and the valuable localization region respectively, \mathcal{L}_{KD} is KD loss [19], \mathcal{C}_{S} and C_T denote the classification head output logits of the student and the teacher, respectively, C^{gt} is the ground truth class label. All the distillation losses will be weighted by the same weight factors according to their types, e.g., LD loss follows the bbox regression and KD loss follows the classification. Also, it is worth mentioning that 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 a separate distillation region manner.

4. 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 largescale MS COCO [31] benchmark.

4.1. Experiment Setup

The train2017 (118K images) is utilized for training and val2017 (5K images) is used for validation. We also obtain the evaluation results on MS COCO test-dev 2019 (20K images) by submitting to the COCO server. The experiments are conducted under mmDetection [6] framework. Unless otherwise stated, we use ResNet [17] with FPN [29] as

our backbone and neck networks, and the FCOS-style [48] 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 [28] 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. More implementation details and more experimental results on PAS-CAL VOC [9] can be found in the supplementary materials.

4.2. Ablation Studies and Analysis

Temperature τ in LD. Our LD introduces a hyperparameter, *i.e.*, the temperature τ . Table 1a 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 1a, 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 [5] is a preliminary attempt to enhance the student on the localization head, *i.e.*, the pseudo bbox regression in Fig. 2, which is represented as:

$$\mathcal{L}_{\text{TBR}} = \lambda \mathcal{L}_{\text{reg}}(\mathcal{B}^s, \mathcal{B}^{gt}), \text{if } \ell_2(\mathcal{B}^s, \mathcal{B}^{gt}) + \varepsilon > \ell_2(\mathcal{B}^t, \mathcal{B}^{gt}),$$
(6)

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, ε is a predefined margin, and \mathcal{L}_{reg} represents the GIoU loss [43]. Here, only the main distillation region is adopted. From Table 1b, TBR loss does yield performance gains (+0.4 AP and +0.7 AP₇₅) when using proper threshold $\varepsilon = 0.1$ in Eqn. (6). However, it uses the coarse bbox representation, which does not contain any localization uncertainty information of the detector, leading to sub-optimal

Table 2. Evaluation of **separate distillation region manner** for KD and our LD. The teacher is ResNet-101 and the student is ResNet-50. "Main" denotes the main distillation region, *i.e.*, the positive locations of label assignment. "VLR" denotes the valuable localization region. The results are reported on MS COCO val2017.

Main KD	Main LD	VLR KD	VLR LD	AP	AP_{50}	AP ₇₅
				40.1	58.2	43.1
\checkmark				40.2	58.6	43.4
	\checkmark			41.1	58.7	44.9
\checkmark	\checkmark			41.4	59.2	45.0
\checkmark		\checkmark		40.4	58.9	43.4
	\checkmark		\checkmark	41.8	59.5	45.4
\checkmark	\checkmark		\checkmark	42.1	60.3	45.6
√	\checkmark	\checkmark	\checkmark	42.0	60.0	45.4

results. On the contrary, our LD directly produces 41.1 AP and 44.9 AP₇₅, since it utilizes the probability distribution of bbox which contains rich localization knowledge.

Various γ **in VLR.** The newly introduced VLR has the parameter γ which controls the range of VLR. As shown in Table 1c, AP is stable when γ ranges from 0 to 0.5. The variation in AP in this range is around 0.1. As γ increases, the VLR gradually shrinks to empty. The performance also gradually drops to 41.1, *i.e.*, conducting LD on the main distillation region only. The sensitivity analysis experiments on the parameter γ indicate that conducting LD on the VLR has a positive effect on performance. In the rest experiments, we set γ to 0.25 for simplicity.

Separate Distillation Region Manner. 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 2, where "Main" indicates 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. It can be seen that conducting "Main KD", "Main LD", and their combination can all improve the student performance by +0.1, +1.0 and +1.3 AP, respectively. 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 distillation on a larger range, *i.e.*, VLR. We can see that "VLR LD" (the 5-th row of Table 2) can further improve AP by +0.7 based on "Main LD" (the 3-rd row). However, we observe that further involving "VLR KD" yields limited improvement (the 2-nd row and the 5-th row of Table 2) or even no improvement (the last two rows of Table 2). This again shows that localization knowledge distillation is more important and efficient than semantic knowledge distillation and our divide-and-conquer distillation scheme, *i.e.*, "Main KD" + "Main LD" + "VLR LD", is complementary to VLR.

Table 3. **Logit Mimicking vs. Feature Imitation.** "Ours" means we use the separate distillation region manner, *i.e.*, conducing KD and LD on the main distillation region, and conducing LD on the VLR. The teacher is ResNet-101 and the student is ResNet-50 [17]. The results are reported on MS COCO val2017.

Method	AP	AP_{50}	AP ₇₅	AP_S	AP_M	AP_L
Baseline (GFocal [28])	40.1	58.2	43.1	23.3	44.4	52.5
FitNets [44]	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 [52]	41.1	58.8	44.8	23.3	45.4	53.1
DeFeat [16]	40.8	58.6	44.2	24.3	44.6	53.7
GI Imitation [8]	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 + 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
. 0.145		(0.55			
E 0.138		Erro	149			
a)		à				
< 0.131 2						
<u>5</u> 0.124						
Se 0.117	_		0.31			
0.110		e,	0.25			
w/o Distillation Fine-Grained GI Imitation Main LD						

Figure 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.

Main LD + VLR LD Main KD + Main LD + VLR LD

Logit Mimicking vs. Feature Imitation. We compare our proposed LD with several state-of-the-art feature imitation methods. We adopt the separate distillation region manner, *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 [29], following previous works [8, 16, 52], we re-implement their methods and impose all the feature imitations on multi-level FPN for a fair comparison. Here, "FitNets" [44] distills the whole feature maps. "DeFeat" [16] means the loss weights of feature imitation outside the GT boxes are larger than those inside GT boxes. "Fine-Grained" [52] distills the deep features on the close anchor box locations. "GI Imitation" [8] selects the distillation regions according to the discriminative predictions of the student and the teacher. "Inside GT Box" means we use the GT 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 3, we can see that distillation within the whole feature maps attains +0.6 AP gains. By setting a



Figure 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 head logits between the teacher and the student at 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 [8], our method (Main LD + VLR LD) can significantly reduce the errors for almost all the locations. Darker is better. Best viewed in color.

larger loss weight for the locations outside the GT boxes (DeFeat [16]), the performance is slightly better than that using the same loss weight for all locations. Fine-Grained [52] 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 [8] 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 separate distillation region manner, which directly produces 42.1 AP. It is worth noting that our method operates on logits instead of features, indicating that logit mimicking is not inferior to feature imitation as long as adopting a proper distillation strategy, like our LD. Moreover, our method is orthogonal to the aforementioned feature imitation methods. Table 3 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 AP₇₅.

We further conduct an experiment to check the average error of classification score and box probability distribution, as shown in Fig. 4. One can see that the Fine-Grained feature imitation [52] and GI imitation [8] reduce the two errors as expected, since the semantic 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 [52] and GI imitation [8] but lower box probability distribution average errors. This indicates that these two settings with

Table 4. Quantitative results of LD for lightweight detectors. The teacher is ResNet-101. The results are reported on MS COCO val2017.

Student	LD	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
DeeNet 19		35.8	53.1	38.2	18.9	38.9	47.9
Resinet-18	\checkmark	37.5	54.7	40.4	20.2	41.2	49.4
DecNet 24		38.9	56.6	42.2	21.5	42.8	51.4
Keshet-54	\checkmark	41.0	58.6	44.6	23.2	45.0	54.2
DecNet 50		40.1	58.2	43.1	23.3	44.4	52.5
Keshet-30	\checkmark	42.1	60.3	45.6	24.5	46.2	54.8

only LD can significantly reduce the box probability distribution distance between the teacher and the student, while it is reasonable that they cannot reduce this error for classification head. If we impose the classification KD on the main distillation region, obtaining "Main KD + Main LD + VLR LD", both classification score average error and 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. In Fig. 5, comparing to "Without Distillation", we can see that the GI imitation [8] 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 reduce this error more observably, and alleviate the localization ambiguity.

LD for Lightweight Detectors. Next, we validate our LD with the separate distillation region manner, *i.e.*, Main KD + Main LD + VLR LD, for lightweight detectors. We select

Table 5. 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	APro	AP ₇₅	APg	APM	AP
		36.0	54.3	30.8	21.2	40.8	18.4
RetinaNet [30]		20.0	54.5	12.4	21.2	42.2	40.4 51.1
	~	39.0	30.4	42.4	23.1	43.2	51.1
ECOS [49]		38.6	57.2	41.5	22.4	42.2	49.8
1003 [40]	\checkmark	40.6	58.4	44.1	24.3	44.1	52.3
ATSS [62]		39.2	57.3	42.4	22.7	43.1	51.5
AI35 [05]	\checkmark	41.6	59.3	45.3	25.2	45.2	53.3

ResNet-101 with 44.7 AP provided by the mmDetection [6] as our teacher to distill a series of lightweight students. As shown in Table 4, our LD can stably improve the students ResNet-18, ResNet-34, ResNet-50 by +1.7, +2.1, +2.0 in AP, and +2.2, +2.4, +2.4 in AP₇₅, respectively. From these results, we can conclude that our LD can stably improve the localization accuracy for all the students.

Extension to Other Dense Object Detectors. Our LD is flexible to incorporate into other dense object detectors, either anchor-based or anchor-free type. We employ LD with the separate distillation region manner to several recently popular detectors such as RetinaNet [30] (anchor-based), FCOS [48] (anchor-free) and ATSS [63] (anchor-based). According to the results in Table 5, LD can consistently improve \sim 2 AP for these dense detectors.

4.3. 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 [27]. For COCO val2017, since most previous works use ResNet-50-FPN backbone with 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 previous work [27], the LD models with $1333 \times [480:960]$ multi-scale $2 \times$ training schedule (24 epochs) are included. The training is carried on a machine node with 8 GPUs using a batch size of 2 per GPU and initial learning rate 0.01 for a fair comparison. During inference, single-scale testing ([1333×800] resolution) is adopted. For different students ResNet-50, ResNet-101 and ResNeXt-101-32x4d-DCN [54, 69], we also choose different networks ResNet-101, ResNet-101-DCN and Res2Net-101-DCN [12] as their teachers, respectively.

As shown in Table 6, our LD improves the AP score of the SOTA GFocalV2 by +1.6 and the AP₇₅ 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 thus can

Table 6. Compare with state-of-the-art methods on COCO *val2017* and *test-dev2019*. **TS**: Traning Schedule. '1×': single-scale training 12 epochs. '2×': multi-scale training 24 epochs.

Method	TS	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	
	ResN	et-50 b	ackbone	on val2	017			
RetinaNet [30]	$1 \times$	36.9	54.3	39.8	21.2	40.8	48.4	
FCOS [48]	1×	38.6	57.2	41.5	22.4	42.2	49.8	
SAPD [67]	1×	38.8	58.7	41.3	22.5	42.6	50.8	
ATSS [63]	1×	39.2	57.3	42.4	22.7	43.1	51.5	
BorderDet [38]	1×	41.4	59.4	44.5	23.6	45.1	54.6	
AutoAssign [66]	1×	40.5	59.8	43.9	23.1	44.7	52.9	
PAA [23]	1×	40.4	58.4	43.9	22.9	44.3	54.0	
OTA [14]	1×	40.7	58.4	44.3	23.2	45.0	53.6	
GFocal [28]	1×	40.1	58.2	43.1	23.3	44.4	52.5	
GFocalV2 [27]	1×	41.1	58.8	44.9	23.5	44.9	53.3	
LD (ours)	1×	42.7	60.2	46.7	25.0	46.4	55.1	
R	esNet-	101 bac	kbone o	n test-de	v 2019			
RetinaNet [30]	$2\times$	39.1	59.1	42.3	21.8	42.7	50.2	
FCOS [48]	2×	41.5	60.7	45.0	24.4	44.8	51.6	
SAPD [67]	$2\times$	43.5	63.6	46.5	24.9	46.8	54.6	
ATSS [63]	2×	43.6	62.1	47.4	26.1	47.0	53.6	
BorderDet [38]	2×	45.4	64.1	48.8	26.7	48.3	56.5	
AutoAssign [66]	2×	44.5	64.3	48.4	25.9	47.4	55.0	
PAA [23]	2×	44.8	63.3	48.7	26.5	48.8	56.3	
OTA [14]	2×	45.3	63.5	49.3	26.9	48.8	56.1	
GFocal [28]	2×	45.0	63.7	48.9	27.2	48.8	54.5	
GFocalV2 [27]	2×	46.0	64.1	50.2	27.6	49.6	56.5	
LD (ours)	2×	47.1	65.0	51.4	28.3	50.9	58.5	
ResNeXt-101-32x4d-DCN backbone on test-dev 2019								
SAPD [67]	$2\times$	46.6	66.6	50.0	27.3	49.7	60.7	
GFocal [28]	2×	48.2	67.4	52.6	29.2	51.7	60.2	
GFocalV2 [27]	2×	49.0	67.6	53.4	29.8	52.3	61.8	
LD (ours)	2×	50.5	69.0	55.3	30.9	54.4	63.4	

guarantee exactly the same inference speed as GFocalV2.

5. Conclusion

In this paper, we propose a flexible localization distillation for dense object detection and design a valuable localization region to distill the student detector in a separate distillation region manner. We show that 1) logit mimicking can be better than feature imitation; and 2) the separate distillation region manner 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 addition, the applications of LD to sparse object detectors (DETR [4] series) and other relevant fields, *e.g.*, instance segmentation, object tracking and 3D object detection, warrant future research.

Acknowledgment This work is funded by the National Key Research and Development Program of China (NO. 2018AAA0100400) and NSFC (NO. 61922046, NO. 62172127).

References

- [1] Ji-Hoon Bae, Doyeob Yeo, Junho Yim, Nae-Soo Kim, Cheol-Sig Pyo, and Junmo Kim. Densely distilled flowbased knowledge transfer in teacher-student framework for image classification. *IEEE Transactions on Image Processing*, 29:5698–5710, 2020.
- [2] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934, 2020.
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: Delving into high quality object detection. In CVPR, 2018.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-toend object detection with transformers. In ECCV, 2020.
- [5] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *NeurIPS*, 2017.
- [6] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155, 2019.
- [7] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian YOLOv3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *ICCV*, 2019.
- [8] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *CVPR*, 2021.
- [9] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal* of Computer Vision, 88(2):303–338, 2010.
- [10] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2009.
- [11] Chengjian Feng, Yujie Zhong, Yu Gao, Matthew R Scott, and Weilin Huang. Tood: Task-aligned one-stage object detection. In *ICCV*, 2021.
- [12] Shang-Hua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip Torr. Res2net: A new multi-scale backbone architecture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2):652–662, 2021.
- [13] Ziteng Gao, Limin Wang, and Gangshan Wu. Mutual supervision for dense object detection. In *ICCV*, 2021.
- [14] Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. OTA: Optimal transport assignment for object detection. In *CVPR*, 2021.
- [15] Spyros Gidaris and Nikos Komodakis. Locnet: Improving localization accuracy for object detection. In CVPR, 2016.

- [16] Jianyuan Guo, Kai Han, Yunhe Wang, Han Wu, Xinghao Chen, Chunjing Xu, and Chang Xu. Distilling object detectors via decoupled features. In *CVPR*, 2021.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [18] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *CVPR*, 2019.
- [19] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015.
- [20] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring R-CNN. In *CVPR*, 2019.
- [21] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yuning Jiang. Acquisition of localization confidence for accurate object detection. In *ECCV*, 2018.
- [22] Zijian Kang, Peizhen Zhang, Xiangyu Zhang, Jian Sun, and Nanning Zheng. Instance-conditional knowledge distillation for object detection. In *NeurIPS*, 2021.
- [23] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. In ECCV, 2020.
- [24] Tao Kong, Fuchun Sun, Chuanqi Tan, Huaping Liu, and Wenbing Huang. Deep feature pyramid reconfiguration for object detection. In *ECCV*, 2018.
- [25] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In CVPR, 2017.
- [26] Wuyang Li, Zhen Chen, Baopu Li, Dingwen Zhang, and Yixuan Yuan. Htd: Heterogeneous task decoupling for two-stage object detection. *IEEE Transactions on Image Processing*, 30:9456–9469, 2021.
- [27] Xiang Li, Wenhai Wang, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In *CVPR*, 2021.
- [28] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized Focal Loss: learning qualified and distributed bounding boxes for dense object detection. In *NeurIPS*, 2020.
- [29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In CVPR, 2017.
- [30] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In ECCV, 2014.
- [32] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. Ssd: Single shot multibox detector. In ECCV, 2016.
- [33] Xin Lu, Buyu Li, Yuxin Yue, Quanquan Li, and Junjie Yan. Grid R-CNN. In CVPR, 2019.

- [34] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In AAAI, 2020.
- [35] Jiangmiao Pang, Kai Chen, Jianping Shi, Huajun Feng, Wanli Ouyang, and Dahua Lin. Libra R-CNN: Towards balanced learning for object detection. In CVPR, 2019.
- [36] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In CVPR, 2019.
- [37] Heqian Qiu, Hongliang Li, Qingbo Wu, and Hengcan Shi. Offset bin classification network for accurate object detection. In *CVPR*, 2020.
- [38] Han Qiu, Yuchen Ma, Zeming Li, Songtao Liu, and Jian Sun. Borderdet: Border feature for dense object detection. In ECCV, 2020.
- [39] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [40] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In CVPR, 2017.
- [41] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [43] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized Intersection over Union: A metric and a loss for bounding box regression. In *CVPR*, 2019.
- [44] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015.
- [45] Wonchul Son, Jaemin Na, Junyong Choi, and Wonjun Hwang. Densely guided knowledge distillation using multiple teacher assistants. In *ICCV*, 2021.
- [46] Guanglu Song, Yu Liu, and Xiaogang Wang. Revisiting the sibling head in object detector. In *CVPR*, 2020.
- [47] Ruoyu Sun, Fuhui Tang, Xiaopeng Zhang, Hongkai Xiong, and Qi Tian. Distilling object detectors with task adaptive regularization. arXiv preprint arXiv:2006.13108, 2020.
- [48] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [49] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *CVPR*, 2019.
- [50] Jiaqi Wang, Wenwei Zhang, Yuhang Cao, Kai Chen, Jiangmiao Pang, Tao Gong, Jianping Shi, Chen Change Loy, and Dahua Lin. Side-aware boundary localization for more precise object detection. In ECCV, 2020.
- [51] Keyang Wang and Lei Zhang. Reconcile prediction consistency for balanced object detection. In *ICCV*, 2021.
- [52] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, 2019.
- [53] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask:

Single shot instance segmentation with polar representation. In *CVPR*, 2020.

- [54] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017.
- [55] Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, Xiaopeng Zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *ICML*, 2021.
- [56] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. Scrdet: Towards more robust detection for small, cluttered and rotated objects. In *ICCV*, 2019.
- [57] Xue Yang, Xiaojiang Yang, Jirui Yang, Qi Ming, Wentao Wang, Qi Tian, and Junchi Yan. Learning high-precision bounding box for rotated object detection via kullbackleibler divergence. In *NeurIPS*, 2021.
- [58] Jiahui Yu, Yuning Jiang, Zhangyang Wang, Zhimin Cao, and Thomas Huang. UnitBox: an advanced object detection network. In ACM MM, 2016.
- [59] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2017.
- [60] Hongkai Zhang, Hong Chang, Bingpeng Ma, Naiyan Wang, and Xilin Chen. Dynamic R-CNN: Towards high quality object detection via dynamic training. In *ECCV*, 2020.
- [61] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sünderhauf. Varifocalnet: An iou-aware dense object detector. In CVPR, 2021.
- [62] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *ICLR*, 2020.
- [63] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z. Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 2020.
- [64] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-IoU Loss: Faster and better learning for bounding box regression. In AAAI, 2020.
- [65] Zhaohui Zheng, Ping Wang, Dongwei Ren, Wei Liu, Rongguang Ye, Qinghua Hu, and Wangmeng Zuo. Enhancing geometric factors in model learning and inference for object detection and instance segmentation. *IEEE Transactions on Cybernetics*, 2021.
- [66] Benjin Zhu, Jianfeng Wang, Zhengkai Jiang, Fuhang Zong, Songtao Liu, Zeming Li, and Jian Sun. Autoassign: Differentiable label assignment for dense object detection. arXiv preprint arXiv:2007.03496, 2020.
- [67] Chenchen Zhu, Fangyi Chen, Zhiqiang Shen, and Marios Savvides. Soft anchor-point object detection. In CVPR, 2020.
- [68] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *CVPR*, 2019.
- [69] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, 2019.