Plug-and-Play Few-shot Object Detection with Meta Strategy and Explicit Localization Inference

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

25-10-2021 / 26-10-2021

CITATION

Huang, Junying; Chen, Fan; Lin, Liang; zhang, dongyu (2021): Plug-and-Play Few-shot Object Detection with Meta Strategy and Explicit Localization Inference. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.16864711.v1>

DOI

[10.36227/techrxiv.16864711.v1](https://doi.org/10.36227/techrxiv.16864711.v1)

Plug-and-Play Few-shot Object Detection with Meta Strategy and Explicit Localization Inference

Junying Huang, Fan Chen, Liang Lin and Dongyu Zhang

Abstract—Aiming at recognizing and localizing the object of novel categories by a few reference samples, few-shot object detection is a quite challenging task. Previous works often depend on the fine-tuning process to transfer their model to the novel category and rarely consider the defect of fine-tuning, resulting in many drawbacks. For example, these methods are far from satisfying in the low-shot or episode-based scenarios since the fine-tuning process in object detection requires much time and high-shot support data. To this end, this paper proposes a plug-and-play few-shot object detection (PnP-FSOD) framework that can accurately and directly detect the objects of novel categories without the fine-tuning process. To accomplish the objective, the PnP-FSOD framework contains two parallel techniques to address the core challenges in the few-shot learning, i.e., across-category task and few-annotation support. Concretely, we first propose two simple but effective meta strategies for the box classifier and RPN module to enable the across-category object detection without fine-tuning. Then, we introduce two explicit inferences into the localization process to reduce its dependence on the annotated data, including explicit localization score and semi-explicit box regression. In addition to the PnP-FSOD framework, we propose a novel one-step tuning method that can avoid the defects in fine-tuning. It is noteworthy that the proposed techniques and tuning method are based on the general object detector without other prior methods, so they are easily compatible with the existing FSOD methods. Extensive experiments show that the PnP-FSOD framework has achieved the state-of-the-art few-shot object detection performance without any tuning method. After applying the one-step tuning method, it further shows a significant lead in both efficiency, precision, and recall, under varied few-shot evaluation protocols.

Index Terms—Meta Learning, Few-Shot Object Detection, Few-Shot Learning, Plug-and-Play, Fine-Tuning.

I. INTRODUCTION

IN the past decade, deep-learning-based methods have achieved remarkable success in various computer vision tasks, such as image classification [13], [5] and object detection [28], [25]. However, these methods are fundamentally dependent on a large amount of annotated data, so their generalization ability towards the open world tasks is limited. This triggers active researches on few-shot learning, which aims to develop models that can be generalized to the unseen (novel) categories with only a few support data with annotations for reference.

By leveraging meta-learning and distance metric learning, remarkable progress has been made in few-shot image classification. However, the works attempting to solve few-shot object detection (FSOD) have encountered setbacks due to the

Fig. 1. Impact of the tuning process on performance, including the improvement achieved and the time required. Averaged results with 10 times running on MS COCO dataset under 10-shot 20-way setting. TFA [42], FSIW [44], and FSCE [31] are inapplicable before fine-tuning. A-RPN [7] can be plug-and-play the same as ours, but it performs poorly before fine-tuning. Our PnP-FSOD has already outperformed these methods before tuning.

complexity of object detection tasks. Most existing methods, whether meta-learning-based or fine-tuning-based, require a cumbersome fine-tuning process on the support data. Otherwise, they will suffer from seriously degraded or even inapplicable. Therefore, these methods still have some obvious drawbacks: (a) Due to the costly fine-tuning time (Figure 1), they are far from satisfying in the episode-based scenario, such as meta-testing that requires evaluating the model in hundreds or thousands of different episodes. (b) They tend to ignore the instances diverse from the support set due to over-fitting to the support instances, resulting in low-recall. (c) These methods perform poorly under the low-shot setting since the effectiveness of fine-tuning depends on the number of support instances.

To remedy these defects, this paper proposes a novel plug-and-play few-shot object detection (PnP-FSOD) framework. Concretely, “plug-and-play” (PnP) means to directly detect the objects of novel categories without fine-tuning. The PnP-FSOD framework is based on Faster R-CNN [28] and built by addressing its defects in few-shot object detection, including the improvements on the box classifier, RPN module, and the localization process in the R-CNN derived model.

First of all, the box multi-classifier can only classify the region proposal into the seen category, which directly leads to many existing few-shot object detectors [42], [31], [44] inapplicable on the novel category before fine-tuning. A-RPN [7] and RepMet [17] attempt to replace the multi-classifier with the comparison-classifier and distance-classifier to recognize the instance of novel categories. However, object detection is a much more complex task, which requires detecting the foreground, localization, and classification. The comparison-

J. Huang, F. Chen, L. Lin, and D. Zhang are with the School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China, (e-mail: huangjy_229@mail2.sysu.edu.cn; chenf99@mail2.sysu.edu.cn; linliang@ieee.org; zhangdy27@mail.sysu.edu.cn).

Fig. 2. Comparisons of the motivation, learnability, and generalization ability between three classifiers. The multi-classifier explicitly learns a hyper-plane for each base category on the feature space, which is inapplicable on the novel category. The comparison-classifier learns a class-agnostic binary classifier on the joint feature space. It can directly recognize the objects of novel categories but still suffers from category bias since its parameters are trained on the base category data. The distance-classifier is non-parameter and performs classification according to the nearest neighbor rule. Although it can't preserve the knowledge from training, it doesn't suffer from category bias.

classifier and distance-classifier perform poorly on learning the complex feature space for object detection since they can't preserve the class-specific knowledge, so the performance of RepMet and A-RPN lags behind the fine-tuning-based methods.

In Figure 2, we compare the multi-classifier, comparison-classifier, and distance-classifier. According to the correlation between parameters and categories, the distance-classifier show the best generalization ability to novel categories, followed by the comparison-classifier. Meanwhile, the distance-classifier perform the worst on learning an appropriate feature space due to the lack of memory for training data, followed by the comparison-classifier. Based on the observation, we propose a meta strategy called the dynamic classifier module that uses different classifiers during training and inference to build a box classifier with both generalization and learnability in the PnP-FSOD framework. While getting rid of fine-tuning, it also significantly improves the few-shot performance.

Second, the region proposal network (RPN) suffers a fatal defect in few-shot object detection. As shown in Figure 3, the positive anchors in the training phase are disjoint with the related anchors for the novel category so that these related anchors are either ignored or regarded as background in the base training. A-RPN [7] proposes to generate class-specific region proposals to avoid RPN focusing only on the base class during testing. But the fundamental problem that only the objects belonging to base categories are treated as positive anchors in base training remains unresolved, which leads to the bottleneck of their model before fine-tuning. FSCE [31] and R-FSOD [32] also point out the defects of the RPN module, but they only adjust the RPN module in the fine-tuning process, increasing the dependence on the fine-tuning process

Different from them, we argue that the RPN module in few-shot object detectors should focus on any potential foreground instances instead of only the object instances belonging to the training category. Therefore, we propose another meta strategy that trains the RPN module with semi-supervised algorithms

Fig. 3. The positive anchors in the training and testing phase are disjoint.

to capture the potential foreground instances belonging to the novel category. Benefited from the two meta strategies, i.e., dynamic classifier and semi-supervised RPN, the PnP-FSOD framework realizes the across-category object detection without fine-tuning. It can already be competitive to the latest few-shot object detectors with fine-tuning [42], [45].

Third, in few-shot learning, only a few annotations for novel categories are supported. However, the localization process in the R-CNN derived model, composed of the localization score and box regression, is enabled by implicit fitting and without logic inference, so it depends on training by a large amount of annotated data. Therefore, we propose to reduce the dependence on the annotated data by introducing explicit logic inference into the localization process. Concretely, we first introduce the pixel-wise spatial contrast into the box classifier to reduce the inconsistency between localization score and classification score. It can generate a higher score for the region proposal with better localization. Then, we propose a semi-explicit box regression by leveraging an explicit regression mechanism. Based on the meta strategies and explicit localization inferences, our PnP-FSOD framework significantly promotes the current upper-performance limit [7] by about 70% under the PnP setting, which also outperforms the state-of-the-art non-PnP framework [31], [44], [15].

Finally, we propose a novel one-step tuning method that only feeds the support instances into the model once. Compared with the existing fine-tuning method, it can improve the performance in almost negligible time (Figure 1) and avoid the defects like low-recall and performance degradation. To evaluate the proposed PnP-FSOD framework and one-step tuning method, we conduct experiments on the one-time FSOD evaluation and standard meta-testing protocols. On the commonly used one-time FSOD evaluation, PnP-FSOD with the one-step tuning outperforms the previous state-of-the-art method by 1.1%-2.7% *AP* under the 1-10 shot settings. On the meta-testing protocols, it achieves 3.7%-5.3% *AP* leads under different settings.

Our main contributions can be summarized as follows:

- 1) We propose two meta strategies that can enable across-category object detection without fine-tuning.
- 2) We introduce two explicit localization inferences to reduce the dependence of the localization process in the R-CNN derived model on the annotated data.

- 3) We present a novel one-step tuning method that can avoid the defects in the existing fine-tuning method.
- 4) Extensive experiments show the proposed PnP-FSOD framework with novel tuning can achieve state-of-the-art results on both efficiency, precision, and recall.

II. RELATED WORK

A. General Object Detection.

General object detection is a fundamental task in computer vision which has attracted lots of attention. Modern object detectors can be divided into two kinds: one-stage detectors and two-stage detectors. One-stage detectors directly predict categories and locations of objects, e.g., YOLO series [25], [26], [27], [1], SSD [22], etc. Two-stage detectors was pioneered by R-CNN [11]. As their name, two-stage detectors first explicitly generate class-agnostic region proposals of potential objects, then further refine the proposals and classify them into special categories [28], [10], [12]. In general, two-stage detectors are slower but achieve better performance than one-stage ones. However, all of these works heavily rely on a huge amount of annotated data with elaborate bounding boxes and thus can not be directly used to solve the few-shot object detection problem.

B. Few-shot Object Detection.

Few-shot object detection was proposed to handle the situation where only a few annotated examples are available. There are mainly two types of methods aiming to address the few-shot object detection problem, i.e., meta-learning-based methods and fine-tuning-based methods.

Meta-learning-based methods attempt to build their few-shot detectors by leveraging various meta-learning techniques to extract the class-agnostic knowledge or transfer the knowledge from base categories to novel categories. Despite this goal, these methods still require the fine-tuning process. Otherwise, they are either inapplicable [16], [45], [40], [39], [44], [15] or lagging behind other methods [17], [7]. FSRW [16] extracts generic meta-features from base categories, then adjusts them using the reweighting features for novel categories. Meta R-CNN [45] proposes to use the reweighting features over RoI features instead of the image feature. MetaDet [39] and GenDet [40] propose to estimate the new parameters in the detector for detecting novel category instances. RepMet [17] incorporates distance metric learning into few-shot detection to help classify the proposals. FSIW [44] performs a similar process as Meta R-CNN [45], but it uses a more complex feature aggregation. A-RPN [7] proposes attention-RPN and multi-relation detector to detect novel objects with a contrastive training strategy. Recent DCNet [15] proposes to fully exploit local information to benefit the detection process and alleviate the scale variation problem by context-aware feature aggregation.

Fine-tuning-based methods adopt the general object detectors and focus on improving the fine-tuning process on the support data to effectively transfer the category-specific model to the novel category. They are once suffered in poor performance, but recent works set the new state-of-the-art. TFA [42] simply fine-tunes the last layer of Faster R-CNN [28] and substantially improves the performance. MPSR [43] proposes to handle the

scale variance issue by multi-scale positive sample refinement, but it needs a manual selection. Recent FSCE [31] builds a strong baseline upon TFA [42] and boosts the performance by large margins. It also integrates contrastive learning to train the detection head and achieves impressive performance.

As can be seen, fine-tuning is essential for the existing method, no matter what type. Figure 1 shows the time in some existing methods [42], [7], [44], [31] required for fine-tuning. Note that it requires at least 15 minutes or even a few hours, which is unacceptable in few-shot learning. By comparison, the fine-tuning process in few-shot image classification often requires less than one second. Thus for the application of few-shot object detection, it is valuable to get rid of fine-tuning or replace it with a time-acceptable tuning method.

C. Problem Definition of FSOD

Given a dataset $\mathcal{D}_b$ with abundant annotated instances of the base (seen) category C_b for base training, the task is to detect all the objects belonging to the novel (unseen) category C_n on the query set $\mathcal{Q}$. C_b and C_n do not intersect. For each novel category, there is also a support set $\mathcal{S}$ with few annotated instances for reference. In more detail, N-way K-shot object detection means C_n contains N categories and the support set contains K annotated instances (usually less than 10), i.e., $\mathcal{S}_c = \{(I_i, b_i), i = 1, \dots, K\}$ where I and b denote the support image and the bounding box of the support instances.

III. PNP-FSOD FRAMEWORK

The overview of PnP-FSOD is shown in Figure 4, which is based on Faster R-CNN [28]. Given a query image and a support set, the goal is to detect the objects belonging to the support category in the query image. Firstly, the backbone extracts the feature maps of the query image and all support images. Then the RPN module predicts the region proposal in the query image. After that, the RoI module, including RoI-pooling and RoI-extractor, extracts the feature maps of region proposals in the query image and the feature maps of all the support instances. Finally, the box classifier and box regressor further predict the category and the box regression of the region proposals by comparing the region features and the support feature that is the average of all support instance feature maps. In addition to the framework, there is also a non-maximum suppression (NMS) algorithm to select the non-overlap prediction box according to the localization score.

In particular, we first propose two meta strategies (Sec. III-A), i.e., dynamic classifier and semi-supervised RPN (SS-RPN), to enable across-category object detection without fine-tuning. Then, we further present two explicit localization inferences (Sec. III-B), i.e., explicit localization score and semi-explicit box regression, to reduce the dependence of the localization process on the annotated data. As shown in Figure 4, we mark the position of the proposed techniques in yellow.

A. Meta Strategies

The key to across-category task lies in the generalization to the novel category. Therefore, we propose two meta strategies for the box classifier and RPN module, respectively.

Fig. 4. The overview of our proposed PnP-FSOD framework. It is built on top of Faster R-CNN. The module without yellow mask is the same as Faster R-CNN. SS-RPN means the proposed semi-supervised RPN. Dynamic Classifier means the dynamic box classifier module.

Dynamic Classifier: As shown in Figure 2, we argue that the three existing commonly used classifiers all have some defects. Therefore, we propose to adopt a dynamic classifier module. Specifically, the box classifier module uses different classifiers in training and inference to realize both the generalization to novel categories and learnability on feature space. In addition, since the distance-classifier is non-parameter, it doesn't require fine-tuning when replacing the trained classifier with the distance-classifier during inference. Table I shows the ablation study of different classifier combinations on the MS COCO dataset under the 10-shot PnP setting and one-time FSOD evaluation protocol. For a fair comparison, both the multi-classifier and comparison-classifier are single connection layers, and the distance-classifier is the sigmoid value of cosine distance between global feature vectors.

As shown in Table I, the multi-classifier is inapplicable on novel categories without fine-tuning. The comparison-classifier performs worse than the distance-classifier since its parameters are still affected by the bias of the base categories. However, the distance-classifier can significantly benefit from models trained by the multi-classifier or comparison-classifier due to their learnability. Based on the results, the PnP-FSOD framework adopts the comparison-classifier in base training and replaces it with the distance-classifier in inference. The details of the two classifiers will be described respectively in Sec. III-B.

TABLE I
ABLATION EXPERIMENTAL RESULTS FOR DYNAMIC CLASSIFIER AND SEMI-SUPERVISED RPN ON MS COCO UNDER 10-SHOT SETTING.

Dynamic Classifier		AP	AP_{50}	AP_{75}
(Training)	(Inference)			
Multi	Multi	inapplicable		
Comparison	Comparison	4.71	8.79	4.41
Distance	Distance	5.94	15.64	2.82
Multi	Distance	6.89	15.28	5.44
Comparison	Distance	8.69	17.38	7.78
+ Semi-supervised RPN		10.54	20.96	9.08

Semi-supervised RPN (SS-RPN): Generating region proposals by the class-agnostic detector (e.g., RPN) is a crucial idea in two-stage detection models, but it has a fatal defect in few-shot object detection. As shown in Figure 3, the positive anchors in the training phase are disjoint with the related anchors for the novel categories so that these related anchors are either

ignored or regarded as background in the base training. Thus the RPN module in few-shot object detection framework is implicitly category-specific to the base categories, making it is hard to capture the anchors related to novel categories.

To address the problem, we propose to adopt semi-supervised algorithms to train the RPN module. Concretely, in the RPN training, all the positive anchors are certain foreground instances with correct labels. But the negative anchors consist of background or potential object instances not belonging to the base categories, so we remark them as unlabeled data. In the PnP-FSOD framework, we adopt a simple but effective semi-supervised algorithm, i.e., Pseudo Label.

In more detail, we first annotate the anchors whose Intersection over Union (IoU) with ground truth box is less than 0.3 as negative, and the anchors whose IoU is greater than 0.7 as positive, following the standard RPN training process. Then, we annotate the negative anchors with RPN prediction probabilities greater than threshold τ as the pseudo positive label and compute positive loss, the same as the positive label. To calculate the balance loss, we keep the ratio of the positive anchors, the negative anchors, and the pseudo positive anchors as 1:1:1. The training process is shown at the top of Figure 5. The selection of threshold τ is shown in Sec. V-E.

Discussion: In the object detection task, it is unrealistic to achieve both high recall and high precision. The RPN with semi-supervised training inevitably leads to more background proposals in the inference. However, we argue that this is worth it because it is possible to eliminate these background proposals by the box classifier in the two-stage detection model. On the contrary, the correct anchors ignored by the RPN module are irreparable. Experiments also validate that the semi-supervised training strategy for RPN can significantly boost plug-and-play performance for the novel category (Table I).

B. Explicit Localization Inferences

In this section, we focus on reducing the dependence of the localization process in the R-CNN model on the annotated data, including the localization score and the box regression.

Explicit Localization Score: R-CNN derived model implicitly evaluates the localization score of the region proposal by the classification score from the box classifier. However, the object bounding box usually contains some low confidence regions, such as background and the low-discrimination parts of the target object. The feature of the complete bounding box will be diluted by low confidence regions, which usually lags behind

Fig. 5. The detail of components in the PnP-FSOD. The input of the comparison-classifier and box regressor are the feature map of a region proposal (blue) and the average feature map of a category support set (red). GAP, FC, and Conv mean the global average pooling, full connection layer, and convolution layer. $\oplus$ means to concatenate the input feature.

the local high confidence region on the classification score (like the yellow and red box in Figure 6 (a)). Thus the localization score and the classification score are inconsistent. Although this inconsistency can be alleviated by the model fully trained on a large amount of annotated data, it is inevitable in the novel category with only a few annotations.

To tackle the problem, we introduce the pixel-wise spatial contrast between feature maps into the box classification, which can evaluate the localization of the region proposal more explicitly. Before the box classification, the RoI module reshapes the feature maps to the same shape, so the pixel-wise spatial contrast can generate a higher score for those proposals whose object's relative positions are more similar to the object's relative positions in the support box. It often indicates better localization. Meanwhile, the contrast between the global features is also essential. Because the pixel-wise spatial contrast may not be accurate when the object's posture and position vary, but the global feature contrast is not affected. Thus, we design the comparison-classifier (in the middle of Figure 5) as the score addition from two lightweight networks, which compare the feature maps and feature vectors, respectively.

For the distance-classifier, we calculate the cosine distance between global feature vectors and the cosine distance between flattened feature maps, then calculate the prediction score by the sharp sigmoid function σ . Concretely, given a region proposal x and a support set of category c , the probability of x belonging to category c is predicted as:

$$Pr(c; x) = \sigma[(1 - \alpha)D(f_x, f_c) + \alpha D(v_x, v_c)] \quad (1)$$

$$D(x, y) = \frac{x^T y}{\|x\| \cdot \|y\|} \quad (2)$$

$$\sigma(x) = \frac{1}{1 + e^{-\lambda x}} \quad (3)$$

where v and f represent the vectors obtained from the feature map by global average pooling and flatten function, respectively.

Fig. 6. The motivations of explicit localization inferences. The yellow box and the red box in (a) mean the bounding box and high confidence region, respectively. (a) shows that the regions with high classification scores and high localization scores are inconsistent. (b) shows an explicit box regression method.

D and σ means the cosine distance and sharp sigmoid function. The selection of α and λ is shown in Sec. V-E.

Semi-explicit Box Regression: General object detectors [28], [25] often implicitly predict the box regression by a lightweight network. It needs to be driven by a large amount of annotated data, which is not applicable in few-shot object detection. To tackle the problem, we propose a semi-explicit box regression method by leveraging an explicit regression mechanism: any two pairs of coordinates between the region proposal and the GT box can provide two regression equations, equivalent to a correct box regression, as shown in Figure 6 (b). Despite the equivalence, this explicit regression is inapplicable since the GT box is unavailable during inference. Therefore, we propose to extract sufficient possible coordinate pairs from the comparison between the region proposal and the support box, and then predict the box regression by these coordinate pairs.

Specifically, we propose to extract and evaluate the coordinate pairs by the pixel-matching contrast. Given the feature map of a region proposal x and the average support feature map of category c (e.g., $F_x, F_c \in R^{d \times r \times r}$), we first reshape them as list of feature vectors (e.g., $\hat{F}_x, \hat{F}_c \in R^{r^2 \times d}$), then compute the distance matrix M between two lists, where

$$M_{i,j} = D(\hat{F}_x^i, \hat{F}_c^j). \quad (4)$$

Then, we flatten the distance matrix to a distance vector $\in R^{r^4}$ and concatenate it with the region proposal feature. Finally, we feed the concatenated feature into a lightweight network to predict the box regression as shown in the bottom of Figure 5.

Discussion: In the distance vector, each index represents a coordinate pair between two feature maps, and the corresponding value indicates the confidence score of the coordinate pair. However, these confidence scores may not be accurate due to the difference between the support instance and the GT instance. Explicitly calculating box regression by these coordinate pairs will suffer from serious errors by inaccurate scores. Thus we still predict the box regression by feeding all confidence scores into a neural network to implicitly synthesize all the equations. This regression method is between the implicit regression in general object detectors and the explicit regression by the regression equation, so we call it semi-explicit.

C. Training Strategy:

Inspired by A-RPN [7], we train our model by the 2-way 10-shot contrastive training strategy. For each training image as a query image, we first randomly select a positive category c_1 that appears in the image and a negative category c_2 that doesn't appear in the image ($c_1, c_2 \in \mathcal{C}_b$) and then collect their support sets ($\mathcal{S}_{c_1}$ and $\mathcal{S}_{c_2}$), both containing ten object instances. After the forward process described in III, we train the box comparison-classifier by the positive loss that matches the same category and the negative loss that distinguishes the different categories. For the box regressor, we only calculate the box regression loss of the region proposal belonging to c_1 . The training of the semi-supervised RPN is the same as Faster R-CNN [28] except for the pseudo-label described above. The final loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{rpn} + \mathcal{L}_{cls} + \mathcal{L}_{reg} \quad (5)$$

where $\mathcal{L}_{rpn}$ consists of the classification loss and regression loss of proposals, $\mathcal{L}_{cls}$ is the binary cross-entropy loss for box classification, $\mathcal{L}_{reg}$ is the smoothed L_1 loss for box regression.

Fig. 7. The visualization of support branch with and without consistency tuning during inference.

IV. CONSISTENCY TUNING

Finally, we present a novel one-step tuning method that can replace the fine-tuning process but avoid defects. While computing the average feature map of a category support set, we collect all the region proposals (R proposals for each support image) extracted from the RPN module, then refine the average support feature. In particular, we propose a consistency optimization (Eq. 6) instead of the classification loss since the box classifier needs to score the localization of the region proposal (as described in Sec. III). Concretely, it aligns the prediction probability of region proposals and their IoU between the GT box. Note that the IoU between other category (background) region proposals and the GT box is zero, so the consistency tuning can also improve the classification performance. Under the K -shot setting, the tuning formula is

$$F_c \leftarrow \arg \min_{F_c} \sum_{i=1}^{KR} \|Pr(c; x_i) - \max_{G_c} IoU(x_i, G_c)\|_2^2 \quad (6)$$

where x is the region proposal, G_c means the GT box belonging to category c , and $Pr(c; x_i)$ means the probability calculated by

the distance-classifier (Eq. 1) with the average support feature F_c . We solve the optimization equation by the gradient descent algorithm, requiring about two seconds when $KR = 5,000$.

Compared with the existing fine-tuning methods, our consistency tuning is one-step that feeds each support instance into the model only once. Thus the time required is almost negligible. In addition, it doesn't update any parameters in the model, so it leads to neither over-fitting to the support instances nor the performance degradation under low-shot settings. Note that, different from the existing FSOD frameworks, the tuning method is not essential for the PnP-FSOD framework. In the pursuit of efficiency, PnP-FSOD can still maintain state-of-the-art performance without the tuning method. As shown in Figure 7, we provide the frameworks with and without the consistency tuning process. In the experiments, we also distinguish the performances whether to use consistency tuning.

V. EXPERIMENTS

A. Experimental Setup

Dataset: In this paper, we conduct experiments on two large and challenging few-shot detection benchmark datasets, MS COCO [20] and FSOD dataset [7], which contains 800K objects belonging to 80 categories and 182K objects belonging to 1000 categories respectively. For MS COCO, we set the 20 categories belonging to PASCAL VOC [6] as the novel categories, and set the remaining 60 categories as the base categories following the existing works [42], [31], [7], [16]. We use the *train2017* with only annotations belonging to base categories for training and evaluate the detection result of the novel categories on the *val2017*. FSOD dataset is specially designed for few-shot object detection, whose training set and test set only contain disjoint 800 base categories and 200 novel categories.

Implement Details: We use the commonly used ResNet-50 [13] as our backbone. Following the existing works [7], [42], [31], the backbone is pretrained on ImageNet [5]. All the network architectures and the most of hyper-parameters remain the same as Faster R-CNN [28] except for the box classifier and box regressor, as described in Sec. III-B. In addition, we halve the number of sampled anchors in RPN and proposals in RoI head used for loss calculation from (512, 256) to (128, 128) during training, and reduce the final generated region proposals during testing from 1000 to 500 for COCO and 100 for FSOD dataset per category per image, respectively. Our model is trained by SGD optimizer on 3 RTX 2080Ti GPUs with batch size of 9 (3 query images per GPU) for 120,000 iterations. The learning rate is initialized as 0.003 with the weight decay of factor 0.1 at 80,000th and 110,000th iteration.

Evaluation protocol: We conduct experiments on the one-time FSOD protocol proposed in [16] and the meta-testing protocol commonly used in the few-shot learning [37]. Given the support sets of all novel categories, the one-time FSOD protocol requires evaluating the performance of detecting these novel categories on the complete test set. The meta-testing protocol requires evaluating the average performance of the detector under multiple random episodes. Each episode consists of a randomly sampled support set and query set.

TABLE II

FEW-SHOT DETECTION RESULTS FOR 20 NOVEL CLASSES ON COCO DATASET. “PNP” MEANS THE MODEL IS PLUG-AND-PLAY AND DOESN’T REQUIRE THE TUNING PROCESS. RED/BLUE INDICATE THE SOTA/SECOND BEST. “FT” AND “CT” MEAN FINE-TUNING AND CONSISTENCY TUNING. + MEANS THE RESULT IS ESTIMATED ACCORDING TO THE DESCRIPTION IN THEIR PAPER.

Model	Backbone	Average Precision						Average Recall						Tuning time
		AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	AR_1	AP_{10}	AR_{100}	AR_S	AR_M	AR_L	
A-RPN [7]	Res-50	7.3	13.2	7.1	4.4	8.7	10.7	17.5	32.3	33.2	10.0	34.7	50.4	PnP
Meta R-CNN [45]	Res-50	8.7	19.1	6.6	2.3	7.7	14.0	12.6	17.8	17.9	7.8	15.6	27.2	5 ⁺ h
MPSR [43]	FPN-101	9.8	17.9	9.7	3.3	9.2	16.1	15.7	21.2	21.2	4.6	19.6	34.3	40 ⁺ min
TFA [42]	FPN-101	9.8	19.7	8.9	2.8	9.2	16.1	14.5	18.6	18.6	5.3	14.8	33.1	16 h
FSCE [31]	FPN-101	11.9	22.3	11.6	2.9	11.1	17.6	17.0	26.6	26.5	6.7	26.3	42.3	3 h
A-RPN+FT[7]	Res-50	12.0	22.4	11.8	2.9	12.2	20.7	18.8	26.4	26.4	3.6	23.6	45.6	15 min
FSIW [44]	Res-50	12.5	27.3	9.8	2.5	13.8	19.9	20.0	25.5	25.7	7.5	27.6	38.9	5 h
DCNet [15]	Res-101	12.8	23.4	11.2	4.3	13.8	21.0	18.1	26.7	25.6	7.9	24.5	36.7	40 ⁺ min
PnP-FSOD	Res-50	12.5	23.6	11.6	6.2	15.6	21.1	19.1	33.5	35.6	11.4	39.4	54.9	PnP
PnP-FSOD+CT	Res-50	13.3	24.7	12.6	5.7	16.8	22.7	19.5	33.6	35.5	11.5	39.3	54.8	40 s

TABLE III

FEW-SHOT DETECTION RESULTS ON COCO DATASET UNDER DIFFERENT SHOT SETTINGS. * MEANS THE RESULT IS RE-IMPLEMENTED. RED/BLUE INDICATE THE SOTA/SECOND BEST. THE RESULTS ARE AVERAGED OVER TEN RANDOM RUNS.

Model	Backbone	1-shot			2-shot			3-shot			5-shot			10-shot		
		AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}
A-RPN* [7]	Res-50	3.6	7.2	3.2	5.1	9.7	4.7	5.6	10.7	5.2	6.3	11.9	5.9	6.7	12.5	5.8
TFA [42]	FPN-101	1.9	3.8	1.7	3.9	7.8	3.6	5.1	9.9	4.8	7.0	13.3	6.5	9.1	17.1	8.8
FSIW [44]	Res-50	3.2	8.9	1.4	4.9	13.3	2.3	6.7	18.6	2.9	8.1	20.1	4.4	10.7	25.6	6.5
A-RPN+FT*[7]	Res-50	-	-	-	4.8	9.2	3.9	5.9	11.6	5.7	8.1	15.5	7.4	10.9	20.5	9.4
FSCE* [31]	FPN-101	2.0	4.9	1.3	4.2	9.5	3.4	5.7	12.0	4.7	7.6	15.6	6.5	11.2	22.3	9.8
PnP-FSOD	Res-50	5.1	10.8	4.3	7.7	15.7	6.6	8.9	17.6	7.9	10.2	20.0	9.2	11.5	22.2	10.6
PnP-FSOD+CT	Res-50	5.2	10.9	4.4	7.8	15.9	6.8	9.3	18.3	8.2	10.7	20.9	9.7	12.3	23.4	11.3

B. One-time FSOD evaluation

MS COCO result: In Table II, we compare our PnP-FSOD with the previous state-of-the-art methods under the 10-shot setting. For a fair comparison, we also report the backbone used in the models and the time requiring for the tuning process. Although the competitive DCNet [15] doesn’t adopt FPN, it still uses multi-scale features to enhance its detector. As shown in the table, PnP-FSOD achieves new state-of-the-art results on both performance and efficiency. Under the PnP setting, we outperform the latest method [7] by about 70% on AP metric. What’s more surprising, PnP-FSOD achieves both high precision and high recall, which is rare in the existing methods. For example, A-RPN [7] before fine-tuning is competitive to our model on the recall, but its precision (AP) is only 55% of ours; DCNet [15] is competitive to our model on the precision, but its recall (AR_{10}) is only 80% of ours. It manifests that PnP-FSOD can not only capture more objects belonging to the novel category but also classify them more accurately.

Then, we conduct further comparison experiments under different shot settings ($K \in \{1, 2, 3, 5, 10\}$). For a fair comparison, we evaluate all the methods over ten random runs. In each run, all the methods adopt the same support set. The support sets are generated from TFA [42]. As shown in Table III, PnP-FSOD with consistency tuning can outperform the

previous state-of-the-art by 1.1%-2.7% AP under different shot settings. Besides PnP-FSOD, A-RPN [7] is the only method that is applicable under the PnP setting. It can outperform other methods under the low-shot settings since the fine-tuning process depends on sufficient support data. However, as the shot number increase, it becomes significantly behind the method with tuning. On the contrary, PnP-FSOD can always stay ahead even it doesn’t adopt tuning, demonstrating the generalized effectiveness of our approach under varied few-shot settings.

FSOD dataset result: Similar to MS COCO, we evaluate all the methods over ten random runs, and all the methods adopt the same support set in each run. The support sets are generated from the code of TFA [42]. The average results under 1/3/5 shot setting are shown in Table IV. As shown in the table, PnP-FSOD achieves state-of-the-art results on 1/3 shots and comparable results on 5-shots without tuning, demonstrating the strong generalization to the various novel categories. It is worth noting that the FSOD dataset has 200 novel classes, so the support set in the 5-shot setting has 1000 object instances. It’s usually hard to obtain so many instances in the practice scenario. Therefore, the detection performance under the low-shot setting is more important. In addition, PnP-FSOD also performs both high precision and high recall on the FSOD dataset, the same as the results on the MS COCO dataset,

indicating that it is not accidental on a particular dataset.

TABLE IV

FEW-SHOT DETECTION RESULTS FOR 200 NOVEL CLASSES ON FSOD DATASET. “PNP” MEANS THE MODEL IS PLUG-AND-PLAY AND DOESN’T REQUIRE THE TUNING PROCESS. “TIME” MEANS THE TUNING TIME. RED/BLUE INDICATE THE SOTA/SECOND BEST. ALL RESULTS ARE RE-IMPLEMENTED AND AVERAGED OVER TEN RANDOM RUNS.

Shot	Method	Backbone	AP	AP_{50}	AP_{75}	AR_{10}	Time
1	A-RPN [7]	Res-50	9.79	16.00	10.25	40.33	PnP
	TFA [42]	FPN-101	7.43	12.07	7.79	14.42	2 h
	FSCE [31]	FPN-101	7.21	12.54	6.92	15.20	1.5 h
	PnP-FSOD	Res-50	10.66	18.41	16.62	38.47	PnP
3	A-RPN [7]	Res-50	14.94	23.68	15.91	50.01	PnP
	TFA [42]	FPN-101	13.21	21.10	14.16	24.87	3 h
	FSCE [31]	FPN-101	14.96	25.85	14.73	29.21	3 h
	PnP-FSOD	Res-50	16.33	27.59	10.65	48.91	PnP
5	A-RPN [7]	Res-50	17.28	27.13	18.49	52.30	PnP
	TFA [42]	FPN-101	15.85	24.89	17.31	27.74	3.5 h
	FSCE [31]	FPN-101	19.58	33.42	19.79	36.10	3.5 h
	PnP-FSOD	Res-50	19.24	32.08	19.75	52.69	PnP

TABLE V

EPISODE-BASED EVALUATION WITH 95% CONFIDENCE INTERVAL ON THE MS-COCO DATASET UNDER 5-WAY SETTING, WITH 1,000 EPISODES. SPE MEANS SECONDS-PER-EPISODE.

K	Method	AP	AP_{50}	AP_{75}	SPE
5	A-RPN [7]	14.27 ± 0.27	26.61 ± 0.45	13.58 ± 0.31	11
	PnP-FSOD	16.59 ± 0.28	31.97 ± 0.47	15.12 ± 0.33	7
	PnP-FSOD+CT	17.94 ± 0.30	33.81 ± 0.48	16.77 ± 0.35	14
10	A-RPN [7]	15.12 ± 0.29	27.74 ± 0.47	14.61 ± 0.32	12
	PnP-FSOD	17.74 ± 0.29	33.59 ± 0.47	16.56 ± 0.34	9
	PnP-FSOD+CT	19.62 ± 0.31	36.16 ± 0.50	18.76 ± 0.39	19

TABLE VI

EPISODE-BASED EVALUATION WITH 95% CONFIDENCE INTERVAL ON THE MS-COCO DATASET UNDER 10-WAY SETTING, WITH 1,000 EPISODES. SPE MEANS SECONDS-PER-EPISODE.

K	Method	AP	AP_{50}	AP_{75}	SPE
5	A-RPN [7]	11.32 ± 0.19	20.84 ± 0.28	10.87 ± 0.20	30
	PnP-FSOD	14.23 ± 0.16	27.39 ± 0.26	13.11 ± 0.18	14
	PnP-FSOD+CT	15.97 ± 0.19	29.79 ± 0.30	15.13 ± 0.24	25
10	A-RPN [7]	12.11 ± 0.15	22.05 ± 0.27	11.78 ± 0.19	36
	PnP-FSOD	15.52 ± 0.14	29.41 ± 0.26	14.55 ± 0.16	16
	PnP-FSOD+CT	17.37 ± 0.19	31.95 ± 0.31	16.62 ± 0.24	39

C. Meta-testing protocol

In this section, we perform the meta-testing protocol on the MS COCO dataset. For an N -way K -shot few-shot object detection, we collect 1,000 episodes, and each episode consists of an N -way K -shot support set and a query set containing ten images for each category. Since the evaluation is performed on each episode independently, including the fine-tuning process

and the inference process, the existing non-PnP models require unacceptable time. Therefore, we only compare our PnP-FSOD with another PnP framework, i.e., A-RPN [7].

Table V and VI report the average results with the 95 % confidence interval and the required time (seconds-per-episode) under different few-shot settings, including $K \in \{5, 10\}$ and $N \in \{5, 10\}$. As shown in the table, our PnP-FSOD achieves a significant lead on both performance and efficiency. Benefited from the proposed consistency tuning that can tune the model in an acceptable time for meta-testing, it can further outperform A-RPN by 3.7%-5.3% AP (about 25%-43% relatively) and only requires comparable time, indicating that our approaches are simpler and more effective.

TABLE VII

ABLATION FOR KEY COMPONENTS PROPOSED IN THIS PAPER, RESULTS FROM 10 SHOT OF COCO DATASET.

Ablation		AP	AP_{50}	AP_{75}
baseline:	Faster R-CNN [28]	inapplicable		
+ Dyn-cls:	(Distance + Distance)	5.94	15.64	2.82
	(Multi + Distance)	6.89	15.28	5.44
	(Comparison + Distance)	8.69	17.38	7.78
+ SS-RPN		10.54	20.96	9.08
+ Reg-PM		10.64	20.59	9.58
+ Cls-PW		10.67	20.92	9.47
+ SS-RPN + Cls-PW		11.31	22.07	10.15
+ SS-RPN + Reg-PM		11.63	22.60	10.34
+ Cls-PW + Reg-PM		11.63	21.61	10.84
+ SS-RPN + Cls-PW + Reg-PM		12.47	23.53	11.62
+ Consistency Tuning		13.29	24.69	12.56

D. Ablation Studies

In this section, we evaluate the effects of the core components in PnP-FSOD. All ablation studies are conducted on the COCO dataset under the 10-shot setting and one-time FSOD evaluation protocol. Our work is built on top of Faster R-CNN [11], which is designed for the general object detection task and is originally inapplicable in PnP-FSOD. As shown in Table VII, we design the ablation experiments in three stages.

In the first stage, we evaluate the effect of the proposed dynamic classifier module, which is the essential strategy to transform the general object detector into a plug-and-play few-shot object detector. Without the dynamic classifier module, the model suffers from low performance or even unavailability due to low learnability or low generalization of the classifier. By introducing the dynamic classifier module into the Faster R-CNN, it has been competitive with Meta R-CNN [45] that adopts the re-weighting strategy and fine-tuning process.

In the second stage, we evaluate the different combinations of three proposed boosted modules, including the semi-supervised RPN (SS-RPN), the box classifier with the pixel-wise contrast (Cls-PW), and the semi-explicit box regressor with the pixel-matching contrast (Reg-PM). As shown in the Table VII, their improvements for the model performance are different

and are all in line with our expectations. Concretely, (a) The semi-supervised RPN mainly achieves the performance improvement on the AP_{50} metric (+1.1%-3.6%), indicating that it successfully captured more potential region proposals belonging to the novel categories; (b) The pixel-matching contrast in the box regressor can significantly improve the result on the AP_{75} metric (+1.3%-1.8%), which shows that it improves the localization accuracy; (c) The pixel-wise contrast in the box classifier can achieve significant improvement on both AP_{50} (+1.1%-3.6%) and AP_{75} (+1.1%-1.7%) metrics, since it can not only improve the classification accuracy but also improve the localization accuracy by generating a higher confidence for the region proposal with better localization.

Finally, we separately evaluate the improvements by consistency tuning at the bottom of Table VII. Similar to the pixel-wise contrast, it can achieve significant improvement on both AP_{50} (+1.1%) and AP_{75} (0.9%) metrics, demonstrating the proposed consistency loss can reflect both the localization the classification performance.

TABLE VIII
HYPER-PARAMETER STUDIES OF τ AND α IN THE PnP-FSOD WITHOUT TUNING, RESULTS FROM 10 SHOT OF COCO DATASET.

τ	AP	AP_{50}	AP_{75}	α	AP	AP_{50}	AP_{75}
-	11.63	21.61	10.84	0.0	11.63	22.60	10.34
0.75	11.99	22.21	11.40	1/3	12.47	23.53	11.62
0.50	12.14	22.75	11.33	1/2	12.41	23.32	11.67
0.25	12.47	23.53	11.62	2/3	12.11	22.75	11.37
0.10	12.26	22.84	11.55	1.0	10.88	20.18	10.37

TABLE IX
HYPER-PARAMETER STUDIES OF λ IN THE PnP-FSOD AND PnP-FSOD WITH TUNING, RESULTS FROM 10 SHOT OF COCO DATASET.

λ	PnP-FSOD			+ Consistency Tuning		
	AP	AP_{50}	AP_{75}	AP	AP_{50}	AP_{75}
1	12.47	23.53	11.62	12.25	22.85	11.57
5	12.47	23.53	11.62	12.53	23.23	11.84
20	12.47	23.53	11.62	13.29	24.69	12.56
50	12.39	23.37	11.61	12.93	24.13	12.34
80	10.76	21.28	9.30	12.24	23.26	11.32

E. Hyper-parameter Studies:

In this section, we study the effect and selection of the hyper-parameters in PnP-FSOD, including the threshold τ in the semi-supervised RPN, as well as the balance weight α and the scaling factor λ in the distance-classifier. For each hyper-parameter, we first select a candidate set by observation and then evaluate their performance on the COCO dataset under the 10-shot setting and one-time FSOD evaluation protocol. The performances at different values of them in Table VIII and IX. The specific analysis is as follow:

Threshold τ : Semi-supervised learning on image classification usually requires higher thresholds to reduce the incorrect pseudo labels. However, in the two-stage detector, we expect RPN to capture all the potential objects as possible and then eliminate the incorrect proposals by the box classifier head.

Therefore, the model performs better when τ is lower and reaches the optimal performance at $\tau = 0.25$.

Balance weight α : Compared with $\alpha = 0.0$ and $\alpha = 1.0$, the performances with other values are improved significantly, indicating that the comparisons between the global features and pixel-wise feature are both essential. PnP-FSOD reaches the optimal performance at $\alpha = \frac{1}{3}$.

Scaling factor λ : In the PnP-FSOD, the performance decrease only when the scale factor is too large since some prediction boxes with low confidence are wiped out. Otherwise, the performance is not affected because the scaling factor doesn't affect the confidence ranking of the predicted box. However, it plays a role in consistency tuning by adjusting the sharpness of the prediction distribution. Therefore we choose $\lambda = 20$ that reaches the optimal performance on both PnP-FSOD with and without consistency tuning.

VI. CONCLUSION

In the few-shot object detection field, the existing methods tend to transfer their model to the detection task by leveraging a fine-tuning process, leading to many drawbacks. Therefore, this paper mainly contributes to few-shot object detection by developing a novel plug-and-play few-shot object detection (PnP-FSOD) framework to accurately and directly apply the few-shot object detector to various detection tasks without fine-tuning. To accomplish the objective, we propose two meta strategies to realize the across-category object detection without fine-tuning and two explicit localization inferences to reduce the dependence of the localization process on the annotated data. In addition to the framework, we present a novel one-step tuning method that can avoid the defects in the existing fine-tuning method. It is noteworthy that our works are built on only Faster R-CNN without other prior methods, so all the approaches are easily compatible with the existing FSOD methods. Extensive experiments show that our methods achieve state-of-the-art results on both efficiency, precision, and recall. We hope our proposed approaches can inspire future works to explore more powerful few-shot object detectors.

REFERENCES

- [1] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [2] Hao Chen, Yali Wang, Guoyou Wang, and Yu Qiao. Lstd: A low-shot transfer detector for object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [3] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *International Conference on Learning Representations*, 2019.
- [4] Yinbo Chen, Xiaolong Wang, Zhuang Liu, Huijuan Xu, and Trevor Darrell. A new meta-baseline for few-shot learning. *arXiv preprint arXiv:2003.04390*, 2020.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [6] Mark Everingham, Luc Gool, Christopher K. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. In *International Journal of Computer Vision*, page 303–338, June 2010.

- [7] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4013–4022, 2020.
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [9] Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *International Conference on Learning Representations*, 2018.
- [10] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [11] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Ruibing Hou, Hong Chang, MA Bingpeng, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. In *Advances in Neural Information Processing Systems*, pages 4003–4014, 2019.
- [15] Hanzhe Hu, Shuai Bai, Aoxue Li, Jinshi Cui, and Liwei Wang. Dense relation distillation with context-aware aggregation for few-shot object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, page TBD, 2021.
- [16] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8420–8429, 2019.
- [17] Leonid Karlinsky, Joseph Shtok, Sivan Harary, Eli Schwartz, Amit Aides, Rogerio Feris, Raja Giryes, and Alex M Bronstein. Repmet: Representative-based metric learning for classification and few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5197–5206, 2019.
- [18] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille, 2015.
- [19] Seong Min Kye, Hae Beom Lee, Hoirin Kim, and Sung Ju Hwang. Transductive few-shot learning with meta-learned confidence. *arXiv preprint arXiv:2002.12017*, 2020.
- [20] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and Larry Zitnick. Microsoft coco: Common objects in context. In *ECCV. European Conference on Computer Vision*, September 2014.
- [21] Jinlu Liu, Liang Song, and Yongqiang Qin. Prototype rectification for few-shot learning. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision*, pages 741–756, Cham, 2020. Springer International Publishing.
- [22] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [23] Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2(2):1, 2018.
- [24] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. *International Conference on Learning Representations*, 2016.
- [25] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [26] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [27] Joseph Redmon and Ali Farhadi. Yolo3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. In *European Conference on Computer Vision*, pages 121–138. Springer, 2020.
- [30] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
- [31] Bo Sun, Banghuai Li, Shengcai Cai, Ye Yuan, and Chi Zhang. Fsce: Few-shot object detection via contrastive proposal encoding. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, page TBD, June 2021.
- [32] Zhibo Fan, Yuchen Ma, Zeming Li, and Jian Sun. Generalized few-shot object detection without forgetting. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, page TBD, June 2021.
- [33] Weilin Zhang, and Yu-Xiong Wang. Hallucination Improves Few-Shot Object Detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, page TBD, June 2021.
- [34] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 403–412, 2019.
- [35] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-verification. *arXiv preprint arXiv:1406.4773*, 2014.
- [36] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1199–1208, 2018.
- [37] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [38] Ting-I Hsieh, Yi-Chen Lo, Hwann-Tzong Chen and Tyng-Luh Liu. One-shot object detection with co-attention and co-excitation. In *Advances in neural information processing systems*, pages 2725–2734, 2019.
- [39] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Meta-learning to detect rare objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9925–9934, 2019.
- [40] Liyang Liu, Bochao Wang, Zhanghui Kuang, Jing-Hao Xue, Yimin Chen, Wenming Yang, Qingmin Liao, and Wayne Zhang. Gendet: Meta learning to generate detectors from few shots. In *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [41] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5265–5274, 2018.
- [42] Xin Wang, Thomas Huang, Joseph Gonzalez, Trevor Darrell, and Fisher Yu. Frustratingly simple few-shot object detection. In *International Conference on Machine Learning*, pages 9919–9928. PMLR, 2020.

- [43] Jiayi Wu, Songtao Liu, Di Huang, and Yunhong Wang. Multi-scale positive sample refinement for few-shot object detection. In *European Conference on Computer Vision*, pages 456–472. Springer, 2020.
- [44] Yang Xiao and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. In *European Conference on Computer Vision (ECCV)*, 2020.
- [45] Xiaopeng Yan, Ziliang Chen, Anni Xu, Xiaoxi Wang, Xiaodan Liang, and Liang Lin. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9577–9586, 2019.
- [46] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [47] Chi Zhang, Yujun Cai, Guosheng Lin, and Chunhua Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12203–12213, 2020.

Junying Huang received the B.Eng degree from the School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China, in 2020. He is currently a postgraduate student in the Human Cyber Physical Intelligence Integration Lab, School of Computer Science and Engineering, Sun Yat-sen University. His research interests include computer vision and machine learning, with a focus on few-shot learning.

Fan Chen received Bachelor Degree from the School of Computer Science and Engineering, Sun Yat-sen University in 2020. He is currently a postgraduate student in the Human Cyber Physical Intelligence Integration Lab, School of Computer Science and Engineering, Sun Yat-sen University. His research interests include computer vision and machine learning, mainly focus on object detection.

Liang Lin is a full professor of Computer Science in Sun Yat-sen University and CEO of DarkerMatter AI. He worked as the Executive Director of the SenseTime Group from 2016 to 2018, leading the R&D teams in developing cutting-edge, deliverable solutions in computer vision, data analysis and mining, and intelligent robotic systems. He has authored or co-authored more than 200 papers in leading academic journals and conferences. He is an associate editor of IEEE Trans. Human-Machine Systems and IET Computer Vision, and he served as the area/session chair for numerous conferences such as CVPR, ICME, ICCV. He was the recipient of Annual Best Paper Award by Pattern Recognition (Elsevier) in 2018, Dimond Award for best paper in IEEE ICME in 2017, ACM NPAR Best Paper Runners-Up Award in 2010, Google Faculty Award in 2012, award for the best student paper in IEEE ICME in 2014, and Hong Kong Scholars Award in 2014. He is a Fellow of IET.

Dongyu Zhang received the Ph. D. from Harbin Institute of Technology in 2010. He is an associate professor of School of Computer and Engineering of Sun Yat-sen University. His research interests include computer vision and machine learning.