

# Breaking Immutable: Information-Coupled Prototype Elaboration for Few-Shot Object Detection

Xiaonan Lu<sup>1,2,3,4</sup>, Wenhui Diao<sup>1,2,3,4\*</sup>, Yongqiang Mao<sup>1,2,3,4</sup>, Junxi Li<sup>1,2,3,4</sup>, Peijin Wang<sup>1,2</sup>,  
Xian Sun<sup>1,2,3,4</sup>, Kun Fu<sup>1,2,3,4</sup>

<sup>1</sup>Aerospace Information Research Institute, Chinese Academy of Sciences

<sup>2</sup>Key Laboratory of Network Information System Technology (NIST), Aerospace Information Research Institute

<sup>3</sup>University of Chinese Academy of Sciences

<sup>4</sup>School of Electronic, Electrical and Communication Engineering, University of Chinese Academy of Sciences  
{luxiaonan19, maoyongqiang19, lijunxi191, wangpeijin17}@mails.ucas.ac.cn  
{diaowh, sunxian}@aircas.ac.cn, kunfuiecas@gmail.com

## Abstract

Few-shot object detection, expecting detectors to detect novel classes with a few instances, has made conspicuous progress. However, the prototypes extracted by existing meta-learning based methods still suffer from insufficient representative information and lack awareness of query images, which cannot be adaptively tailored to different query images. Firstly, only the support images are involved for extracting prototypes, resulting in scarce perceptual information of query images. Secondly, all pixels of all support images are treated equally when aggregating features into prototype vectors, thus the salient objects are overwhelmed by the cluttered background. In this paper, we propose an Information-Coupled Prototype Elaboration (ICPE) method to generate specific and representative prototypes for each query image. Concretely, a conditional information coupling module is introduced to couple information from the query branch to the support branch, strengthening the query-perceptual information in support features. Besides, we design a prototype dynamic aggregation module that dynamically adjusts intra-image and inter-image aggregation weights to highlight the salient information useful for detecting query images. Experimental results on both Pascal VOC and MS COCO demonstrate that our method achieves state-of-the-art performance in almost all settings. Code will be available at: <https://github.com/lxn96/ICPE>.

## Introduction

Object detection has witnessed significant progress by deep learning (Ren et al. 2015; Redmon et al. 2016; Lin et al. 2017b; Carion et al. 2020). However, most detectors require abundant labeled data for training and impracticably generalize to new tasks with a few images. In contrast, humans can observe novel objects with limited instances. Thus, few-shot object detection (FSOD) is proposed to bridge the gap between deep learning and human learning systems.

FSOD assists in improving the performance on novel classes with a few labeled images by learning knowledge on base classes with abundant data. Some methods (Chen et al.

\*Corresponding author.

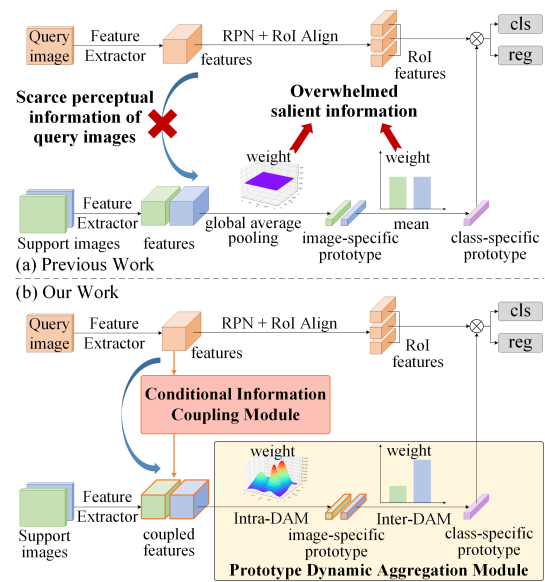


Figure 1: Prototype extraction in meta-learning methods. (a) Prototypes extracted in previous works lack perceptual information of query images. And averagely aggregating features makes the salient information overwhelmed. (b) In our work, a conditional information coupling module strengthens query-perceptual information. And we design a prototype dynamic aggregation module consisting of intra- and inter-image dynamic aggregation mechanisms (Intra-DAM, Inter-DAM) to consolidate salient information of prototypes.

2018; Wang et al. 2020) transfer the knowledge from base dataset to novel dataset. Still, they are sensitive to the difference between datasets and may produce negative transfer when the difference is obvious. Other methods (Hsieh et al. 2019; Yan et al. 2019; Li et al. 2021) achieve remarkable performance based on meta-learning. They simulate few-shot scenes and construct episodes consisting of support sets and query sets. The annotated support sets are used to generate prototypes, then the objects in query sets are detected based

on the generated prototypes. The quality of prototypes directly affects the performance on query images. However, the prototypes extracted in most meta-learning based methods are immutable, suffering from insufficient query-aware and representative information. They are unable to be adaptively tailored to different query images. As Figure 1(a) shows, it mainly manifests in the following two aspects:

- 1) *Scarce perceptual information of query images.* Existing methods extract prototypes only from support images, lacking the perception of query images. The prototypes extracted from the same support images are the same, which may be not suitable for predicting all query images. It is inevitable that the detection performance of different query images varies greatly, especially when the extreme perspective span occurs.
- 2) *Overwhelmed salient information of support images.* Majority works averagely aggregate support features into prototypes. All pixels of all images contribute equally to prototypes. On the one hand, using averaging operation within an image (intra-image) causes that the cluttered background weakens the interested objects. On the other hand, the similarities between support images and query images are different. Averagely aggregating multiple images (inter-image) results in the determinative images being overwhelmed by other non-significant images.

To alleviate the above problems, we propose a novel Information-Coupled Prototype Elaboration (ICPE) network for FSOD. As Figure 1(b) shows, it generates representative and specific prototypes for each query image. Specifically, when encountering novel objects, humans intuitively apply the immediate information to search for the relevant knowledge they have seen for identifying the objects. Thus, a conditional information coupling module is proposed. It couples the information of query images into support images to generate coupled features with query-aware information, which are used to obtain tailored prototypes for each query image. In addition, the averaging operation might blur discriminative details that are crucial for object detection (Gao, Wang, and Wu 2019). To address the issue, we propose a prototype dynamic aggregation module. It dynamically adjusts intra- and inter-image aggregation weights to highlight salient and useful information for predicting query images. On the one hand, an intra-image dynamic aggregation mechanism is designed to highlight the important local information within each support image by building local-to-global dependencies. On the other hand, based on the law of large numbers, an inter-image dynamic aggregation mechanism is proposed to learn the implicit similarities between support images and query images, with emphasis on highlighting the key support images. Through the above two modules, our model generates representative and tailored prototypes for each query image and effectively improves the performance.

The major contributions of this work are as follows: (1) A conditional information coupling module is proposed to couple query features into support features. It enhances support features with query-aware information. (2) We design a prototype dynamic aggregation module to emphasize the salient information both within each support image and

cross multiple images when aggregating the enhanced support features into prototypes. (3) Extensive experiments illustrate that our method outperforms state-of-the-art methods in most experimental settings.

## Related Work

**Object Detection.** Modern detectors are divided into one-stage and two-stage methods. The former (Redmon et al. 2016; Liu et al. 2016; Lin et al. 2017b) directly predicts objects with a simple convolutional neural network, while the latter (Girshick et al. 2014; Girshick 2015; Ren et al. 2015; He et al. 2017) extracts proposals on different receptive fields (Lin et al. 2017a; Mao et al. 2022a) and then makes refined predictions. Two-stage methods are generally more accurate than one-stage methods with slightly longer time. These detectors all rely on sufficient labeled data and suffer from over-fitting when facing limited data.

**Few-Shot Learning (FSL).** FSL is divided into augmentation, metric learning, and meta-learning methods. Augmentation methods (Hariharan and Girshick 2017; Wang et al. 2018b) perform data or feature augmentation to meet the data requirements for training. Metric learning methods (Koch et al. 2015; Scott, Ridgeway, and Mozer 2018; Mao et al. 2022b) calculate the distance between test and known samples to find adjacent classes. Meta-learning methods (Andrychowicz et al. 2016; Finn, Abbeel, and Levine 2017; Jamal and Qi 2019) construct few-shot tasks and learn meta-knowledge to adapt for new tasks quickly. These methods only focus on classification without locating objects.

**Few-Shot Object Detection (FSOD).** FSOD consists of data augmentation based, transfer learning based, and meta-learning based methods. The methods based on data augmentation (Yang et al. 2020; Wu et al. 2020, 2021; Xu et al. 2021; Zhang and Wang 2021) alleviate data scarcity by expanding data or extracting various features from limited data. The transfer learning based methods (Chen et al. 2018; Karlinsky et al. 2019; Fan et al. 2020, 2021) utilize fine-tuning or metric learning to transfer the knowledge learned on a data-rich base dataset to a few-shot novel dataset. TFA (Wang et al. 2020) only fine-tunes the last layers of prediction heads on novel classes while freezing all other layers. The meta-learning based methods (Xiao and Marlet 2020; Hu et al. 2021; Li et al. 2021; Li and Li 2021; Han et al. 2022a) intuitively and explicitly simulate few-shot scenarios including query sets and support sets, and apply a query branch and a support branch to process the two sets, respectively. Based on YOLOv2 (Redmon and Farhadi 2017), Meta YOLO (Kang et al. 2019) extracts reweighting vectors from support images and applies them as channel-wise attention on query images to detect objects. Meta R-CNN (Yan et al. 2019) extracts prototypes and uses them to the region of interest (RoI) features of query images for detecting objects, which is selected as our baseline.

Although existing meta-learning methods have achieved prominent performance, the prototypes generated by them are immutable and lack query-aware information. We propose a method to generate unique prototypes for each query image and improve the quality of prototypes.

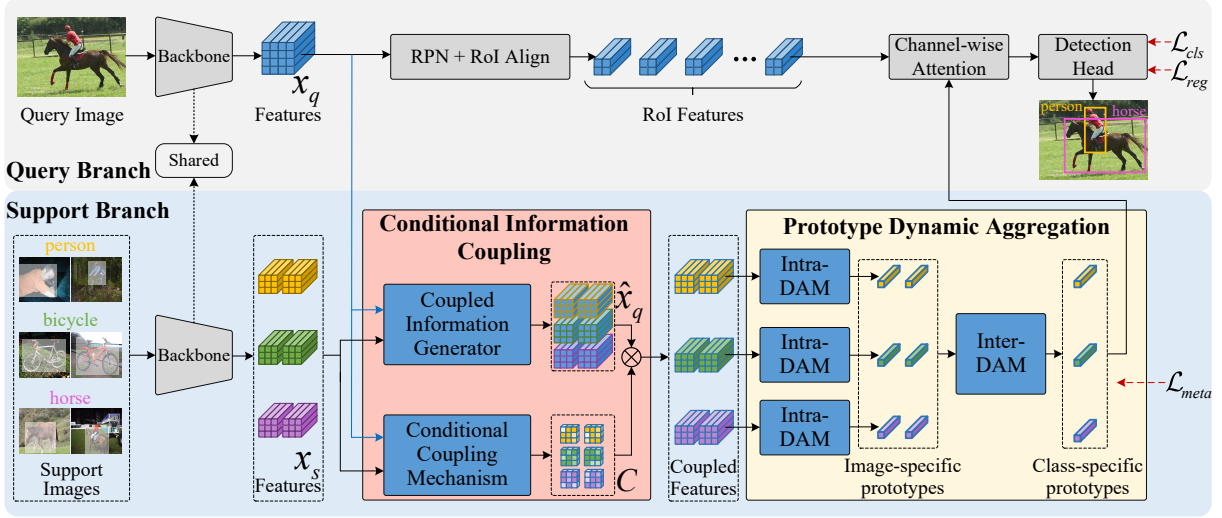


Figure 2: Architecture of ICPE. The conditional information coupling module integrates support features and query features to generate coupled features containing query-aware information. Through the prototype dynamic aggregation module, the coupled features are aggregated into representative prototypes which facilitate the query branch to detect objects accurately.

## Preliminaries

**Problem Definition.** Given a base class set  $C_b$  with a data-rich dataset  $\mathcal{D}_b$  and a novel class set  $C_n$  ( $C_n \cap C_b = \emptyset$ ) with a dataset  $\mathcal{D}_n$  which comprises  $k$  instances of each category, FSOD expects that detectors can detect objects of both  $C_b$  and  $C_n$  simultaneously. As with most progresses of FSOD (Li et al. 2021; Wu et al. 2020; Xiao and Marlet 2020), we follow the episodic meta-learning paradigm (Kang et al. 2019; Yan et al. 2019). It constructs abundant few-shot tasks with support sets and query sets in  $\mathcal{D}_b$  and  $\mathcal{D}_n$ , respectively, namely episodes. And a two-step training strategy is applied. Firstly, the detector is trained on the episodes of  $\mathcal{D}_b$ , called meta-training step. Then, in meta-finetuning step, the detector is fine-tuned on the episodes of  $\mathcal{D}_n$  and a subset of  $\mathcal{D}_b$ .

**Baseline Method.** Meta R-CNN (Yan et al. 2019), as a typical meta-learning based FSOD method, is selected as our baseline. Based on Faster R-CNN (Ren et al. 2015), it contains a support branch and a query branch. In the support branch, images and object masks of support sets are sent to the backbone to obtain features, and global average pooling is applied to get feature vectors. Prototypes are obtained by averaging vectors belonging to the same class. The query branch receives query images and extracts RoI features. Then, given the generated prototypes, the channel-wise attention is applied for all RoIs to strengthen the category-related representations, promoting the detection head to detect objects. The prototypes extracted in Meta R-CNN are immutable and unrepresentative, while our model generates representative and tailored prototypes for query images.

## Method

In this section, we first briefly introduce the architecture and the training strategy of our method. Then, we separately elaborate on the two proposed modules.

## Overview

The purpose of our model is definite: Coupling query-perceptual information into support features and emphasizing salient details to generate adaptive high-quality prototypes. As Figure 2 shows, based on Meta R-CNN (Yan et al. 2019), our model has a query branch and a support branch that share a common backbone, and introduces a conditional information coupling module and a prototype dynamic aggregation module for refining prototypes to be representative and tailored for each query image.

Given a query image and  $k$  support instances of each category ( $k$ -shot), the features extracted from the backbone of the query branch and the support branch are sent to the conditional information coupling module. It effectively interacts the information of the two features to polish up support features with query-perceptual information. Then, the prototype dynamic aggregation module is introduced to highlight the salient information when aggregating the coupled features into class-specific prototypes. By introducing perceptual information of the query image and highlighting salient information, our model generates high-quality and representative prototypes tailored to the query image. Finally, the refined prototypes are used as the channel-wise attention for RoI features in the query branch, stimulating the detection head to accurately predict objects in the query image.

Like Meta R-CNN (Yan et al. 2019), in addition to the typical detection loss consisting of classification loss  $\mathcal{L}_{cls}$  implemented by Cross-Entropy (CE) loss and regression loss  $\mathcal{L}_{reg}$  implemented by L1 loss, a meta loss  $\mathcal{L}_{meta}$  implemented by CE loss is also applied on the class-specific prototypes to make prototypes distinguishable and avoid prediction ambiguity. Thus, the overall loss of our ICPE is:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{reg} + \lambda \mathcal{L}_{meta} \quad (1)$$

where  $\lambda$  is a loss weight for the balance between detection loss and meta loss, which is set to 1 in our experiments.

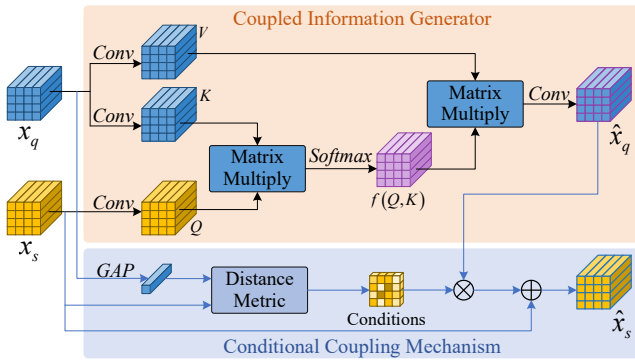


Figure 3: Architecture of the conditional information coupling module.

### Conditional Information Coupling

To make prototypes have query-perceptual information and be consistent with the objects in query images, QA-FewDet (Han et al. 2021) applies heterogeneous graph networks to obtain query-adaptive prototypes. CoAE (Hsieh et al. 2019) enriches query-aware information in support features through non-local (Wang et al. 2018a). In order to avoid redundant and useless information flow when directly using non-local interaction, it is necessary to calculate the coupling conditions to guide the way of information coupling. Thus, as Figure 3 shows, we design a conditional information coupling module. It consists of a coupled information generator and a conditional coupling mechanism.

**Coupled information generator** polishes query features to be appropriate for coupling to support features. The query features are reorganized to obtain customized perceptual information for support features. For comprehensive information interaction, the coupled information generator is designed based on non-local, which measures similarities between the two features and realizes the recombination of query features. It receives features extracted from the query branch and the support branch, denoted as  $x_q$  and  $x_s$ , respectively.  $x_q$  is represented as key-value pairs, and  $x_s$  is represented as query. The output is the restructured query features implemented as a weighted sum of values  $V$ , where the weights are represented by the affinity relationship between  $x_q$  and  $x_s$ , which is calculated by a compatibility function (Vaswani et al. 2017) of query  $Q$  and key  $K$ :

$$f(Q, K) = \text{Softmax}(Q^T \cdot K) \quad (2)$$

Then, the values  $V$  of query features are assembled by the weights  $f(Q, K)$  and passed through a convolutional layer to produce the customized query features  $\hat{x}_q$  for coupling into support features:

$$\hat{x}_q = \text{Conv}(f(Q, K) \cdot V) \quad (3)$$

**Conditional coupling mechanism** calculates the conditions for information coupling based on  $x_q$  and  $x_s$ . It predicts the regions of support features that are highly consistent with query features as coupling conditions, which guide the restructured query features  $\hat{x}_q$  to flow into support features  $x_s$  efficiently. Taking the global representation of  $x_q$ ,

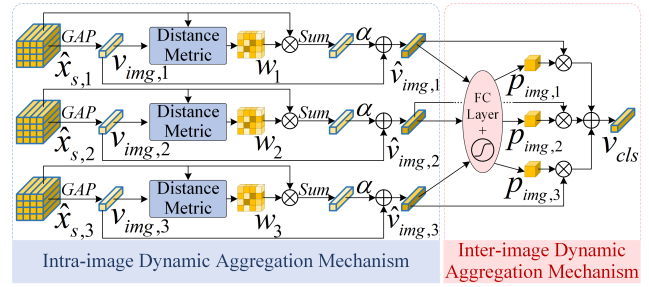


Figure 4: Architecture of the prototype dynamic aggregation module.

the conditional coupling mechanism adequately delves the areas of  $x_s$  with high similarities of  $x_q$ . The conditions  $C$  are calculated as follows:

$$C = \mathcal{M}(GAP(x_q), x_s) \quad (4)$$

where  $GAP$  denotes the global average pooling, and  $\mathcal{M}$  represents the distance measurement implemented with cosine similarity. Then,  $C$  is used as a mask to activate the areas of  $\hat{x}_q$  for information coupling. And the information coupling from query features to support features is realized by a residual connection (He et al. 2016). Thus, the coupled features  $\hat{x}_s$  with query-perceptual information are expressed as follows:

$$\hat{x}_s = C \circ \hat{x}_q + x_s \quad (5)$$

where  $\circ$  represents Hadamard product. Compared with the original support features,  $\hat{x}_s$  contains the perceptual information of query features through the deep and efficient interaction between the two features.

Different from DCNet (Hu et al. 2021) which replaces channel-wise attention in the query branch with cross-attention to densely match support features (prototypes) into query features, our coupling module performs cross-attention to improve the quality of prototypes. It enhances query-aware information in support features which are used to obtain unique and tailored prototypes for each query image subsequently. Furthermore, our module extends cross-attention with a conditional coupling mechanism for efficient information coupling and reduced noise introduction.

### Prototype Dynamic Aggregation

For dynamic aggregation of support features, CrossTransformers (Doersch, Gupta, and Zisserman 2020) aggregates information in a spatially-aware way. DANa (Chen et al. 2021) views features as a wave along the channel dimension and maintains the wave patterns of foreground features. To generate high-quality and representative class-specific prototypes from  $\hat{x}_s$ , a prototype dynamic aggregation module is proposed. As Figure 4 shows, it incorporates intra-image and inter-image dynamic aggregation mechanisms. They replace average aggregations both within a single image and across multiple images, preserving the discriminative and salient details when aggregating features into prototype vectors.

**Intra-image dynamic aggregation mechanism (Intra-DAM)** is designed to preserve prominent features within

each coupled feature by building local-to-global dependencies when generating image-specific prototypes. Unlike generating prototypes with GAP, Intra-DAM reinforces the contributions of salient regions. It conducts the global representation of  $\hat{x}_s$ , and then calculates the global dependency of each local pixel to obtain the dynamic aggregation weights within each coupled feature. The weight  $w$  signifies the amount of the global information representing the whole features contained in each pixel, which is defined as:

$$w = \mathcal{N}(GAP(\hat{x}_s), \hat{x}_s) \quad (6)$$

where  $\mathcal{N}$  means a function that conducts the dependency between two features, which is implemented with the cosine similarity. After yielding the aggregation weight of each position, the pixels of  $\hat{x}_s$  are aggregated to generate a refined and representative image-specific prototype  $\hat{v}_{img}$ :

$$\hat{v}_{img} = v_{img} + \frac{\alpha}{N} \sum w \circ \hat{x}_s \quad (7)$$

where  $v_{img} = GAP(\hat{x}_s)$  is a popular way to generate prototypes.  $\sum$  denotes the sum operation of vectors.  $N$  means the number of vectors for normalization, and  $\alpha$  is a balance factor set to 1. The generated  $\hat{v}_{img}$  emphasizes salient pixels within  $\hat{x}_s$  on the basis of treating all pixels equally.

**Inter-image dynamic aggregation mechanism (Inter-DAM)** achieves dynamic aggregation between multiple image-specific prototypes of the same class to generate class-specific prototypes. For a query image, the discriminative information carried in different support images is diverse. Inter-DAM consolidates the dominance of the coupled features which contain relatively rich salient and useful information for detection. Given the image-specific prototypes of each class, it models the implicit similarities between multiple support images and the query image through a fully connected (FC) layer to generate contribution probabilities  $p_{img}$ , that is, the proportion of the useful information carried in each support image, which is defined as:

$$p_{img} = Sigmoid(FC(\hat{v}_{img})) \quad (8)$$

Based on the contribution probabilities  $p_{img}$  of support images, the weighted summation is applied to  $\hat{v}_{img}$  belonging to the same category to get the expectation of them, which is denoted as the class-specific prototype:

$$v_{cls} = E(\hat{v}_{img}) = \sum_{i=1}^k p_{img,i} \hat{v}_{img,i} \quad (9)$$

where  $k$  means the number of support instances for each class, i.e.,  $k$ -shot. As noted by the Law of Large Numbers, the barycenter of samples converges to the expectation as the number of samples becomes infinity. Hence, with few-shot images, the class-specific prototypes represented by the expectations of image-specific prototypes can simulate the class information in the presence of sufficient data. The generated class-specific prototypes  $v_{cls}$  adequately carry the representative and salient information of categories.

## Experiments

### Experimental Setup

**Dataset.** Same as Meta YOLO (Kang et al. 2019) and Meta R-CNN (Yan et al. 2019), our method is evaluated on

Pascal VOC (Everingham et al. 2010, 2015) and MS COCO (Lin et al. 2014) datasets. For Pascal VOC, our model is trained on the trainval sets of VOC 2007 and 2012, and tested on VOC 2007 test set. The dataset is partitioned into three different splits, where 5 categories are selected as novel classes and the other 15 categories are used as base classes. There are  $k$  annotated instances of each category for meta-finetuning.  $k$  is set to 1, 2, 3, 5, and 10. For MS COCO, the model is trained on a modified dataset consisting of 80k training images and 35k validation images. The remaining 5k images are used for test. The 20 classes included in Pascal VOC are used as novel classes, and the other 60 categories are as base classes.  $k$  is set to 10 and 30.

**Implementation Details.** Same as previous work (Wang et al. 2020; Qiao et al. 2021), for query images, we use multiple scale images for training and a single scale for testing. The shorter sides are 480 ~ 800 pixels while the longer sides are less than 1333 pixels. The support images are resized to 224 × 224 pixels. Only random horizontal flipping and normalization are used for training. The backbone is ResNet-101 (He et al. 2016) with the pretrained weight on ImageNet (Russakovsky et al. 2015). We set batch size as 4 and use stochastic gradient descent (SGD) optimizer with the momentum of 0.9 and the weight decay of 0.0001. The training iterations and learning rates are the same as Meta R-CNN (Yan et al. 2019). Consistent with TFA (Wang et al. 2020), the present results of our method are averaged over multiple random runs. We implement our experiments in Py-Torch (Paszke et al. 2019) on 2 V100 GPUs.

### Comparisons with State-of-the-art Methods

We conduct experiments on Pascal VOC and MS COCO, and compare our method with state-of-the-art (SOTA) methods. The results demonstrate the effectiveness of our ICPE.

**Pascal VOC.** In Table 1, we report the results of our method and existing SOTA methods on Pascal VOC. Our ICPE achieves state-of-the-art performance on almost all three splits with different shot settings. For the average mAP of three splits, our method achieves 46.2%, 50.9%, 55.5%, 59.3%, and 59.6% in 1, 2, 3, 5, and 10-shot. As the number of annotated instances increases, the performance of the model gradually improves. Especially in extremely low-shot settings, the performance is improved significantly as the number of shots increases. For example, when only increasing one instance from 1-shot to 2-shot and from 2-shot to 3-shot, the average performance of all splits is improved by 4.7% and 4.6%, respectively. This is because when the number of samples is tiny, adding one sample can provide relatively rich information for prototypes.

**MS COCO.** We also evaluate our method on a challenging dataset MS COCO. The results with the standard COCO metrics are reported in Table 2. Compared with Meta R-CNN (2019), our method improves  $AP$  by more than 10% in both 10 and 30-shot settings, reaching 19.3% and 23.1%, respectively. Besides, our method outperforms SOTA methods in most evaluation metrics. It shows that our ICPE is still effective on the challenging and complex datasets.

Method / Shot	Novel Set 1					Novel Set 2					Novel Set 3				
	1	2	3	5	10	1	2	3	5	10	1	2	3	5	10
LSTD (Chen et al. 2018)	8.2	1.0	12.4	29.1	38.5	11.4	3.8	5.0	15.7	31.0	12.6	8.5	15.0	27.3	36.3
Meta YOLO (Kang et al. 2019)	14.8	15.5	26.7	33.9	47.2	15.7	15.3	22.7	30.1	40.5	21.3	25.6	28.4	42.8	45.9
Meta R-CNN (Yan et al. 2019)	19.9	25.5	35.0	45.7	51.5	10.4	19.4	29.6	34.8	45.4	14.3	18.2	27.5	41.2	48.1
Viewpoint (Xiao and Marlet 2020)	24.2	35.3	42.2	49.1	57.4	21.6	24.6	31.9	37.0	45.7	21.2	30.0	37.2	43.8	49.6
TFA w/fc (Wang et al. 2020)	36.8	29.1	43.6	55.7	57.0	18.2	29.0	33.4	35.5	39.0	27.7	33.6	42.5	48.7	50.2
TFA w/cos (Wang et al. 2020)	39.8	36.1	44.7	55.7	56.0	23.5	26.9	34.1	35.1	39.1	30.8	34.8	42.8	49.5	49.8
MPSR (Wu et al. 2020)	41.7	42.5	51.4	52.2	61.8	24.4	29.3	39.2	39.9	47.8	35.6	41.8	42.3	48.0	49.7
SRR-FSD (Zhu et al. 2021)	47.8	50.5	51.3	55.2	56.8	32.5	35.3	39.1	40.8	43.8	40.1	41.5	44.3	46.9	46.4
FSCE (Sun et al. 2021)	44.2	43.8	51.4	61.9	63.4	27.3	29.5	43.5	44.5	50.2	37.2	41.9	47.5	54.6	58.5
DCNet (Hu et al. 2021)	33.9	37.4	43.7	51.1	59.6	23.2	24.8	30.6	36.7	46.6	32.3	34.9	39.7	42.6	50.7
CME (Li et al. 2021)	41.5	47.5	50.4	58.2	60.9	27.2	30.2	41.4	42.5	46.8	34.3	39.6	45.1	48.3	51.5
DeFRCN (Qiao et al. 2021)	<u>53.6</u>	<u>57.5</u>	<u>61.5</u>	64.1	60.8	30.1	<u>38.1</u>	<u>47.0</u>	<b>53.3</b>	47.9	<u>48.4</u>	<u>50.9</u>	52.3	54.9	57.4
Meta FRCNN (Han et al. 2022a)	43.0	54.5	60.6	<b>66.1</b>	65.4	27.7	35.5	46.1	47.8	<u>51.4</u>	40.6	46.4	<u>53.4</u>	<u>59.9</u>	<u>58.6</u>
KFSOD (Zhang et al. 2022)	44.6	-	54.4	60.9	<u>65.8</u>	<b>37.8</b>	-	43.1	48.1	50.4	34.8	-	44.1	52.7	53.9
FCT (Han et al. 2022b)	38.5	49.6	53.5	59.8	64.3	25.9	34.2	40.1	44.9	47.4	34.7	43.9	49.3	53.1	56.3
<b>ICPE(ours)</b>	<b>54.3</b>	<b>59.5</b>	<b>62.4</b>	<u>65.7</u>	<b>66.2</b>	<u>33.5</u>	<b>40.1</b>	<b>48.7</b>	<u>51.7</u>	<b>52.5</b>	<b>50.9</b>	<b>53.1</b>	<b>55.3</b>	<b>60.6</b>	<b>60.1</b>

Table 1: FSOD performance in terms of mAP on Pascal VOC dataset. Best results are in **bold** and second best are underlined.

Method / Shot	10			30		
	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>
LSTD (2018)	3.2	8.1	2.1	6.7	15.8	5.1
Meta YOLO (2019)	5.6	12.3	4.6	9.1	19.0	7.6
Meta R-CNN (2019)	8.7	19.1	6.6	12.4	25.3	10.8
Viewpoint (2020)	12.5	<u>27.3</u>	9.8	14.7	30.6	12.2
TFA w/fc (2020)	10.0	19.2	9.2	13.4	24.7	13.2
TFA w/cos (2020)	10.0	19.1	9.3	13.7	24.9	13.4
MPSR (2020)	9.8	17.9	9.7	14.1	25.4	14.2
SRR-FSD (2021)	11.3	23.0	9.8	14.7	29.2	13.5
FSCE (2021)	11.1	-	9.8	15.3	-	14.2
DCNet (2021)	12.8	23.4	11.2	18.6	<u>32.6</u>	17.5
CME (2021)	15.1	24.6	16.4	16.9	28.0	<u>17.8</u>
DeFRCN (2021)	<u>18.5</u>	-	-	<u>22.6</u>	-	-
Meta FRCNN (2022a)	12.7	25.7	10.8	16.6	31.8	15.8
KFSOD (2022)	<u>18.5</u>	26.3	<b>18.7</b>	-	-	-
FCT (2022b)	15.3	-	-	20.2	-	-
<b>ICPE(ours)</b>	<b>19.3</b>	<b>27.9</b>	<u>18.0</u>	<b>23.1</b>	<b>32.9</b>	<b>19.2</b>

Table 2: FSOD performance on MS COCO dataset.

## Ablation Study

To analyze the effectiveness of the two components proposed in ICPE, we conduct comprehensive ablation studies on the Novel Set 1 of Pascal VOC. And the baseline is reproduced to verify the superiority of our method convincingly.

**Effects of Different Modules.** Table 3 shows the effectiveness of the proposed modules in ICPE. The conditional information coupling module improves the average perfor-

#	Method		Shot					Avg.	$\Delta$ Avg.
	CIC	PDA	1	2	3	5	10		
1			44.7	49.3	55.5	58.3	61.2	53.8	
2	✓		50.7	57.2	59.6	63.2	64.1	59.0	+5.2
3		✓	49.3	55.9	57.3	61.4	63.8	57.5	+3.7
4	✓	✓	<b>54.3</b>	<b>59.5</b>	<b>62.4</b>	<b>65.7</b>	<b>66.2</b>	<b>61.6</b>	<b>+7.8</b>

Table 3: Ablation study of different modules. CIC means the conditional information coupling module. PDA means the prototype dynamic aggregation module. Avg. represents the average performance of all shot settings.

mance of all shot settings by 5.2%. It augments query-aware information in support features. Besides, the average performance gain is 3.7% when simply introducing the prototype dynamic aggregation module to highlight useful and salient information. With both conditional information coupling and prototype dynamic aggregation, the performance gain increases to 7.8%. Thus, our ICPE efficiently strengthens the query-aware and salient information when generating high-quality prototypes tailored to query images. It achieves significant improvement over the baseline.

**Effects of Conditional Information Coupling.** Figure 5(a) verifies the advancement of the conditional coupling mechanism. When only using the coupled information generator (CIC w/o CCM), the mAPs of five settings are 47.7%, 54.2%, 57.7%, 61.4%, and 62.5%, respectively. The average performance is 2.9% higher than the baseline. Further adding the conditional coupling mechanism, as #2 in Table 3 shows, the mAP is improved by 5.2%. Besides, Figure 5(b) reflects the effects of the number of channels in  $f(Q, K)$ .




Figure 5: Ablation study of Conditional Information Coupling (CIC). (a) Effects of conditional coupling mechanism (CCM). (b) Effects of the number of channels in  $f(Q, K)$ .

Method	Shot					Avg.	$\Delta Avg.$
	1	2	3	5	10		
baseline	44.7	49.3	55.5	58.3	61.2	53.8	
+GMP	45.9	52.2	55.8	58.7	61.9	54.9	+1.1
+LIP	47.5	52.8	56.0	59.2	61.9	55.5	+1.7
+Intra-DAM	49.3	54.4	56.4	59.9	62.6	56.5	+2.7
+Intra-DAM & Inter-DAM	<b>49.3</b>	<b>55.9</b>	<b>57.3</b>	<b>61.4</b>	<b>63.8</b>	<b>57.5</b>	<b>+3.7</b>

Table 4: Ablation study of prototype dynamic aggregation. +GMP means that image-specific prototypes are implemented by the sum of GAP and global max pooling (GMP). +LIP means generating image-specific prototypes via Local Importance-based Pooling (Gao, Wang, and Wu 2019).

Method	Parameters(MB)	FPS(img/s)	FLOPs(GB)
baseline	45.9	13.4	1752.7
ICPE(ours)	50.1	9.6	2401.1

Table 5: Ablation study of computational cost.

It reveals that  $f(Q, K)$  with 1024 channels reports the best results. 256 and 512 channels are insufficient because of the depression of feature representation, while 2048 channels face performance degradation due to over-fitting.

**Effects of Prototype Dynamic Aggregation.** To verify the effectiveness of Intra-DAM, we conduct experiments with two other approaches to generate image-specific prototypes: one is to use the sum of GAP and global max pooling (GMP), and the other is to use LIP (Gao, Wang, and Wu 2019). As Table 4 shows, the performance with GMP retaining the most discriminative features is increased by 1.1%. And LIP, strengthening the local importance, improves the performance by 1.7%. Our Intra-DAM improves the performance by 2.7%. Further adding Inter-DAM, the average performance is improved by 3.7%. It is concluded that the prototype dynamic aggregation module effectively highlights salient and useful information when generating prototypes.

**Computational Cost.** We analyze the computational cost between our ICPE and Meta R-CNN on 10-shot settings. As

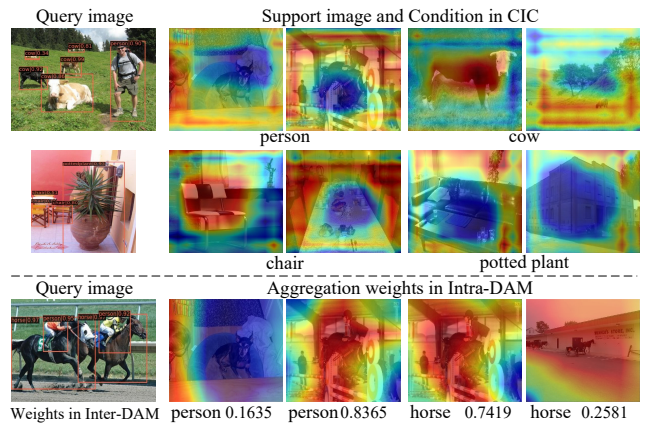


Figure 6: Visualization of components. Top: Conditions in conditional information coupling module. Bottom: Aggregation weights of prototype dynamic aggregation module.




Figure 7: Detection results of the baseline and our method.

Table 5 shows, compared with the baseline, our model only increases a slight number of parameters and computations.

## Qualitative Results

In Figure 6, the conditions predicted by the conditional coupling mechanism are almost the same as the areas where the objects are located in. It guides the perceptual information of query images to be coupled accurately to the effective regions of support images and reduces redundant information flow. Besides, for prototype dynamic aggregation module, Intra-DAM pays attention to the areas with comprehensive information in each support image, and Inter-DAM assigns significant weights to the images with high similarities for query images. We also present the detection results of our method and the baseline in Figure 7. Our ICPE achieves superior performance and effectively reduces false detection.

## Conclusion

In this work, we propose an Information-Coupled Prototype Elaboration network to generate high-quality prototypes tailored to each query image. A conditional information coupling module is proposed to enhance query-perceptual information in support features. Besides, we design a prototype dynamic aggregation mechanism to highlight the salient and useful information both within a single image and across multiple images when generating prototypes. Extensive experiments validate the effectiveness of our method which achieves state-of-the-art performance in most settings.

## Acknowledgments

This work was supported by the National Key R&D Program of China (Grant No. 2021YFB390050\*).

## References

- Andrychowicz, M.; Denil, M.; Gomez, S.; Hoffman, M. W.; Pfau, D.; Schaul, T.; Shillingford, B.; and De Freitas, N. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in neural information processing systems*, 3981–3989.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229. Springer.
- Chen, H.; Wang, Y.; Wang, G.; and Qiao, Y. 2018. Lstd: A low-shot transfer detector for object detection. In *Proceedings of the AAAI conference on artificial intelligence*, 2836–2843.
- Chen, T.-I.; Liu, Y.-C.; Su, H.-T.; Chang, Y.-C.; Lin, Y.-H.; Yeh, J.-F.; Chen, W.-C.; and Hsu, W. 2021. Dual-awareness attention for few-shot object detection. *IEEE Transactions on Multimedia*.
- Doersch, C.; Gupta, A.; and Zisserman, A. 2020. Crosstransformers: spatially-aware few-shot transfer. *Advances in Neural Information Processing Systems*, 33: 21981–21993.
- Everingham, M.; Eslami, S.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1): 98–136.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2): 303–338.
- Fan, Q.; Zhuo, W.; Tang, C.-K.; and Tai, Y.-W. 2020. Few-shot object detection with attention-RPN and multi-relation detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4013–4022.
- Fan, Z.; Ma, Y.; Li, Z.; and Sun, J. 2021. Generalized Few-Shot Object Detection without Forgetting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4527–4536.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 1126–1135. PMLR.
- Gao, Z.; Wang, L.; and Wu, G. 2019. Lip: Local importance-based pooling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3355–3364.
- Girshick, R. 2015. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 1440–1448.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 580–587.
- Han, G.; He, Y.; Huang, S.; Ma, J.; and Chang, S.-F. 2021. Query adaptive few-shot object detection with heterogeneous graph convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3263–3272.
- Han, G.; Huang, S.; Ma, J.; He, Y.; and Chang, S.-F. 2022a. Meta faster r-cnn: Towards accurate few-shot object detection with attentive feature alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 780–789.
- Han, G.; Ma, J.; Huang, S.; Chen, L.; and Chang, S.-F. 2022b. Few-shot object detection with fully cross-transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5321–5330.
- Hariharan, B.; and Girshick, R. 2017. Low-shot visual recognition by shrinking and hallucinating features. In *Proceedings of the IEEE International Conference on Computer Vision*, 3018–3027.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hsieh, T.-I.; Lo, Y.-C.; Chen, H.-T.; and Liu, T.-L. 2019. One-shot object detection with co-attention and co-excitation. *Advances in neural information processing systems*, 32.
- Hu, H.; Bai, S.; Li, A.; Cui, J.; and Wang, L. 2021. Dense Relation Distillation with Context-aware Aggregation for Few-Shot Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10185–10194.
- Jamal, M. A.; and Qi, G.-J. 2019. Task agnostic meta-learning for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11719–11727.
- Kang, B.; Liu, Z.; Wang, X.; Yu, F.; Feng, J.; and Darrell, T. 2019. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8420–8429.
- Karlinsky, L.; Shtok, J.; Harary, S.; Schwartz, E.; Aides, A.; Feris, R.; Giryas, R.; and Bronstein, A. M. 2019. 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*, 5197–5206.
- Koch, G.; Zemel, R.; Salakhutdinov, R.; et al. 2015. Siamese neural networks for one-shot image recognition. In *ICML deep learning workshop*, volume 2. Lille.
- Li, A.; and Li, Z. 2021. Transformation Invariant Few-Shot Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3094–3102.
- Li, B.; Yang, B.; Liu, C.; Liu, F.; Ji, R.; and Ye, Q. 2021. Beyond Max-Margin: Class Margin Equilibrium for Few-shot



- Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7363–7372.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *European conference on computer vision*, 21–37. Springer.
- Mao, Y.; Chen, K.; Diao, W.; Sun, X.; Lu, X.; Fu, K.; and Weinmann, M. 2022a. Beyond single receptive field: A receptive field fusion-and-stratification network for airborne laser scanning point cloud classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 188: 45–61.
- Mao, Y.; Guo, Z.; Lu, X.; Yuan, Z.; and Guo, H. 2022b. Bidirectional Feature Globalization for Few-shot Semantic Segmentation of 3D Point Cloud Scenes. *arXiv preprint arXiv:2208.06671*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Qiao, L.; Zhao, Y.; Li, Z.; Qiu, X.; Wu, J.; and Zhang, C. 2021. Defrcn: Decoupled faster r-cnn for few-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8681–8690.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Redmon, J.; and Farhadi, A. 2017. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7263–7271.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28: 91–99.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.
- Scott, T. R.; Ridgeway, K.; and Mozer, M. C. 2018. Adapted Deep Embeddings: A Synthesis of Methods for  $k$ -Shot Inductive Transfer Learning. *arXiv preprint arXiv:1805.08402*.
- Sun, B.; Li, B.; Cai, S.; Yuan, Y.; and Zhang, C. 2021. Fsce: Few-shot object detection via contrastive proposal encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7352–7362.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018a. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7794–7803.
- Wang, X.; Huang, T. E.; Darrell, T.; Gonzalez, J. E.; and Yu, F. 2020. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*.
- Wang, Y.-X.; Girshick, R.; Hebert, M.; and Hariharan, B. 2018b. Low-shot learning from imaginary data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7278–7286.
- Wu, A.; Han, Y.; Zhu, L.; Yang, Y.; and Deng, C. 2021. Universal-Prototype Augmentation for Few-Shot Object Detection. *arXiv preprint arXiv:2103.01077*.
- Wu, J.; Liu, S.; Huang, D.; and Wang, Y. 2020. Multi-scale positive sample refinement for few-shot object detection. In *European Conference on Computer Vision*, 456–472. Springer.
- Xiao, Y.; and Marlet, R. 2020. Few-shot object detection and viewpoint estimation for objects in the wild. In *European Conference on Computer Vision*, 192–210. Springer.
- Xu, H.; Wang, X.; Shao, F.; Duan, B.; and Zhang, P. 2021. Few-Shot Object Detection via Sample Processing. *IEEE Access*, 9: 29207–29221.
- Yan, X.; Chen, Z.; Xu, A.; Wang, X.; Liang, X.; and Lin, L. 2019. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9577–9586.
- Yang, Z.; Wang, Y.; Chen, X.; Liu, J.; and Qiao, Y. 2020. Context-transformer: tackling object confusion for few-shot detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12653–12660.
- Zhang, S.; Wang, L.; Murray, N.; and Koniusz, P. 2022. Kernelized Few-Shot Object Detection With Efficient Integral Aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19207–19216.
- Zhang, W.; and Wang, Y.-X. 2021. Hallucination Improves Few-Shot Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13008–13017.
- Zhu, C.; Chen, F.; Ahmed, U.; Shen, Z.; and Savvides, M. 2021. Semantic relation reasoning for shot-stable few-shot object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8782–8791.