

Fine-Grained Prototypes Distillation for Few-Shot Object Detection

Zichen Wang, Bo Yang*, Haonan Yue, Zhenghao Ma

School of Automation, Northwestern Polytechnical University, Xi'an, China
 {wangchen1801, hnyue, mazh0819}@mail.nwpu.edu.cn, byang@nwpu.edu.cn

Abstract

Few-shot object detection (FSOD) aims at extending a generic detector for novel object detection with only a few training examples. It attracts great concerns recently due to the practical meanings. Meta-learning has been demonstrated to be an effective paradigm for this task. In general, methods based on meta-learning employ an additional support branch to encode novel examples (a.k.a. support images) into class prototypes, which are then fused with query branch to facilitate the model prediction. However, the class-level prototypes are difficult to precisely generate, and they also lack detailed information, leading to instability in performance. New methods are required to capture the distinctive local context for more robust novel object detection. To this end, we propose to distill the most representative support features into fine-grained prototypes. These prototypes are then assigned into query feature maps based on the matching results, modeling the detailed feature relations between two branches. This process is realized by our Fine-Grained Feature Aggregation (FFA) module. Moreover, in terms of high-level feature fusion, we propose Balanced Class-Agnostic Sampling (B-CAS) strategy and Non-Linear Fusion (NLF) module from different perspectives. They are complementary to each other and depict the high-level feature relations more effectively. Extensive experiments on PASCAL VOC and MS COCO benchmarks show that our method sets a new state-of-the-art performance in most settings. Our code is available at <https://github.com/wangchen1801/FPD>.

Introduction

Object detection is a fundamental task in computer vision and the methods based on deep learning have been well established over the past few years (Redmon et al. 2016; Ren et al. 2017; Carion et al. 2020; Liu et al. 2016). While remarkable achievements have been made, most of them require a large amount of labeled data to obtain a satisfactory performance, otherwise they are prone to overfitting and hardly generalize to the unknown data.

Few-shot object detection (FSOD) is a more challenging task to detect object specially in data-scarce scenarios. FSOD assumes that there are sufficient amount of examples for base classes while only k-shot examples for each novel class.

*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

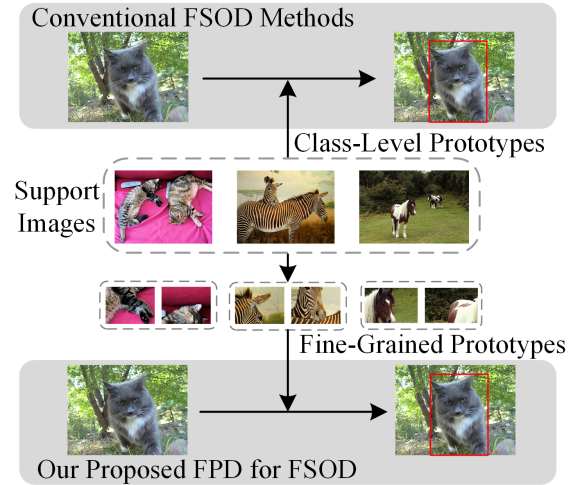


Figure 1: Overview of the proposed method, which we denote as FPD. In addition to class-level prototypes, we distill representative detailed features into fine-grained prototypes, enabling more robust novel object detection.

Therefore, the key question is how to transfer the knowledge learnt from base classes to the novel classes. **Transfer learning** based methods (Wang et al. 2020; Cao et al. 2021; Qiao et al. 2021) focus on fine-tuning the model more effectively. They use the same architecture as generic object detection, additionally with advanced techniques such as parameter freezing and gradient decoupling to improve performance. **Meta-learning** based methods (Kang et al. 2019; Wang, Ramanan, and Hebert 2019; Yan et al. 2019; Han et al. 2023), instead, follow the idea: learn how to learn the new tasks rapidly. As illustrated in Figure 2, an additional support branch is incorporated to encode support images into class-level prototypes, which function as dynamic parameters to interact with the query branch. In this way, the connections between novel examples and the model predictions are enhanced, thereby improving the generalization ability and learning the new tasks more quickly.

This work studies the meta-learning based FSOD and aims at realizing a more effective method. In general, features from the two branches are fused on top of the framework to make

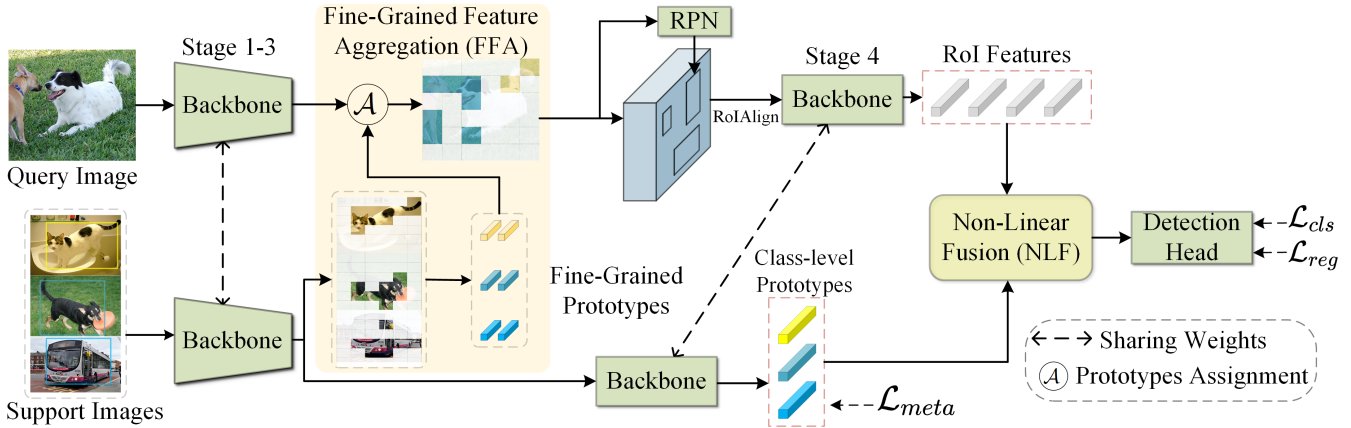


Figure 2: The overall architecture of our method. FFA and NLF are proposed to improve the performance.

the final prediction (Kang et al. 2019; Yan et al. 2019; Xiao and Marlet 2020), while most of the layers are separated and do not exchange information. This hinders the model from learning the correlations among detailed features especially in data-scarce scenarios.

DCNet (Hu et al. 2021) proposes to directly match the mid-level support features into query features in a pixel-wise manner, which enables the relation modeling of detailed local context. However, this approach has its limitations in terms of effect and implementation. **First**, the mid-level features with an extensive range of patterns are intricate and complex, thus the model might struggle to capture the most critical details. **Second**, directly matching between dense feature maps is inefficiency and will cost more computational resources. **Third**, this approach has difficulty in transitioning seamlessly from the training phase to the testing phase, as it can not integrate the mid-level support features across different shots to boost the performance.

To address the aforementioned issues, we propose a novel Fine-Grained Feature Aggregation (FFA) module to aggregate the mid-level features. As illustrated in Figure 3, different from DCNet, we propose to distill features into fine-grained prototypes. These prototypes, which reside in a highly refined and reduced feature space, embody the most distinctive and representative details of the support images. Specifically, we employ a set of embeddings following the object queries in DETR (Carion et al. 2020) to distill prototypes. Rather than being encoded with positional information and representing specific objects, the embeddings here function within the feature space and thereby are denoted as feature queries. We give each class a unique set of feature queries to distill prototypes independently. It can avoid confusion and is a key factor for our method to work. The distilled prototypes are then assigned into query feature map based on the matching results, modeling the fine-grained relations and highlighting the features with similar details.

The proposed FFA enables a more effective feature aggregation by focusing on the key information encapsulated within prototypes. This method also reduces the computational complexity by avoiding the directly matching between

dense feature maps. Furthermore, it can naturally transition into the testing phase through a weighted sum of prototypes across different shots, preserving the full potential derived from the training phase.

In terms of high-level feature aggregation, we revisit the previous methods and propose two improvements from different perspectives. **First**, we propose Balanced Class-Agnostic Sampling (B-CAS) strategy to control the ratio of support classes aggregated with query features. Meta R-CNN (Yan et al. 2019) adopts a simple class-specific aggregation scheme where only the features having the same classes are aggregated. While VFA (Han et al. 2023) proposes a class-agnostic aggregation scheme which randomly selects the support classes to reduce class bias. Our insight is that different support classes are served as positive and negative samples, thereby the balanced sampling is required to keep the most important positive samples from being overwhelmed. **Second**, many works (Kang et al. 2019; Yan et al. 2019; Han et al. 2023) employ element-wise multiplication to explore the relations within the same classes. However, it is not compatible with our proposed B-CAS which incorporates the feature aggregation between different classes. To solve this issue, we propose a stronger Non-Linear Fusion (NLF) module motivated by (Han et al. 2022a; Xiao and Marlet 2020) to fuse features more effectively. Our contributions can be summarized as follows:

- We propose to distill support features into fine-grained prototypes before being integrated into query feature maps, which can help the model grasp the key information. They are implemented in the Fine-Grained Feature Aggregation (FFA) module.
- We propose Balanced Class-Agnostic Sampling (B-CAS) strategy and Non-Linear Fusion (NLF) module. They are complementary to each other and can fuse high-level features more effectively.
- Extensive experiments illustrate that our method significantly improves the performance and achieves state-of-the-art results on the two widely used FSOD benchmarks.

Related Works

General Object Detection

Deep learning based object detection has been extensively studied in recent years. The well-established object detectors can be categorized into one-stage and two-stage methods. One-stage detectors (Redmon et al. 2016; Liu et al. 2016; Lin et al. 2017b) directly make predictions upon the CNN feature maps. While two-stage detectors (Ren et al. 2017; He et al. 2017) additionally employ a Region Proposal Network (RPN) to generate object proposals, which will be further refined into the final predictions. Both of them require the predefined dense anchors to generate candidates.

Recently, anchor-free detectors DETR (Carion et al. 2020) and Deformable DETR (Zhu et al. 2020) have been developed and are drawing more attention. They use a CNN backbone combining with Transformer encoder-decoders (Vaswani et al. 2017) for end-to-end object detection. A set of object queries are proposed to replace the anchor boxes. They will be refined into the detected objects layer by layer through Transformer decoders.

We employ the two-stage Faster R-CNN (Ren et al. 2017) framework to build our FSOD detector, and draw inspirations from DETR (Carion et al. 2020) into our approach.

Few-Shot Object Detection

Few-Shot Object Detection (FSOD), which studies the detection task in data-scarce situations, has been attracting an increased interest recently. LSTD (Chen et al. 2018) first proposes a transfer learning based approach to detect novel objects in a FSOD data setting. TFA (Wang et al. 2020) utilizes a cosine similarity based classifier and only fine-tunes the last layer with novel examples, achieving a comparable results with other complex methods. DeFRCN (Qiao et al. 2021) employs advanced gradient decoupling technique into the Faster R-CNN framework and integrates an offline prototypical calibration block to refine the classification results, which achieves an impressive performance.

The meta-learning is also a promising paradigm for FSOD. FSRW (Kang et al. 2019) proposes to re-weight the YOLOv2 feature maps along channel dimension using proposed reweighting vectors, which can highlight the relevant features. Meta R-CNN (Yan et al. 2019) adopts the Faster R-CNN framework to build a two-branch based siamese network. It processes query and support images in parallel to produce the Region of Interest (RoI) features and class prototypes, which are then fused to make predictions. Instead of learning a softmax-based classifier for all classes, (Han et al. 2022a) constructs a meta-classifier through feature alignment and non-linear matching. It calculates the similarity between query-support feature maps, producing binary classification results for novel classes. VFA (Han et al. 2023) introduces variational feature learning into Meta R-CNN, further boosting its performance. Recently, there are some works incorporate meta-learning into other advanced frameworks. Meta-DETR (Zhang et al. 2022) employs Deformable DETR (Zhu et al. 2020) to build a few-shot detector. (Han et al. 2022b) utilizes PVT (Wang et al. 2021) to construct

a fully cross transformer for few-shot detection. They all achieve remarkable results.

A two stage training paradigm has been widely adopted in both transfer learning and meta-learning based methods due to its effectiveness. At the base training stage, the model is trained on abundant base class examples. While at the fine-tuning stage, the model is fine-tuned only with K -shot examples for each base and novel class.

Our approach is based on Meta R-CNN and we propose to distill fine-grained prototypes for effectively exploiting the relations between detailed features.

Our Approach

In this section, we first introduce the task definition and the overall architecture of our model. Then we will elaborate the fine-grained and high-level feature aggregation.

Task Definition

We adopt the standard FSOD setting following (Kang et al. 2019; Wang et al. 2020). Specifically, given a dataset \mathcal{D} with two sets of classes C_{base} and C_{novel} , where each class in C_{base} has abundant training data while each class in C_{novel} has only K -shot annotated objects, FSOD aims at detecting the objects of $C_{base} \cup C_{novel}$ using the detector trained on \mathcal{D} . Please note that $C_{base} \cap C_{novel} = \emptyset$.

The Model Architecture

As illustrated in Figure 2, our model is based on Meta R-CNN, which is a siamese network with query branch and support branch that share a same backbone. Typically, we use the first three stages of ResNet-50/101 backbone (He et al. 2016) to extract mid-level features for both query images and support images. Then our proposed FFA module is employed to distill the fine-grained prototypes and assign them into the query branch. Subsequently, we use the last stage (i.e. stage four) of the backbone to extract high-level features for both branches, which produces RoI features and class-level prototypes, respectively. They are further processed by the proposed NLF module, following by the detection head to make the final prediction. We would like to mention that the RPN is fed with the query features which have already interacted with the support branch. It gives the RPN more ability learning to identify the new instances.

Fine-Grained Feature Aggregation

The Fine-Grained Feature Aggregation (FFA) module is the key component of our proposed method, which is a class-agnostic aggregator that matches all classes of support features into query features. It models inter-class relations in the early stage of the detection framework where the features are low-level and have more detailed information. Instead of directly performing feature matching, we propose to distill the representative support features into fine-grained prototypes. These prototypes are then assigned into query feature maps based on the matching results. FFA can help the model distinguish foreground from background and learn the similarities and differences between object classes. We will elaborate the

prototypes distillation and feature assignment in the following subsections. We also discuss our strategy to transfer this method to novel classes, as well as test-time natural integration of prototypes across different shots.

Prototypes Distillation Inspired by DETR, we incorporate a new component which is a set of learnable embeddings to distill prototypes. Different from object queries in DETR, which are encoded with positional information and are refined into a specific instance layer by layer, the embeddings here work as a guidance to refine the entire support feature space into a set of representative features. It can filter out the noise and ease the training. We refer to these embeddings as feature queries.

We employ the cross-attention mechanism to perform the prototypes distillation. Specifically, given a support feature map $X_s \in \mathbb{R}^{hw \times d}$ and a set of feature queries $q \in \mathbb{R}^{n \times d'}$, where hw denote the height and width, d and d' is the feature dimension, and n is the number of feature queries, the affinity matrix is calculated through a matching operation:

$$A = \text{softmax}\left(\frac{q(X_s W)^T}{\sqrt{d'}}\right) \quad (1)$$

where W is a linear projection to project X_s into the latent space with dimensionality d' , and the *softmax* function is performed along hw dimension. Subsequently, the fine-grained prototypes can be distilled from X_s via:

$$p = AX_s + E_{cls} \quad (2)$$

where the affinity matrix is applied directly on the support feature map. We do not project X_s to keep feature space the same. An additional class embedding E_{cls} is added to retain the class information.

We would like to mention that each class has its exclusive feature queries. This is different from object queries in DETR and is a crucial factor for our method to work. It means that q is the feature queries of one class and is part of $Q \in \mathbb{R}^{nc \times d'}$, where Q denotes the feature queries of all classes. This setting makes feature queries class-relevant and avoids them getting overwhelmed and confused by too many object classes.

Prototypes Assignment We densely match the fine-grained prototypes into query feature map to achieve the prototypes assignment. Considering that the background area should not be matched to any prototypes that represent salient object features, we incorporate a set of embeddings to serve as background prototypes. We also use the cross-attention mechanism to assign prototypes. Specifically, given a query feature map $X_q \in \mathbb{R}^{HW \times d}$, prototypes assignment is performed via:

$$A' = \text{softmax}\left(\frac{(X_q W')(PW')^T}{\sqrt{d'}}\right) \quad (3)$$

$$P = \text{concat}(p_1, p_2, \dots, p_c, p_{bg}) \quad (4)$$

$$X'_q = X_q + \alpha \cdot A'P \quad (5)$$

where $P \in \mathbb{R}^{(nc+n_{bg}) \times d}$ is the prototypes of c support classes with additional n_{bg} background classes, and W' is a linear projection shared by X_q and P which projects them into the

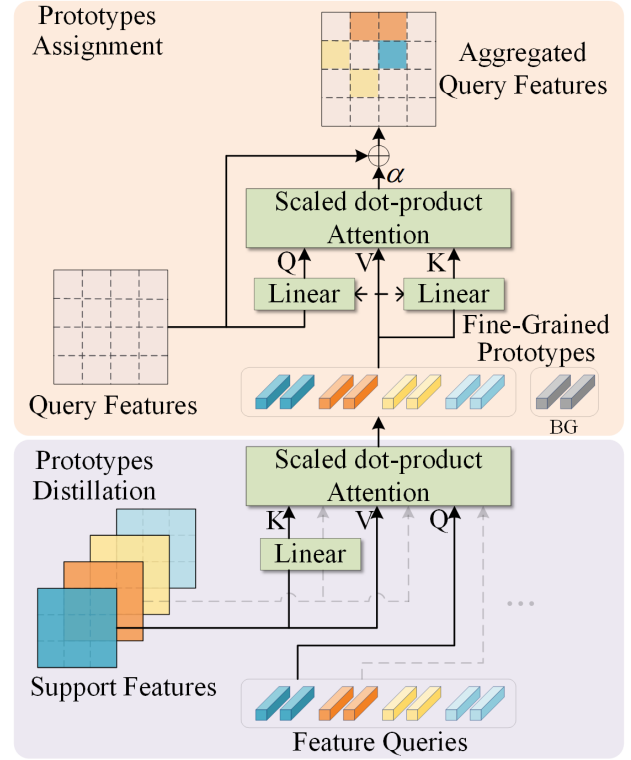


Figure 3: The architecture of the Fine-Grained Feature Aggregation (FFA) module. It can be divided into Prototypes Distillation and Prototypes Assignment.

same latent space. The prototypes are assigned into query feature map based on the affinity matrix A' , which produces the aggregated query features. The α is a learnable parameter initialized as zero to help stabilize the training.

Transferring to Novel Classes At the base training stage, the feature queries of base classes are randomly initialized and well trained. However, at the fine-tuning stage, training the feature queries from scratch becomes challenging due to the limited novel class examples, which means that an effective knowledge transfer method is required. To address this issue, we propose to duplicate the most compatible feature queries from the base classes to serve as those in the novel classes. To be specific, given feature queries of base classes $Q \in \mathbb{R}^{nc \times d'}$ and support feature map of a novel class $X_{ns} \in \mathbb{R}^{hw \times d}$, the compatibility matrix and the *weight* of each feature query can be obtained via:

$$C = \text{topk}(Q(X_{ns}W)^T) \quad (6)$$

$$\text{weight}_i = \sum_{j=0}^k C_{ij}, i = 1, 2, \dots, nc \quad (7)$$

where *topk* is performed along hw dimension to filter out irrelevant locations. We select n feature queries for each novel class based on the largest *weight*. Instead of sharing the same feature queries with base classes, they are created as a duplicate and can be trained independently.

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
<i>Single run results:</i>															
FSRW (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
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	55.2	61.8	24.4	29.3	39.2	39.9	47.8	35.6	41.8	42.3	48.0	49.7
Retentive (Fan et al. 2021)	42.4	45.8	45.9	53.7	56.1	21.7	27.8	35.2	37.0	40.3	30.2	37.6	43.0	49.7	50.1
FSCE (Sun et al. 2021)	44.2	43.8	51.4	61.9	63.4	27.3	29.5	43.5	44.2	50.2	37.2	41.9	47.5	54.6	58.5
Meta FR-CNN (Han et al. 2022a)	43.0	54.5	60.6	66.1	65.4	27.7	35.5	46.1	47.8	51.4	40.6	46.4	53.4	<u>59.9</u>	58.6
Meta-DETR (Zhang et al. 2022)	40.6	51.4	58.0	59.2	63.6	<u>37.0</u>	36.6	43.7	49.1	54.6	41.6	45.9	52.7	58.9	<u>60.6</u>
FCT (Han et al. 2022b)	49.9	57.1	57.9	63.2	67.1	27.6	34.5	43.7	49.2	51.2	39.5	<u>54.7</u>	52.3	57.0	58.7
VFA (Han et al. 2023)	57.7	64.6	64.7	<u>67.2</u>	<u>67.4</u>	41.4	46.2	51.1	<u>51.8</u>	51.6	48.9	54.8	<u>56.6</u>	59.0	58.9
FPD(Ours)	48.1	<u>62.2</u>	<u>64.0</u>	67.6	68.4	29.8	<u>43.2</u>	<u>47.7</u>	52.0	<u>53.9</u>	<u>44.9</u>	53.8	58.1	61.6	62.9
<i>Average results over multiple runs:</i>															
FSDetView (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
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
Meta-DETR (Zhang et al. 2022)	35.1	49.0	53.2	57.4	62.0	27.9	32.3	38.4	43.2	51.8	34.9	41.8	47.1	54.1	58.2
DeFRCN (Qiao et al. 2021)	40.2	<u>53.6</u>	58.2	63.6	<u>66.5</u>	<u>29.5</u>	39.7	43.4	48.1	<u>52.8</u>	<u>35.0</u>	38.3	52.9	57.7	60.8
FCT (Han et al. 2022b)	38.5	<u>49.6</u>	53.5	59.8	<u>64.3</u>	25.9	34.2	40.1	44.9	47.4	34.7	43.9	49.3	53.1	56.3
VFA (Han et al. 2023)	47.4	54.4	58.5	<u>64.5</u>	<u>66.5</u>	33.7	38.2	<u>43.5</u>	<u>48.3</u>	52.4	43.8	48.9	<u>53.3</u>	<u>58.1</u>	60.0
FPD(Ours)	<u>41.5</u>	52.8	<u>58.4</u>	64.9	67.1	28.2	<u>38.7</u>	43.8	50.3	53.6	34.9	<u>48.6</u>	54.0	58.4	61.5

Table 1: FSOD results (AP50) on the three splits of Pascal VOC dataset. We report both single run and multiple run results. Bold and Underline indicate the best and the second best results.

Test-Time Natural Integration A simple method to integrate fine-grained prototypes across different shots is to take the average. However, the detailed features represented by a feature query may not appear in some support images. Directly averaging might hurt the performance. Therefore, we compute a weighted sum using the aforementioned *weight*. Specifically, given K shot support images in a class, which produces K prototypes, the integration is performed via:

$$p_{avg} = \sum_{s=1}^K weight_s^* \cdot p_s \quad (8)$$

where $weight^*$ denote the *weight* after the *softmax* operation across different shot, p_{avg} is the integrated prototypes. This approach effectively filters out the prototypes that are not compatible with the current feature query, improving the robustness of our detector.

High-Level Feature Aggregation

Feature aggregation between RoI features and class-level prototypes is a crucial step for meta-learning based FSOD, where the high-level semantic information is aligned to make the final prediction. We revisit the conventional methods and propose two improvements from different perspectives.

Balanced Class-Agnostic Sampling Meta R-CNN adopts a simple class-specific aggregation scheme where the RoI features are aggregated only with the prototypes of the same class. While VFA proposes a class-agnostic aggregation scheme which aggregates RoI features with randomly selected class prototypes to reduce class bias. Nonetheless, we argue that the completely random sampling might disturb

the model from focusing on the most crucial positive prototypes and thus hurt the performance. Instead, we propose a balanced sampling strategy named B-CAS which selects a pair of positive and negative prototypes to aggregate with RoI features in parallel. The B-CAS not only enables the relation modeling between different classes but also keeps the positive prototype from being overwhelmed by too many negative examples, and therefore can learn the high-level semantic relations more effectively.

(Fan et al. 2020) employs a more complex training strategy which divides training pairs into three types and maintains a ratio of 1:2:1. Additionally, a matching loss is computed to align RoI features with prototypes. However, we find it instead hurts the performance. A plausible reason is that FFA introduces the asymmetry upon two branches, making the matching loss no longer beneficial. Consequently, a simple yet effective method B-CAS is adopted in our experiments.

Non-Linear Fusion Module Many previous meta-learning based methods use element-wise multiplication to handle the feature fusion. We argue that while this approach learns the similarities within the same class effectively, it struggles to capture the class differences. Therefore it is not compatible with the proposed B-CAS. To solve this problem, we employ a novel non-linear fusion network following (Han et al. 2022a; Xiao and Marlet 2020) with modifications.

Specifically, features after element-wise multiplication, subtraction and concatenation are processed independently to refine their relation to the new feature. Then they are concatenated with the vanilla RoI features and further refined before fed into the detection head. Given RoI feature $f_{roi} \in \mathbb{R}^{1 \times 2d}$ and class prototype $p_{cls} \in \mathbb{R}^{1 \times 2d}$, the aggregation

	Method	Framework	shot	
			10	30
<i>Single run results:</i>				
T	TFA w/ cos (Wang et al. 2020)	FR-CNN	10.0	13.7
	Retentive (Fan et al. 2021)	FR-CNN	10.5	13.8
	FSCE (Sun et al. 2021)	FR-CNN	11.9	16.4
	FADI (Cao et al. 2021)	FR-CNN	12.2	16.1
	DeFRCN (Qiao et al. 2021)	FR-CNN	18.5	22.6
M*	FCT (Han et al. 2022b)	Transformer	17.1	21.4
M	FSRW (Kang et al. 2019)	YOLOv2	5.6	9.1
	Meta R-CNN (Yan et al. 2019)	FR-CNN	8.7	12.4
	FSDetView (Xiao and Marlet 2020)	FR-CNN	12.5	14.7
	Meta FR-CNN (Han et al. 2022a)	FR-CNN	12.7	16.6
	VFA (Han et al. 2023)	FR-CNN	16.2	18.9
	FPD(ours)	FR-CNN	16.5	20.1
<i>Average results over multiple runs:</i>				
T	TFA w/ cos (Wang et al. 2020)	FR-CNN	9.1	12.1
	DeFRCN (Qiao et al. 2021)	FR-CNN	16.8	21.2
M*	FCT (Han et al. 2022b)	Transformer	15.3	20.2
	Meta-DETR (Zhang et al. 2022)	Def DETR	19.0	22.2
M	FSDetView (Xiao and Marlet 2020)	FR-CNN	10.7	15.9
	DCNet (Hu et al. 2021)	FR-CNN	12.8	18.6
	VFA (Han et al. 2023)	FR-CNN	15.9	18.4
	FPD(ours)	FR-CNN	15.9	19.3

Table 2: FSOD results (AP) on the MS COCO dataset. T: Transfer-learning based methods. M: Meta-learning based methods. M*: Meta-learning with advanced framework.

can be formulated as:

$$f' = [\mathcal{F}_1(f_{roi} \odot p_{cls}), \mathcal{F}_2(f_{roi} - p_{cls}), \mathcal{F}_3[f_{roi}, p_{cls}], f_{roi}] \quad (9)$$

$$f = \mathcal{F}_{agg}(f') \quad (10)$$

where \mathcal{F}_1 , \mathcal{F}_2 and \mathcal{F}_3 represent independent fully-connected layer followed by ReLU activation function, and \mathcal{F}_{agg} denote a pure fully-connected layer. This formulation provides a stronger capability to thoroughly explore the relations between high-level features. In addition, an exclusive path for RoI features is reserved to propagate the original RoI information, which reduces the noise introduced by random prototypes and can be used to regress the object location.

Experiments

Benchmarks

We evaluate our method on two widely-used FSOD benchmarks PASCAL VOC (Everingham et al. 2010) and MS COCO (Lin et al. 2014), using exactly the same class partitions and few-shot examples as in (Wang et al. 2020).

PASCAL VOC. The 20 PASCAL VOC classes are split into 15 base classes and 5 novel classes. There are three different class partitions for a more comprehensive evaluation. The VOC07 and VOC12 train/val sets are used for training and the VOC07 test set is used for evaluation. The Mean Average Precision at IoU=0.5 (AP50) is reported under $K=\{1, 2, 3, 5, 10\}$ shot settings.

	B-CAS	NLF	FFA	shot		
				3	5	10
Baseline				56.7	58.3	61.4
Ours	✓			61.2	64.7	64.9
	✓	✓		62.8	67.1	66.3
	✓	✓	✓	64.0	67.6	68.4

Table 3: Ablation study of different components.

MS COCO. For MS COCO, the 20 PASCAL VOC classes are used as novel classes, the other 60 classes are used as base classes. The 5k images from COCO2017 val are used for evaluation and the rest are used for training. We report the AP at IoU=0.5:0.95 under $K=\{10, 30\}$ shot settings.

Implementation Details

Our method is implemented with MMDetection (Chen et al. 2019). We adopt ResNet-101 (He et al. 2016) pretrained on ImageNet (Russakovsky et al. 2015) as the backbone. The single scale feature map is used for detection without FPN (Lin et al. 2017a). We resize the query images to a maximum of 1333x800 pixels, and the cropped instances from support images are resized to 224x224 pixels.

Our model is trained on 2x3090 Nvidia GPUs with a total batch size of 8, using the SGD optimizer. In the base training stage, the model is trained on VOC and COCO datasets for 20k/110k iterations. The learning rate is set to 0.004 and decayed at 17k/92k iteration by a factor of 0.1. In the fine-tuning stage, the learning rate is set to 0.001. We use exactly the same loss functions with Meta R-CNN.

Comparison with the State-of-the-Art Methods

PASCAL VOC. We show both the single run results and the average results over multiple runs of PASCAL VOC in Table 1. It can be seen that FPD significantly outperforms previous methods, achieving the state-of-the-art performance in most settings. Specifically, FPD outperforms previous method by 1.5%, 4.4%, and 6.8% on the three data splits under $K=10$ shot setting, respectively. We notice that under $K=\{1, 2\}$ shot settings, our method is less effective than VFA, which is a strong FSOD detector utilizing a variational autoencoder to estimate class distributions. Our analysis suggests that in extremely data-scarce scenarios, it is more challenging for the FFA to capture the representative and common features across different shots, therefore it fails to achieve the expected effect under $K=\{1, 2\}$ shot settings.

MS COCO. Table 2 shows the results of MS COCO. It can be seen that FPD outperforms all of the meta-learning based methods adopting the Faster R-CNN framework. For example, FPD improves performance by 6.3% compared to previous best result under $K=30$ shot setting. FPD ranks fourth among all the methods. Please note that our method focuses on the three proposed components, without using advanced frameworks or techniques such as DETR, Transformer or gradient decoupled layer. Given the challenging nature of the MS COCO dataset, we believe that the performance can be further improved with more refinements.



Figure 4: Visualization of the detection results on novel classes.

Method	Directly Match	FFA	shot		
			3	5	10
Baseline*			62.8	67.1	66.3
Ours	✓	✓	63.2	67.0	67.5
			64.0	67.6	68.4

Table 4: Comparison with directly matching.

Ablation Study

We conduct comprehensive experiments on the Novel Set 1 of PASCAL VOC under $K=\{3, 5, 10\}$ shot settings, which demonstrates the effectiveness of our proposed method.

Effect of Different Components. We show the results with different components in Table 3. It can be seen that B-CAS and NLF together improve the performance by about 10% over the baseline. Based on this, our FFA can further boost the results, achieving the state-of-the-art performance.

Effect of the FFA. FFA differs from DCNet in that it distills the fine-grained prototypes to aggregate with query branch. To demonstrate the superiority of this method, we re-implement the DRD module following DCNet to directly match dense feature maps for aggregation. We show the experimental results in Table 4. It can be seen that FFA consistently achieves better performance than directly matching, which validates the effectiveness of our method.

Effect of Feature Queries. We assign each class a set of feature queries, which are the key guidance to distill fine-grained prototypes. The number of feature queries for a class is set to 5 by default. Figure 5 shows the effect of this number.

Moreover, to explore the fundamental working mechanism, we visualize the attention heatmap of feature queries on support images. As shown in Figure 6, two feature queries from *person* category are listed. They are prone to focus on the specific details, e.g., *head* and *hand*, which conforms to our expectations. Please note that the generated heat maps has a resolution of 14x14. It is not absolutely aligned with the original images.

Visualize Detection Results

We show the detection results in Figure 4. The model is trained on the Novel Set 3 of PASCAL VOC under 10 shot

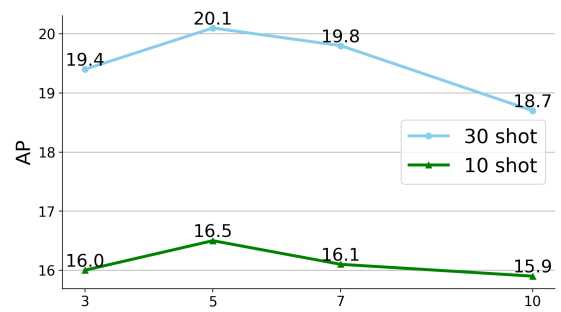


Figure 5: Ablation study on the number of feature queries.



Figure 6: Attention heatmap of feature queries. Please find more discussion and results in Appendix.

setting and tested on the VOC07 test set. It can be seen that many of the novel instances are effectively detected, even though the detected bboxes are not perfectly aligned. This results demonstrate the promising potential of our method.

Conclusion

This paper studies the meta-learning based FSOD. We propose a novel FFA module which can distill fine-grained prototypes in addition to class-level ones. It enables more robust novel object detection by focusing on the detailed features. We also propose B-CAS strategy and NLF module to aggregate high-level features more effectively. Both quantitative and qualitative results demonstrate the effectiveness of our method and the promising prospect of FSOD.

Acknowledgments

This work was supported in part by the Overseas Students Science and Technology Activities Project (No. 2018024), by the National Natural Science Foundation of China (No. 61502389), by the Natural Science Basic Research Program of Shaanxi Province, China (No. 2023-JC-YB-508).

References

- Cao, Y.; Wang, J.; Jin, Y.; Wu, T.; Chen, K.; Liu, Z.; and Lin, D. 2021. Few-Shot Object Detection via Association and Discrimination. In *NeurIPS*, 16570–16581.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*, 213–229.
- Chen, H.; Wang, Y.; Wang, G.; and Qiao, Y. 2018. LSTD: A low-shot transfer detector for object detection. In *AAAI*, 6066–6073.
- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; Zhang, Z.; Cheng, D.; Zhu, C.; Cheng, T.; Zhao, Q.; Li, B.; Lu, X.; Zhu, R.; Wu, Y.; Dai, J.; Wang, J.; Shi, J.; Ouyang, W.; Loy, C. C.; and Lin, D. 2019. MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv preprint arXiv:1906.07155*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *IJCV*, 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 *CVPR*, 4013–4022.
- Fan, Z.; Ma, Y.; Li, Z.; and Sun, J. 2021. Generalized few-shot object detection without forgetting. In *CVPR*, 4527–4536.
- 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 *AAAI*, 780–789.
- Han, G.; Ma, J.; Huang, S.; Chen, L.; and Chang, S.-F. 2022b. Few-shot object detection with fully cross-transformer. In *CVPR*, 5321–5330.
- Han, J.; Ren, Y.; Ding, J.; Yan, K.; and Xia, G.-S. 2023. Few-Shot Object Detection via Variational Feature Aggregation. In *AAAI*, 755–763.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- 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 *CVPR*, 10185–10194.
- Kang, B.; Liu, Z.; Wang, X.; Yu, F.; Feng, J.; and Darrell, T. 2019. Few-shot object detection via feature reweighting. In *ICCV*, 8420–8429.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal Loss for Dense Object Detection. In *ICCV*, 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 *ECCV*, 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 *ECCV*, 21–37.
- 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 *ICCV*, 8681–8690.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *CVPR*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2017. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE TPAMI*, 39(6): 1137–1149.
- 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. *IJCV*, 115(3): 211–252.
- Sun, B.; Li, B.; Cai, S.; Yuan, Y.; and Zhang, C. 2021. Fscf: Few-shot object detection via contrastive proposal encoding. In *CVPR*, 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. In *NeurIPS*, 5998–6008.
- Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 568–578.
- 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.; Ramanan, D.; and Hebert, M. 2019. Meta-learning to detect rare objects. In *ICCV*, 9925–9934.
- Wu, J.; Liu, S.; Huang, D.; and Wang, Y. 2020. Multi-scale positive sample refinement for few-shot object detection. In *ECCV*, 456–472.
- Xiao, Y.; and Marlet, R. 2020. Few-shot object detection and viewpoint estimation for objects in the wild. In *ECCV*, 192–210. Springer.
- 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 *ICCV*, 9577–9586.
- Zhang, G.; Luo, Z.; Cui, K.; Lu, S.; and Xing, E. P. 2022. Meta-DETR: Image-Level Few-Shot Detection with Inter-Class Correlation Exploitation. *IEEE TPAMI*, 45(11): 12832–12843.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.