

Query Adaptive Few-Shot Object Detection with Heterogeneous Graph Convolutional Networks

Guangxing Han, Yicheng He, Shiyuan Huang, Jiawei Ma, Shih-Fu Chang
Columbia University

{gh2561, yh3330, sh3813, jiawei.m, sc250}@columbia.edu

Abstract

Few-shot object detection (FSOD) aims to detect never-seen objects using few examples. This field sees recent improvement owing to the meta-learning techniques by learning how to match between the query image and few-shot class examples, such that the learned model can generalize to few-shot novel classes. However, currently, most of the meta-learning-based methods perform pairwise matching between query image regions (usually proposals) and novel classes separately, therefore failing to take into account multiple relationships among them. In this paper, we propose a novel FSOD model using heterogeneous graph convolutional networks. Through efficient message passing among all the proposal and class nodes with three different types of edges, we could obtain context-aware proposal features and query-adaptive, multiclass-enhanced prototype representations for each class, which could help promote the pairwise matching and improve final FSOD accuracy. Extensive experimental results show that our proposed model, denoted as QA-FewDet, outperforms the current state-of-the-art approaches on the PASCAL VOC and MSCOCO FSOD benchmarks under different shots and evaluation metrics.

1. Introduction

With abundant annotated training examples of objects, deep neural networks are tailored to extract commonalities and detect object instances accordingly. However, such methods tend to over-fit when there are only a few examples available. On the other hand, having seen many similar objects, humans can recognize a novel object when shown only a few examples of it. Inspired by humans' outstanding ability to generalize knowledge, few-shot object detection aims to detect novel object instances in an image given a few examples of the novel object (a.k.a a novel class) and abundant examples of other objects (a.k.a a base classes).

Although object detection methods like Faster-RCNN [31] work extremely well on data-abundant base classes, it

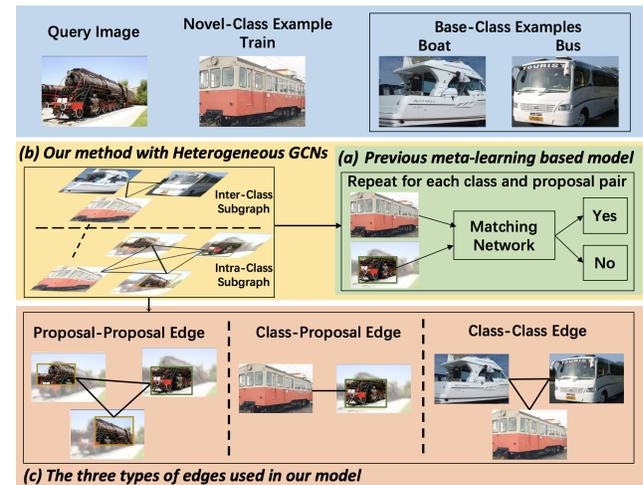


Figure 1. Overview of our proposed few-shot object detection (FSOD) method. (a) Previous meta-learning-based FSOD method. These methods aim to learn how to match between the regions (usually proposals) in the query image and few-shot class examples. Note that the key module, the matching network, will be repeatedly applied on each proposal and class pair. (b) Built upon the meta-learning-based approach, our method proposes a novel heterogeneous GCNs module with three different types of edges, allowing efficient message passing among all the nodes. (c) Three types of edges used in our model.

is non-trivial to adapt the model to few-shot novel classes. This is mainly because the softmax classifier used in R-CNN is tasked to perform classification among foreground classes and reject background regions at the same time. However, the notion of background changes when we adapt the model to novel classes. Typically, there are two approaches addressing this problem. One approach [37, 40] is to use long-tailed learning methods for training on the unbalanced dataset, but the generalization ability of such model is still limited. The other approach [3, 29, 39, 14] is to use meta-learning method to learn a class-agnostic few-shot detector on base classes, which can be easily adapted to novel classes without additional training. The simple framework and high detection

accuracy make the meta-learning-based approach a promising choice for FSOD.

The key of these meta-learning approaches is to learn how to match between the regions (usually RPN proposals [31]) in the query image and few-shot class examples. This is achieved by learning a class-agnostic matching network with a binary classifier. The two inputs are the proposal feature extracted from an image [6] and the prototype representation [32] of a few-shot class. The matching network will then be repeatedly applied on every proposal and class pair. However, there are three potential limits in this approach. **Firstly**, this kind of method turns out to be a ‘single-class’ detector, without modeling multi-class relations. This is important especially when there are base classes similar to the novel class as we could ‘borrow’ robust features from these classes. **Secondly**, since the class prototype is extracted only from the few-shot examples, there could potentially be a huge discrepancy between the extracted prototype and the proposal feature, considering the wide variety of objects and their different image statistics. **Thirdly**, proposals could be noisy and may not contain complete objects. Current methods do not consider the contextual information in the image for matching.

To address these three challenges, we propose a novel GCN-based FSOD model, denoted as QA-FewDet (**Q**uery-**A**daptive **F**ew-**S**hot **O**bject **D**etection), which utilizes graph propagation to learn context-aware proposal features and query-adaptive, multiclass-enhanced class prototypes. As shown in Fig. 1, we construct a graph among the proposal and class nodes for efficient class-class, class-proposal, and proposal-proposal communications. Firstly, by connecting different classes (including base classes) through **the class-class edge**, our method can model multi-class relations and enhance novel-class prototypes with prototypes from other similar classes. Secondly, **the class-proposal edge** provides mutual adaptation between class prototypes and proposal features and therefore reduces the distribution discrepancy between the two features. Meanwhile, it could provide additional examples of the class from proposals belonging to that class. Thirdly, **the proposal-proposal edge** provides both local and global contextual information to help classification and bounding box localization.

The naive way of graph construction is to include all proposals and classes in a single graph. However, such a graph is memory-expensive and inefficient for message passing. To better incorporate these three types of edges in our model, as shown in Fig. 1, we propose a novel heterogeneous graph consisting of a query-agnostic Inter-Class Subgraph, and multiple class-specific Intra-Class Subgraphs for each query image. The two subgraphs are processed sequentially for efficient message passing among all the nodes.

The entire network can be learned end-to-end using episode-based training on the abundant base-class data. To

show the effectiveness of our model, we conduct comprehensive experiments on two widely-used FSOD benchmarks. The results show that our QA-FewDet surpasses the current SOTAs by a huge margin under different shots and metrics.

Our contributions are 1) To our knowledge, we are the first to propose a graph model that considers class-class, class-proposal, and proposal-proposal relations in few-shot object detection, 2) we propose a novel heterogeneous graph structure that allows efficient message passing among all the nodes, 3) our model achieves significantly better results than current state-of-the-art methods on the PASCAL VOC and MSCOCO FSOD benchmarks across various settings.

2. Related Works

Object Detection. Current DCNNs-based object detection methods can mainly be grouped into two categories: proposal-based methods and proposal-free methods. Proposal-based methods [31, 11, 8, 10] divide object detection into two sequential stages by firstly generating a set of region proposals and then performing classification and bounding box regression for each proposal. Proposal-free methods [30, 23, 34, 9] directly predict the bounding boxes and the corresponding class labels on top of CNN features. We choose to use one of the most representative proposal-based methods, Faster R-CNN [31], in our model as it usually has better detection accuracy than proposal-free detectors due to the cascade design and improving detection accuracy is still the top priority for FSOD.

Few-Shot Learning and Meta-Learning. Few-shot learning aims to recognize novel classes using only few examples. Meta-learning has been demonstrated a promising learning paradigm for few-shot learning tasks by transferring meta-knowledge learned from data-abundant base classes to data-scarce novel classes. Current meta-learning based few-shot learning methods can be roughly divided into three categories: optimization-based methods [4], parameter-generation-based methods [5] and metric-learning-based methods [36, 32, 33, 25]. Most of the few-shot learning methods are developed for the image classification task.

Few-Shot Object Detection. Different from few-shot image classification, few-shot object detection needs to not only recognize objects with an arbitrary appearance, pose and scale using few-shot examples as references, but also localize (multiple) objects in the image and reject numerous background regions. Existing works can be mainly grouped into the following two categories: (1) Long-tailed-learning-based methods [37, 40]. These methods attempt to learn object detection by using training data from both data-abundant base classes and data-scarce novel classes. To deal with the unbalanced training set, re-sampling [37] and re-weighting [21] are the two main strategies [17]. However, the models trained on the joint dataset are inflexible for adding never-seen few-shot classes. (2) Meta-learning-based methods

[18, 16]. Meta-learner [16, 3, 29, 43] is introduced to acquire class-level meta knowledge via feature re-weighting and helps the model to generalize to novel classes. Meta-learning based methods [39, 18, 41, 3, 44, 7, 14, 28] has been demonstrated to be successful for FSOD. Moreover, meta-learning based methods can be efficient for incrementally adding new few-shot classes during network inference. Our method belongs to this category.

Graph Convolutional Networks. First proposed by Kipf et al. [19], graph convolutional networks (GCNs) and their variants like Graph Attentional Networks (GAT) [35], have seen massive applications in Computer Vision, including modeling video proposal relations in action localization [27, 45], object relations in visual relation reasoning [26], joint relations in skeleton-based action recognition [42], and object proposal relations in object detection [1]. Liu et al. [24] inject a GNN into a Faster R-CNN framework to contextualize the features of region proposals before the R-CNN classifier. This improves the results but it only shows results under traditional many-shot setting. Some methods [38] utilize GNNs on graphs that represent the ontology of concepts, which could enable generalization to unseen concepts by considering their relations with frequently-seen concepts. Different from previous works, we propose a novel heterogeneous GCNs in this paper that considers various relations within and between proposals and classes for FSOD.

3. Task Formulation

In few-shot object detection (FSOD), we split object classes C into C_{base} and C_{novel} such that $C = C_{base} \cup C_{novel}$ and $C_{base} \cap C_{novel} = \emptyset$. For each class $c \in C$, its annotations T_c contain the object instances' labels and bounding boxes inside the image. Formally, $T_c = \{(c, u, I) | u \in U, I \in \mathbb{R}^{H_I * W_I * 3}\}$, and $U \subseteq \mathbb{R}^4 = \{(x_i, y_i, w_i, h_i)\}$ represents bounding boxes in the image.

For each class $c \in C_{base}$, we have plenty of annotated instances. For each class $c \in C_{novel}$, we only have limited K examples, also known as support images, (e.g., $|T_c| = K$ for $K = 1, 5, 10$). FSOD aims to detect novel-class object instances with few annotated object instances. Formally, given a query image $I_q \in \mathbb{R}^{H_q * W_q * 3}$, FSOD outputs a set of detections $S_q = \{(c, u) | c \in C_{novel}, u \in U\}$.

4. The Baseline FSOD Model

As shown in Fig. 2, we mainly follow Fan et al.'s work [3] to build our baseline FSOD model, which adopts a siamese Faster R-CNN (with ResNet [12]) with two branches. In one branch, a query image I_q is fed into the feature extraction network to extract its feature $r(I_q)$ (output of *res4* block) for the detection head. Similarly, the other branch extracts the feature $r(I_s)$ given an input support image I_s .

Then the Attention-RPN [3] is used to produce N

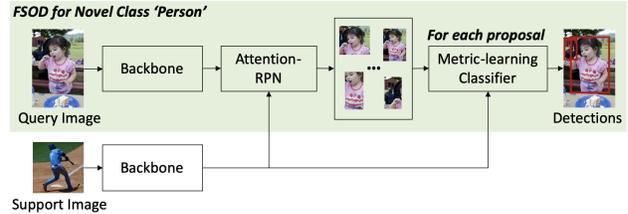


Figure 2. The baseline FSOD model [3].

($N = 100$ by default following [3]) class-specific proposals $P_c = \{\{p_i^c\}_{i=1}^N, p_i^c \in U\}$ for the novel class $c \in C_{novel}$. After that, we use the *res5* block and RoIpooling [6] to extract the feature $f(p_i^c) \in \mathbb{R}^{H \times W \times C}$ ($H = W = 7, C = 2048$) for proposal p_i^c from the query image feature $r(I_q)$. We apply the same layers to $r(I_s)$, and take the average feature of all support images belonging to the novel class c as the class prototype $f(c)$. After that, a multi-relation network [3] is used to calculate similarity score between the proposal feature and the class prototype, and then produce the final detection results for class c following [31]. The above process will be applied to each novel class independently.

The reason we use [3] as our baseline model is as follows. Firstly, the whole framework is simple and elegant, and is a natural extension of the original Faster R-CNN to the few-shot setting. Secondly, as shown in the Section 6.3, it has achieved SOTA accuracy in major FSOD benchmarks.

Although [3] has been demonstrated to be a promising FSOD model, there are three potential limits as discussed in the Section 1, and the primary cause is the separate classification of each proposal and class pair. To deal with this problem, we propose a novel FSOD model with heterogeneous GCNs in the following section.

5. Our Method with Heterogeneous GCNs

The ultimate goal of our proposed heterogeneous GCNs is to enable efficient message passing among all the proposals and classes before pairwise classification. To this end, we first generate class-specific proposals and extract proposal features and class prototypes following the baseline model. Then we establish a novel heterogeneous graph using the generated proposal and class nodes, and use GCN layers to update features for each node in our graph in a sequential manner. After that, we use the updated features for the final pairwise classification.

5.1. Overview of Our Heterogeneous Graph

We aim to construct a graph to capture various types of relations within and between proposals and classes. One way to construct the graph is to include all proposals and classes in a single graph. However, such a construction is memory-expensive as the number of proposals increases significantly with the number of classes, and the edge between cross-class

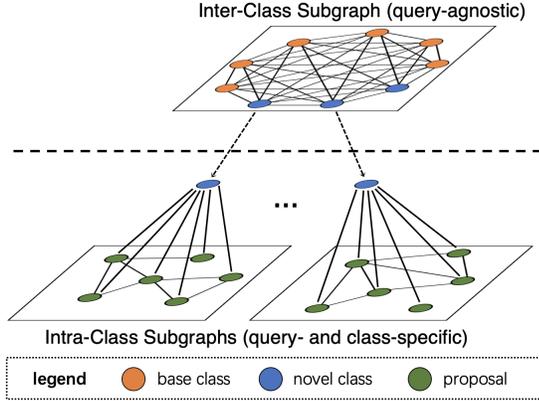


Figure 3. Our proposed heterogeneous graphs.

proposals incurs redundant and noisy information.

Considering the class-specific proposals generated in our model, we build a heterogeneous graph G that consists of two types of nodes, three types of edges, and two types of subgraphs as illustrated in Fig. 3. Specifically, we build a graph $G = (V, E)$, where V and E denote the node sets and edge sets respectively. The two types of nodes in V are the proposal nodes $V_p = \cup_{c \in C_{novel}} \{P_c\}$, namely the class-specific proposals for each novel class, and the class nodes $V_c = C$ including all novel and base classes. E has three components: the class-class edges E_{c-c} , the proposal-proposal edges E_{p-p} , and the class-proposal edges E_{c-p} . Our heterogeneous graph G has two types of subgraphs: a query-agnostic Inter-Class Subgraph G_{inter} shared among all query images, and multiple Intra-Class Subgraphs $G_{intra} = \{G_c, c \in C_{novel}\}$ for each query image, where G_c represents the class-specific subgraph for novel class c .

5.2. The Inter-Class Subgraph

Inspired by previous works [13, 5], modeling multi-class relationships could help enhance the prototype representation of few-shot classes and perform robust classification. In this paper, we first establish an Inter-Class Subgraph among all classes. Given the set of all classes $C = \{c_1, c_2, \dots, c_{|C|}\}$, we construct $G_{inter} = (V_{inter}, E_{inter})$, where $V_{inter} = C$ and $E_{inter} = \{(c_i, c_j) : 1 \leq i, j \leq |C|\}$. A_{inter} is the adjacency matrix of G_{inter} with edge weights for E_{inter} . G_{inter} is a graph with only class nodes, and considers relations among all base classes and novel classes for enhancing novel class prototypes.

The Class-Class Edge. To efficiently aggregate valuable information from other classes, we compute the class-wise correlation between every class pair. The key idea is that if two classes are similar, it would be meaningful to update one class’s prototype using that of the other class. Therefore, the edge weight between the two should increase correspondingly. Formally, given a class pair (c_i, c_j) , we use cosine

similarity to estimate their correlation by ¹

$$e(c_i, c_j) = \frac{f(c_i)^T f(c_j)}{\|f(c_i)\|_2 \cdot \|f(c_j)\|_2} \quad (1)$$

We then apply softmax for each class c_i to normalize its pairwise correlation

$$A_{inter}^{ij} = \frac{\exp(e(c_i, c_j))}{\sum_{k \in C} \exp(e(c_i, c_k))}, \quad (2)$$

5.3. The Intra-Class Subgraph

Considering the class-specific proposals generated in our model, we build an Intra-Class Subgraph for each novel class. Each subgraph consists of one class node and the corresponding class-specific proposals. Formally, for each novel class c , we construct $G_c = (V_c, E_c)$, where $V_c = \hat{P}_c \cup \{c\}$, and $\hat{P}_c = P_c \cup \{g\}$. Here, g denotes the ‘proposal’ containing the entire image I , and we obtain its feature $f(g)$ by performing RoI-pooling from the whole image feature. A_c is the adjacency matrix of G_c with edge weights for E_c .

To model the different types of relations within an Intra-Class Subgraph, we break E_c into two components:

The Proposal-Proposal Edge. Contextual information has been demonstrated to be crucial for traditional many-shot object detection [45, 15, 24]. In this paper, we apply this idea to the few-shot setting to obtain context-aware proposal features. To be specific, we establish an edge between proposal p_i ² and proposal p_j if $IoU(p_i, p_j) > \theta$, where θ is a fixed threshold for determining meaningful overlaps ($\theta = 0.7$ by default following [45]), and IoU is defined as

$$IoU(p_i, p_j) = \frac{\cap(p_i, p_j)}{\cup(p_i, p_j)}, \quad p_i, p_j \in P_c \quad (3)$$

At the same time, we also provide image-level contextual information by connecting the aforementioned global image node g to each proposal.

By building edges with nearby proposals and with the global scene-context, the proposal-proposal edges enrich the original proposal features with both local and global contextual information, leading to better classification accuracy and more precise bounding box locations. The weight of the proposal-proposal edge in A_c is calculated and normalized as in Eq. 1 and Eq. 2.

The Class-Proposal Edge. Traditional methods [3, 41, 43] usually use the K -shot support images to extract the prototype of a novel class. However, the extracted prototype cannot represent the novel class well using very few examples. Moreover, it is challenging to extract robust novel-class prototypes using the feature backbone trained only on

¹As in section 4, we use $f(c)$ to represent the prototype of class c . The same rule applies for the proposal feature.

²We use novel class c as an example in the rest of this section, and remove the class label c in p_i^c for simplicity reason.

base classes. To calibrate the different statistical distribution between the proposal feature and the class prototype, we introduce the class-proposal edge for dynamic mutual adaptations. Moreover, if the query image contains instances of the novel class, the class-specific proposals should cover these regions, and thus contributing extra-shots from the query image. This could help extract accurate novel-class prototype that is more suitable for the query image.

In practice, we establish bidirectional edges between the class node c and all the class-specific proposal nodes P_c . We show in Section 6.2 that mutual adaptation is better than adaption with only a single direction. Furthermore, since we only want to connect relevant proposals to the class node, we compute the cosine similarity as in Eq. 1 to estimate the correlation between the class-proposal pair and filter out noisy relations. Then, for the class node c , we normalize the weights of all the incoming class-proposal edges similar to Eq. 2. For the proposal nodes, we use the original cosine similarity as the weight of the incoming class-proposal edge.

5.4. Our Heterogeneous GCNs

After building our heterogeneous graph, we first perform message passing on the query-agnostic Inter-Class Subgraph for enhancing novel-class prototypes before processing any query image. With the enhanced novel-class prototypes, we then build multiple Intra-Class Subgraphs for each query image, and apply message passing for facilitating the communication among the proposal and class nodes.

For each graph, we sequentially perform L -layer GCNs ($L = 1$ as shown in section 6.2), which take in feature X_0 and output feature X_L of the same size ($\mathbb{R}^{7 \times 7 \times 2048}$). In practice, we implement the l^{th} GCN layer ($1 \leq l \leq L$) by

$$X^l = AX^{l-1}W^l \quad (4)$$

In this equation, $X^l \in \mathbb{R}^{H \times W \times d_l}$ is the output feature of the l^{th} GCN layer. A_{inter} and A_c are the respective adjacency matrix of the Inter-Class and class c 's Intra-Class Subgraph. $W^l \in \mathbb{R}^{d_{l-1} \times d_l}$ is a learnable parameter matrix. Each layer is followed by a residual block.

Note that to keep a consistent feature space between the class and proposal nodes, we apply the same number of learnable transformation layers to both of them as a siamese network. Therefore, we do not use W^l in any GCN layer in the Inter-Class Subgraph. A detailed discussion about this can be found in section 6.2. For each class c_i , the effect of each layer is equivalent to a weighted sum of other class prototypes by ³

$$\tilde{f}(c_i) = \sum_{j \in C} A_{inter}^{ji} \cdot f(c_j) + f(c_i) \quad (5)$$

³For simplicity reason, we use $f(c_i)$ and $\tilde{f}(c_i)$ to denote the input and output of the l^{th} layer GCN respectively. The similar strategy applies for Eq. 6 and Eq. 7.

In the Intra-Class Subgraph, for each proposal p_i of novel class c , let $P_{p_i} = \{g\} \cup \{p_j \in P_c \mid IoU(p_i, p_j) > \theta\}$ denote the set of its overlapping proposals and the global 'proposal'. We aggregate p_i 's feature from the enhanced class prototype and proposal features of P_{p_i} using a GCN layer by

$$\tilde{f}(p_i) = (A_c^{cp_i} \cdot f(c) + \sum_{p \in P_{p_i}} A_c^{pp_i} \cdot f(p))W + f(p_i), \quad (6)$$

$A_c^{cp_i}$ is the edge weight of class c and proposal p_i in G_c . $A_c^{pp_i}$ is defined similarly. We update class c 's prototype by

$$\tilde{f}(c) = (f(c) + \sum_{p_k \in P_c} A_c^{pkc} \cdot f(p_k))W + f(c), \quad (7)$$

After obtaining the updated proposal features and the class prototype, we feed them into the pairwise matching network [3] for the final classification.

5.5. Training Framework

To transfer knowledge from the base classes to the novel classes, we adopt a two-stage training strategy.

Meta-learning with Base Classes. With a pretrained feature extractor, we perform episode-based training on base classes. To simulate the few-shot scenario, each episode consists of one annotated query image and K randomly sampled shots for each base class. The whole model is supervised under a binary cross-entropy loss for classification and a smooth $L1$ loss for bounding box regression. During meta-testing, we can adapt our model to novel classes by simply calculating their prototype representations.

Fine-tuning with Novel Classes (Optional). We can further fine-tune the class-agnostic few-shot detector on novel classes following previous works [3, 41, 43, 37]. The difference between fine-tuning and meta-learning-only is that during fine-tuning, we use positive and negative proposals generated from the original novel class images to train our few-shot detector, while in meta-learning-only there is no training over novel classes. We study the performance of our model both with and without fine-tuning in section 6.2.

6. Experimental Results

6.1. Datasets

We use two widely-used few-shot object detection benchmarks MSCOCO 2014 [22] and PASCAL VOC 2007 and 2012 [2] for model evaluation, and follow the same FSOD settings as previous works [16, 37] for fair comparison.

MSCOCO. We set the 20 PASCAL VOC categories as novel classes and the remaining 60 categories as base classes. We use the same few-shot support images as [16]. We report detection accuracy with AP, AP50, and AP75 under shots

1, 2, 3, 5, 10 and 30. Note that 30-shot is considered as few-shot in MSCOCO dataset because the accuracy still largely falls behind the many-shot setting [16]. We use the MSCOCO dataset for ablation study in Section 6.2.

PASCAL VOC, The 20 PASCAL VOC categories are split into 15 base classes and 5 novel classes. We follow [16] and use the same base/novel splits and support images. We report AP50 results under shots 1, 2, 3, 5, and 10.

More implementation details are included in the Supplementary file.

6.2. Ablation Study

How do the graph convolutional layers help for FSOD? As shown in the Table 1 and Figure 4, we analyze the impact of each component in our model. We first verify the effectiveness of the GCN layers. To this end, we replace the GCN layers with fully-connected layers (MLP). Specifically, if we only consider a one-layer GCN, Eq. 4 becomes $Y = AXW$, where A is the adjacency matrix, and W is the learnable parameter. Notice that the MLP baseline shares the same structure as GCN except that we remove the adjacency matrix A . In other words, the MLP can be formulated as $Y = XW$. Compared with the GCN layer, the MLP baseline only uses self-connected edges in the graph, and as a result, each node updates its features independently. Comparing the MLP baseline in Table 1 (b) with the vanilla baseline model in Table 1 (a) and our heterogeneous GCNs in Table 1 (g), we can conclude that the additional learnable modules are useful, and the message passing among different nodes in the graph is crucial for the final performance.

How do the three types of edges help for FSOD? We then analyze the roles of the three types of edges in our heterogeneous graph. Firstly, we experiment on using only one type of edges as shown in Table 1 (c-e). We notice that all three types of edges can improve the baseline model’s performance in Table 1 (b). This demonstrates the effectiveness of multi-class modeling, class-proposal mutual adaptation, and learning context-aware proposal features in our model. Among all three types of edges, we observe that the class-proposal edge is the most important. This is because the objective of FSOD is to calculate the similarity score between the proposal feature and class prototype, thereby preferring that both sides adapt to each other. By further adding proposal-proposal edges, we obtain the full Intra-Class Subgraph, which, as shown in Table 1 (f), further improves the performance owing to context-aware proposal features. Finally, our full model, shown in Table 1 (g), achieves the best results after introducing the Inter-Class Subgraph.

The effectiveness of meta-learning and fine-tuning. We show the comparison results between meta-training-only and fine-tuning in Table 1 (g) and (h). We find that fine-tuning improves the performance in 10/30 shot settings.

However, when examples are extremely scarce, e.g., 2-shot as in Table 1, the performance hardly improves as fine-tuning tends to over-fit with small samples. This demonstrates the strong generalization ability of our meta-learning-based model, and fine-tuning needs large number of samples to perform well.

The roles of local and global context in the proposal-proposal edge. We show the ablation study of using local and global context in Table 2. We find that both local and global contextual information contributes to the model’s performance. Local context can provide missing features and help refine bounding boxes, especially when proposals are not accurate. Global scene-context, on the other hand, can provide complementary information from the global view. Using both context produces the best results.

The efficacy of the bidirectional class-proposal edge. We show in Table 3 the results of using bidirectional class-proposal edges against using solely uni-directional edges. We observe that the model with mutual adaptation achieves better results than with any of the two types of uni-directional edges.

The comparison between the class-proposal edge and non-local attention in [14]. We compare our proposed method with *Hsieh et al.*’s work [14] in Table 3. *Hsieh et al.* [14] proposes to use non-local attention (a.k.a co-attention) between the query image and the support image for feature enhancement. To compare its performance with our class-proposal edge, we use the official codebase of [14] and perform training/testing in our FSOD pipeline. Our model outperforms *Hsieh et al.*’s method [14] significantly. The main difference is that in [14], each ‘pixel’ in the feature maps is regarded as a basic unit for co-attention. In contrast, the category-specific proposal is used as the basic node in our model. Compared with the ‘pixels’ in the feature maps used in [14], the category-specific proposals could provide richer semantics related to the target object in the query image, and therefore are more suitable for mutual adaptation.

The advantage of using base class memory in the Inter-Class Subgraph. We show in Table 6 the results of using different numbers of base classes in our Inter-Class Subgraph. If we only use novel classes, the model gains little advantage from multi-class modeling. However, the Inter-Class Subgraph sees massive improvement when introduced with all base-class prototypes, which could enhance novel-class prototypes with more robust features from base classes. In practice, we deploy all 60 base classes by default in our Inter-Class Subgraph.

The analysis on different numbers of GCN layers in the Inter-Class and the Intra-Class Subgraphs. We show in Table 4 and 5 the results of applying different numbers of GCN layers in the two Subgraphs. (1) We first emphasize that proposal and class nodes should go through the same number of transformation layers in a siamese network before

Table 1. Ablation study on each component in the our model.

	Class-class	Class-proposal	Proposal-proposal	2-shot			10-shot			30-shot		
				AP	AP50	AP75	AP	AP50	AP75	AP	AP50	AP75
Meta-training the model on base classes, and meta-testing on novel classes												
(a)			w/o heterogeneous GCNs	5.4	11.6	4.6	7.6	15.4	6.8	8.9	17.8	8.0
(b)			Using MLP instead of GCN layers (only self-connected edges)	5.9	12.5	5.1	8.4	17.0	7.6	9.8	20.3	8.8
(c)	✓			6.3	13.3	5.5	9.0	17.7	8.1	10.6	20.9	9.6
(d)		✓		7.6	16.0	6.5	9.8	19.7	8.8	11.2	22.8	10.1
(e)			✓	6.7	14.0	5.8	9.3	18.5	8.3	10.8	21.5	9.7
(f)		✓		7.6	16.2	6.5	10.0	20.1	8.9	11.3	23.1	10.1
(g)	✓	✓	✓	7.8	16.4	6.6	10.2	20.4	9.0	11.5	23.4	10.3
Fine-tuning the model on novel classes, and testing on novel classes												
(h)	✓	✓	✓	7.6	16.1	6.2	11.6	23.9	9.8	16.5	31.9	15.5

Table 2. Ablation study on the proposal-proposal edge.

Local	Global	2-shot			10-shot		
		AP	AP50	AP75	AP	AP50	AP75
✓		6.6	13.8	5.7	9.0	17.9	8.1
	✓	6.4	13.2	5.7	9.1	17.9	8.2
✓	✓	6.7	14.0	5.8	9.3	18.5	8.3

Table 3. Ablation study on the class-proposal edge.

Model	2-shot			10-shot		
	AP	AP50	AP75	AP	AP50	AP75
class→proposal	6.4	13.2	5.9	8.6	17.1	7.9
class←proposal	7.2	15.7	5.7	9.4	19.1	8.6
class↔proposal	7.6	16.0	6.5	9.8	19.7	8.8
Non-local attention [14]	6.0	12.8	5.3	8.3	17.3	7.4

Table 4. Ablation study on the number of GCN layers for the Inter-Class Subgraph.

#GCN Layer	2-shot			10-shot		
	AP	AP50	AP75	AP	AP50	AP75
1 w/ W	5.0	10.4	4.3	7.2	14.2	6.5
1 w/o W	6.3	13.3	5.5	9.0	17.7	8.1
2 w/o W	6.1	13.0	5.3	8.8	17.6	8.0
3 w/o W	5.8	12.2	4.9	8.6	17.3	7.8

Table 5. Ablation study on the number of GCN layers for the Intra-Class Subgraph.

#GCN Layer	2-shot			10-shot		
	AP	AP50	AP75	AP	AP50	AP75
1 w/o W	7.2	15.6	6.0	9.4	19.3	8.4
1 w/ W	7.6	16.2	6.5	10.0	20.1	8.9
2 w/ W	7.4	15.9	6.3	9.8	19.7	8.8
3 w/ W	6.9	14.5	5.5	9.2	18.6	8.1

the final pairwise matching. If we learn W for the inter-class subgraph, the few-shot classes would have one more learnable layer than the proposals, and we show in Table 4 that the performance is worse. But for intra-class subgraphs, the GCN layers are applied to both proposals and few-shot

Table 6. Ablation study on base class memory in the Inter-Class Subgraph.

#Base Classes	2-shot			10-shot		
	AP	AP50	AP75	AP	AP50	AP75
0	6.0	12.8	5.1	8.5	17.0	7.7
20	6.1	13.1	5.2	8.6	17.3	7.7
40	6.3	13.4	5.4	8.9	17.5	8.0
60	6.3	13.3	5.5	9.0	17.7	8.1

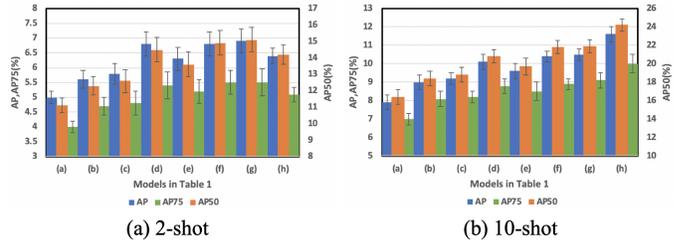


Figure 4. The average accuracy and standard deviation result of models in Table 1 over 10 runs.

classes. As shown in Table 5, learning W does not violate the siamese structure, and could help improve performance. (2) As shown in both Table 4 and 5, using one GCN layer is sufficient in both subgraphs as we already connect edges to all neighbors that a node needs in our model. Using more GCN layers are not helpful due to the over-smoothing problem [20] in GCNs.

6.3. Comparison with State-of-the-arts

As shown in Table 7 and 8, we compare our QA-FewDet with the STOA on PASCAL VOC and MSCOCO FSOD benchmarks. We draw the following three conclusion. (1) Our final model significantly outperforms previous STOA by more than 4.0% on AP50 in most of the shots and metrics of the PASCAL VOC. We achieve similar improvement on the MSCOCO. (2) Fine-tuning does not help too much in extreme few-shot settings because it is prone to over-fitting with very few samples (e.g., 1-shot in the PASCAL VOC,

Table 7. Few-shot object detection performance (AP50) on the PASCAL VOC dataset. †We re-evaluate the methods following the standard procedure in [16, 37]. Our approach with only meta-learning could achieve competitive results compared with other methods on extreme few-shot setting. After finetuning our model consistently outperforms other state-of-the-art methods in almost all settings.

Method / Shot	Venue	Backbone	Novel Set 1					Novel Set 2					Novel Set 3				
			1	2	3	5	10	1	2	3	5	10	1	2	3	5	10
Meta-training the model on base classes, and meta-testing on novel classes																	
Fan et al. [3]†	CVPR 2020	ResNet-101	32.4	22.1	23.1	31.7	35.7	14.8	18.1	24.4	18.6	19.5	25.8	20.9	23.9	27.8	29.0
QA-FewDet (Ours)	This work	ResNet-101	41.0	33.2	35.3	47.5	52.0	23.5	29.4	37.9	35.9	37.1	33.2	29.4	37.6	39.8	41.5
Fine-tuning the model on novel classes, and testing on novel classes																	
FSRW [16]	ICCV 2019	YOLOv2	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
MetaDet [39]	ICCV 2019	VGG16	18.9	20.6	30.2	36.8	49.6	21.8	23.1	27.8	31.7	43.0	20.6	23.9	29.4	43.9	44.1
Meta R-CNN [43]	ICCV 2019	ResNet-101	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/ fc [37]	ICML 2020	ResNet-101	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 [37]	ICML 2020	ResNet-101	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
Xiao et al. [41]	ECCV 2020	ResNet-101	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
MPSR [40]	ECCV 2020	ResNet-101	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
Fan et al. [3] †	CVPR 2020	ResNet-101	37.8	43.6	51.6	56.5	58.6	22.5	30.6	40.7	43.1	47.6	31.0	37.9	43.7	51.3	49.8
QA-FewDet (Ours)	This work	ResNet-101	42.4	51.9	55.7	62.6	63.4	25.9	37.8	46.6	48.9	51.1	35.2	42.9	47.8	54.8	53.5

Table 8. Few-shot detection performance on COCO novel classes. †We re-evaluate the methods following the standard procedure in [16, 37]. ‡The authors show these results at <https://github.com/YoungXIAO13/FewShotDetection>. Our method consistently outperforms state-of-the-art methods in most of the shots and metrics.

Method	1-shot			2-shot			3-shot			5-shot			10-shot			30-shot		
	AP	AP50	AP75	AP	AP50	AP75	AP	AP50	AP75									
Meta-training the model on base classes, and meta-testing on novel classes																		
Fan et al. [3]†	4.0	8.5	3.5	5.4	11.6	4.6	5.9	12.5	5.0	6.9	14.3	6.0	7.6	15.4	6.8	8.9	17.8	8.0
QA-FewDet (Ours)	5.1	10.5	4.5	7.8	16.4	6.6	8.6	17.7	7.5	9.5	19.3	8.5	10.2	20.4	9.0	11.5	23.4	10.3
Fine-tuning the model on novel classes, and testing on novel classes																		
FSRW [16]	-	-	-	-	-	-	-	-	-	-	-	-	5.6	12.3	4.6	9.1	19.0	7.6
MetaDet [39]	-	-	-	-	-	-	-	-	-	-	-	-	7.1	14.6	6.1	11.3	21.7	8.1
Meta R-CNN [43]	-	-	-	-	-	-	-	-	-	-	-	-	8.7	19.1	6.6	12.4	25.3	10.8
TFA w/ fc [37]	2.9	5.7	2.8	4.3	8.5	4.1	6.7	12.6	6.6	8.4	16.0	8.4	10.0	19.2	9.2	13.4	24.7	13.2
TFA w/ cos [37]	3.4	5.8	3.8	4.6	8.3	4.8	6.6	12.1	6.5	8.3	15.3	8.0	10.0	19.1	9.3	13.7	24.9	13.4
Xiao et al. [41]‡	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	15.9	31.7	15.1
MPSR [40] †	2.3	4.1	2.3	3.5	6.3	3.4	5.2	9.5	5.1	6.7	12.6	6.4	9.8	17.9	9.7	14.1	25.4	14.2
Fan et al. [3]†	4.2	9.1	3.0	5.6	14.0	3.9	6.6	15.9	4.9	8.0	18.5	6.3	9.6	20.7	7.7	13.5	28.5	11.7
QA-FewDet (Ours)	4.9	10.3	4.4	7.6	16.1	6.2	8.4	18.0	7.3	9.7	20.3	8.6	11.6	23.9	9.8	16.5	31.9	15.5

and 1/2/3-shot in the MSCOCO.), but could help in larger shot settings. (3) Our meta-learning-only model improves significantly compared with the strong baseline model [3], and outperforms or at least attains comparable results compared with other SOTAs using fine-tuning on 1/2-shot in the PASCAL VOC and on 1/2/3/5/10-shot in the MSCOCO.

7. Conclusion

In this paper, we introduce a novel heterogeneous GCNs that consider multi-relations among the proposal and class nodes for FSOD. The Inter-Class Subgraph enhances novel-class prototype representation via modeling multi-class relations. The Intra-Class Subgraph provides query-adaptive class prototypes and context-aware proposal features to facilitate pairwise matching. Our experiments show that our

model, QA-FewDet, with only meta-learning, can outperform or achieve competitive results especially on extreme few-shot settings. After finetuning, our model outperforms current SOTAs by a large margin across various settings.

Acknowledgment: This research is based upon work supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number D17PC00345. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied of IARPA, DOI/IBC or the U.S. Government.

References

- [1] Jintai Chen, Biwen Lei, Qingyu Song, Haochao Ying, Danny Z Chen, and Jian Wu. A hierarchical graph network for 3d object detection on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 392–401, 2020. 3
- [2] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 5
- [3] 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. 1, 3, 4, 5, 8
- [4] 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, 2017. 2
- [5] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018. 2, 4
- [6] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2, 3
- [7] Guangxing Han, Shiyuan Huang, Jiawei Ma, Yicheng He, and Shih-Fu Chang. Meta faster r-cnn: Towards accurate few-shot object detection with attentive feature alignment. *arXiv preprint arXiv:2104.07719*, 2021. 3
- [8] Guangxing Han, Xuan Zhang, and Chongrong Li. Revisiting faster r-cnn: A deeper look at region proposal network. In *ICONIP*, pages 14–24, 2017. 2
- [9] Guangxing Han, Xuan Zhang, and Chongrong Li. Single shot object detection with top-down refinement. In *ICIP*, pages 3360–3364, 2017. 2
- [10] Guangxing Han, Xuan Zhang, and Chongrong Li. Semi-supervised dff: Decoupling detection and feature flow for video object detectors. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1811–1819, 2018. 2
- [11] 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. 2
- [12] 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. 3
- [13] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2014. 4
- [14] 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. 1, 3, 6, 7
- [15] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3588–3597, 2018. 4
- [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 International Conference on Computer Vision*, pages 8420–8429, 2019. 3, 5, 6, 8
- [17] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *Eighth International Conference on Learning Representations (ICLR)*, 2020. 2
- [18] 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 Conference on Computer Vision and Pattern Recognition*, pages 5197–5206, 2019. 3
- [19] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017. 3
- [20] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 7
- [21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 2
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [23] 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. 2
- [24] Yong Liu, Ruiping Wang, Shiguang Shan, and Xilin Chen. Structure inference net: Object detection using scene-level context and instance-level relationships. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6985–6994, 2018. 3, 4
- [25] Jiawei Ma, Hanchen Xie, Guangxing Han, Shih-Fu Chang, Aram Galstyan, and Wael Abd-Almageed. Partner-assisted learning for few-shot image classification. *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 2
- [26] Li Mi and Zhenzhong Chen. Hierarchical graph attention network for visual relationship detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13886–13895, 2020. 3
- [27] Megha Nawhal and Greg Mori. Activity graph transformer for temporal action localization. *arXiv preprint arXiv:2101.08540*, 2021. 3
- [28] Anton Osokin, Denis Sumin, and Vasily Lomakin. Os2d: One-stage one-shot object detection by matching anchor features. In *European Conference on Computer Vision*, 2020. 3

- [29] Juan-Manuel Perez-Rua, Xiatian Zhu, Timothy M Hospedales, and Tao Xiang. Incremental few-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13846–13855, 2020. 1, 3
- [30] 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. 2
- [31] 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. 1, 2, 3
- [32] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 2
- [33] 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. 2
- [34] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 2
- [35] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018. 3
- [36] 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. 2
- [37] Xin Wang, Thomas E. Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. In *International Conference on Machine Learning (ICML)*, July 2020. 1, 2, 5, 8
- [38] Xiaolong Wang, Yufei Ye, and Abhinav Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6857–6866, 2018. 3
- [39] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Meta-learning to detect rare objects. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9925–9934, 2019. 1, 3, 8
- [40] 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. 1, 2, 8
- [41] Yang Xiao and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. In *European Conference on Computer Vision*, 2020. 3, 4, 5, 8
- [42] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 3
- [43] 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 International Conference on Computer Vision*, pages 9577–9586, 2019. 3, 4, 5, 8
- [44] Yukuan Yang, Fangyu Wei, Miaoqing Shi, and Guoqi Li. Restoring negative information in few-shot object detection. In *Advances in neural information processing systems*, 2020. 3
- [45] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7094–7103, 2019. 3, 4