

Boosting Few-Shot Open-Set Object Detection via Prompt Learning and Robust Decision Boundary

Zhaowei Wu¹, Binyi Su^{2,3}, Qichuan Geng⁴, Hua Zhang⁵ and Zhong Zhou^{1,6,*}

¹State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, China

²The School of Artificial Intelligence and Data Science, Hebei University of Technology, China

³China Xiongan Group Digital City Technology Co., Ltd.

⁴Information Engineering College, Capital Normal University, China

⁵Institute of Information Engineering, Chinese Academy of Sciences

⁶Zhongguancun Laboratory, Beijing, China

{wuzhaowei, zz}@buaa.edu.cn, subinyi@hebut.edu.cn, gengqichuan1989@cnu.edu.cn, zhanghua@iie.ac.cn

Abstract

Few-shot Open-set Object Detection (FOOD) poses a challenge in many open-world scenarios. It aims to train an open-set detector to detect known objects while rejecting unknowns with scarce training samples. Existing FOOD methods are subject to limited visual information, and often exhibit an ambiguous decision boundary between known and unknown classes. To address these limitations, we propose the first prompt-based few-shot open-set object detection framework, which exploits additional textual information and delves into constructing a robust decision boundary for unknown rejection. Specifically, as no available training data for unknown classes, we select pseudo-unknown samples with Attribution-Gradient based Pseudo-unknown Mining (AGPM), which leverages the discrepancy in attribution gradients to quantify uncertainty. Subsequently, we propose Conditional Evidence Decoupling (CED) to decouple and extract distinct knowledge from selected pseudo-unknown samples by eliminating opposing evidence. This optimization process can enhance the discrimination between known and unknown classes. To further regularize the model and form a robust decision boundary for unknown rejection, we introduce Abnormal Distribution Calibration (ADC) to calibrate the output probability distribution of local abnormal features in pseudo-unknown samples. Our method achieves superior performance over previous state-of-the-art approaches, improving the average recall of unknown class by 7.24% across all shots in VOC10-5-5 dataset settings and 1.38% in VOC-COCO dataset settings. Our source code is available at https://gitee.com/VR_NAVE/ced-food.

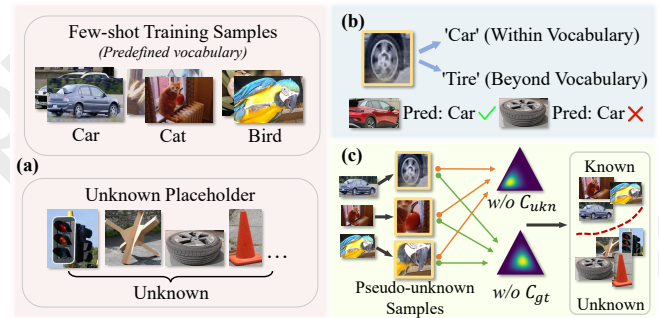


Figure 1: The challenge of unknown classes in the image-text joint space and our solution. (a) There are still numerous unknown objects beyond the predefined vocabulary in real-world scenarios. (b) Our intuition is that region proposals with high uncertainty (yellow border) consist of features from both known and unknown classes. (c) Our method decouples and learns distinct information from pseudo-unknown samples to construct a discriminative decision boundary.

1 Introduction

Object detection [Ren *et al.*, 2015; Redmon *et al.*, 2016; Lin *et al.*, 2017] has achieved remarkable success, enabling downstream tasks by leveraging large-scale labeled training data under the close-set setting, where training and testing sets share the same classes. However, in real-world scenarios such as safe autonomous driving, training data typically follows a long-tail distribution and many unknown objects may lack clear definitions, which could potentially lead to serious safety risks with close-set detector. To address these limitations, Few-shot Open-set Object Detection (FOOD) [Su *et al.*, 2024] enables the detection of known classes and rejection of unknowns with only few-shot close-set training data, which breaks the conventional assumption of identical class labels in training and testing. Despite its potential, this task is highly challenging due to limited training data and the absence of labels for unknown objects, resulting in weak generalization for unknown rejection and low recall rates.

Existing FOOD methods have utilized weight sparsifica-

*Corresponding author.



Figure 2: The detector misidentifies the zebra as horses (**left**). The detector misses the Czech hedgehog (**middle**). The detector successfully detects the known and rejects the unknown (**right**).

tion [Su *et al.*, 2024] or moving weight average [Su *et al.*, 2023] to facilitate the generalization for unknown rejection. However, the decision boundary between known and unknown classes is still ambiguous due to limited visual information. This ambiguity often leads to the misclassification of unknown classes as known ones with high confidence, resulting in a relatively low recall rate. Recently, the vision-language models [Radford *et al.*, 2021] introduce textual modality as complementary for downstream tasks through prompt learning [Zhou *et al.*, 2022b; Zhou *et al.*, 2022a], shedding light on few-shot out-of-distribution image classification [Miyai *et al.*, 2023; Li *et al.*, 2024] and open vocabulary detection [Zhong *et al.*, 2022; Minderer *et al.*, 2022]. However, there are still countless objects that remain outside the scope of the predefined vocabulary in the real world, as depicted in Fig. 1(a). For instance, assuming the current label consists of the vocabulary [‘car’, ‘bird’, ‘horse’], applying it to the two images in Fig. 2 may lead to false or missed detection. Ideally, the model should distinguish between “*what is known*” and “*what is unknown*”, as shown in Fig. 2 (right). Therefore, combining prompt learning with the construction of a robust decision boundary could enhance the FOOD detector for unknown rejection.

In this paper, we develop the first prompt-based FOOD framework for unknown rejection. As no training data is available for unknown classes, we select pseudo-unknown samples from region proposals for unknown class optimization. Drawing inspiration from the gradient-based attribution method [Chen *et al.*, 2023] for uncertainty estimation, we propose an Attribution-Gradient based Pseudo-unknown Mining (AGPM) method to select pseudo-unknown samples. It benefits from the diverse interpretative capabilities of different texts for the same image, which is reflected in the attribution gradient differences within the large model. However, these pseudo-unknown samples often contain a mixture of known and unknown properties, which cannot fit the real unknown distribution. As shown in Fig. 1(b), the region proposal could contain features of both the car within the vocabulary and the unknown tire, causing an ambiguous decision boundary between knowns and unknowns. To alleviate this problem, the proposed Conditional Evidence Decoupling (CED) method decouples known and unknown properties from pseudo-unknown samples. This approach is derived from the uncertainty mining property of Evidential Deep Learning [Sensoy *et al.*, 2018] while removing the evidence influence of the ground truth class, as shown in Fig. 1(c). Furthermore, to reduce the impact of local abnormal features on the final decision, the proposed Abnormal Distribution Calibration (ADC) method adjusts the output probability distribution of local features to regularize the model

and strengthen the unknown decision boundary. Experimental results demonstrate the superiority of our method on both known and unknown classes. We summarize our main contributions as follows:

- To the best of our knowledge, this is the first work to employ prompt learning for the FOOD task. We propose an Attribution-Gradient based Pseudo-unknown Mining method by innovatively quantifying the text interpretative capabilities in the image-text joint space to mine high uncertainty region proposals.
- We propose Conditional Evidence Decoupling to extract distinct information from pseudo-unknown samples. To enhance robustness, we introduce Abnormality Distribution Calibration to regularize the model and shape a discriminative unknown decision boundary.
- Extensive experiments demonstrate the effectiveness of our FOOD method, which outperforms previous vision-only frameworks and achieves superior unknown rejection performance.

2 Related Work

Vision Language Models (VLMs) [Radford *et al.*, 2021; Zhong *et al.*, 2022] leverage extensive image-text pre-training data, enabling zero-shot detection of desired classes based on given vocabulary, and they could quickly adapt to downstream tasks with prompt learning [Zhou *et al.*, 2022b; Zhou *et al.*, 2022a]. VLMs use predefined vocabulary to extend their understanding of the real world. However, real-world data often follows a long-tail distribution, where many rare or non-typical objects are difficult to define and cannot be effectively addressed by Open-Vocabulary Detection (OVD) [Zareian *et al.*, 2021; Zhong *et al.*, 2022; Minderer *et al.*, 2022; Liu *et al.*, 2023]. We tackle this challenge by FOOD, which ensures accurate recognition of known classes in vocabulary while preventing false detections. It is valuable for advancing the robustness and versatility of vision-language models in handling complex and diverse real-world scenarios.

Pseudo-unknown Sample Mining. There are no training samples for unknown classes, thus, the goal of pseudo-unknown sample mining is to select highly uncertain samples from foreground and background proposals for subsequent optimization of unknown classes. [Han *et al.*, 2022] used a maximum entropy for pseudo-unknown sample mining in open-set object detection. [Su *et al.*, 2024] employed maximum conditional energy in few-shot open-set object detection, and in FOODv2 [Su *et al.*, 2023], they selected proposals with high evidence uncertainty as pseudo-unknown samples. While these methods all operated within the visual feature space, we explore pseudo-unknown sample mining with additional textual information.

Few-Shot Open-Set Object Detection. Compared with Few-Shot Open-Set Recognition (FSOSR) [Liu *et al.*, 2020; Jeong *et al.*, 2021; Wang *et al.*, 2023a; Boudiaf *et al.*, 2023; Nag *et al.*, 2023], the task of few-shot open-set object detection (FOOD) required fine-grained, region-level representations and cannot overlook the impact of background region

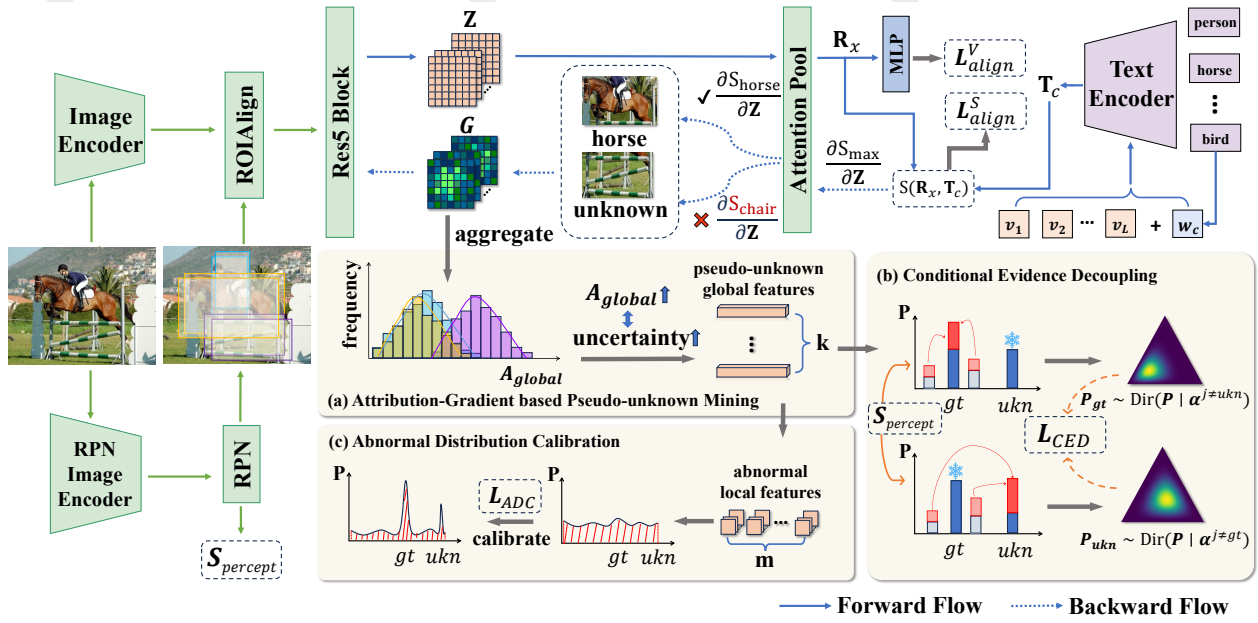


Figure 3: The overview architecture of our method, which is a two-stage detector with (a) Attribution-Gradient based Pseudo-unknown Mining, (b) Conditional Evidence Decoupling For Unknown Optimization, (c) Abnormal Distribution Calibration For Robust Decision Boundary.

proposals for the unknown rejection. [Su *et al.*, 2024] initially established a benchmark for the FOOD task. To enhance generalization for unknown classes, [Su *et al.*, 2023] proposed a Hilbert-Schmidt Independence Criterion (HSIC) based moving weight averaging technique to regularize the updating of model parameters. In this paper, we are dedicated to decoupling known and unknown information in pseudo-unknown samples with evidential deep learning to establish a robust decision boundary between known and unknown classes.

3 Method

Our method is a prompt-based few-shot open-set object detection framework including Attribution-Gradient based Pseudo-unknown Mining for uncertainty quantification, Conditional Evidence Decoupling for unknown optimization, and Abnormal Distribution Calibration for robust decision boundary. An overview of our method is shown in Fig. 3. RegionCLIP is adopted [Zhong *et al.*, 2022] as the base framework, composed of two image encoders, a separately trained region proposal network (RPN), and a text encoder. To alleviate the issue of traditional RPNs falsely being class-agnostic [Wang *et al.*, 2023b; Sarkar *et al.*, 2024], we train the RPN with a parallel branch to compute the centerness score [Tian *et al.*, 2020], which provides robust localization ability from object position and shape.

3.1 Preliminary

We formalize the FOOD task based on previous research [Su *et al.*, 2024] as training a detector with a class-imbalanced training dataset to accurately classify the K known classes $C_{Known} = C_{Base} \cup C_{Novel}$, reject all unknown classes C_{Ukn} , and distinguish between foreground and background according to C_{BG} .

Conventionally, the prompt learning method CoOp [Zhou *et al.*, 2022b] aligns the region feature \mathbf{R} with class-specific prompts \mathbf{T} , where context words of prompt templates (e.g., “a photo of a”) are replaced with continuously learnable parameters, denoted as $\mathbf{t}_c = \{v_1, v_2, \dots, v_L, w_c\}$. The semantic alignment loss is defined as:

$$L_{align}^S = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{K+2} y_{ij} \log \frac{\exp(S(\mathbf{R}_i, \mathbf{T}_j)/\tau)}{\sum_{c=1}^K \exp((S(\mathbf{R}_i, \mathbf{T}_c))/\tau)}, \quad (1)$$

where $S(\cdot, \cdot)$ represents the cosine similarity and τ denotes the temperature parameter, y_{ij} is an indicator (0 or 1) of sample i belonging to category j in the ground truth label.

Following [Han *et al.*, 2022], we implement enqueue/dequeue operations based on the memory bank and regularize the model with the following visual alignment loss to align the region feature with the same class:

$$L_{align}^V = \frac{1}{N} \sum_{i=1}^N L_{align}^V(\mathbf{z}_i), \quad (2)$$

$$L_{align}^V(\mathbf{z}_i) = \frac{1}{|Q(\mathbf{c}_i)|} \sum_{\mathbf{z}_j \in Q(\mathbf{c}_i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_j/\varepsilon)}{\sum_{\mathbf{z}_k \in Q \setminus Q_{\mathbf{c}_i}} \exp(\mathbf{z}_i \cdot \mathbf{z}_k/\varepsilon)}, \quad (3)$$

where \mathbf{z} is 128-dimensional latent embeddings mapped from \mathbf{R} , \mathbf{c}_i is the class label for the i -th proposal, ε is a hyperparameter, and $Q(\mathbf{c}_i)$ represents the embedding queue for class \mathbf{c}_i .

3.2 Attribution-Gradient based Pseudo-unknown Mining

Due to the lack of training data for unknown classes, it is difficult to establish a clear unknown decision boundary. To




Figure 4: Distribution of A_{global} across known, background, and unknown classes. Proposals are sampled from 500 randomly selected images in the VOC10-5-5 and the VOC-COCO testing set.

tackle the above issue, we select a subset of known proposals as pseudo-unknown samples, which may exhibit features of unknown classes. Inspired by the gradient-based attribution method, which was first introduced in the sensitivity analysis (SA) [Simonyan *et al.*, 2013], it evaluates the sensitivity of a particular input feature on the final prediction output for visual interpretability [Selvaraju *et al.*, 2017]. We propose a novel Attribution-Gradient based Pseudo-unknown Mining (AGPM) method to mine high-uncertainty pseudo-unknown samples, which are then employed to construct an unknown decision boundary. Specifically, we take the intermediate feature layer \mathbf{Z} (in Fig. 3) as the target layer. For a given proposal feature \mathbf{R}_x , we obtain the attribution gradient at \mathbf{Z}_{ij}^k corresponding to the maximum text-image matching score:

$$G_{ij}^k = \frac{\partial \max_{c=1 \dots K} S(\mathbf{R}_x, \mathbf{T}_c)}{\partial \mathbf{Z}_{ij}^k}, \quad (4)$$

where i , j , and k represent the indices of height, width, and channel, respectively. We can obtain the attribution gradient map \mathbf{G} corresponding to different proposals. Consequently, we perform global aggregation of attribution gradients as follows:

$$A_{global} = \frac{1}{C} \sum_k \left(\sum_i \sum_j \gamma_{ij}^k \right) \cdot \left(\sum_i \sum_j |G_{ij}^k| \right), \quad (5)$$

where γ_{ij}^k is an indicator function such that $\gamma_{ij}^k = 1$ if $G_{ij}^k \neq 0$ and $\gamma_{ij}^k = 0$ if $G_{ij}^k = 0$, and $|\cdot|$ denotes the absolute function, resulting in a scalar aggregated outcome. This result could serve as a metric for quantifying uncertainty and assessing the differences between known and unknown classes.

We then analyze the distributions of A_{global} for known, background, and unknown classes with all labels available, identifying distinct distribution patterns, as shown in Fig. 4. Under the premise of having only known class labels, a higher A_{global} aligns more closely with the distribution characteristic of unknown classes. Therefore, we select the proposals corresponding to the top- k highest A_{global} from the foreground and background proposals as pseudo-unknown samples with sampling ratio $S_{fg/bg}$.

3.3 Conditional Evidence Decoupling For Unknown Optimization

For FOOD, the unknown objects are easily misclassified into known ones with a high confidence score, which could be at-

tributed to its coupling of known and unknown information. To decouple and learn distinct information from the pseudo-unknown samples, we reserve a placeholder beyond the vocabulary for unknown classes and model the relationship between known and unknown classes based on conditional evidence. Specifically, we employ Evidential Deep Learning (EDL) [Sensoy *et al.*, 2018] based on the evidence framework of Dempster-Shafer Theory (DST) [Sentz and Ferson, 2002] and the subjective logic (SL) [Jøsang, 2016] to estimate uncertainty. By assuming that output probabilities \mathbf{P} follow a Dirichlet distribution, denoted as $\mathbf{P} \sim \text{Dir}(\mathbf{P} | \boldsymbol{\alpha})$, EDL constructs distribution of distributions for uncertainty modeling. Drawing on the DST and SL theory, for a classifier with $K + 2$ classes, we denote $\exp(l_i^j)$ as the evidence output for the j -th class from the i -th proposal, where $l_i^j = S(\mathbf{R}_i, \mathbf{T}_j) / \tau$. Consequently, this allows deriving the parameters for the Dirichlet distribution:

$$\alpha_i^j = \exp(l_i^j) + 1. \quad (6)$$

To extract distinct knowledge from identical features, we optimize evidence for known and unknown classes separately. Furthermore, we eliminate the evidence of the ground-truth class while optimizing for the unknown class, and conversely for known classes. In this case, we can alleviate the performance degradation caused by the contradictory evidence of decoupled classes. We formalize this as a conditional EDL loss in the following form:

$$\mathbf{L}_i^{ukn} = \psi \left(\sum_{j=1, j \neq gt}^{K+2} \alpha_i^j \right) - \psi(\alpha_i^{ukn}), \quad (7)$$

$$\mathbf{L}_i^{gt} = \psi \left(\sum_{j=1, j \neq unk}^{K+2} \alpha_i^j \right) - \psi(\alpha_i^{gt}), \quad (8)$$

where $\psi(\cdot)$ represents the digamma function, \mathbf{L}^{ukn} and \mathbf{L}^{gt} optimize the evidence for the known and unknown classes, respectively. Subsequently, we use the object perception score as a weight factor to balance the optimization between known and unknown classes, which is derived from the geometric mean of the original objectness score S_{obj} and the centerness score S_{center} from RPN:

$$S_{percept} = \sqrt{S_{obj} \cdot S_{center}}. \quad (9)$$

Intuitively, a higher $S_{percept}$ of foreground proposals indicates more known information, thereby increasing the weight for optimizing known classes, conversely for background proposals. Consequently, we derive the following foreground and background conditional evidence decoupling losses:

$$\mathbf{L}_{CED}^{fg} = \frac{1}{N} \sum_{i=1}^N (1 - S_{percept_i}) \cdot \mathbf{L}_i^{ukn} + S_{percept_i} \cdot \mathbf{L}_i^{gt}, \quad (10)$$

$$\mathbf{L}_{CED}^{bg} = \frac{1}{N} \sum_{i=1}^N S_{percept_i} \cdot \mathbf{L}_i^{ukn} + (1 - S_{percept_i}) \cdot \mathbf{L}_i^{gt}. \quad (11)$$

To this end, the final loss expression is as follows:

$$\mathbf{L}_{CED} = \mathbf{L}_{CED}^{fg} + \mathbf{L}_{CED}^{bg}. \quad (12)$$




Figure 5: Scatter plots of known and unknown classes on the VOC07+12trainval dataset. Each point represents a local feature \mathbf{Z}_{xy} from intermediate output \mathbf{Z} . Unknown classes show twice as many local aggregated gradient outliers (threshold > 0.0002) as known and background classes, with 100 selected proposals per plot.

3.4 Abnormal Distribution Calibration For Robust Decision Boundary

By employing conditional evidence decoupling, the detector can distinguish known and unknown classes using global features. However, certain local anomalous features \mathbf{Z}_{xy} still pose a disruption to the decision-making process. Therefore, we delve into the reasons for the differences in global attribution gradient distributions by aggregating local attribution gradients. We observe that, compared to known classes, unknown classes exhibit a greater number of outliers in locally aggregated attribution gradients, as shown in Fig. 5. For the attribution gradient map G , we perform aggregation along the channel dimension, resulting in local aggregation as follows:

$$A_{local} = \frac{1}{C} \sum_k |G_{xy}^k|, \quad (13)$$

for each local position (x, y) , A_{local} is a scalar, C is the total number of channels. We posit that outlier gradients reflect uncertain local features that hinder global discrimination between known and unknown classes. To address this, we recalibrate the probability distribution by reducing non-ground-truth logits to mitigate overconfidence. Specifically, we project the pseudo-unknown local features \mathbf{Z}_{xy} into the image-text joint space. The match scores between local and textual features are then computed to obtain the local output logits l' . These logits are then adjusted using the following abnormality distribution calibration loss to recalibrate the local output distribution:

$$\begin{aligned} L_{ADC} = & -\frac{1}{M} \sum_i \left(\sum_{j=1, j \neq gt}^K \log \frac{\exp(-l_i^j)}{1 + \exp(-l_i^j)} \right. \\ & \left. + H_{norm}(\mathbf{p}') \cdot \log \frac{1}{1 + \exp(-l_i^{u_{kn}})} \right), \end{aligned} \quad (14)$$

where $H_{norm}(\mathbf{p}) = -\sum_c p'_c \log p'_c / \log(K)$ represents the normalized entropy, serving as a weighting factor to constrain the learning of the unknown class. For each pseudo-unknown sample, we select the top- m highest A_{local} local features to recalibrate the output probability distribution,

3.5 Overall Optimization

We adopt a two-stage fine-tuning strategy [Wang *et al.*, 2020] to train the few-shot open-set detector, for the base training

stage:

$$\mathbf{L}_{base} = \mathbf{L}_{reg} + \mathbf{L}_{align}^S + \gamma_t \mathbf{L}_{align}^V, \quad (15)$$

and for the few-shot fine-tuning stage:

$$\mathbf{L}_{novel} = \mathbf{L}_{reg} + \mathbf{L}_{align}^S + \gamma_t \mathbf{L}_{align}^V + \lambda_t (\mathbf{L}_{CED} + \beta \mathbf{L}_{ADC}), \quad (16)$$

where \mathbf{L}_{reg} is smooth L1 loss for box regression, γ_t is a step-wise decreasing weight strategy similar to [Han *et al.*, 2022], β is a hyperparameter and $\lambda_t = \exp(\log(\lambda) \cdot (1 - t/T)) \in [\lambda, 1]$ denotes the weight that changes exponentially with the current iteration (t) and the total iteration (T), of which intention is to first learn well-defined semantic clusters and then gradually establish decision boundary between known and unknown classes.

4 Experiments

4.1 Experimental Details

1) *Datasets*: Following [Su *et al.*, 2024], the data splits VOC10-5-5, VOC-COCO, and COCO-RoadAnomaly [Lis *et al.*, 2019] are used for performance evaluation. **VOC10-5-5** includes 10 base, 5 novel, and 5 unknown classes from PASCAL VOC [Everingham *et al.*, 2010]. **VOC-COCO** has 20 base classes from PASCAL VOC, 20 novel classes from non-overlapping MS COCO [Lin *et al.*, 2014], and 40 unknown classes. **COCO-RoadAnomaly** evaluates model generalization in open-set road scenes.

2) *Evaluation Metrics*: For the FOOD task, the mean Average Precision (mAP) of known classes (mAP_K) and novel classes (mAP_N) are adopted as known class metrics. For unknown class metrics, the recall (R_U) and average recall (AR_U) of unknown classes are reported as in [Su *et al.*, 2023]. Furthermore, we report Wilderness Impact (WI) under a recall level of 0.8 to measure the degree of unknown objects misclassified to known ones: $WI = \frac{P_K}{P_{K \cup U}} - 1$, and Absolute Open-Set Error ($AOSE$) to count the number of misclassified unknown objects as in [Han *et al.*, 2022].

3) *Implementation Details*: We employ RegionCLIP [Zhong *et al.*, 2022] as the image encoder, and ResNet-50 [He *et al.*, 2016] pre-trained on ImageNet as the RPN image encoder. Class-specific prompt training follows CoOp [Zhou *et al.*, 2022b] with a context length of 16, using a two-stage training strategy [Wang *et al.*, 2020], where the base training phase trains a base detector using C_{base} , and the few-shot fine-tuning phase refines the model with a small, balanced set comprising both C_{base} and C_{novel} . We employ SGD with 0.9 momentum, 5e-5 weight decay, and a batch size of 1 on a GTX 1080 Ti GPU. The learning rate is 0.0002 for base training and 0.0001 for fine-tuning. Visual alignment loss settings follow the setting in [Han *et al.*, 2022]. Other hyperparameters τ , ε , λ , and β are 0.01, 0.1, 1e-4, and 1.0, respectively.

4) *Baselines*: We compare vision only open-set framework including: ORE [Joseph *et al.*, 2021], PROSER [Zhou *et al.*, 2021], OPENDET [Han *et al.*, 2022], FOOD [Su *et al.*, 2024], and FOODv2 [Su *et al.*, 2023]. Furthermore, we employ OPENDET and FOODv2 within our open-set detection framework, denoted by OPENDET(+Ours) and FOODv2(+Ours) to conduct a fair comparison. We employ max entropy and max evidential uncertainty as pseudo-unknown sampling methods respectively, with unknown

Method	1-shot			3-shot		
	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$
ORE [Joseph <i>et al.</i> , 2021]	43.25 / 8.62	18.25 / –	9.54 / 930.30	45.88 / 14.52	22.23 / –	9.88 / 1058.70
PROSER [Zhou <i>et al.</i> , 2021]	41.64 / 8.49	30.95 / 15.41	11.15 / 994.60	43.30 / 15.16	32.30 / 16.17	10.45 / 1021.70
OPENDET [Han <i>et al.</i> , 2022]	43.45 / 8.27	33.64 / 17.28	10.47 / 867.30	46.47 / 14.09	30.62 / 15.89	9.27 / 954.50
FOOD [Su <i>et al.</i> , 2024]	43.97 / 8.95	43.72 / 23.51	6.96 / 598.60	48.48 / 16.83	44.52 / 23.58	7.83 / 859.00
FOODv2 [Su <i>et al.</i> , 2023]	45.12 / 11.56	60.03 / 31.19	– / –	48.90 / 18.96	61.21 / 32.02	– / –
OPENDET(+Ours)	50.28 / 18.40	<u>78.56 / 36.76</u>	5.89 / 781.60	<u>55.61 / 33.03</u>	78.87 / 38.43	<u>4.75 / 547.20</u>
FOODv2(+Ours)	53.71 / 22.62	77.28 / 34.70	<u>5.65 / 1042.20</u>	56.53 / 35.65	80.19 / 36.80	5.52 / 949.80
Ours	<u>51.94 / 21.43</u>	79.88 / 38.12	4.12 / 459.60	53.09 / 31.70	80.55 / 39.53	3.72 / 451.20

Method	5-shot			10-shot		
	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$
ORE [Joseph <i>et al.</i> , 2021]	46.29 / 18.49	23.01 / –	10.16 / 1019.70	48.17 / 25.40	23.48 / –	9.65 / 1063.70
PROSER [Zhou <i>et al.</i> , 2021]	45.12 / 20.08	32.68 / 16.48	10.65 / 1009.80	48.35 / 25.13	32.61 / 17.01	10.29 / 956.70
OPENDET [Han <i>et al.</i> , 2022]	47.56 / 17.90	32.13 / 16.72	9.01 / 1031.50	50.95 / 25.14	36.30 / 18.89	8.50 / 1021.40
FOOD [Su <i>et al.</i> , 2024]	50.18 / 23.10	45.65 / 23.61	7.59 / 908.00	53.23 / 28.60	45.84 / 23.86	6.99 / 900.20
FOODv2 [Su <i>et al.</i> , 2023]	52.55 / 27.31	62.02 / 32.79	– / –	57.24 / 32.63	62.14 / 32.80	– / –
OPENDET(+Ours)	56.01 / 36.57	79.70 / 39.42	4.53 / 519.40	58.70 / 42.69	74.60 / 37.16	4.90 / 530.60
FOODv2(+Ours)	<u>55.13 / 38.28</u>	<u>80.62 / 37.05</u>	4.98 / 1185.60	60.84 / 45.56	79.45 / 37.17	<u>4.12 / 953.30</u>
Ours	54.35 / <u>36.67</u>	81.37 / 40.32	3.78 / 512.20	58.55 / <u>43.52</u>	<u>79.39 / 39.79</u>	3.43 / 546.30

Table 1: Few-Shot Open-Set Object Detection results on VOC10-5-5. ‘(+Ours)’ indicates the implementation with our prompt-based FOOD framework while ‘Ours’ denotes our framework with all of our methods. **Bold** indicates the best, underlined indicates the second best.

Method	1-shot			5-shot		
	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$
ORE [Joseph <i>et al.</i> , 2021]	14.14 / 2.18	4.59 / –	12.08 / 1087.00	16.21 / 6.29	4.99 / –	12.30 / 1344.00
PROSER [Zhou <i>et al.</i> , 2021]	13.58 / 2.32	7.53 / 3.07	11.68 / 925.30	15.67 / 6.40	9.59 / 4.08	12.56 / 1165.90
OPENDET [Han <i>et al.</i> , 2022]	16.01 / 2.29	7.24 / 3.14	9.82 / 690.90	17.16 / 6.56	11.49 / 5.21	9.55 / 1176.90
FOOD [Su <i>et al.</i> , 2024]	15.83 / 2.26	15.76 / 7.20	6.78 / 485.00	18.08 / 6.69	20.02 / 9.45	7.37 / 859.00
FOODv2 [Su <i>et al.</i> , 2023]	18.54 / 4.33	30.87 / 14.13	– / –	19.88 / 11.95	32.53 / 15.74	– / –
OPENDET(+Ours)	18.42 / 4.42	<u>36.70 / 16.17</u>	5.42 / 796.80	20.42 / 12.23	<u>39.10 / 17.89</u>	4.83 / 742.40
FOODv2(+Ours)	20.44 / 5.69	36.25 / 15.74	5.14 / 945.40	<u>21.12 / 12.47</u>	39.05 / 16.72	4.70 / 835.90
Ours	<u>19.49 / 5.41</u>	38.53 / 16.68	<u>4.51 / 638.70</u>	21.46 / 13.24	40.52 / 17.91	2.99 / 808.90

Method	10-shot			30-shot		
	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$	$mAP_K / mAP_N \uparrow$	$R_U / AR_U \uparrow$	$WI / AOSE \downarrow$
ORE [Joseph <i>et al.</i> , 2021]	17.98 / 8.75	5.13 / –	11.65 / 1463.20	23.07 / 15.17	5.51 / –	11.22 / 1867.00
PROSER [Zhou <i>et al.</i> , 2021]	17.00 / 8.75	10.06 / 4.89	12.47 / 1160.00	21.44 / 14.30	12.06 / 5.98	12.00 / 1561.60
OPENDET [Han <i>et al.</i> , 2022]	18.53 / 8.70	13.89 / 6.32	9.83 / 1400.60	22.93 / 14.02	18.07 / 8.76	9.02 / 1818.00
FOOD [Su <i>et al.</i> , 2024]	20.17 / 9.48	21.48 / 9.56	7.59 / 1099.30	23.90 / 14.17	23.17 / 11.45	8.13 / 1480.00
FOODv2 [Su <i>et al.</i> , 2023]	22.64 / 13.82	32.78 / 16.52	– / –	23.71 / 17.67	35.74 / 17.26	– / –
OPENDET(+Ours)	22.74 / 15.34	<u>38.12 / 17.72</u>	5.12 / 934.30	25.34 / <u>21.56</u>	38.78 / 16.68	4.97 / 1463.6
FOODv2(+Ours)	24.42 / 16.83	37.33 / 16.04	4.38 / 1046.60	26.70 / 22.73	39.46 / 17.24	<u>4.23 / 1442.40</u>
Ours	<u>23.75 / 16.77</u>	38.69 / 17.06	2.58 / 856.40	<u>25.72 / 21.16</u>	<u>39.43 / 17.52</u>	2.46 / 1339.30

Table 2: Few-Shot Open-Set Object Detection results on VOC-COCO. ‘(+Ours)’ indicates the implementation with our prompt-based FOOD framework while ‘Ours’ denotes our framework with all of our methods. **Bold** indicates the best, underlined indicates the second best.

probability loss and IoU loss as optimization strategies for unknown classes respectively.

4.2 Main results

1) *Experiments on VOC10-5-5*: Table 1 presents the FOOD results on VOC10-5-5, where we report the results of fine-tuning on 1, 3, 5, and 10 shots, averaging ten runs per setting for a fairer comparison. Based on our framework, both OPENDET(+Ours) and FOODv2(+Ours) show significant improvements compared to their original versions, demonstrating the advantages of our open-set framework. Compared to previous state-of-the-art methods with traditional open-set framework, our approach (with $k = 3, S_{fg:bg} = 1 : 3, m = 1$) achieves significant improvement on unknown

class metrics, surpassing the second best (mean results) by **18.95% R_U , 7.24% AR_U , 3.58 WI and 324.13 $AOSE$** . Our method consistently achieves the state-of-the-art performance on most metrics related to unknown classes, despite a slight decrease in accuracy for known classes.

2) *Experiments on VOC-COCO*: Table 2 displays the FOOD results on VOC-COCO, which is a more challenging dataset. We report results of fine-tuning on 1, 5, 10, and 30 shots, averaging ten runs per shot setting to ensure a fair comparison. Both OPENDET(+Ours) and FOODv2(+Ours) also show significant improvements compared to the original frameworks. Compared to prior state-of-the-art methods with the traditional open-set framework, our approach (with $k = 3, S_{fg:bg} = 1 : 1, m = 1$) shows a marked improve-

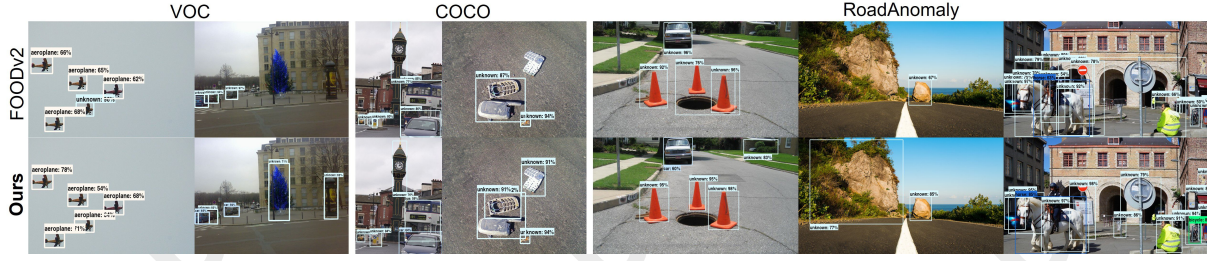


Figure 6: The visualized results on VOC, COCO, and RoadAnomaly datasets under 10-shot VOC10-5-5 setting. Our method recalls more unknown objects and better distinguishes between the knowns and the unknowns.

L_{align}^S	L_{align}^V	L_{CED}	L_{ADC}	$WI \downarrow$	$AOSE \downarrow$	$mAP_K \uparrow$	$R_U \uparrow$	$AR_U \uparrow$
✓				7.06	2314.10	57.79	0.00	0.00
✓	✓			6.83	2169.10	58.66	0.00	0.00
✓		✓		5.17	615.80	49.76	79.63	37.96
✓		✓	✓	4.95	567.60	51.93	79.38	37.94
✓	✓	✓		4.69	539.10	52.02	79.93	38.02
✓	✓	✓	✓	4.12	459.60	51.94	79.88	38.12

Table 3: The ablation study of proposed components.

Method	$WI \downarrow$	$AOSE \downarrow$	$mAP_K \uparrow$	$R_U \uparrow$	$AR_U \uparrow$
OPENDET [Han <i>et al.</i> , 2022]	10.47	867.30	43.45	33.64	17.28
OPENDET(+Ours)	<u>5.89</u>	781.60	<u>50.28</u>	<u>78.56</u>	<u>36.76</u>
FOODv2 [Su <i>et al.</i> , 2023]	-	-	45.12	60.03	31.19
FOODv2(+Ours)	5.65	1042.20	53.71	77.28	34.70
FOODv1(+Ours)(P) [Su <i>et al.</i> , 2024]	4.57	896.00	52.39	78.90	35.54
FOODv2(+Ours)(P)	4.87	<u>884.00</u>	<u>52.95</u>	78.79	34.35
GAIA-Z(+Ours) [Chen <i>et al.</i> , 2023]	4.47	948.60	53.69	79.64	35.86
Ours(AGPM)	4.12	459.60	51.94	79.88	38.12
FOODv2(+Ours)(U)	5.34	706.10	52.58	79.32	37.35
Non-Decoupled	4.22	424.20	50.67	79.38	37.87
Non-Conditional	3.73	538.40	48.98	82.12	38.45
Original L_{CED}	4.12	459.60	51.94	79.88	38.12

Table 4: The main contributions analysis.

ment, outperforming the second best (mean results) by **6.31%** R_U , **1.38%** AR_U , **4.42** WI and **70.00** $AOSE$. The 1-shot $AOSE$ performance does not surpass previous benchmarks, likely due to the strong learning capability of prompt-based methods with limited samples, which incurs overfitting of known classes. Our method generally achieves higher AR_U and lower WI and $AOSE$, effectively rejecting unknown objects rather than misclassifying them as known.

4.3 Ablation Studies

We ablate the proposed losses with the 1-shot VOC 10-5-5 setting in Tab. 3. By employing attribution gradients to filter pseudo-unknown samples, the proposed L_{CED} establishes a discriminative decision boundary between known and unknown classes through decoupled evidential learning. The regularization with L_{ADC} yields improved WI and $AOSE$ without adversely affecting mAP_K , indicating its facilitation for the decision boundary construction.

We conduct a more detailed ablation analysis in Tab. 4, ‘(+Ours)’ denotes our open-set detection framework, while ‘(P)’ and ‘(U)’ represent the Pseudo-unknown sample mining

method and the Unknown class optimization method from the corresponding paper, respectively. The top section demonstrates the advantages of our framework. The middle section reports the experimental results of different pseudo-unknown sample mining methods. Our AGPM shows better overall performance. The bottom section presents the results of various optimization methods for unknown classes. When we apply only Eq. 7 to optimize unknown classes (Non-Decoupled), it results in a degradation of accuracy for known classes. Subsequently, we remove the conditions $j \neq gt$ in Eq. 7 and $j \neq unkn$ in Eq. 8 (Non-Conditional), which significantly reduces mAP_K . By retaining the conditions, we achieve a better trade-off between known and unknown metrics. For details on the parameter ablation study, please refer to the appendix.

4.4 Visualized results

We conduct visual comparisons between FOODv2 [Su *et al.*, 2023] and our proposed method in Fig. 6 under 10-shot VOC10-5-5 experimental setup. It reveals that our method successfully recalls more unknown objects across three open-set datasets and makes more accurate distinctions between known and unknown objects. For the VOC dataset, the two images on the left illustrate that FOODv2 mistakenly classifies an airplane as an unknown object, while our method correctly identifies it. The right demonstrates that our approach successfully detects the car within the vocabulary and rejects the unknown classes (tree and billboard). This suggests that our approach enables enhanced perception of object presence and facilitates the unknown rejection.

5 Conclusion

In this paper, we propose a novel approach for addressing the challenging few-shot open-set detection problem. By introducing prompt learning to the FOOD task and incorporating an unknown class placeholder, our framework captures information beyond the known vocabulary. To address the unavailable training data of unknown classes, an attribution-gradient based method is proposed to mine high-uncertainty samples as pseudo-unknowns. Additionally, we propose a conditional evidence decoupling loss and a local abnormal distribution calibration loss to improve unknown class optimization and establish a robust decision boundary for unknown rejection. Extensive experiments demonstrate that our method achieves state-of-the-art performance, surpassing existing approaches.

Acknowledgments

This paper is supported by the National Key Research and Development Program of China (Grant No. 2024YFB3310904), the Science and Technology Project of Hainan Provincial Department of Transportation (Grant No. HNJTT-KXC-2024-3-22-02), the National Natural Science Foundation of China (Grant No. 62206184, 62272018, 62372448), the Basal Research Fund of Central Public Research Institute of China (Grant No.246Z4306G) and Zhongguancun Laboratory.

References

- [Boudiaf *et al.*, 2023] Malik Boudiaf, Etienne Bennequin, Myriam Tami, Antoine Toubhans, Pablo Piantanida, Celine Hudelot, and Ismail Ben Ayed. Open-set likelihood maximization for few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24007–24016, 2023.
- [Chen *et al.*, 2023] Jinggang Chen, Junjie Li, Xiaoyang Qu, Jianzong Wang, Jiguang Wan, and Jing Xiao. Gaia: Delving into gradient-based attribution abnormality for out-of-distribution detection. *Advances in Neural Information Processing Systems*, 36, 2023.
- [Everingham *et al.*, 2010] 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*, 88:303–338, 2010.
- [Han *et al.*, 2022] Jiaming Han, Yuqiang Ren, Jian Ding, Xingjia Pan, Ke Yan, and Gui-Song Xia. Expanding low-density latent regions for open-set object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9591–9600, 2022.
- [He *et al.*, 2016] 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.
- [Jeong *et al.*, 2021] Minki Jeong, Seokeon Choi, and Chang-ick Kim. Few-shot open-set recognition by transformation consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12566–12575, 2021.
- [Jøsang, 2016] Audun Jøsang. *Subjective logic*, volume 3. Springer, 2016.
- [Joseph *et al.*, 2021] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5830–5840, 2021.
- [Li *et al.*, 2024] Tianqi Li, Guansong Pang, Xiao Bai, Wenjun Miao, and Jin Zheng. Learning transferable negative prompts for out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17584–17594, 2024.
- [Lin *et al.*, 2014] 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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [Lin *et al.*, 2017] 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.
- [Lis *et al.*, 2019] Krzysztof Lis, Krishna Nakka, Pascal Fua, and Mathieu Salzmann. Detecting the unexpected via image resynthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2152–2161, 2019.
- [Liu *et al.*, 2020] Bo Liu, Hao Kang, Haoxiang Li, Gang Hua, and Nuno Vasconcelos. Few-shot open-set recognition using meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2020.
- [Liu *et al.*, 2023] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023.
- [Minderer *et al.*, 2022] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European Conference on Computer Vision*, pages 728–755. Springer, 2022.
- [Miyai *et al.*, 2023] Atsuyuki Miyai, Qing Yu, Go Irie, and Kiyoharu Aizawa. Locoop: Few-shot out-of-distribution detection via prompt learning. *Advances in Neural Information Processing Systems*, 36, 2023.
- [Nag *et al.*, 2023] Sayak Nag, Dripta S Raychaudhuri, Sujoy Paul, and Amit K Roy-Chowdhury. Reconstruction guided meta-learning for few shot open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [Redmon *et al.*, 2016] 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.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.

- [Sarkar *et al.*, 2024] Hiran Sarkar, Vishal Chudasama, Naoyuki Onoe, Pankaj Wasnik, and Vineeth N Balasubramanian. Open-set object detection by aligning known class representations. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 219–228, 2024.
- [Selvaraju *et al.*, 2017] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [Sensoy *et al.*, 2018] Murat Sensoy, Lance Kaplan, and Melih Kandemir. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31, 2018.
- [Sentz and Ferson, 2002] Kari Sentz and Scott Ferson. Combination of evidence in dempster-shafer theory. *Sandia National Lab.(SNL-NM), Albuquerque, NM (United States)*, 2002.
- [Simonyan *et al.*, 2013] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [Su *et al.*, 2023] Binyi Su, Hua Zhang, and Zhong Zhou. Hsic-based moving weight averaging for few-shot open-set object detection. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5358–5369, 2023.
- [Su *et al.*, 2024] Binyi Su, Hua Zhang, Jingzhi Li, and Zhong Zhou. Toward generalized few-shot open-set object detection. *IEEE Transactions on Image Processing*, 33:1389–1402, 2024.
- [Tian *et al.*, 2020] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: A simple and strong anchor-free object detector. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):1922–1933, 2020.
- [Wang *et al.*, 2020] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*, 2020.
- [Wang *et al.*, 2023a] Haoyu Wang, Guansong Pang, Peng Wang, Lei Zhang, Wei Wei, and Yanning Zhang. Glocal energy-based learning for few-shot open-set recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7507–7516, 2023.
- [Wang *et al.*, 2023b] Yanghao Wang, Zhongqi Yue, Xian-Sheng Hua, and Hanwang Zhang. Random boxes are open-world object detectors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6233–6243, 2023.
- [Zareian *et al.*, 2021] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021.
- [Zhong *et al.*, 2022] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Region-clip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022.
- [Zhou *et al.*, 2021] Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Learning placeholders for open-set recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2021.
- [Zhou *et al.*, 2022a] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825, 2022.
- [Zhou *et al.*, 2022b] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022.