

# Cross-Domain Few-Shot Object Detection via Enhanced Open-Set Object Detector

Yuqian Fu<sup>\*1,2,3</sup>, Yu Wang<sup>\*1</sup>, Yixuan Pan<sup>\*4</sup>, Lian Huai<sup>†5</sup>, Xingyu Qiu<sup>1</sup>, Zeyu Shangguan<sup>5</sup>, Tong Liu<sup>5</sup>, Yanwei Fu<sup>1</sup>, Luc Van Gool<sup>2,3</sup>, Xingqun Jiang<sup>5</sup>

<sup>1</sup>Fudan University, <sup>2</sup>ETH Zürich, <sup>3</sup>INSAIT, <sup>4</sup>Southeast University, <sup>5</sup>BOE Technology  
yuqian.fu@insait.ai, yu\_w13@fudan.edu.cn, yixuanpan@seu.edu.cn

**Abstract.** This paper studies the challenging cross-domain few-shot object detection (CD-FSOD), aiming to develop an accurate object detector for novel domains with minimal labeled examples. While transformer-based open-set detectors, such as DE-ViT, show promise in traditional few-shot object detection, their generalization to CD-FSOD remains unclear: 1) can such open-set detection methods easily generalize to CD-FSOD? 2) If not, how can models be enhanced when facing huge domain gaps? To answer the first question, we employ measures including style, inter-class variance (ICV), and indefinable boundaries (IB) to understand the domain gap. Based on these measures, we establish a new benchmark named CD-FSOD to evaluate object detection methods, revealing that most of the current approaches fail to generalize across domains. Technically, we observe that the performance decline is associated with our proposed measures: style, ICV, and IB. Consequently, we propose several novel modules to address these issues. First, the learnable instance features align initial fixed instances with target categories, enhancing feature distinctiveness. Second, the instance reweighting module assigns higher importance to high-quality instances with slight IB. Third, the domain prompter encourages features resilient to different styles by synthesizing imaginary domains without altering semantic contents. These techniques collectively contribute to the development of the Cross-Domain Vision Transformer for CD-FSOD (CD-ViTO), significantly improving upon the base DE-ViT. Experimental results validate the efficacy of our model. Datasets and codes are available at <http://yuqianfu.com/CDFSOD-benchmark>.

**Keywords:** Cross-Domain Few-Shot Learning · Few-Shot Object Detection · Open-Set Detector

## 1 Introduction

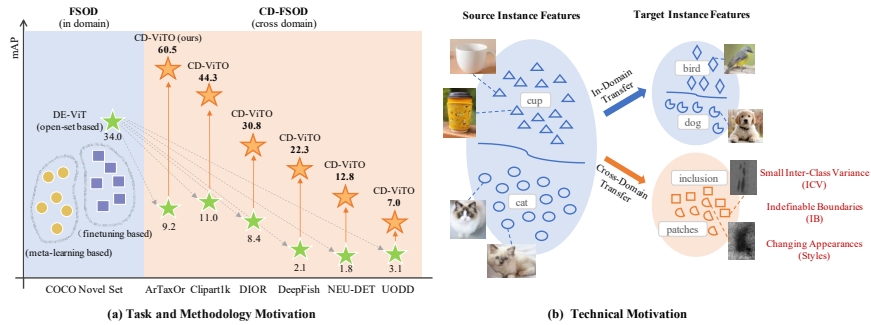
Cross-domain Few-Shot Learning (CD-FSL) [12] deals with learning across different domains, transferring knowledge from a source dataset to new target domains with few labeled data. While many methods focus on the classification

---

\*: equal contributions.

†: corresponding author.

Part of this work commenced during Dr. Yuqian Fu’s PhD at Fudan University.



**Fig. 1:** (a) Our motivation: The DE-ViT open-set detector excels in FSOD but struggles in CD-FSOD, inspiring our creation of CD-ViTO. (b) Technical motivation: FSOD models face challenges when dealing with cross-domain targets, such as small inter-class variance (ICV), indefinable boundaries (IB), and varying appearances (styles).

task in CD-FSL, they often overlook *object detection*. This paper delves into object detection tasks in CD-FSL, also known as Cross-Domain Few-Shot Object Detection (CD-FSOD).

CD-FSOD is derived from Few-Shot Object Detection (FSOD) [23], aiming to detect new objects with limited instances. Much effort has been dedicated to this objective. Current FSOD methods fall into two main categories: meta-learning based methods [13, 15, 48], which train a meta-learner adaptable to new objects, and finetuning based approaches [22, 31, 33, 39, 45], involving initial training followed by finetuning on few labeled data. Notably, DE-ViT [52], an open-set object detector, shows exceptional performance in FSOD, surpassing other methods as depicted in Fig. 1(a). This motivates us to explore if the detector is effective for the CD-FSOD task, posing two main questions: ① Can open-set detection methods generalize to cross-domain target datasets without performance degradation? ② If not, how can the open-set methods be improved in the presence of significant domain gaps?

To address question ①, we establish a new benchmark for CD-FSOD by rearranging existing object detection datasets. Our benchmark utilizes COCO [28] as source training set and incorporates six additional datasets, namely ArTaxOr [2], Clipart1K [18], DIOR [26], DeepFish [36], NEU-DET [38], and UODD [20], as novel target datasets. We also introduce metrics such as style, Inter-Class Variance (ICV), and Indefinable Boundaries (IB) to measure domain differences. Style represents different visual appearances, ICV measures dissimilarity between semantic labels, and IB reflects confusion between targets and backgrounds. Our target datasets exhibit significant variations in style, ICV, and IB, making our benchmark a comprehensive testbed for evaluating CD-FSOD methods. We evaluate various object detectors on this benchmark. As shown in Fig. 1(a), we observe obvious performance degradation even for the State-Of-The-Art (SOTA) FSOD detector, DE-ViT. This indicates the challenge of tackling CD-FSOD even with well-designed open-set detectors.

To tackle the performance degradation issue, i.e., question ②, we aim to upgrade the vanilla DE-ViT to improve the open-set detector’s performance. Our technical motivation, illustrated in Fig. 1(b), highlights that FSOD model degradation correlates with domain metrics we mentioned earlier: style, ICV, and IB. Specifically, a model trained on instances with coarse semantic concepts (large ICV) and clear foreground-background boundaries (slight IB), like "cup" and "cat", can easily transfer to similar in-domain concepts like "bird" and "dog". However, cross-domain cases, such as industrial defect types like "inclusion" and "patches", pose challenges with smaller ICV, more significant IB, and different styles. This insight inspires us to design a method that addresses issues of small ICV, significant IB, and changing styles in cross-domain scenarios.

To achieve this, we propose several novel modules: learnable instance features, an instance reweighting module, and a domain prompter. (1) Initially, to counter the negative effects of small ICV, we aim to increase feature discriminability making the target classes more distinguishable. This is achieved by our learnable instance features which adapt fixed initial instance features using target labels, aligning the features with semantic concepts. With explicit categories as supervision, the optimized features naturally exhibit greater dispersion than directly extracted ones. (2) Next, we introduce the instance reweighting module to assign different weights to instances, prioritizing those with slight IB. By emphasizing high-quality instances in forming class prototypes, we mitigate challenges posed by IB. (3) Further, to boost model robustness across domains, we propose a domain prompter to synthesize virtual "domains" and ensure prototype features remain stable under domain perturbations. To encourage the uniqueness of each synthesized domain, a domain diversity loss is proposed. Simultaneously, a contrastive learning loss and a classification loss are incorporated to maintain semantic content consistency in the presence of domain perturbations.

Furthermore, we also explore the finetuning technique. Although finetuning has been studied in cross-domain tasks [6, 12], we provide a detailed analysis of its effectiveness on our benchmark. By integrating these novel modules into DE-ViT, we introduce a new vision transformer based open-set detector called CD-ViTO for CD-FSOD. CD-ViTO substantially boosts DE-ViT’s performance across all target datasets, as shown in Fig. 1(a) (highlighted by the orange stars).

We summarize our contributions. 1) We propose a comprehensive benchmark for CD-FSOD that spans diverse target domains. Additionally, we introduce several metrics related to domain gap analysis, including style, inter-class variance, and indefinable boundaries. 2) Through extensive evaluations, we uncover the challenges posed by the domain gap issue in existing open-set object detectors and other object detection methods for CD-FSL. 3) Leveraging the existing DE-ViT, we present CD-ViTO, a novel method that incorporates finetuning, learnable instance features, an instance reweighting module, and a domain prompter. Notably, these new modules introduce minimal parameters and require only a few target support instances for training, ensuring high efficiency. 4) Our results demonstrate that CD-ViTO significantly improves upon the base DE-ViT, achieving new state-of-the-art performance in CD-FSOD.

## 2 Related Works

**Few-Shot Object Detection.** FSOD [23] identifies objects with few labeled data. Existing FSOD methods can be roughly divided into meta-learning based [13, 14, 15, 21, 48] and transfer learning based [4, 11, 22, 31, 33, 39, 45]. Meta-learning trains model across tasks for quick adaptation. Meta-RCNN [48] is a typical work of this type. Transfer learning finetunes pretrained model on limited data for few-shot tasks. Flagship works include TFA [45], FSCE [39], and DeFRCN [33]. Notably, a recently proposed method, DE-ViT [52], unifying the problem of FSOD and open-vocabulary object detection [49] and constructing an open-set detector for both tasks, shows great advantage. However, almost all of the methods mentioned above ideally assume that the training and testing sets belong to the same domain. Exceptions include Distill-cdfsod [47] and MoF-SOD [24]. Distill-cdfsod [47] introduces several datasets and a distillation-based method, but both the benchmark and method have room for improvement. MoF-SOD [24] also explores CD-FSOD and studies various model architectures, finetuning, and pre-trained datasets. Compared with them, apart from the datasets, we for the first time propose benchmark metrics and study the open-set detectors for CD-FSOD. Note that AcroFOD [9] and AsyFOD [8] also explore FSOD across domains. However, they share the same class set between source and target domains, while we handle entirely novel target categories.

**Cross-Domain Few-Shot Learning.** CD-FSL aims to improve few-shot learning (FSL) [29, 37, 40, 42, 51] across diverse domains. While there are several existing works [5, 6, 12, 17, 41, 43, 57] addressing classification tasks, there is a noticeable gap in tackling domain gap challenges across various few-shot vision tasks, including object detection [13, 48], semantic segmentation [44, 46], and action recognition [7, 50]. This paper specifically explores the area of CD-FSOD, seeking to develop a robust object detector. Despite recent advancements [9, 24, 47] in CD-FSOD, certain crucial aspects, such as domain gap metrics, comprehensive baseline evaluations, and technical improvements, remain less explored. Thus, we meticulously organize the benchmark, conduct an in-depth study on various existing methods, and introduce CD-ViTO as a novel method. Our goal is to propel the field of CD-FSOD forward, shedding light on unexplored aspects.

**Open-Vocabulary Object Detection (OVD).** Traditional object detection models like Fast R-CNN [10], YOLO [35], and ViTDet [27] recognize a predefined set of object categories. OVD [49] aims to extend this recognition to novel categories by leveraging the shared space between images and category text. This shared space can be established by the method itself [49] or by utilizing existing vision-language frameworks like CLIP [34], as shown in RegionCLIP [54], VL-PLM [53], and Detic [56]. Particularly, DE-ViT [52] is an open-set detector<sup>1</sup> capable of detecting objects of arbitrary classes. Despite DE-ViT relying solely on the visual modality without using the text modality in practice, its motiva-

<sup>1</sup> Despite varying definitions, we adopt the one of [52] that “both open-vocabulary and few-shot belong to open-set except their category representations”.

tion, method design, and experimental validations align with the OVD scope. Here, we made an exploration of open-set object detectors for CD-FSOD.

### 3 CD-FSOD: Setup, Metrics, and Benchmark

#### 3.1 Task Configurations

**Task Formulation.** Formally, given a source dataset denoted as  $\mathcal{D}_S = \{I, y\}, y \in \mathcal{C}_S$  with distribution  $\mathcal{P}_S$  and a novel target dataset  $\mathcal{D}_T = \{I, y\}, y \in \mathcal{C}_T$  with distribution  $\mathcal{P}_T$ , consistent with FSOD, CD-FSOD supposes the labeled data for the source classes  $\mathcal{C}_S$  is sufficient while each novel target class in  $\mathcal{C}_T$  has only a few labeled instances. All target classes must be novel to model i.e.,  $\mathcal{C}_S \cap \mathcal{C}_T = \emptyset$ . In addition, different from the FSOD which assumes the distribution of source  $\mathcal{P}_S$  equals that of the target  $\mathcal{P}_T$ , CD-FSOD tackles a more realistic scenario where  $\mathcal{P}_S \neq \mathcal{P}_T$ . Models are required to train on  $\mathcal{D}_S$  and then test on  $\mathcal{D}_T$ .

**$N$ -way  $K$ -shot Protocol.** To evaluate the FSL ability of CD-FSOD models, we follow the  $N$ -way  $K$ -shot evaluation protocol. Specifically, for each novel class in  $\mathcal{C}_T$ ,  $K$  labeled instances are provided which is also termed as support set  $S$ , and other unlabeled instances are used as query set  $Q$ . Formally, we have  $|S| = N \times K, N = |\mathcal{C}_T|$ .

**Training Strategy.** We adopt a "pretrain, finetune, and testing" pipeline, as used in many previous CD-FSL methods [12, 17]. This pipeline involves training methods on the source dataset  $\mathcal{D}_S$ , followed by finetuning the trainable parameters using the few labeled support instances  $S$  from target dataset  $\mathcal{D}_T$ , and finally testing on the novel query set  $Q$ . Typical objective functions for training and finetuning include the box regression loss and the classification loss.

#### 3.2 Metrics for Domain Difference


We aim to comprehensively capture challenges in cross-domain scenarios by evaluating the datasets across the following three dimensions.

**Style:** We recognize style’s pivotal role in various domain-related tasks like domain adaptation [30], domain generalization [55], and cross-domain few-shot learning [6]. Common styles include photorealistic, cartoon, sketch, etc.

**Inter-class variance (ICV):** ICV, a widely used metric in learning, measures the dissimilarity between classes. Higher ICV values indicate easier recognition of semantic labels. Coarser datasets like COCO usually have higher ICV, whereas finer-grained datasets exhibit smaller ICV values.

**Indefinable Boundaries (IB):** IB, borrowed from camouflaged object detection [3], reflects the confusion between the target object and its background. Greater confusion poses a challenge for object detectors. For example, detecting a person against a clean background is relatively straightforward, but identifying a fish in a coral reef is much more challenging.

We categorize ICV levels as large, medium, and small, and IB levels as slight, moderate, and significant. Specifics regarding the measurement of ICV and IB levels for datasets are provided in the supplementary materials.



**Fig. 2:** CD-FSOD benchmark: COCO serves as the training source data, while six datasets are utilized as novel testing target datasets. These datasets exhibit variations in styles, inter-class variance (ICV), and indefinable boundaries (IB).


### 3.3 Benchmarks of Different Domains

Based on the style, ICV, and IB metrics, we carefully review and reorganize existing object detection datasets to construct the CD-FSOD benchmark. Our benchmark comprises seven datasets. COCO [28], the widely used dataset for object detection, which provides a diverse set of object categories such as people, animals, vehicles, and various everyday objects serve as the  $\mathcal{D}_S$ . The other six datasets including ArTaxOr [2], Clipart1k [18], DIOR [26], DeepFish [36], NEU-DET [38], and UODD [20] function as  $\mathcal{D}_T$ . Examples of these datasets as well as the corresponding style, ICV, and IB are depicted in Fig. 2. Note that DeepFish has only one class and thus hasn't an ICV value. More details of the datasets are attached to the supplementary materials.

## 4 Methodology

### 4.1 Overview of CD-ViTO

**Preliminary.** DE-ViT [52] builds an open-set detector by using the visual features from large pretrained model and disentangling the localization and classification tasks. The pipeline for the base DE-ViT is represented by blue arrows in Fig. 3(a). It mainly contains a pretrained DINOv2 ViT [32], a region proposal network ( $M_{RPN}$ ), an ROI align module ( $M_{ROI}$ ), a detection head ( $M_{DET}$ ), and a one-vs-rest classification head ( $M_{CLS}$ ). Particularly, given a query image  $q$  and a set of support instances  $S$ , DE-ViT first uses the DINOv2 to extract the instance features  $F_{ins} = \{F_{ins}^{ob}, F_{ins}^{bg}\}$ , where the  $F_{ins}^{ob}$  represents features for foreground objects from  $S$  and the  $F_{ins}^{bg}$  denotes features for background instances.



**Fig. 3:** (a) Overall framework of our CD-ViT. We build our method upon an open-set detector (DE-ViT). Modules in blue are inherited from DE-ViT while modules in orange are proposed by us. New improvements include learnable instance features, instance reweighting, domain prompter, and finetuning; (b) Illustration of our modules.

Then, the class prototypes  $F_{pro} = \{F_{pro}^{ob}, F_{ins}^{bg}\}$  are obtained by averaging  $F_{ins}^{ob}$  on object categories while keeping  $F_{ins}^{bg}$  unchanged (as in blue dashed arrows). For the query image  $q$ , DE-ViT applies the DINOv2,  $M_{RPN}$ , and  $M_{ROI}$  to generate the region proposals  $R_q$ , visual features  $F_q$ , and ROI features  $F_{q_{roi}}$ . After that, the  $M_{DET}$  takes the  $R_q$ ,  $F_q$ , and  $F_{pro}$  as input for the localization task producing  $\mathcal{L}_{loc}$ . In the meantime, the  $M_{CLS}$  performs the classification task based on  $F_{q_{roi}}$  and  $F_{pro}$  resulting in  $\mathcal{L}_{cls}$ . The network is optimized by  $\mathcal{L}_{loc}$  and  $\mathcal{L}_{cls}$ .

Despite DE-ViT’s purported ability to detect any novel objects without finetuning, our experiments reveal its failure to generalize to cross-domain target datasets (Sec. 5.2). This motivates us to improve the cross-domain ability of DE-ViT and form a new method. Note that the basic components of DE-ViT including the DINOv2,  $M_{RPN}$ ,  $M_{ROI}$ ,  $M_{DET}$ ,  $M_{CLS}$ ,  $\mathcal{L}_{loc}$ , and  $\mathcal{L}_{cls}$ , validated on FSOD tasks, are directly inherited in our framework. Our method, CD-ViT, is outlined in Fig. 3(a) with new enhancements highlighted in orange. Recall that in Sec. 1, we point out that the domain gap of CD-FSOD is highly related to the ICV, IB, and style of the target datasets. Therefore, we are inspired to design our CD-ViT by tackling these factors. As a result, the novel *learnable instance features* ( $M_{LIF}$ ), the *instance reweighting* ( $M_{IR}$ ), the *domain prompter* ( $M_{DP}$ ) are proposed to tackle the small ICV, significant IB, and changing styles, respectively. In addition, though being explored in related cross-domain tasks [6, 12], the finetuning (FT) is also explored in our benchmark.

Specifically, given  $S$  and  $q$  as input, we perform the same steps as DE-ViT to obtain instance features  $F_{ins} = \{F_{ins}^{ob}, F_{ins}^{bg}\}$ , region proposals  $R_q$ , visual features  $F_q$ , and ROI features  $F_{q_{roi}}$ . But different from directly averaging  $F_{ins}^{ob}$  to obtain class prototypes  $F_{pro}^{ob}$  for objects, we first utilize our  $M_{LIF}$  to make the precalculated  $F_{ins}^{ob}$  learnable and adapt them as  $F_{ins}^{lea}$  by performing the disentangled detection tasks i.e., the localization task supervised by  $\mathcal{L}_{loc}$  and the classification task supervised by  $\mathcal{L}_{cls}$ . Since the training data is labeled with class categories, such an optimization would align the instance features to their

corresponding categories thus increasing the discriminability of features and addressing the small ICV issue. Then, considering the qualities especially the IB of object instances varies, we apply our  $M_{IR}$  to reweight the weights of objects in the  $F_{ins}^{lea}$ , this step forms the new prototypes  $F'_{pro}$ . Generally,  $M_{IR}$  is expected to assign higher values to the high-quality instances with slight IB. In this way, the high-quality instances could contribute more to its class prototype. The  $M_{IR}$  is inserted after the  $M_{LIF}$  and is optimized use the same  $\mathcal{L}_{loc}$ ,  $\mathcal{L}_{cls}$ . In addition, to make the model more robust to the changing styles, our proposed  $M_{DP}$  is applied to synthesize several virtual "domains"  $F_{domain}$ . The goal of the  $M_{DP}$  is to introduce such domain perturbations and force our features to be consistent with different domains. To ensure the diversity of the introduced domains and constrain the domains would not affect the semantics of features, an extra loss  $\mathcal{L}_{dp}$  is introduced. The  $M_{DP}$  is build upon the prototypes reweighed by  $M_{IR}$  i.e.,  $F'_{pro}$ , and the  $\mathcal{L}_{dp}$  works with the  $\mathcal{L}_{loc}$  and  $\mathcal{L}_{cls}$  together.

As for finetuning, we finetune the top two modules i.e., the  $M_{DET}$  and  $M_{CLS}$ . We highlight that despite its simplicity, it is still worth exploring in the context of using open-set detectors for CD-FSOD. Experimental results (Sec. 5.2) also demonstrate the necessity of our approach. As a whole, our CD-ViT0 tackles the cross-domain issues by setting instance features learnable, reweighting instances, adding domain prompter as useful perturbations, and finetuning key components. All the introduced parameters (indicated by the fire icons in Fig. 3(a)) are very lightweight and we optimize them via the object function  $\mathcal{L}$  as:

$$\mathcal{L} = \mathcal{L}_{loc} + \mathcal{L}_{cls} + \mathcal{L}_{dp}. \quad (1)$$

## 4.2 Key Components of Our CD-ViT0

**Learnable Instance Features  $M_{LIF}$ .** The implementation of our learnable instance features  $M_{LIF}$  is quiet easy and efficient. Given the precalculated instance features  $F_{ins} \in \mathcal{R}^{(N \times K + N_{bg}) \times D}$ , where  $N$ ,  $K$ ,  $N_{bg}$ ,  $D$  represents the number of object classes, the number of shots, the number of background instances, and the dimension of visual features,  $M_{LIF}$  defines a new learnable matrix  $F_{ins}^{lea}$  with the same size of  $F_{ins}$  and use the  $F_{ins}$  for initialization. Formally, we have  $F_{ins}^{lea} \in \mathcal{R}^{(N \times K + N_{bg}) \times D}$ ,  $F_{ins}^{lea} = F_{ins}$ . By optimizing the values of the  $F_{ins}^{lea}$ , we obtain the instance features better aligned to the target semantics. The idea of learnable instance features aligns with the visual prompt learning [19] where the aim is to modify the input to better suit downstream tasks. In our case, we propose to modify the key instance features by setting them as learnable parameters. Most importantly, our module serves as the solution to making features more discriminative thus better tackling target data with small ICV.

**Instance Reweighting Module  $M_{IR}$ .** We propose  $M_{IR}$  to enable high-quality object instances e.g., objects with slight IB contribute more to the class prototype. The architecture of  $M_{IR}$  is shown in Fig. 3(b). Given the learnable instance features  $F_{ins}^{lea}$  from  $M_{LIF}$ , we first obtain the object instance features from  $F_{ins}^{lea}$  and denote it as  $F_{ins}^0 \in \mathcal{R}^{N \times K \times D}$ . Our  $M_{IR}$  has two paths connected in a residual way. More specifically, the lower path feeds  $F_{ins}^0$  into an MLP module to

get the weighted score which is further used to perform the weighted sum on the initial  $F_{ins}^0$  resulting in the deformed prototypes  $F_{pro}^{att} \in \mathcal{R}^{N \times D}$ . The upper path performs the average operation on  $F_{ins}^0$  getting  $F_{pro}^{avg} \in \mathcal{R}^{N \times D}$ . The final object prototypes  $F_{pro}^{ob'}$  are represented as  $\alpha fc(F_{pro}^{att}) + (1 - \alpha)F_{pro}^{avg}$ , where  $fc(\cdot)$  is a fully connected layer and  $\alpha$  is a hyperparameter. With  $F_{pro}^{ob'}$  and background features, we form new prototype  $F'_{pro}$ .

**Domain Prompter  $M_{DP}$ .** To make our object prototype features  $F_{pro}^{ob'}$  robust to various domains, we have to solve two goals: 1) synthesize several "domains" with diversities; 2) prevent semantic drift when perturbing features with such "domains". To achieve these, our  $M_{DP}$  first introduces several learnable vectors  $F_{domain} \in \mathcal{R}^{N_{dom} \times D}$  as virtual domains and then proposes several loss functions as supervision. The concept of "prompter" is also related but differs from the prompt tuning [19, 25]. Notably, both approaches employ trainable parameters as prompters, but general prompt tuning facilitates the transfer of pretrained models through efficient tuning, whereas we introduce prompters as perturbations. The illustration of  $M_{DP}$  is demonstrated in Fig. 3(b).  $N_{dom}$  is a hyperparameter.

Concretely, given  $F_{domain}$  and  $F_{pro}^{ob'}$ , we first propose the domain diversity loss  $\mathcal{L}_{domain}$  to force the different domains e.g.,  $f^{d_i}$  and  $f^{d_j}$  to be far from each other. Secondly, for a prototype  $f_{p_i}$  in  $F_{pro}^{ob'}$ , we randomly perturbs it with two various domains  $f^{d_k}$  and  $f^{d_m}$  sampled from  $F_{domain}$  forming two perturbed prototypes  $f_{p_i}^{d_k}, f_{p_i}^{d_m}$ . The perturbation is done by adding the features e.g.,  $f_{p_i}^{d_k} = f_{p_i} + f^{d_k}$ . The domain  $f^{d_m}$  is also added to other object prototypes forming  $f_{p_j}^{d_m}$ , where  $j = 1, \dots, N$ . We constrain that  $f_{p_i}^{d_k}$  and  $f_{p_i}^{d_m}$  should be close to each other while being distant from perturbed prototypes generated from different class prototypes. This constitutes our prototype consistency loss  $\mathcal{L}_{proto}$ . Both the  $\mathcal{L}_{domain}$  and  $\mathcal{L}_{proto}$  are implemented by the InfoNCE loss as follows:

$$\mathcal{L}_{domain} = -\frac{1}{N_{dom}} \sum_{i=1}^{N_{dom}} \left( \log \frac{\exp(f^{d_i} \cdot f^{d_i} / \tau)}{\sum_{j=1}^{N_{dom}} \exp(f^{d_i} \cdot f^{d_j} / \tau)} \right), \quad (2)$$

$$\mathcal{L}_{proto} = -\frac{1}{N} \sum_{i=1}^N \left( \log \frac{\exp(f_{p_i}^{d_k} \cdot f_{p_i}^{d_m} / \tau)}{\sum_{j=1}^N \exp(f_{p_i}^{d_k} \cdot f_{p_j}^{d_m} / \tau)} \right), \quad (3)$$

where the  $\tau$  is a temperature hyperparameter. Besides, we also classify the perturbed features e.g.,  $f_{p_i}^{d_k}$  into the corresponding class labels producing a new cross-entropy loss  $\mathcal{L}_{proto_{cls}}$ . Based on  $\mathcal{L}_{domain}$ ,  $\mathcal{L}_{proto}$ , and  $\mathcal{L}_{proto_{cls}}$ , we have:

$$\mathcal{L}_{dp} = \mathcal{L}_{domain} + \mathcal{L}_{proto} + \mathcal{L}_{proto_{cls}}. \quad (4)$$

## 5 Experiments

**Datasets.** We perform CD-FSOD experiments using the benchmark introduced in Sec. 3.3. In summary, the model is trained on COCO and evaluated on six

other datasets. For ArTaxOr, DIOR, and UODD, we adhere to the splits used in Distill-cdfsod [47]. For Clipart1k, DeepFish, and NEU-DET, we create new splits for CD-FSOD evaluation.

**Network Modules.** All the modules inherited from the base DE-ViT including the pretrained DINOv2 ViT,  $M_{RPN}$ ,  $M_{ROI}$ ,  $M_{DET}$ , and  $M_{CLS}$  remain exactly the same as DE-ViT. For details of these modules, please refer to DE-ViT [52]. Architectures of our novel modules, i.e.,  $M_{LIF}$ ,  $M_{IR}$ ,  $M_{DP}$  are stated in Sec. 4. The MLP contained in the  $M_{IR}$  module is a fully connected layer and a softmax.

**Implementation Details.** The mAPs for 1/5/10 shots are reported. Besides the vanilla OVD methods that detect novel objects directly, we adopt the "pre-train, finetune, and test" pipeline for all evaluated methods as in Sec. 3.1. In the case of our CD-ViT, the base DE-ViT model pretrained on COCO is taken, then the  $M_{DET}$ ,  $M_{CLS}$ ,  $M_{LIF}$ ,  $M_{IR}$ , and  $M_{DP}$  are tuned on novel support set  $S$  using the loss as in Eq. 1. Finetuning details e.g., hyperparameters, optimizer, learning rate, and epochs will be attached in the supplementary materials.

## 5.1 Evaluated Methods

To conduct thorough experiments and offer valuable insights, we explore various methods across four types in the newly proposed CD-FSOD benchmark.

**Typical FSOD Methods.** Several FSOD methods, including Meta-RCNN [48], TFA [45], FSCE [39], and DeFRCN [33], are included to examine whether these traditional algorithms can handle the huge domain gap posed by CD-FSOD.

**CD-FSOD Methods.** Distill-cdfsod [47] and MoFSOD [24] are two works for CD-FSOD. Distill-cdfsod is taken as a competitor. However, MoFSOD explores factors like finetuning and pretrained datasets, but doesn't propose a specific method, thus precluding numerical comparisons.

**ViT-based OD Methods.** Since most of the typical FSOD methods adopt the ResNet50 [16] as the backbone, we would like to study how the large-scale pretrained ViT [1] models can benefit CD-FSOD. To that end, we adapt the ViTDeT [27], a flagship model that uses ViT for standard object detection, to our CD-FSOD. For clarity, we refer to the adapted ViTDeT as "ViTDeT-FT".

**Open-Set FSOD/OD Methods.** We study DE-ViT [52], which excels in FSOD benchmarks. Considering the straightforward nature of finetuning, we also report results using the best finetuning strategy on DE-ViT, referred to as "DE-ViT-FT" (details in supplementary materials). Besides the purely vision-based DE-ViT, we evaluate another widely known vision-language matching OVD detector, namely Detic [56], and finetune it, further resulting in "Detic-FT".

## 5.2 Main Results on CD-FSOD

The results for 1/5/10 shot scenarios on six novel target datasets are summarized in Tab. 1. What can be apparently drawn from the table is the remarkable superiority of our CD-ViT over the base DE-ViT across all target datasets. Moreover, it outperforms other competitors on the majority of these datasets.

**Table 1:** Main Results (mAP) on CD-FSOD benchmark. Four types of object detection methods including typical FSOD, CD-FSOD, ViT-based OD, and open-set based OD/FSOD methods are evaluated and compared. The  $\circ$  denotes results from Distill-cdfsod [47];  $\dagger$  denotes that the methods are developed by us or the results are reported by us. The best results are highlighted in blue. "Avg." means the average result.

	Method	Backbone	ArTaxOr	Clipart1k	DIOR	DeepFish	NEU-DET	UODD	Avg.
1-shot	Meta-RCNN $\circ$ [48]	ResNet50	2.8	-	7.8	-	-	3.6	/
	TFA w/cos $\circ$ [45]	ResNet50	3.1	-	8.0	-	-	4.4	/
	FSCE $\circ$ [39]	ResNet50	3.7	-	8.6	-	-	3.9	/
	DeFRCN $\circ$ [33]	ResNet50	3.6	-	9.3	-	-	4.5	/
	Distill-cdfsod $\circ$ [47]	ResNet50	5.1	7.6	10.5	nan	nan	5.9	/
	ViTDeT-FT $\dagger$ [27]	ViT-B/14	5.9	6.1	12.9	0.9	2.4	4.0	5.4
	Detic $\dagger$ [56]	ViT-L/14	0.6	11.4	0.1	0.9	0.0	0.0	2.2
	Detic-FT $\dagger$ [56]	ViT-L/14	3.2	15.1	4.1	9.0	3.8	4.2	6.6
	DE-ViT $\dagger$ [52]	ViT-L/14	0.4	0.5	2.7	0.4	0.4	1.5	1.0
	DE-ViT-FT $\dagger$ [52]	ViT-L/14	10.5	13.0	14.7	19.3	0.6	2.4	10.1
<b>CD-ViT0 (ours)</b>	ViT-L/14	21.0	17.7	17.8	20.3	3.6	3.1	13.9	
5-shot	Meta-RCNN $\circ$ [48]	ResNet50	8.5	-	17.7	-	-	8.8	/
	TFA w/cos $\circ$ [45]	ResNet50	8.8	-	18.1	-	-	8.7	/
	FSCE $\circ$ [39]	ResNet50	10.2	-	18.7	-	-	9.6	/
	DeFRCN $\circ$ [33]	ResNet50	9.9	-	18.9	-	-	9.9	/
	Distill-cdfsod $\circ$ [47]	ResNet50	12.5	23.3	19.1	15.5	16.0	12.2	16.4
	ViTDeT-FT $\dagger$ [27]	ViT-B/14	20.9	23.3	23.3	9.0	13.5	11.1	16.9
	Detic $\dagger$ [56]	ViT-L/14	0.6	11.4	0.1	0.9	0.0	0.0	2.2
	Detic-FT $\dagger$ [56]	ViT-L/14	8.7	20.2	12.1	14.3	14.1	10.4	13.3
	DE-ViT $\dagger$ [52]	ViT-L/14	10.1	5.5	7.8	2.5	1.5	3.1	5.1
	DE-ViT-FT $\dagger$ [52]	ViT-L/14	38.0	38.1	23.4	21.2	7.8	5.0	22.3
<b>CD-ViT0 (ours)</b>	ViT-L/14	47.9	41.1	26.9	22.3	11.4	6.8	26.1	
10-shot	Meta-RCNN $\circ$ [48]	ResNet50	14.0	-	20.6	-	-	11.2	/
	TFA w/cos $\circ$ [45]	ResNet50	14.8	-	20.5	-	-	11.8	/
	FSCE $\circ$ [39]	ResNet50	15.9	-	21.9	-	-	12.0	/
	DeFRCN $\circ$ [33]	ResNet50	15.5	-	22.9	-	-	12.1	/
	Distill-cdfsod $\circ$ [47]	ResNet50	18.1	27.3	26.5	15.5	21.1	14.5	20.5
	ViTDeT-FT $\dagger$ [27]	ViT-B/14	23.4	25.6	29.4	6.5	15.8	15.6	19.4
	Detic $\dagger$ [56]	ViT-L/14	0.6	11.4	0.1	0.9	0.0	0.0	2.2
	Detic-FT $\dagger$ [56]	ViT-L/14	12.0	22.3	15.4	17.9	16.8	14.4	16.5
	DE-ViT $\dagger$ [52]	ViT-L/14	9.2	11.0	8.4	2.1	1.8	3.1	5.9
	DE-ViT-FT $\dagger$ [52]	ViT-L/14	49.2	40.8	25.6	21.3	8.8	5.4	25.2
<b>CD-ViT0 (ours)</b>	ViT-L/14	60.5	44.3	30.8	22.3	12.8	7.0	29.6	

For instance, under the 10-shot setting on ArTaxOr, CD-ViT0 surpasses Meta-RCNN by 332.1% (from 14.0 to 60.5), ViTDeT-FT by 158.5% (from 23.4 to 60.5), and Detic-FT by 404.2% (from 12.0 to 60.5). These results strongly highlight the effectiveness of our CD-ViT0 method in the context of FSOD across domains. Beyond performance improvement, we also delve into the following aspects:

① **Does the domain gap pose challenges for FSOD models?** Yes, it does. In examining the results of typical FSOD methods, we observe relatively low performance on the target datasets. Consider DeFRCN, for instance, as reported in its paper [33], which achieves 9.3 mAP on COCO’s novel testing set under the 1-shot setting. However, on ArTaxOr and UODD, it drops significantly to 3.6 (a decrease of 61.3%) and 4.5 (a decrease of 51.6%), respectively. Even the Distill-cdfsod, specifically designed for CD-FSOD, fails to surpass 9.3 on the

majority of target domain datasets. These observations indicate the importance of advancing CD-FSOD methods.

② **Are ViT-based models superior to ResNet-based ones?** The introduction of ViTDeT-FT aims to assess whether the ViT backbone can strengthen the performance of CD-FSOD. The comparison between ViTDeT-FT and ResNet50-based methods reveals that the effectiveness depends on the specific target datasets. ViTDeT-FT shows strong performance on ArTaxOr, Clipart1k, and DIOR, but it is less effective on DeepFish, NEU-DET, and UODD. Despite the absence of clear patterns, it can be observed that the choice of backbone is not a decisive factor in constructing accurate CD-FSOD models.

③ **Can open-set models directly address the challenges of CD-FSOD?** Unfortunately, the results of Detic and DE-ViT clearly show that the answer is no. Despite being designed to recognize any novel category, open-set detectors still face challenges posed by domain gaps. It is noteworthy that Detic performs inference using category names and the query image, essentially adopting a zero-shot approach. Consequently, the results for 1/5/10 shots remain the same. On the other hand, DE-ViT requires support instances to construct prototypes, leading to variations in results. These findings highlight the nontrivial nature of solving the CD-FSOD by directly applying existing open-set models.

④ **Can open-set models be significantly improved?** Absolutely! While the performance of vanilla open-set models may be initially low, we have demonstrated that it can be substantially enhanced. This conclusion becomes evident when comparing Detic-FT to Detic and our CD-ViTO to DE-ViT. The success of Detic-FT and CD-ViTO underscores the necessity of finetuning for adapting open-set models to CD-FSOD. Coupled with finetuning and our other novel components, CD-ViTO achieves very impressive results, transforming the initial shortcomings of open-set models into a very promising solution for CD-FSOD.

⑤ **Do different open-set models exhibit distinct performances?** Yes. We delve into the performance of DE-ViT and Detic, two models that achieve open-set learning through entirely different approaches. DE-ViT relies solely on vision features, while Detic relies on the relationship between visual features and category text. Upon analyzing the results of DE-ViT-FT and Detic-FT, a clear trend emerges. DE-ViT-FT demonstrates a decisive advantage on ArTaxOr, Clipart1k, DIOR, and DeepFish, but falls behind Detic-FT on NEU-DET and UODD. We observe that when the Instance Boundary (IB) degree of the target dataset is slight or moderate, the vision-based DE-ViT-FT performs better. Conversely, in scenarios where boundaries are significantly indefinable, indicating strong similarity between foreground and background, methods like Detic-FT, which leverages both visual and text information, prove more competent.

⑥ **How do style, inter-class variance (ICV), and indefinable boundaries (IB) impact the domain gap?** We meticulously analyze the influence of style, ICV, and IB on domain gaps based on experimental results and benchmark analysis in Sec. 3.3. 1) Style, often considered strongly related to the domain, has a relatively smaller impact on CD-FSOD compared to ICV and IB. For instance, DIOR shares relatively similar ICV and IB with COCO but has a

different style. Despite this, the performance of typical FSOD methods does not significantly drop on DIOR. For example, DeFRCN achieves 9.3 mAP on both COCO’s testing set and DIOR under the 1-shot setting. Similarly, the cartoon-style Clipart1k, distinct from the photorealistic COCO, also receives relatively high mAPs. 2) ICV can cause domain issues but appears manageable. ArTaxOr, with the same style and similar IB as COCO, but with fine-grained categories (arthropods) compared to the coarse-level semantic concepts in COCO. This difference in ICV leads to a significant drop in performance for typical FSOD methods on ArTaxOr, such as from 9.1 to 3.6 on 1-shot using DeFRCN. However, we observe that this negative effect can be mitigated by utilizing better vision features, as demonstrated by ViTDeT-FT, DE-ViT-FT, and our CD-ViTO. 3) The change in IB poses the greatest challenge for CD-FSOD. Nearly all types of methods struggle to achieve satisfactory results on NEU-DET and UODD when significant IB exists. Blurred boundaries make it challenging to detect correct bounding boxes, resulting in poor overall performance.

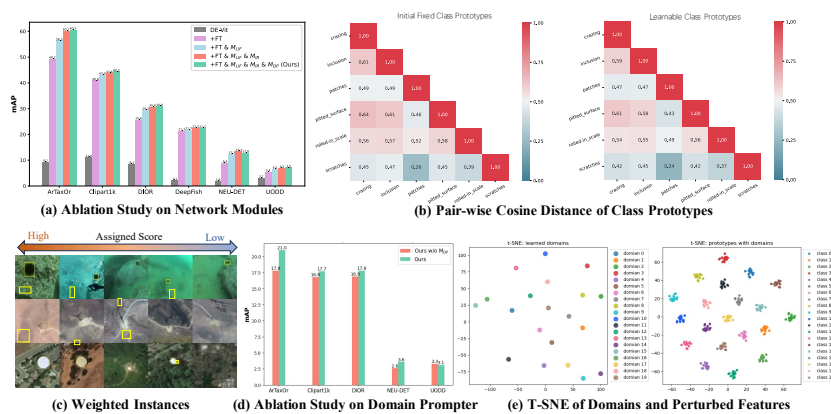


Fig. 4: Ablation study on our proposed modules.

### 5.3 Analysis on the Proposed Modules

To demonstrate the effectiveness of our proposed components, we conduct ablation studies by incrementally adding them to the base DE-ViT. The 10-shot results are presented in Fig. 4(a). Overall, we observe that: 1) By simply finetuning top heads, the performance is significantly improved. The results reinforce the observed phenomenon and analysis from Tab. 1 that finetuning is crucial in CD-FSOD. 2) The learnable instance features further boost performance clearly, indicating that we do make the features better by optimizing them. 3) Although not as pronounced as in finetuning and  $M_{LIF}$ , the results still demonstrate the effectiveness of our  $M_{IR}$  and  $M_{DP}$  modules.

**More Discussions on Learnable Instance Features.** We propose the learnable instance features  $M_{LIF}$  to tackle the small ICV issue hoping to make the visual features more discriminative. To verify whether the  $M_{LIF}$  module meets expectations, we calculate the pair-wise cosine distances of class prototypes and compare the results of our learnable class prototypes produced by  $M_{LIF}$  with that of the initial fixed ones. Results on 5-shot NEU-DET are given in Fig. 4(b), we observe that the cosine distances between classes are consistently declining. This indicates that our  $M_{LIF}$  enlarges the ICV.

**More Discussions on Instance Reweighting.** The instance reweighting module  $M_{IR}$  is presented to assign higher weights to the instances with slight IB values. Thus, we visualize several weighted instances ordered by weights from high to low. 5-shot results from three datasets are demonstrated in Fig. 4(c). Results show that instances with relatively clear boundaries are given higher scores; conversely, those with confusing boundaries receive lower scores.

**More Discussions on Domain Prompter.** The  $M_{DP}$  is designed to generate virtual domains as useful perturbations. Different from the relatively marginal improvements in 10-shot (Fig. 4(a)), the 1-shot results on Fig. 4(d) show significant gains. Note that we don't apply  $M_{DP}$  on DeepFish since it only contains one object category and the slight drop in UODD indicates that training the new parameters introduced by  $M_{DP}$  in extremely data-limited cases (3 classes for UODD) may pose challenges. Besides quantified results, in Fig. 4(e), we provide the t-SNE of our learned domains and the perturbed features by adding domains to the class prototypes. The 1-shot results on Clipart1K are visualized. The results reveal that: 1) the learned domains exhibit diversity; 2) adding domains to the class prototypes generates new data but keeps the semantic concepts maintained. Finally, we term these learnable perturbations as "domain" prompts rather than "style" prompts since learning meaningful styles without explicit style information is challenging. The term "domain" is used broadly to encompass factors, including style, that could induce changes in data distribution.

## 6 Conclusion

In conclusion, this paper presents several key contributions to the field of cross-domain few-shot object detection (CD-FSOD). Firstly, we meticulously reviewed and reorganized object detection datasets, establishing a novel benchmark specifically tailored for CD-FSOD. Secondly, we conducted a comprehensive study involving various methods, including typical FSOD models, CD-FSOD models, ViT-based OD models, and open-set-based OD/FSOD models, based on this benchmark. Through careful analysis, we provided detailed insights, marking the first exploration of open-set based detectors for CD-FSOD. Thirdly, leveraging the state-of-the-art DE-ViT, we introduced the novel CD-ViTO method, enhancing the open-set detector with our novel techniques, i.e., finetuning, learnable instance features, instance reweighting, and domain prompter. The effectiveness of CD-ViTO and its modules are thoroughly validated on target domains. We believe our contributions will propel advancements for CD-FSOD.

## Bibliography

- [1] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
- [2] Drange, G.: Arthropod taxonomy orders object detection dataset. In: <https://doi.org/10.34740/kaggle/dsv/1240192> (2019)
- [3] Fan, D.P., Ji, G.P., Cheng, M.M., Shao, L.: Concealed object detection. TPAMI (2021)
- [4] Fan, Z., Ma, Y., Li, Z., Sun, J.: Generalized few-shot object detection without forgetting. In: CVPR (2021)
- [5] Fu, Y., Fu, Y., Jiang, Y.G.: Meta-fdmixup: Cross-domain few-shot learning guided by labeled target data. In: ACM MM (2021)
- [6] Fu, Y., Xie, Y., Fu, Y., Jiang, Y.G.: Styleadv: Meta style adversarial training for cross-domain few-shot learning. In: CVPR (2023)
- [7] Fu, Y., Zhang, L., Wang, J., Fu, Y., Jiang, Y.G.: Depth guided adaptive meta-fusion network for few-shot video recognition. In: ACM MM (2020)
- [8] Gao, Y., Lin, K.Y., Yan, J., Wang, Y., Zheng, W.S.: Asyfod: An asymmetric adaptation paradigm for few-shot domain adaptive object detection. In: CVPR (2023)
- [9] Gao, Y., Yang, L., Huang, Y., Xie, S., Li, S., Zheng, W.S.: Acrofofod: An adaptive method for cross-domain few-shot object detection. In: ECCV (2022)
- [10] Girshick, R.: Fast r-cnn. In: ICCV (2015)
- [11] Guirguis, K., Meier, J., Eskandar, G., Kayser, M., Yang, B., Beyerer, J.: Niff: Alleviating forgetting in generalized few-shot object detection via neural instance feature forging. In: CVPR (2023)
- [12] Guo, Y., Codella, N.C., Karlinsky, L., Codella, J.V., Smith, J.R., Saenko, K., Rosing, T., Feris, R.: A broader study of cross-domain few-shot learning. In: ECCV (2020)
- [13] Han, G., He, Y., Huang, S., Ma, J., Chang, S.F.: Query adaptive few-shot object detection with heterogeneous graph convolutional networks. In: ICCV (2021)
- [14] Han, G., Huang, S., Ma, J., He, Y., Chang, S.F.: Meta faster r-cnn: Towards accurate few-shot object detection with attentive feature alignment. In: AAAI (2022)
- [15] Han, G., Ma, J., Huang, S., Chen, L., Chang, S.F.: Few-shot object detection with fully cross-transformer. In: CVPR (2022)
- [16] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- [17] Hu, S.X., Li, D., Stühmer, J., Kim, M., Hospedales, T.M.: Pushing the limits of simple pipelines for few-shot learning: External data and fine-tuning make a difference. In: CVPR (2022)

- [18] Inoue, N., Furuta, R., Yamasaki, T., Aizawa, K.: Cross-domain weakly-supervised object detection through progressive domain adaptation. In: CVPR (2018)
- [19] Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: ECCV (2022)
- [20] Jiang, L., Wang, Y., Jia, Q., Xu, S., Liu, Y., Fan, X., Li, H., Liu, R., Xue, X., Wang, R.: Underwater species detection using channel sharpening attention. In: ACM MM (2021)
- [21] Kang, B., Liu, Z., Wang, X., Yu, F., Feng, J., Darrell, T.: Few-shot object detection via feature reweighting. In: ICCV (2019)
- [22] Kaul, P., Xie, W., Zisserman, A.: Label, verify, correct: A simple few shot object detection method. In: CVPR (2022)
- [23] Köhler, M., Eisenbach, M., Gross, H.M.: Few-shot object detection: A comprehensive survey. arXiv preprint arXiv:2112.11699 (2021)
- [24] Lee, K., Yang, H., Chakraborty, S., Cai, Z., Swaminathan, G., Ravichandran, A., Dabeer, O.: Rethinking few-shot object detection on a multi-domain benchmark. In: ECCV (2022)
- [25] Lester, B., Al-Rfou, R., Constant, N.: The power of scale for parameter-efficient prompt tuning. arXiv preprint arXiv:2104.08691 (2021)
- [26] Li, K., Wan, G., Cheng, G., Meng, L., Han, J.: Object detection in optical remote sensing images: A survey and a new benchmark. ISPRS (2020)
- [27] Li, Y., Mao, H., Girshick, R., He, K.: Exploring plain vision transformer backbones for object detection. In: ECCV (2022)
- [28] Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
- [29] Luo, X., Wu, H., Zhang, J., Gao, L., Xu, J., Song, J.: A closer look at few-shot classification again. In: ICML (2023)
- [30] Luo, Y., Liu, P., Guan, T., Yu, J., Yang, Y.: Adversarial style mining for one-shot unsupervised domain adaptation. In: NeurIPS (2020)
- [31] Ma, J., Niu, Y., Xu, J., Huang, S., Han, G., Chang, S.F.: Digeo: Discriminative geometry-aware learning for generalized few-shot object detection. In: CVPR (2023)
- [32] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: DINOv2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
- [33] Qiao, L., Zhao, Y., Li, Z., Qiu, X., Wu, J., Zhang, C.: Defrcn: Decoupled faster r-cnn for few-shot object detection. In: ICCV (2021)
- [34] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
- [35] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR (2016)
- [36] Saleh, A., Laradji, I.H., Konovalov, D.A., Bradley, M., Vazquez, D., Sheaves, M.: A realistic fish-habitat dataset to evaluate algorithms for underwater visual analysis. Scientific Reports (2020)

- [37] Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: *NeurIPS* (2017)
- [38] Song, K., Yan, Y.: A noise robust method based on completed local binary patterns for hot-rolled steel strip surface defects. *Applied Surface Science* (2013)
- [39] Sun, B., Li, B., Cai, S., Yuan, Y., Zhang, C.: Fscse: Few-shot object detection via contrastive proposal encoding. In: *CVPR* (2021)
- [40] Tang, H., Yuan, C., Li, Z., Tang, J.: Learning attention-guided pyramidal features for few-shot fine-grained recognition. *Pattern Recognit.* (2022)
- [41] Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. In: *ICLR* (2020)
- [42] Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: *NeurIPS* (2016)
- [43] Wang, H., Deng, Z.H.: Cross-domain few-shot classification via adversarial task augmentation. *arXiv preprint* (2021)
- [44] Wang, K., Liew, J.H., Zou, Y., Zhou, D., Feng, J.: Panet: Few-shot image semantic segmentation with prototype alignment. In: *ICCV* (2019)
- [45] Wang, X., Huang, T.E., Darrell, T., Gonzalez, J.E., Yu, F.: Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957* (2020)
- [46] Xie, G.S., Xiong, H., Liu, J., Yao, Y., Shao, L.: Few-shot semantic segmentation with cyclic memory network. In: *ICCV* (2021)
- [47] Xiong, W.: Cd-fsod: A benchmark for cross-domain few-shot object detection. In: *ICASSP* (2023)
- [48] Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., Lin, L.: Meta r-cnn: Towards general solver for instance-level low-shot learning. In: *ICCV* (2019)
- [49] Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: *CVPR* (2021)
- [50] Zhang, H., Zhang, L., Qi, X., Li, H., Torr, P.H., Koniusz, P.: Few-shot action recognition with permutation-invariant attention. In: *ECCV* (2020)
- [51] Zhang, J., Gao, Lianliand Luo, X., Shen, H., Song, J.: Deta: Denoised task adaptation for few-shot learning. In: *ICCV* (2023)
- [52] Zhang, X., Wang, Y., Boularias, A.: Detect every thing with few examples. *arXiv preprint arXiv:2309.12969* (2023)
- [53] Zhao, S., Zhang, Z., Schuler, S., Zhao, L., Vijay Kumar, B., Stathopoulos, A., Chandraker, M., Metaxas, D.N.: Exploiting unlabeled data with vision and language models for object detection. In: *ECCV* (2022)
- [54] Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. In: *CVPR* (2022)
- [55] Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain generalization with mixstyle. In: *ICLR* (2021)
- [56] Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. In: *ECCV* (2022)
- [57] Zhuo, L., Fu, Y., Chen, J., Cao, Y., Jiang, Y.G.: Tgdm: Target guided dynamic mixup for cross-domain few-shot learning. In: *ACM MM* (2022)