

Open-Set Semi-Supervised Object Detection

Yen-Cheng Liu^{1*}, Chih-Yao Ma², Xiaoliang Dai², Junjiao Tian¹,
Peter Vajda², Zijian He², and Zsolt Kira¹

¹Georgia Tech, ²Meta

Abstract. Recent developments for Semi-Supervised Object Detection (SSOD) have shown the promise of leveraging unlabeled data to improve an object detector. However, thus far these methods have assumed that the unlabeled data does not contain out-of-distribution (OOD) classes, which is unrealistic with larger-scale unlabeled datasets. In this paper, we consider a more practical yet challenging problem, Open-Set Semi-Supervised Object Detection (OSSOD). We first find the existing SSOD method obtains a lower performance gain in open-set conditions, and this is caused by the semantic expansion, where the distracting OOD objects are mispredicted as in-distribution pseudo-labels for the semi-supervised training. To address this problem, we consider online and offline OOD detection modules, which are integrated with SSOD methods. With the extensive studies, we found that leveraging an offline OOD detector based on a self-supervised vision transformer performs favorably against online OOD detectors due to its robustness to the interference of pseudo-labeling. In the experiment, our proposed framework effectively addresses the semantic expansion issue and shows consistent improvements on many OSSOD benchmarks, including large-scale COCO-OpenImages. We also verify the effectiveness of our framework under different OSSOD conditions, including varying numbers of in-distribution classes, different degrees of supervision, and different combinations of unlabeled sets.

1 Introduction

The success of deep neural networks relies on large collections of labeled data, although creating such large-scale datasets is expensive and time-consuming. The recent development of successful Semi-Supervised Learning (SSL) methods alleviates this requirement by making better use of unlabeled data, narrowing the performance gap between the SSL models and the fully-supervised model.

Inspired by SSL methods for image classification [1,20,35,39,48,46,10,13,36], recent works on semi-supervised object detection (SSOD) [37,25,49,38] have applied the self-training method, which trains the model with the pseudo-labels of the unlabeled data. These works often consider a scenario where the labeled set is randomly sampled from a dataset (*e.g.*, MS-COCO [23]) and use the remaining images as the unlabeled set. This implicitly assumes the label spaces of labeled and unlabeled data are identical. However, this closed-set assumption is unlikely

* Work done partially while interning at Meta.

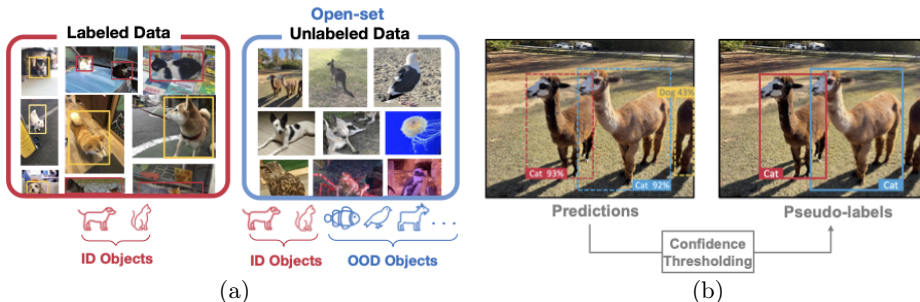


Fig. 1. (a) Open-Set Semi-Supervised Object Detection aims to learn an in-distribution object detector with a number of labeled images and another set of unconstrained/open-set unlabeled images. The objects appearing in the labeled data are defined as in-distribution (ID) objects, while some objects in the open-set unlabeled data are out-of-distribution (OOD) objects, which are unseen, unknown, and undefined in the labeled set. (b) While most of the self-training methods [36,25] rely on thresholding for filtering unreliable pseudo-labels, some OOD objects in unlabeled images are mispredicted as inlier objects with high confidence. These OOD objects are then mis-labeled, trained on, and lead to worse performance.

to happen in real-world situations, where unlabeled images collected in the wild might contain out-of-distribution (OOD) objects, which are unseen, undefined, and unknown in the available labeled set.

We are thus interested in a more practical yet challenging problem, Open-Set Semi-Supervised Object Detection (OSSOD), which aims to leverage the **unconstrained** unlabeled images (*i.e.*, images containing unseen OOD objects), to improve an object detector trained with the available labeled data as shown in Figure 1a. When adding unlabeled data containing open-set categories, we observe that the existing successful SSOD method leads to a lower performance gain or even *degraded* results. This is different from the common belief that SSL methods can benefit from using more unlabeled data. We attribute the above phenomena to the *semantic expansion* issue, where OOD objects are mispredicted as in-distribution objects with high confidence and misused as pseudo-labels with confidence thresholding (Figure 1b).

To eliminate the detrimental effect of OOD samples, we propose to add an additional OOD filtering process into the existing SSOD training pipeline. More concretely, we first consider **online** OOD detectors to perform OOD filtering. An online OOD detector is a prediction head/branch we straightforwardly add on the object detector using existing OOD methods [12,21,14,41]. However, we find that such methods cannot produce satisfactory results due to interference with other tasks, *e.g.*, bounding box localization and box classification (See Section 4.2 for further discussion). In order to address this, we propose a simple but effective strategy that uses an **offline** OOD detection module, which is disentangled from the architecture of the object detector. This OOD detector is based on a self-supervised DINO [2] model, and it provides several advantages. First, the pre-training of DINO does not require label annotations, so it is suitable for

the low-label setting and alleviates the concern of limited amounts of labels for OOD detection tasks. Secondly, it is more effective in detecting OOD objects in the pseudo-labels compared with other (online) OOD methods as shown in Section 5.2, and this eliminates the detrimental effect of OOD samples for OSSOD tasks. Lastly, since the architecture of OOD detector and object detector are independent, the training of the two models can be done separately and prevent interference as we observed in online OOD detectors.

In our experiments, we first provide a systematic analysis between several OOD detection methods, and our results suggest that offline methods show consistent improvements over the online detection methods. We also show that using the offline OOD detector can filter the OOD objects in pseudo-labels and consistently improve against the existing SSOD methods under different open-set scenarios, including under different combination of unlabeled sets, varying number of in-distribution (ID) classes, and different number of images. We also find that using the pseudo-labels generated from our framework is even more effective than using the ground-truth labels provided in OpenImages (see Section 5.3).

We highlight the contributions of this paper as follows:

- To the best of our knowledge, we are the *first* to address the open-set semi-supervised object detection tasks, and we analyze the limitation of existing SSOD methods on OSSOD tasks.
- We identify the challenges in designing OOD detection in OSSOD tasks, present online and offline OOD detectors, and provide a systematic comparison between these two modules.
- Through our extensive experiments, we demonstrate that an offline OOD detector can effectively remove OOD objects in pseudo-labels, and this leads to a significant improvements under different OSSOD scenarios, including a varying number of ID/OOD classes, different degrees of supervision, and different scales of datasets.

2 Related Work

Open-set semi-supervised object detection focuses on improving an in-distribution object detector with in-the-wild unlabeled data and preventing the detrimental effects caused by distracting OOD objects in a more practical semi-supervised setup. This task is different from the existing works on open-vocabulary/open-world object detection [47,9,33,50,18,17,16], where the goal is to improve the recall and accuracy of novel/OOD objects.

Semi-Supervised Object Detection. Recent successful methods on semi-supervised learning for image classification have applied various data augmentation and consistency regularization on unlabeled data [1,20,35,39,48,46,10,13,36]. These techniques have also promoted the development of semi-supervised learning in object detection tasks. For example, several works [37,25,49,38] apply the self-training method, which is based on the Teacher-Student framework. For example, STAC [37] uses the labeled data to train the Teacher, which generates pseudo-labels to supervise the Student model. To refine the quality of pseudo-labels,

existing works propose several techniques, including exponential moving average (EMA) [49,43,25,44], co-rectify mechanism [49], replacing hard pseudo-labels with soft pseudo-labels [38], and addressing class imbalance in pseudo-labels [25]. While promising results have been made, existing SSOD works usually experiment on datasets where the labeled and unlabeled sets have the same object categories. We are interested in a more challenging scenario, where some objects in the unlabeled set are novel and never appear in the labeled set.

Out-of-Distribution Detection. Out-of-distribution (OOD) classes are object categories that do not appear in the training label space and are not known *a-priori*, and OOD detection, which plays a vital role in OSSOD tasks, is a binary classification problem that decides whether a sample is an ID or OOD. Early works use different scoring functions to estimate the likelihood of a sample being OOD, including maximum softmax probability (MSP) [11,22] and Mahalanobis distance [21]. Another line of works [12,24,28] exploit large sets of OOD samples during training and they generally achieve better performance on detecting the outlier objects. There are also several generative-based methods [52,30,32,29] for OOD detection tasks, though it is difficult and prohibitively challenging to apply them on large-scale semi-supervised object detection tasks. Among the existing OOD works, Fort *et al.* [7] shows state-of-the-art performance on OOD detection by simply fine-tuning the pretrained vision transformer with the available inlier data. While promising results have been made on OOD detection tasks, most of the existing works only experiment on small-scale image classification tasks, and their extension to the large-scale object detection tasks is not verified. The only work on OOD detection for object detection is VOS [6], but the interference with pseudo-labeling in SSOD tasks and computation cost of virtual negative sampling make it incompatible with the existing SSOD methods.

Open-Set Semi-Supervised Learning. Existing works on open-set semi-supervised learning focus on image classification tasks [45,34,15,26] or image generation [8]. For example, MTC [45] estimates the OOD score for each unlabeled image, and both the network parameters and OOD scores are updated alternately. Based on the estimated OOD scores, MTC eliminates the OOD samples with low OOD scores and improves the image classifier in a semi-supervised scenario. OpenMatch [34] applies a consistency regularization on a one-vs-all classifier, which is used as an OOD detector to filter the OOD samples during semi-supervised learning. Despite the promising results, no prior work has addressed open-set semi-supervised learning for object detection tasks, which have more challenges than image classification tasks. For instance, in image classification, each image only contains one single object/class, whereas, in object detection, each image may contain an arbitrary number of ID/OOD objects, making pseudo-labeling integrated with OOD detection much more challenging. Additionally, we also observed that balancing object detection losses against OOD detector is difficult as we showed in Section 4.2.

3 Revisiting Semi-Supervised Object Detection

Semi-supervised object detection (SSOD) aims to learn an object detector by using a set of labeled images $\mathbf{D}_s = \{\mathbf{x}_i^s, \mathbf{y}_i^s\}_{i=1}^{N_s}$ and unlabeled images $\mathbf{D}_u = \{\mathbf{x}_i^u\}_{i=1}^{N_u}$, where $N_s \ll N_u$. Existing SSOD works [37,25,49,43] apply the self-training method and have shown significant improvements in this semi-supervised scenario. The typical strategy of such methods is to generate pseudo-labels (*i.e.*, pseudo-boxes and the corresponding class labels) of unlabeled images and then train the object detector using both labeled data and unlabeled data with pseudo-labels.

These works adopt the Teacher-Student framework, where the **Teacher** model generates pseudo-labels to train the **Student** model. Specifically, the **Teacher** model takes the weakly-augmented unlabeled data as input and generates the box predictions, and the bounding boxes with confidence higher than a pre-defined threshold τ are selected as the pseudo-boxes. The **Student** then takes as input the same images with stronger augmentation, and we enforce the consistency loss \mathcal{L}_{unsup} (*i.e.*, unsupervised loss) between its predictions and the generated pseudo-labels. To train the **Student**, we combine both the supervised loss \mathcal{L}_{sup} and the unsupervised loss \mathcal{L}_{unsup} .

$$\mathcal{L}_{ssod} = \mathcal{L}_{sup} + \lambda \mathcal{L}_{unsup} = \sum_i \mathcal{L}(\mathbf{x}_i^s, \mathbf{y}_i^s) + \lambda \sum_i \mathcal{L}(\mathbf{x}_i^u, \hat{\mathbf{y}}_i^u), \quad (1)$$

where λ is the unsupervised loss weight and $\hat{\mathbf{y}}_i^u = \delta(\tilde{\mathbf{y}}_i^u; \tau)$ represents the pseudo-labels, which are derived from the bounding box prediction $\tilde{\mathbf{y}}_i^u$ after the confidence thresholding function $\delta(\cdot)$ with the pre-defined threshold τ .

To further refine the quality of pseudo-labels, some existing works [38,25,43] update the **Teacher** model (θ_t) by exploiting the model weights of **Student** (θ_s) via the exponential moving average (*i.e.*, $\theta_t \leftarrow \alpha \theta_t + (1 - \alpha) \theta_s$).

To verify the effectiveness of the models, existing works often experiment on a setting where the labeled and unlabeled sets are randomly sampled from the same dataset. This implicitly assumes the set of object categories are the same across the labeled and unlabeled sets, and we formalize this setup as the **closed-set SSOD**.

4 Open-Set Semi-Supervised Object Detection

While the above self-training methods show promising results in closed-set SSOD, it restricts the set of object categories in the unlabeled set to be the same as the labeled set. This setting, however, is less practical since unlabeled datasets collected in the wild might contain images with novel objects beyond what were presented in the labeled dataset. We are thus interested in OSSOD, a more practical yet challenging problem setup, where novel and undefined objects appear in the unconstrained unlabeled set and are not available in the labeled set with a limited amount of data.

4.1 Semantic Expansion in Open-Set Scenarios

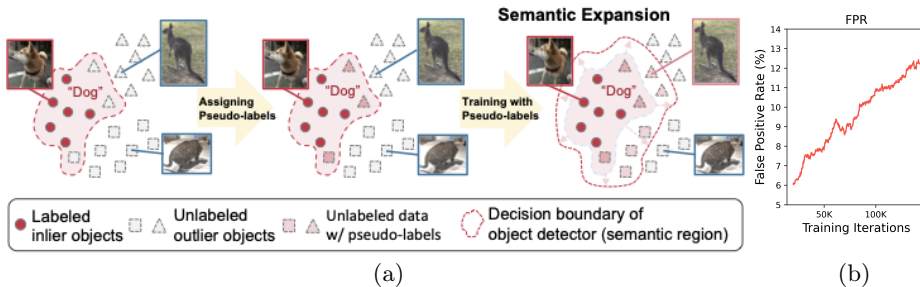


Fig. 2. Illustration of **semantic expansion**. (a) With the object detector trained with the labeled data, some OOD objects in unlabeled data are predicted as ID objects, and solely using the confidence thresholding based on the box score cannot effectively suppress these OOD objects in pseudo-labels. Using the noisy pseudo-labels in self-training methods makes the open-set issue become more severe after several training iterations, and (b) thus the false-positive rate (*i.e.*, percentage of OOD objects predicted as ID) increases over time.

We first analyze the existing successful SSOD work¹ on OSSOD tasks. Interestingly, we found it shows limited performance gain in the open-set scenario compared with closed-set scenario as shown in Figure 3. We attribute this lower performance gain to *semantic expansion* as presented in Figure 2a. To be more specific, when OOD/novel objects exist in the unlabeled set, the closed-set classifier (*e.g.*, ROIhead classifier in Faster-RCNN [31]) might mispredict these OOD instances as ID objects with high confidence (a similar observation is also made recently in prior works [4,27]). These over-confident and incorrect OOD instances are prone to be selected as the pseudo-labels of ID objects

even after confidence thresholding, and using these incorrect OOD pseudo-labels in the training makes the EMA-updated Teacher mispredict more OOD objects as ID objects. After the iterative EMA updates between the Teacher and Student, the semantic discrimination of ID objects is *enlarged* and incorrectly covers more OOD objects. This causes the false-positive rate of the OOD pseudo-labels to increase as the model is trained longer, as presented in Figure 2b.

To address the above issue and improve the performance on OSSOD tasks, we propose an integrated framework to detect the OOD instances in the pseudo-boxes and eliminate their detrimental effect during semi-supervised learning. With the




Fig. 3. Performance comparison between closed-set and open-set SSOD. The performance gain of the existing SSOD method [25] is much smaller in the open-set conditions. Note that we randomly select 0.5%/1% data as the labeled set for both conditions.

¹ As all SoTA SSOD methods [25,38,49,43] use the Teacher-student mechanism and pseudo-labeling method, we choose UT [25] as an example while our framework is not restricted to it.

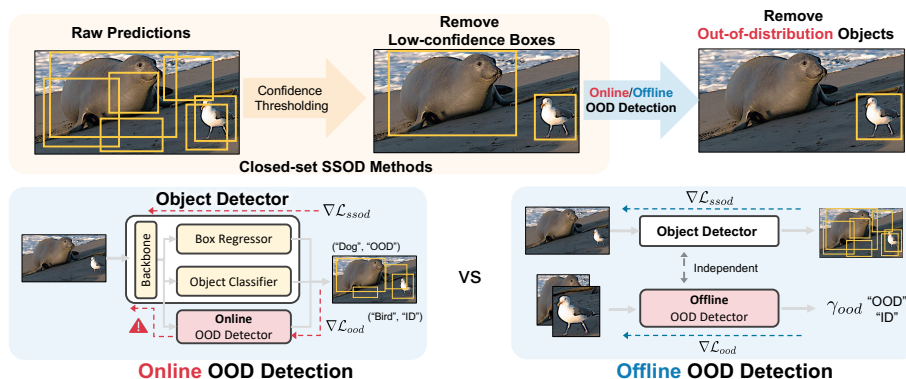


Fig. 4. Illustration of online and offline OOD detection frameworks for OSSOD tasks. Combined with the existing closed-set SSOD methods, OOD detection aims to remove the OOD objects in pseudo-labels and prevent *semantic expansion* in OSSOD. For the online OOD detector, we add an additional branch on the object detector, but it suffers from the instability of pseudo-labels used for SSOD and leads to degraded performance of the object detector. In contrast, the offline OOD detector does not have these limitations and produce better results on OOD detection due to its nature of independent model architecture.

goal of deciding whether a sample is an ID or OOD data (*i.e.*, binary classification), prior works have proposed several OOD detection methods [11,21,24,12,34], although most of these works only verify the effectiveness of the methods on the image-level vision tasks. To further advance to OOD detection for instance-level tasks (*e.g.*, object detection), we consider **online** and **offline** OOD detectors in the following sections.

4.2 Online OOD detection

As illustrated in Figure 4, online OOD detection expands the architecture of the existing object detector so that it can perform OOD detection in addition to bounding box classification and localization. One simple way to detect the OOD objects is to use the maximum softmax probability (MSP) [11] of ROIhead classifier as an indicator for identifying OOD samples. In a similar spirit, another way is to apply a range of existing OOD detection methods [11,21,24,12,34,41] and add an additional OOD branch on the ROIhead. We combine the semi-supervised object detection loss \mathcal{L}_{ssod} as mentioned in Eq. 3 and OOD detection loss \mathcal{L}_{ood} defined in the original works, and we train the model in an end-to-end manner.

However, these online OOD detectors have several limitations in OSSOD tasks. On the one hand, as presented in Figure 5a, pseudo-labeling used for SSOD significantly degrades the performance of online OOD detection (we use MSP [11] as an example), and such a trend is hypothetically caused by the instability of pseudo-labels (*i.e.*, classification noise in pseudo-labels). On the other hand, bundling two different tasks (*e.g.*, OOD detection and object detection) in a

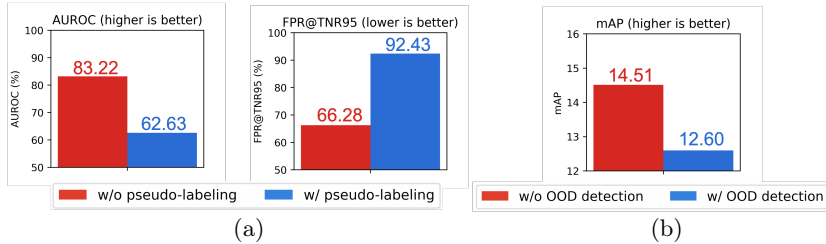


Fig. 5. (a) Pseudo-labeling used in SSOD degrades the performance of online OOD detection. (b) Accuracy of an object detector is degraded by adding online OOD detection, where we create another branch on ROIhead and exploit outlier exposure [12] to detect OOD samples. Both AUROC and FPR@TNR95 are standard evaluation metrics for OOD detection (defined in Section 5.2).

shared architecture leads to sub-optimal performance of object detection task as shown in Figure 5b, making the online OOD detection less favorable to be applied for OSSOD tasks.

4.3 Offline OOD detection.

To ameliorate the above issues, as shown in Figure 4, we alternatively propose to use **offline** OOD detection, where we construct an OOD detector that is disentangled from the architecture of the object detector and *not* jointly trained with the object detector. This framework design provides two advantages: 1) the offline OOD detector is compatible with *any* existing SSOD methods [37,49,25,51], since the offline OOD detector is modularized and independent from the object detector. 2) Such a framework design alleviates the concern about competing task objectives between OOD detection and object detection.

To design the offline OOD detector, we exploit the current state-of-the-art OOD method, which is simple yet effective on image-level OOD tasks. Specifically, as pointed out by Fort *et al.* [7], one highly successful method for OOD detection is to simply fine-tune a pretrained ViT-B [5] with available ID data. Despite its promising results, pre-training a ViT-B requires ImageNet-21k images and the corresponding large amount of annotated labels, which is not suitable for our semi-supervised setting and would result in an unfair comparison to prior works. **DINO as an offline OOD detector.** We therefore utilize the self-supervised DINO [2], which is only trained with ImageNet-1k images without using any label annotations². To perform the OOD detection, we fine-tune DINO with multi-class training and then compute different OOD scores (described below) to decide whether an object is an ID or an OOD sample. To be more concrete, the OOD detector is trained as a simple classifier with $K + 1$ output classes (K is the number of ID object classes), where the additional node is for the OOD and novel samples. Unlike prior works [28,12,24] that used a large amount of

² The backbone of object detectors is usually pretrained on the ImageNet1k classification [42,3], and both images and ground-truth labels are required for pretraining of the model weights. A systematic comparison between using DINO and ViT will be discussed in supplementary material.

additional OOD data during training, we do not assume the availability of such data. We therefore only use the available ID labeled data and propose to regard the background proposal patches as the OOD samples. Specifically, we crop the image patches according to the ground-truth boxes, and the background patches are randomly sampled from the proposal boxes labeled with the low IoU to the ground-truth boxes. We compute the OOD detection loss \mathcal{L}_{ood} with the cross-entropy which enforces the ID patches as the corresponding ID foreground labels and enforces the background patches as the OOD class. While it is sub-optimal to use background instances as OOD class, they still provide sufficient negative gradients for distinguishing the ID objects when the novel OOD objects are not available in the labeled set.

After the fine-tuning of the DINO model, we freeze the DINO model and experiment the following common scores for OOD detection:

1. **Mahalanobis Distance [21]:**

$$\gamma_{ood} = \max_k -(f(x) - \tilde{\mu}_k)^\top \tilde{\Sigma}^{-1} (f(x) - \tilde{\mu}_k). \quad (2)$$

where $f(x)$ represents the DINO intermediate feature vector of patch x , $\tilde{\mu}_k$ is the estimated mean vector of class k , and $\tilde{\Sigma}$ is the estimated covariance matrix. Both $\tilde{\mu}_k$ and $\tilde{\Sigma}$ are estimated by using the available labeled data.

2. **Inverse Abstaining Confidence [28,40]:**

$$\gamma_{ood} = 1 - \tilde{p}_{K+1}, \quad (3)$$

where \tilde{p}_{K+1} is the prediction confidence of $K + 1$ -th class (*i.e.*, abstention class) from the DINO classifier.

3. **Energy Score [24]:**

$$\gamma_{ood} = -T \log \sum_i^K e^{f_i(x)/T}, \quad (4)$$

where T is the temperature value, K is the number of classes, $f_i(x)$ represents the i -th index of the logit corresponding to the class i . We also consider **Shannon Entropy** and **Euclidean Distance** for the OOD scores. To decide whether an object is an ID or OOD class, we set a threshold δ_{ood} on the OOD score, and the objects with OOD scores lower than the threshold are regarded as OOD samples.

Integration with Semi-Supervised Object Detection. With the goal of filtering the OOD objects in pseudo-labels, we integrate the trained OOD detector with the SSOD method by inserting the OOD filtering after the confidence thresholding as shown in Figure 4. In other words, the pseudo-boxes are derived by sequentially applying both confidence thresholding and OOD filtering on predicted boxes from Teacher model, and they are then used to compute the unsupervised loss defined in Eq. 3 to train the Student. Note that our OOD detector is complementary to the existing semi-supervised object detection works [37,49,25,51] and can also be combined with these SSOD methods to address the open-set issue.

5 Experiments

5.1 Experimental Setting and Datasets

COCO-Open. COCO2017-**train** contains 117k images with box-level labels from 80 object categories. We randomly sample 20/40/60 classes as the ID classes and the remaining classes as the OOD classes. Specifically, since each image contains multiple objects, COCO2017-**train** is divided into the pure-ID, mixed, and pure-OOD image sets. Each image in the pure-ID set has at least one ID object and no OOD object, and each image in the pure-OOD set has at least one OOD object and no ID object. As for the mixed set, each image has at least one ID object and at least one OOD object. To make sure that OOD objects are not in the labeled set, we randomly sample the pure-ID images as a labeled set and the remaining data as the unlabeled set. Thus, the unlabeled set contains images from the pure-ID, pure-OOD, and mixed sets.

COCO-OpenImages. To examine all methods in a large-scale OSSOD experiment, we use COCO2017-**train** as the labeled set and OpenImagesv5 [19] as the unlabeled set. OpenImagesv5 contains 1.7M images with 601 object categories, and most object categories do not exist in COCO (and are hence OOD).

We also include more experiments in the supplementary material.

5.2 Comparison between online and offline OOD detection

We first present a comparison between online and offline detectors on **OOD detection** tasks (object detection results shown in Section 5.3), and we evaluate them with the standard OOD detection metrics, including area under the ROC curve (AUROC) and false-positive rate under the true-false rate with $x\%$ (FPR@TNR x). We consider several OOD detection methods [11,21,24,12,34,34,14,41]³. For a fair and thorough comparison, we adapt these methods from image classification to object detection tasks and implement them in a unified framework with the same implementation details such as batch size and optimization. To better understand the efficacy of each OOD method and prevent the interference of pseudo-labeling from affecting the performance, we do not use the pseudo-labels of unlabeled data and only use the available labeled data to train all OOD detectors (adding pseudo-labels worsen online OOD detection methods as we discussed in Section 4.2).

As presented in Table 1, we list our observations as follows,

1) *offline OOD detectors are significantly better than online OOD detectors on all evaluation metrics*, and this trend also supports that the disentanglement between OOD detector and object detector can alleviate the interference between OOD detection and object detection. It is worth repeating that the feature backbone of Faster-RCNN is pretrained on ImageNet1k with label annotations, so the superiority of the DINO on OOD detection does not come from using extra

³ As the problem setting of GODIN [14] and GSD [41] is OOD detection without OOD samples, we adapt them to our setup and enhance the performance by using an linear classifier to predict background samples.

Table 1. Evaluation of OOD detection for object detection. We sample 20 classes from COCO as ID objects, and $4k$ pure-ID images are selected as labeled images.

OOD Models	Methods	OoD Scores γ_{ood}	AUROC \uparrow	FPR75 \downarrow	FPR95 \downarrow
Online (ROIhead)	Vanilla	MSP [11]	67.0	58.4	92.3
		Energy [24]	75.5	36.8	83.6
		Entropy	75.9	38.5	83.1
		Mahalanobis [21]	50.2	83.0	98.1
		Euclidean	56.3	74.3	96.1
	OE [12]	MSP	67.0	55.0	89.1
	OVA [34]	MSP	73.0	45.7	90.0
	GODIN [14]	Cosine $h(x)$	77.8	33.8	77.4
	GSD [41]	Feat. angle	78.7	32.1	73.9
	Offline (DINO)	Ours	IAC [40]	83.6	22.4
Energy			89.6	12.2	47.5
Entropy			88.9	12.6	51.1
Mahalanobis [21]			81.8	25.6	57.6
Euclidean			90.8	10.7	38.6

images and labels. In addition, as mentioned in Section 4.3, another merit of the offline OOD detector is its robustness to the instability to pseudo-labeling used in SSOD due to the isolated training of OOD detectors and object detectors.

2) Among the OOD scores used in offline detectors, simply applying the Euclidean distance performs the bests on all metrics. The popular Mahalanobis distance cannot lead to satisfactory results due to the inaccurate covariance matrix estimated limited amounts of labeled data.

Additionally, in the supplementary material, we also list other in-depth analyses and discussions on OOD detection, including using supervised ViT-B as an offline OOD detector, and other offline OOD detection methods with different ID/OOD class splits and different degrees of supervision. Owing to the superior performance on OOD detection, we apply the offline OOD detectors in the following OSSOD experiments.

5.3 Experiments on Open-set Semi-Supervised Object Detection

In addition to the OOD detection performance shown in the previous section, we further advance to OSSOD tasks and assess the object detection performance under different OSSOD conditions, including different combinations of unlabeled sets, varying numbers of ID classes, different degrees of supervision, and larger-scale OSSOD tasks (COCO-additional and COCO-GOI).

Effect of different combinations of unlabeled sets. We experiment with the existing SSOD method (UT [25]) and our proposed OOD detectors by using different combinations of unlabeled data, including the pure-ID, the mixed, and the pure-OOD sets. The purpose of this experiment is to verify the existence of open-set issue/semantic expansion and investigate whether the OOD samples affect the SSOD methods.




Fig. 6. Using more unlabeled images containing OOD objects (*i.e.*, Mix and OOD) cannot lead to a better result for UT, while applying our OOD filtering can alleviate the performance drop and improve the accuracy. We experiment on COCO-Open with 4k images from 40 ID classes.

Table 2. Mean average precision of COCO-Open when **varying the number of ID objects**. We experiment on COCO-Open with 4k labeled images from 20/40/60 ID classes. We run each method 3 times and report the standard deviation.

Num. of ID/OOD objects	20/60	40/40	60/20
Label-only	16.89±2.6	15.98±0.49	16.64±0.59
UT	18.37±1.67 (+1.48)	20.28±0.85 (+4.29)	23.09±0.25 (+6.45)
UT + OF-DINO	23.43±2.19 (+6.54)	22.91±0.28 (+6.93)	24.89±0.34 (+8.25)

As shown in Figure 6, UT achieves 22.10 mAP when the pure-ID images are used in the unlabeled set, while the performance drops when the mixed and pure-OOD sets, containing OOD objects, are added. This shows that OOD objects in the unlabeled set do affect the effectiveness of the existing SS-OD method, and the trend is also contrary to the common belief that semi-supervised methods can derive more performance gains by using more unlabeled data. Therefore, with our proposed OOD filtering (OF-DINO), the performance of the object detector does not degrade but instead improves by using the mixed set, and this reflects the importance of eliminating the detrimental effect of OOD objects.

Varying numbers of ID classes. We also experiment by varying the number of ID classes in COCO-Open, and we first find that UT shows less performance gain against label-only baseline when there are fewer ID classes (see Table 2). This is because the semantic expansion/open-set issue becomes more severe when more OOD objects appear in the unlabeled set, and such an issue can be alleviated by the OOD filtering. With the OOD filtering, we can obtain a larger improvement gain against UT, and it improves 6.54 mAP against the labeled-only baseline and 5.06 mAP against UT in the case of 20 ID classes. This verifies the efficacy of the OOD filtering using the different number of ID/OOD classes.

Different degrees of supervision. In addition, we consider varying the number of labeled images in COCO-Open as shown in Table 3, and we find that using the OOD detector can also consistently improve against the label-only baseline and UT under different numbers of labeled images. When fewer labeled images are used, we could obtain more improvement gain by filtering OOD samples.

Analysis on pseudo-labels. To better understand how our OOD detector helps the learning of the SSOD method in the open-set condition, we use the ground-truth labels to measure the number of ID/OOD boxes in pseudo-labels

Table 3. Mean average precision of COCO-open under **different degree of supervision**. We experiment on COCO-Open with 1/2/4k labeled images from 20 selected ID classes. We run each method 3 times and report the standard deviation.

Num. of Labeled Images	1,000	2,000	4,000
Label-only	10.20± 0.34	11.84± 0.33	16.35± 0.28
UT	11.77±0.38 (+1.57)	13.87±0.68 (+2.03)	18.23±0.47 (+1.88)
UT + OF-DINO	16.80±0.53 (+6.60)	18.10±0.71 (+6.26)	22.56±0.51 (+6.21)




Fig. 7. Using OOD detector can (a) suppress the OOD objects while (b) maintain the ID objects in pseudo-labels, and this leads to (c) the improvement on OSSOD. We experiment on COCO-Open with 500 labeled images from 20 ID classes.

(note that we only use ground-truth labels for the analysis and do not use them in the training). As illustrated in Figure 7, our OOD detector can effectively suppress OOD objects in pseudo-labels without sacrificing the ID objects, and thus this alleviates the semantic expansion issue mentioned in Section 4.1 and leads to improvement for the OSSOD task.

COCO-GOI. We also consider a more challenging experiment, where we use unlabeled `OpenImagesv5` to improve the object detector trained on `COCO2017-train`. Although `OpenImagesv5` is substantially larger than COCO datasets, the object categories in `OpenImagesv5` (601 classes) are more diverse than the object categories in MSCOCO (80 classes). This implies there are more distracting OOD objects that can affect the performance of existing SSOD methods. In addition to UT, we also consider a fully-supervised baseline, which uses ground-truth labels from *both* `COCO2017-train` and `OpenImagesv5` to train an object detector. To be more specific, we first manually label the correspondence between 601 classes in `OpenImagesv5` and 80 classes in `COCO2017-train`, and we only use labels of 80 COCO-classes in `OpenImagesv5`. We present the details of the correspondence in our supplementary material.

As presented in Table 4, with our OOD filtering mechanism (DINO) we can improve UT from 41.81 mAP to 43.14 mAP. Another interesting result is that our framework can perform *even better than the model with the ground-truth labels from OpenImages*, and this indicates that using pseudo-labels generated from our framework might be more effective than the ground-truth labels provided in `OpenImages`. This is potentially caused by noisy labels from human annotations. In addition, we also find that some COCO object categories are rare

Table 4. Experimental results of COCO-OpenImages.

Methods	Labeled	Unlabeled	mAP
Fully-supervised	COCO	-	40.90
Fully-supervised	COCO+OpenImage	-	42.91
Unbiased Teacher [25]	COCO	OpenImage	41.81
Unbiased Teacher+OF-DINO	COCO	OpenImage	43.14

in **OpenImages**, and this potentially limits further improvements by exploiting **OpenImages** as an unlabeled set (more discussion in the supplementary material). **Limitations and future works.** While addressing OSSOD, we do not address other issues such as covariate shift and mismatch in object category distributions between datasets. The offline OOD detector is an individual module from the object detector, so it requires more computational resources in the training stage. However, this concern does not exist as we remove the offline OOD detector and only use the object detector in the inference stage. Our key message is that combining an offline OOD detection module and an SSOD method is a simple yet effective solution to address OSSOD tasks. Based on this integrated framework, there will be more advanced techniques for both SSOD and OOD detection methods, which can potentially improve the performance on OSSOD tasks.

6 Conclusion

Unlike traditional closed-set Semi-Supervised Object Detection, where the majority of the unlabeled images are from in-distribution datasets, we aim to train object detector with *unconstrained* unlabeled images – Open-Set Semi-Supervised Object Detection (OSSOD). In our proposed setting, unlabeled training images can contain object classes that are out-of-distribution, *i.e.*, classes not included in the labeled set. To the best of our knowledge, this is the first work investigating OSSOD. We present comprehensive studies to understand the challenges for OSSOD and analyze why the performance of existing SoTA methods degrade even when trained with more unlabeled images. To overcome the challenges, we consider both online and offline OOD detectors and show that a simple yet effective offline OOD detector can further improve the SSOD models when training with large-scale unlabeled datasets, *e.g.*, COCO-OpenImage. We achieved state-of-the-art results across existing and new benchmark settings, and some interesting directions such as more accurate and efficient OOD detectors are worth exploring for future research.

Acknowledgments. We thank Ross Girshick for many helpful suggestions and valuable discussions. Yen-Cheng Liu and Zolt Kira were partly supported by DARPA’s Learning with Less Labels (LwLL) program under agreement HR0011-18-S-0044, as part of their affiliation with Georgia Tech.

References

1. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: *Advances in Neural Information Processing Systems (NeurIPS)*. pp. 5049–5059 (2019) [1](#), [3](#)
2. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294* (2021) [2](#), [8](#)
3. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155* (2019) [8](#)
4. Dhamija, A., Gunther, M., Ventura, J., Boulton, T.: The overlooked elephant of object detection: Open set. In: *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)* (2020) [6](#)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2021) [8](#)
6. Du, X., Wang, Z., Cai, M., Li, Y.: Vos: Learning what you don't know by virtual outlier synthesis. *arXiv preprint arXiv:2202.01197* (2022) [4](#)
7. Fort, S., Ren, J., Lakshminarayanan, B.: Exploring the limits of out-of-distribution detection. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021) [4](#), [8](#)
8. Girish, S., Suri, S., Rambhatla, S.S., Shrivastava, A.: Towards discovery and attribution of open-world gan generated images. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 14094–14103 (2021) [4](#)
9. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2022) [3](#)
10. Guo, H., Mao, Y., Zhang, R.: Mixup as locally linear out-of-manifold regularization. In: *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. vol. 33, pp. 3714–3722 (2019) [1](#), [3](#)
11. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2017) [4](#), [7](#), [10](#), [11](#)
12. Hendrycks, D., Mazeika, M., Dietterich, T.: Deep anomaly detection with outlier exposure. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2019) [2](#), [4](#), [7](#), [8](#), [10](#), [11](#)
13. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)* (2020) [1](#), [3](#)
14. Hsu, Y.C., Shen, Y., Jin, H., Kira, Z.: Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020) [2](#), [10](#), [11](#)
15. Huang, J., Fang, C., Chen, W., Chai, Z., Wei, X., Wei, P., Lin, L., Li, G.: Trash to treasure: Harvesting ood data with cross-modal matching for open-set semi-supervised learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8310–8319 (2021) [4](#)

16. Huynh, D., Kuen, J., Lin, Z., Gu, J., Elhamifar, E.: Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. arXiv preprint arXiv:2111.12698 (2021) [3](#)
17. Joseph, K., Khan, S., Khan, F.S., Balasubramanian, V.N.: Towards open world object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [3](#)
18. Kim, D., Lin, T.Y., Angelova, A., Kweon, I.S., Kuo, W.: Learning open-world object proposals without learning to classify. IEEE Robotics and Automation Letters (2022) [3](#)
19. Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Kamali, S., Mallocci, M., Pont-Tuset, J., Veit, A., Belongie, S., Gomes, V., Gupta, A., Sun, C., Chechik, G., Cai, D., Feng, Z., Narayanan, D., Murphy, K.: Openimages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from <https://storage.googleapis.com/openimages/web/index.html> (2017) [10](#)
20. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: Proceedings of the International Conference on Learning Representations (ICLR) (2017) [1](#), [3](#)
21. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) [2](#), [4](#), [7](#), [9](#), [10](#), [11](#)
22. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018) [4](#)
23. 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: Proceedings of the European Conference on Computer Vision (ECCV) (2014) [1](#)
24. Liu, W., Wang, X., Owens, J.D., Li, Y.: Energy-based out-of-distribution detection. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) [4](#), [7](#), [8](#), [9](#), [10](#), [11](#)
25. Liu, Y.C., Ma, C.Y., He, Z., Kuo, C.W., Chen, K., Zhang, P., Wu, B., Kira, Z., Vajda, P.: Unbiased teacher for semi-supervised object detection. In: Proceedings of the International Conference on Learning Representations (ICLR) (2021) [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [9](#), [11](#), [14](#)
26. Luo, H., Cheng, H., Gao, Y., Li, K., Zhang, M., Meng, F., Guo, X., Huang, F., Sun, X.: On the consistency training for open-set semi-supervised learning. arXiv preprint arXiv:2101.08237 (2021) [4](#)
27. Miller, D., Sünderhauf, N., Milford, M., Dayoub, F.: Uncertainty for identifying open-set errors in visual object detection. arXiv preprint arXiv:2104.01328 (2021) [6](#)
28. Mohseni, S., Pitale, M., Yadawa, J., Wang, Z.: Self-supervised learning for generalizable out-of-distribution detection. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI) (2020) [4](#), [8](#), [9](#)
29. Nalisnick, E., Matsukawa, A., Teh, Y.W., Gorur, D., Lakshminarayanan, B.: Do deep generative models know what they don't know? In: Proceedings of the International Conference on Learning Representations (ICLR) (2019) [4](#)
30. Pidhorskyi, S., Almoheisen, R., Adjero, D.A., Doretto, G.: Generative probabilistic novelty detection with adversarial autoencoders. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) [4](#)
31. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems (NeurIPS). pp. 91–99 (2015) [6](#)

32. Sabokrou, M., Khaloeei, M., Fathy, M., Adeli, E.: Adversarially learned one-class classifier for novelty detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [4](#)
33. Saito, K., Hu, P., Darrell, T., Saenko, K.: Learning to detect every thing in an open world. arXiv preprint arXiv:2112.01698 (2021) [3](#)
34. Saito, K., Kim, D., Saenko, K.: Openmatch: Open-set consistency regularization for semi-supervised learning with outliers. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) [4](#), [7](#), [10](#), [11](#)
35. Sajjadi, M., Javanmardi, M., Tasdizen, T.: Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 1163–1171 (2016) [1](#), [3](#)
36. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) [1](#), [2](#), [3](#)
37. Sohn, K., Zhang, Z., Li, C.L., Zhang, H., Lee, C.Y., Pfister, T.: A simple semi-supervised learning framework for object detection. arXiv preprint arXiv:2005.04757 (2020) [1](#), [3](#), [5](#), [8](#), [9](#)
38. Tang, Y., Chen, W., Luo, Y., Zhang, Y.: Humble teachers teach better students for semi-supervised object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3132–3141 (2021) [1](#), [3](#), [4](#), [5](#), [6](#)
39. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in neural information processing systems (NeurIPS). pp. 1195–1204 (2017) [1](#), [3](#)
40. Thulasidasan, S., Thapa, S., Dhaubhadel, S., Chennupati, G., Bhattacharya, T., Bilmes, J.: An effective baseline for robustness to distributional shift. arXiv preprint arXiv:2105.07107 (2021) [9](#), [11](#)
41. Tian, J., Yung, D., Hsu, Y.C., Kira, Z.: A geometric perspective towards neural calibration via sensitivity decomposition. In: Advances in Neural Information Processing Systems (NeurIPS) (2021) [2](#), [7](#), [10](#), [11](#)
42. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019) [8](#)
43. Xu, M., Zhang, Z., Hu, H., Wang, J., Wang, L., Wei, F., Bai, X., Liu, Z.: End-to-end semi-supervised object detection with soft teacher. arXiv preprint arXiv:2106.09018 (2021) [4](#), [5](#), [6](#)
44. Yang, Q., Wei, X., Wang, B., Hua, X.S., Zhang, L.: Interactive self-training with mean teachers for semi-supervised object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
45. Yu, Q., Ikami, D., Irie, G., Aizawa, K.: Multi-task curriculum framework for open-set semi-supervised learning. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) [4](#)
46. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 6023–6032 (2019) [1](#), [3](#)
47. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [3](#)
48. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: Proc. International Conference on Learning Representations (ICLR) (2018) [1](#), [3](#)

49. Zhou, Q., Yu, C., Wang, Z., Qian, Q., Li, H.: Instant-teaching: An end-to-end semi-supervised object detection framework. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [1](#), [3](#), [4](#), [5](#), [6](#), [8](#), [9](#)
50. Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. arXiv preprint arXiv:2201.02605 (2022) [3](#)
51. Zhu, C., Chen, F., Shen, Z., Savvides, M.: Soft anchor-point object detection. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) [8](#), [9](#)
52. Zong, B., Song, Q., Min, M.R., Cheng, W., Lumezanu, C., Cho, D., Chen, H.: Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018) [4](#)