

Random Boxes Are Open-world Object Detectors

Yanghao Wang¹, Zhongqi Yue^{1,2}, Xian-Sheng Hua³, Hanwang Zhang¹

¹Nanyang Technological University, ²Damo Academy, Alibaba Group, ³Terminus Group

scuwyh2000@gmail.com, yuez0003@ntu.edu.sg, huaxiansheng@gmail.com, hanwangzhang@ntu.edu.sg

Abstract

We show that classifiers trained with random region proposals achieve state-of-the-art Open-world Object Detection (OWOD): they can not only maintain the accuracy of the known objects (w/ training labels), but also considerably improve the recall of unknown ones (w/o training labels). Specifically, we propose RandBox, a Fast R-CNN based architecture trained on random proposals at each training iteration, surpassing existing Faster R-CNN and Transformer based OWOD. Its effectiveness stems from the following two benefits introduced by randomness. First, as the randomization is independent of the distribution of the limited known objects, the random proposals become the instrumental variable that prevents the training from being confounded by the known objects. Second, the unbiased training encourages more proposal explorations by using our proposed matching score that does not penalize the random proposals whose prediction scores do not match the known objects. On two benchmarks: Pascal-VOC/MS-COCO and LVIS, RandBox significantly outperforms the previous state-of-the-art in all metrics. We also detail the ablations on randomization and loss designs. Codes are available at <https://github.com/scuwyh2000/RandBox>.

1. Introduction

Different objects may share common visual features. For example, the detection of “horse” helps that of “cow” as they are both “four-legged” cattle. Therefore, it is possible to generalize the detector trained from a limited object class inventory to detect visually similar objects whose classes are out of the inventory. As shown in the right column of Figure 1, besides the conventional detection of a close-world object inventory, Open-world Object Detection [23] (OWOD) can detect out-of-inventory objects and label them as “unknown”.

To exploit the feature transfer from “known” to “unknown”, the training loss of existing OWOD methods [23, 20, 55] has the following break-downs (Figure 1):

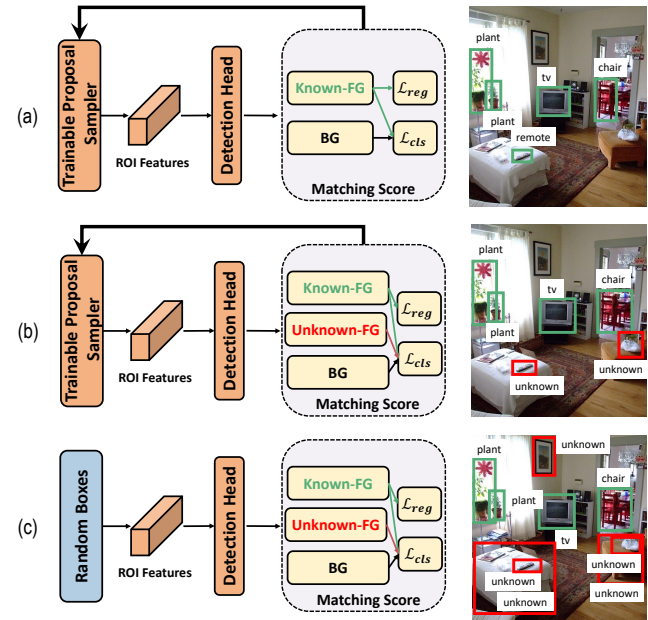


Figure 1: The paradigm of (a) closed-world object detection, (b) existing open-world object detection, and (c) our RandBox that uses random proposals.

Known-FG (foreground): like conventional closed-world detection (Figure 1a), we calculate the matching scores between the predicted region proposal labels and the ground-truth labels of known objects, where the label is a pair of class and bbox (bounding box). In particular, for two-stage detectors like Faster R-CNN [42], the “proposal sampler” and “matching score” are region proposal network and bbox IoU (intersection over union); for end-to-end detectors like DETR [9], they are query transformer decoder and bipartite matching score. For the matched proposals, the loss is class cross-entropy \mathcal{L}_{cls} and bbox regression \mathcal{L}_{reg} .

BG (background): for conventional detection, all the mismatched proposals are considered background, whose loss is only \mathcal{L}_{cls} with the label “BG”; for OWOD, we only use the mismatched proposals excluding the Unknown-FG proposals below.

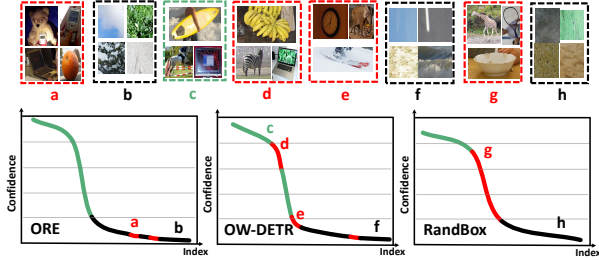


Figure 2: The sorted foreground confidence assigned by two baselines (ORE [23], OW-DETR [20]) and our RandBox to the generated proposals in training. Green, red and black correspond to “known”, “unknown” and “BG”, respectively, with examples showing on the top row.

Unknown-FG: this is the key difference from conventional detection. Its design heuristic is that, for the mismatched proposals with class confidence larger than a threshold—probably unknown objects—the loss is only \mathcal{L}_{cls} with a special label “unknown” w/o bbox regression. Note that this heuristic is based on the assumption of feature transfer [56, 10]—visually similar proposals have similar class confidences (see Section 3.1 for the review of the loss implementation in two-stage and end-to-end detection)—the more unknown objects the Unknown-FG collects, the better the OWOD.

However, in practice, Unknown-FG usually has a low recall of unknown objects. As shown in Figure 2, in existing methods, most proposals of high confidence are matched as Known-FG; the mismatched proposals, mixed with both unknown objects and BG, all have low confidence, losing a gradual confidence decrease to differentiate Unknown-FG from BG (the red and black colors are inseparable). One may easily understand that this failure is predictable because, yet, we don’t have the ground-truth labels for unknown objects, so the proposal sampler trained on known objects will be inevitably biased towards them. In fact, such bias is ubiquitously in any bootstrapping learning [43] with limited supervision, such as semi-supervised learning [34, 49] and reinforcement learning [24, 53]. Unfortunately, existing OWOD methods do not properly address the bias.

In this paper, we point out that the bias is caused by the confounding effect [37, 39] of the training data with limited known object labels that mislead the feature transfer. The effect can be understood in this way: if we only let the training data determine what region proposals to be used to extract the region features for detector training (Figure 1b), we cannot avoid the trivial loss minimization from only detecting the known objects. To this end, we propose to remove the confounding effect by introducing a random Instrumental Variable [5]—the proposed RandBox—that is independent of the training data. Intuitively, RandBox sim-

ulates a Randomized Controlled Experiment (RCT) [8] to learn a causal detector from region features to object labels, by inviting an independent 3rd-party proposal sampler to explore more possible locations of unknown objects. As shown in Figure 2 RandBox, the recall of unknown objects is greatly improved by RandBox, due to the fact that the confidence score change of the proposals from known to unknown is much more separable. We will detail the theory behind RandBox in Section 4. To get a qualitative sense of open-world detection task, see Figures 3.

RandBox is simple yet effective (Section 3.2). For each training image in every iteration, we sample 500 random boxes as proposals from a 4-dimensional Gaussian noise (center_x, center_y, height, width). For these random boxes, we crop their RoI (region of interest) features from a backbone feature map, and feed them into a detection head to obtain \mathcal{L}_{cls} and \mathcal{L}_{reg} as in (Figure 1c), where the proposed matching score that encourages exploration is detailed in Section 3.2. At inference, we remove the randomness and grid-sample 10,000 boxes per image by 10 scales, 10 shapes, and 100 spatial locations. We use NMS [35] (Non-maximum Suppression) to prune the 10,000 boxes for final detection. Overall, our contributions include:

- Our proposed RandBox is a Fast R-CNN based architecture for OWOD, which leverages random proposals to remove the confounding effect from the known objects, and uses an improved matching score to encourage proposal exploration.
- By removing the confounding effect, we learn an unbiased detector that improves both known and unknown detection accuracy. In particular, we achieve new state-of-the-art results on Pascal-VOC/MS-COCO [29].
- We show that our RandBox detector trained on Pascal-VOC/MS-COCO remains robust when transferring to long-tailed datasets such as LVIS [19].

2. Related Work

Open-World Object Detection (OWOD). Object detection systems are significantly advanced recently [41, 57, 42, 18, 48]. Yet they are based on the conventional close-world paradigm, which assumes that all objects of interests are annotated. In practice, an object detector deployed in real life is constantly challenged by the vast open and dynamic visual world [33, 23, 25]. To bridge the gap, open-set detection [33, 32, 14] aims to reject out-of-inventory objects as “unknown”, and other works focus on incrementally learning [40, 1, 31] a detector to recognize new object classes without forgetting old ones. Inspired by the open-world setting proposed in image classification [6, 7], recent works [23, 20, 55, 58] unify open-set detection and incremental learning as OWOD, which are based on two stage detectors like Faster R-CNN [42], or end-to-end ones like



Figure 3: **Example results.** We show the ground-truth (GT) locations of unknown objects, those predicted by OW-DETR [20], and those predicted by our RandBox. RandBox not only detects the annotated unknown objects (first 2 images), but even discovers valid unknown ones without annotations (e.g., *camp* in image 3, *bracelet* in image 5, *banner* in image 6).

DETR [57]. We systematically analyze the failure of these existing works—low recall of unknown objects—which is caused by the confounding effect [37, 36], and remove it using random proposals to learn a causal detector.

Debiasing by Randomization. Randomization is an effective tool to debias across many tasks and domains. For example, using randomization as an instrument variable [12, 4, 5] is a classic tool to identify causal effect in social studies [15]. In reinforcement learning [53], randomization is used to encourage exploration [54, 27], such that the agent is not biased to exploiting spurious shortcuts. Several works [3, 50] in adversarial training [17] leverages randomization to build a robust model not biased to feature that can be manipulated in adversarial attack, e.g., by image perturbation [59]. In fact, data augmentation that is prevailing in vision [45, 21, 47] and language [46, 44] also leverages randomness to remove augmentation-related context bias in training (e.g., using green grass to classify cow). In our work, we present a novel detection framework that leverages random proposals to learn causal effect in OWO.

3. Method

In OWO, the task is to detect a gradually expanding set of known inventory classes K , while simultaneously rejecting the out-of-inventory classes set U as “unknown”. The training dataset contains images where known objects are labelled. Specifically, for a training image containing n known objects, it is labelled with a set $\{(b_i, y_i)\}_{i=1}^n$, where b_i specifies the bounding box (bbox) of the i -th object (i.e., center_x, center_y, height, width), and $y_i \in K$ is its class label. We drop the subscript i for simplicity when the context is clear. The goal in OWO is to train a detector, whose input is an image, and output is a set of m predictions $P = \{(\hat{b}_i, \hat{y}_i)\}_{i=1}^m$, where \hat{b}_i is the i -th bounding box prediction (predicted bbox), and $\hat{y}_i \in \mathbb{R}^{|K|+2}$ denotes the clas-

sification logit for each class in $K \cup \{\text{“unknown”}, \text{“BG”}\}$.

3.1. Preliminaries: Existing OWO Methods

Existing methods adopt either the two-stage Faster R-CNN [42] or the end-to-end DETR [9] as the detector. In training, the predictions P generated by the detector are partitioned into three subsets: Known-FG, Unknown-FG and BG, which are used to calculate the classification loss and bbox regression loss. We detail each part below:

Detector. Given an image, a detector generates m proposals, extracts ROI features x_i for each proposal, and finally computes bounding box \hat{b}_i and logits \hat{y}_i using a classifier to form P . The two detector frameworks mainly differ in the first step: 1) Faster R-CNN leverages a region proposal network (RPN) pre-trained on the labelled data. It outputs the top- m proposals with the largest *objectness* (i.e., probability of containing a known object). 2) DETR is an end-to-end detector. It first transforms an image into tokens using a combination of CNN backbone and transformer encoder. Then a transformer decoder uses the image tokens to generate the proposals. More details are in Appendix.

Known-FG is selected from all predictions P by matching with the labels $\{(b_i, y_i)\}_{i=1}^n$ based on a Known-FG score. To compute the score for a prediction $(\hat{b}, \hat{y}) \in P$ with a ground-truth (b, y) , Faster R-CNN uses the IoU between \hat{b} and b , and DETR uses the bipartite matching score, which additionally considers if \hat{y} has a large value on the class y . As the process is the same as that in closed-set detection, we leave other details in Appendix. After matching, we denote the Known-FG subset as P^K . For $(\hat{b}, \hat{y}) \in P^K$, we denote its matched ground-truth as (b^*, y^*) .

Unknown-FG & BG are selected from the unmatched predictions $P \setminus P^K$. Specifically, current methods compute an Unknown-FG score for each prediction, and select 5 predictions with the largest score as the Unknown-FG. For the score, methods based on Faster R-CNN [42, 23] use object-

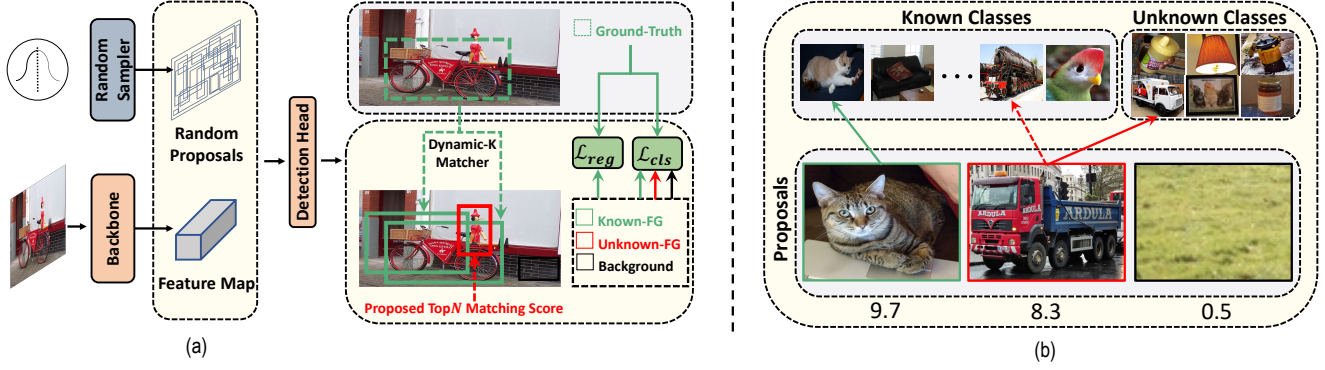


Figure 4: (a) **RandBox Pipeline**. Our proposals are randomly generated. We obtain the Known-FG subset using a dynamic- k matcher, and the Unknown-FG subset based on our proposed matching score. Our training losses follow existing methods. (b) **Proposed Matching Score** is the sum of likelihood of belonging to each foreground classes (including “unknown”). The proposals of known (green) and unknown (red) objects both have large matching score, *i.e.*, high likelihood on the ground-truth class (solid arrows) or by feature transfer (dashed arrow), while BG proposals (black) have low score.

ness, and those based on DETR [20] use the average value of the ROI feature. We denote the Unknown-FG subset as P^U . Finally, the rest predictions become the BG subset P^B .

Training Loss.

We calculate the the loss on each of the three subsets. The loss on Known-FG \mathcal{L}^K has the same form with closed-set detection:

$$\mathcal{L}^K = \sum_{i=1}^{|P^K|} \mathcal{L}_{cls}(\hat{y}_i, y_i^*) + \lambda \mathcal{L}_{reg}(\hat{b}_i, b_i^*), \quad (1)$$

where \mathcal{L}_{cls} is the cross-entropy loss or focal loss [28] to handle class-imbalance, λ is a balancing weight, and \mathcal{L}_{reg} is the smooth L1 regression loss [18] to rectify the predicted bbox. Details of \mathcal{L}_{cls} and \mathcal{L}_{reg} are in Appendix. As Unknown-FG and BG have no annotated bbox, we only calculate \mathcal{L}_{cls} on them:

$$\mathcal{L}^{UB} = \sum_{i=1}^{|P^U|} \mathcal{L}_{cls}(\hat{y}_i, \text{“unknown”}) + \sum_{i=1}^{|P^B|} \mathcal{L}_{cls}(\hat{y}_i, \text{“BG”}), \quad (2)$$

Overall, the detector parameters θ are optimized with:

$$\min_{\theta} \mathcal{L}^K + \beta \mathcal{L}^{UB}, \quad (3)$$

where β is a weight with ablation in Section 5.2.

OWOD Pipeline. The detector is first randomly initialized. Then it is trained with the labelled data on K , and evaluated with data on both K and U (Task 1). Next, additional labelled data on a subset of classes $K' \subset U$ is provided, *i.e.*, the known classes set expands as $K \leftarrow K \cup K'$, and the unknown one becomes $U \leftarrow U \setminus K'$. The detector trained on Task 1 is fine-tuned and evaluated with the updated K and U (Task 2). This continues for the following tasks.

3.2. Proposed RandBox

Existing methods suffer from the low recall of the unknown objects in Unknown-FG. To tackle this, our proposed RandBox has two improvements: 1) We propose to generate region proposals randomly, instead of using a proposal sampler trained on the known objects, which explores more possible unknown object locations. 2) We propose a matching score without penalizing proposals that are not matched with the known objects, which explores more possible unknown object proposals. We detail the two improvements below, and explain why our RandBox solves the existing problems from a causal view in Section 4.

Detector with Random Proposals. In training, we randomly generate 500 bboxes as region proposals on each image (Figure 4a). For each bbox, we sample 4 random real numbers in $[0, 1]$, where each number is first drawn from the standard Gaussian distribution, truncated to $[-2, 2]$ and then scaled linearly to $[0, 1]$. The 4 numbers correspond to the horizontal and vertical coordinate of the bbox center, as well as the width and height of the bbox, respectively (proportional to image size). The width and height of out-of-boundary bboxes are adjusted to stay in the image. In testing, to remove the prediction randomness, we use 10,000 pre-defined bboxes as proposals on each image, which cover 10 scales, 10 aspect ratios and 100 spatial locations (details in Appendix). We use non-maximum suppression to prune the bboxes as per standard practice [35]. We highlight the following points:

- Our region proposals are randomly sampled for every image in each iteration throughout the training.
- While we use random sampling, the number of our random proposals is in fact *comparable* with that used in the existing methods, *e.g.*, ORE [23] uses 1024 proposals.

Known-FG. We adopted the dynamic- k matcher [11] as shown in Figure 4a. Specifically, each ground-truth bbox b is matched with k proposals based on the bipartite matching score, where k is dynamically selected as the sum of the IoU between each proposal and b .

Unknown-FG. We propose a novel **matching score** to form Unknown-FG. Specifically, we select the top N proposals with the largest matching score (ablation of N in Section 5.2). Among the selected proposals, those already in Known-FG are removed to form our Unknown-FG subset. We use the following equation to compute our matching score for a proposal with the classification logits \hat{y} :

$$s(\hat{y}) = \sum_{i=1}^{|K|+1} \frac{1}{1 + \exp(-\hat{y}_i)}, \quad (4)$$

where we compute and sum the sigmoid of the $|K|$ known classes and “+1” unknown class, and \hat{y}_i denotes the logit for the class i . Note that we implement \mathcal{L}_{cls} with the BCE loss [28]. Hence our score essentially computes how likely \hat{y} corresponds to the foreground. Specifically, the top- N (largest matching score) proposals that are not in Known-FG are pseudo-labeled as “unknown” to train the Unknown-FG logit (Eq. 2 first loss). This design is based on the assumption of feature transfer [38, 52]—similarity between Unknown-FG and Known-FG is larger than that between BG and Known-FG. As shown in Figure 4b, an unknown object proposal can have non-trivial likelihood of belonging to a known class due to shared feature (dashed arrow), while BG proposal almost shares no feature with known. Hence future unknown object proposals can have large likelihood of belonging to “unknown” (solid red arrow), increasing the unknown recall. We have the following comparisons with existing methods:

- Methods based on Faster R-CNN use the objectness predicted by RPN as the matching score to select Unknown-FG. However, RPN is trained to produce large scores *only* on the proposals matched with the known objects, thus penalizing those unmatched ones.
- Without RPN, methods based on DETR use the mean activation of the ROI feature instead. However, this heuristic is extremely unreliable (Figure 3b), and unknown object proposals are often mistakenly placed in the BG subset and penalized as “background”.
- In contrast, our matching score fairly evaluates each proposal by checking both known classes (first $|K|$ logits), as well as “unknown” ($|K|+1$ -th logit). Hence unknown object proposals can be reliably selected in the Unknown-FG subset with high likelihood of belonging to “unknown”.

Other Details. Our training objective is the same as Eq. (3). We used ResNet-50 [22] as backbone. Our RandBox is Fast R-CNN [18] based architecture and detection head is borrowed from Sparse R-CNN [48].



Figure 5: **Causal graph** of (a) existing OWOD methods, and (b) our RandBox.

4. Theory

We use a Structural Causal Model (SCM) to analyze the causalities between the key components in OWOD: training data D , region proposal R , ROI feature X , and label Y on $K \cup \{\text{“unknown”}, \text{“BG”}\}$. Note that D only contains labelled data of the known classes K . The SCM is illustrated with a causal graph in Figure 5, where each directed link denotes a causal relation between the two connected nodes. $R \rightarrow X \rightarrow Y$ represents the desired causal effect from region proposal R to label Y , as label Y describes what a region R (or its ROI feature X) contains. A detector is unbiased if it identifies $R \rightarrow X \rightarrow Y$, *i.e.*, given any region, the detector can fairly evaluate what it contains.

$R \rightarrow X \leftarrow D$. The ROI feature X is jointly determined by what the region contains ($R \rightarrow X$), as well as the feature representation learned on the training data ($D \rightarrow X$).

$D \rightarrow Y$ is because the detector leverages the training knowledge obtained from D to predict Y , *e.g.*, identifying the BG subset with prediction “BG” (Section 3.1).

$D \rightarrow R$ exists in current OWOD methods as shown in Figure 5a. This is because the region proposal network in Faster R-CNN and the transformer decoder in DETR are both trained on D to generate proposal R .

Confounding Effect. In Figure 5a, $R \leftarrow D \rightarrow Y$ is known as a backdoor path, which leads to the undesired confounding effect [37]. Specifically, as the proposal sampler is supervised by the labelled data of known objects D , the generated proposals R are inevitably biased to them ($D \rightarrow R$), *i.e.*, R only faithfully reflects the known object locations. Consequently, the training is dominated by the loss \mathcal{L}^K on Known-FG, and the detector recklessly exploits *any* discriminative feature in D towards the known objects ($D \rightarrow Y$), which may not transfer to the unknown ones. This causes the low recall of Unknown-FG in Figure 2.

Removing Confounding Effect. In our RandBox, R is randomly generated with no influence from D , *i.e.*, $D \not\rightarrow R$ in Figure 5b. This makes R an Instrument Variable [5], which removes the confounding effect by simulating a Randomized Controlled Experiment [8]. We provide a three-fold intuitive explanation: 1) The backdoor path is cut, *i.e.*, $R \not\leftarrow D \rightarrow Y$; 2) The path $R \rightarrow X \leftarrow D \rightarrow Y$ is blocked due to the collider [37] X (*i.e.*, with two incoming arrows). 3) This makes $R \rightarrow X \rightarrow Y$ the only unblocked path from R to Y . Hence when learning to predict Y from R , our de-

Method	Task 1				Task 2				Task 3				Task 4
	K-mAP	U-R	WI	A-OSE	K-mAP	U-R	WI	A-OSE	K-mAP	U-R	WI	A-OSE	K-mAP
Faster-RCNN [42]	56.4	-	0.0699	13,396	15.2	-	0.0371	12,291	6.7	-	0.0213	9,174	4.2
Faster-RCNN+Finetuning	56.4	-	0.0699	13,396	38.0	-	0.0375	12,497	30.0	-	0.0279	9,622	25.6
DDETR [57]	60.3	-	0.0608	33,270	17.9	-	0.0368	18,115	8.5	-	0.0197	9,392	6.0
DDETR+Finetuning	60.3	-	0.0608	33,270	44.8	-	0.0337	17,834	33.3	-	0.0195	10,095	29.4
ORE [23]	56.0	4.9	0.0621	10,459	39.4	2.9	0.0282	10,445	29.7	3.9	0.0156	6,803	25.3
OST [55]	56.2	-	0.0417	4,889	39.9	-	0.0213	2,546	29.6	-	0.0146	2,120	25.9
OW-DETR [20]	59.3	7.5	0.0571	10,240	42.9	6.2	0.0278	8,441	30.8	5.7	0.0156	6,803	27.8
RandBox	61.8	10.6	0.0240	4,498	45.3	6.3	0.0078	1,880	39.4	7.8	0.0054	1,452	35.4

Table 1: **OWOD Evaluation** results and comparison with existing methods. We can also achieve SOTA on previously/currently known scores, more details are in Appendix.

tor captures the causal effect free from the confounding one. The formal definition and proof are in Appendix.

5. Experiments

We introduce important experiment setups below, and leave other details in Appendix.

Datasets. We extensively evaluated RandBox on three tasks: OWOD, transfer learning and open-set detection. For **OWOD**, we used Pascal VOC [16] and MS-COCO [29] following the standard setups [23, 20]. In Task 1, the known set K contains the 20 classes in Pascal VOC, and the unknown set U contains the 60 classes unique to MS-COCO. For each of the Task 2-4, 20 classes in U are added to K , *i.e.*, Task 4 corresponds to closed-set detection with 80 classes. The training and testing images are selected from the respective split of the two datasets. For **transfer learning**, we evaluated models trained on Task 1 on the 4,809 validation images of LVIS v1.0 [19], which are from over 1000 classes that exhibits long-tailed distribution. For **open-set detection**, we followed [33] to train models on Pascal VOC train split, and evaluate them on WR1 [33].

Evaluation Metrics: On known classes, we used the standard Known-class mAP (K-mAP). On unknown classes, we used three metrics: 1) Unknown class Recall (U-R) is the recall of unknown classes at IoU threshold of 0.5. 2) Wilderness Impact (WI) [14] is computed by $(\text{closed-set precision})/(\text{open-set precision}) - 1$. A small WI is desired in OWOD, *i.e.*, precision does not drop when unknown objects are introduced in evaluation. 3) Absolute Open-Set Error (A-OSE) [33] is the predictions that are misclassified as one of the known classes.

Implementation Details: In training, we used the ResNet-50 [22] pre-trained on ImageNet [13] as our backbone. We trained RandBox using AdamW [30] optimizer with learning rate as $2.5e^{-5}$ and weight decay as $1e^{-4}$. We used the same methodology as prior work [23, 2, 51, 26] to mitigate forgetting, *i.e.*, by storing a balanced set of exemplars for model fine-tuning after each incremental step. In inference, after getting the predictions on the 10,000 pre-defined bboxes, we used Non-Maximum Suppression (NMS) [35]

to prune redundant bboxes. After pruning, the final predictions are the bboxes whose probability of belonging to a class in $K \cup \{\text{"unknown"}\}$ are at least 0.1.

5.1. Main Results

Table 1 presents the results of OWOD evaluation. We have the following observations:

1) For the conventional framework Faster-RCNN [42] and Deformable-DETR (DDETR) [57], their K-mAP drops significantly in Task 2-4, showing that they suffer from catastrophic forgetting when new classes are incrementally added. By using data replay [23] (+Finetuning), their K-mAP results improve significantly (*e.g.*, 6.0 to 29.4 on Task 4 for DDETR). Note that they are not capable to detect unknown objects (*e.g.*, U-R is not applicable).

2) By comparing existing OWOD methods with the conventional frameworks, we notice that *all* of them have reduced K-mAP (*e.g.*, the DDETR-based OW-DETR reduced 1.6 mAP on Task 4 compared with DDETR+Finetuning). This means that their capabilities of unknown detection are at the cost of known detection accuracy.

3) By comparing RandBox with existing OWOD methods on U-R and A-OSE, we observe significant improvements, *e.g.*, 3.1% U-R improvement and 391 A-OSE reduction on Task 1. This suggests that our training scheme with random boxes and the proposed matching score encourage proposal explorations, which detects and learn on more unknown objects to improve their evaluation performance.

4) In particular, by comparing K-mAP and WI of RandBox and OWOD methods, we see that RandBox also achieves the state-of-the-art (SOTA) performance, *e.g.*, absolute gains ranging from 2.4% to 8.6% in terms of K-mAP and relatively reduces ranging from 42% to 64% in terms of WI. Remarkably, our RandBox even outperforms the conventional detection frameworks, *e.g.*, Faster-RCNN and DDETR. This means that instead of sacrificing “known” to detect “unknown”, RandBox learns “unknown” to benefit “known”. This strongly validates that our approach removes confounding bias by simulating Randomized Controlled Experiment with random boxes, where our detector captures the causal effect from *any* region proposal (known,

Proposals	K-mAP	U-R	WI	A-OSE
Selective search	57.3	9.1	0.029	5,433
RPN	59.0	8.6	0.031	5,128
Query based	59.7	9.1	0.043	4,679
Random (Ours)	61.8	10.6	0.024	4,498

(a) Ablation on proposal generation method.

Match	Rand	K-mAP	U-R	WI	A-OSE
		51.8	-	0.042	7,652
Ours		59.7	9.1	0.043	4,679
Mean activation	✓	60.3	7.1	0.024	5,929
Ours	✓	61.8	10.6	0.024	4,498

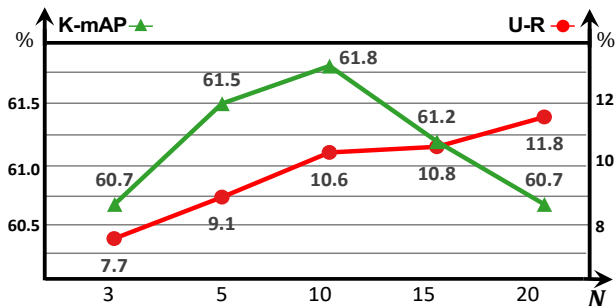
(c) Ablation on components.

distribution	num	K-mAP	U-R	WI	A-OSE
Uniform	100	60.2	8.7	0.027	4,877
Gaussian	100	60.7	9.3	0.026	4,732
Uniform	500	62.0	10.1	0.026	4,588
Gaussian	500	61.8	10.6	0.024	4,498
Uniform	1,000	61.7	10.4	0.028	4,975
Gaussian	1,000	61.8	10.6	0.030	4,763

(b) Strategy for Random Proposals.

ratio	scale	location	K-mAP	U-R	WI	A-OSE
5	5	20	60.2	8.2	0.041	3,001
10	10	20	61.6	9.6	0.032	3,715
10	10	40	61.3	10.2	0.029	4,098
10	10	100	61.8	10.6	0.024	4,498

(d) Ablation on inference bboxes.

Table 2: **Ablation results.** Models are trained on Task 1. We adopted the best performing setting as our default (in gray).Figure 6: Ablation on N by training RandBox on Task 1.

unknown or BG) to its label (Section 4). Hence we observe unconditional improvements overall.

5.2. Ablation Experiments

Choice of N . Recall that in training, we identify the top N proposals with the largest matching score to form Unknown-FG (by removing those already in Known-FG). Hence N is directly correlated with the size of Unknown-FG. The ablation results on N are given in Figure 6. Initially, increasing N benefits both known detection (K-mAP) and unknown recall (U-R). This is because 1) With more unknown proposals for optimizing Eq. (3), the detector improves in recognizing “unknown” (increased U-R). 2) To differentiate “unknown” from “known”, the detector also avoids spurious correlations of known classes in training (*i.e.*, captures causal effect), which may not generalize to known objects in testing (increased K-mAP). However, further increasing N leads to the drop in K-mAP. This is because increasing the size of Unknown-FG by large N essentially increases the weight of the classification loss on “unknown” in Eq. (2). Hence the detector learning on known classes is adversely impacted when the weight is too large. We used $N = 10$ in all our experiments.

Proposal generation methods Its ablation is in Table 2a. We use our matching score Eq. 4 and fix the same proposal number (500) to ablate the proposal generation methods. We compared our Random Proposals with Selective search, RPN and Query based proposals. We can find that our Random proposals surpass all other methods which means the randomness can make more explorations on more image regions, without biased towards the training data.

Strategy for Random Proposals. When generating random proposals in training, we experimented with different distribution (for drawing random numbers) and the number of random proposals (num), and the results are in Table 2b. Using the same number of proposals, the results on the two distributions are both competitive. Regarding the number of proposals, we observe that increasing it beyond 500 does not bring any improvement. This shows that the image locations are already sufficiently explored using 500 random bboxes. Overall, we used 500 random proposals drawn from Gaussian distribution.

Components. The main components in RandBox are the use of random proposals (Rand), and the matching process for Known-FG and Unknown-FG (Match). Their ablations are given in Table 2c. The first line corresponds to the baseline Sparse R-CNN [48]. Comparing the second line, by combining our matching process with NMS, our matching score can improve all metrics and already achieved SOTA (*e.g.*, by comparing with other OWOD methods in Table 1 Task 1). In the third and forth lines, we introduce the Random proposals and compare two kinds of matching score, *i.e.*, Mean activation and our matching score. We can find that our matching score improve all metrics significantly (*e.g.*, 1.5% on K-mAP and 3.5% on U-R) compared with the Mean activation. This strongly validates the effectiveness of each component in RandBox.

Training schedule. We show the change of U-R and K-mAP over the training and compare RandBox with ORE

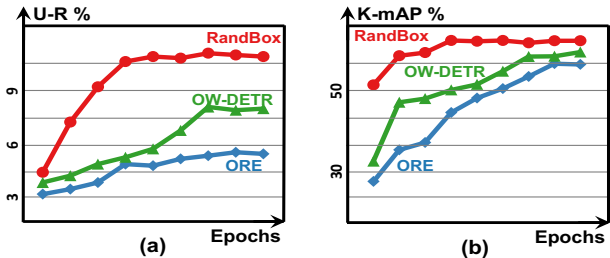


Figure 7: Training schedule compared with existing works.

Method	K-mAP	U-R	WI	A-OSE
ORE [23]	23.2	3.5	0.1743	23,225
OW-DETR [20]	25.1	6.3	0.1548	19,672
RandBox	27.7	8.4	0.0735	7,842

Table 3: Transfer evaluation results on LVIS v1.0 [19].

Evaluated on →	Pascal-VOC	Open-Set (WR1)
Faster R-CNN [42]	81.8	77.1
RetinaNet [28]	79.2	73.8
Dropout Sampling [33]	78.1	73.8
ORE [23]	81.3	71.7
OW-DETR [20]	82.1	78.6
RandBox	84.4	82.0

Table 4: Results on Pascal-VOC [16] (closed-set) and WR1 [33] (open-set).

and OW-DETR in Figure 7. Our RandBox converges quickly and trains stably, outperforming the two methods.

Inference Bboxes. We ablation on the number of pre-defined bbox ratios, scales and locations in Table 2d. Increasing the bbox number generally improves the performance due to the increased recall on both known and unknown classes. However, it also increases the number of mis-classified unknown predictions, leading to a larger A-OSE. As the performance tends to saturate with 10,000 bboxes (last row), we use this configuration for efficiency.

5.3. Transfer and Open-Set Evaluations

Transfer to LVIS v1.0. We evaluated the models trained on Task 1 on LVIS 1.0. [19] We observe significant improvements on all metrics compared to ORE and OW-DETR in Table 3 This validates the robustness of RandBox over the existing works under distribution shift.

Open-Set Object Detection. Detectors in OWO can naturally tackle open-set detection. We followed [33] and evaluated models trained on Task 1 on Pascal-VOC test split (closed-set) and WR1 (open-set). We significantly improve existing SOTAs on both evaluations (Table 4).

5.4. Additional Qualitative Results

As shown in Figure 8 first column, OW-DETR misses the known objects “mouse” and “chair” due to confusing

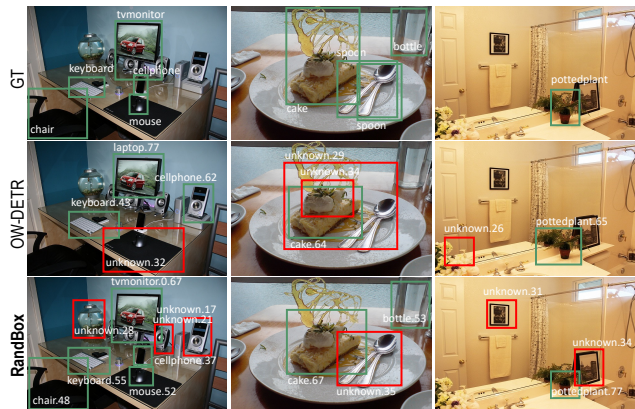


Figure 8: Additional qualitative results. We show the ground-truth labels of known objects in the first row, and compare the predictions from OW-DETR and our RandBox. **Green:** known. **Red:** unknown.

background and partial view, respectively. Our RandBox not only detects all known objects, but also identify many valid unknown objects. In the second column, RandBox predicts the overlapping spoons as “unknown”, which is indeed plausible without other prior knowledge. The third column shows that our RandBox detects the unknown object “picture frame” both in the real world and in the mirror.

6. Conclusion

We presented a novel Open-World Object Detection (OWOD) method called RandBox, trained on random proposals and surpassing existing Faster R-CNN and Transformer based OWOD. We show from a causality theoretical perspective that RandBox simulates a Randomized Controlled Experiment, which removes the confounding effect of the training data with limited known object labels. In addition, we propose a matching score without penalizing random proposals that are not matched with the known objects, which further encourages proposal explorations. Extensive evaluations of our proposed RandBox show that it significantly improves both the accuracy of the known objects and the recall of unknown ones. As future direction, we will seek additional inductive bias to remove the confounding effect and extend our approach to other detection tasks.

7. Acknowledgements

The authors would like to thank all reviewers and ACs for their constructive suggestions. This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-022) and Alibaba-NTU Singapore Joint Research Institute.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, pages 139–154, 2018. [2](#)
- [2] Takeshi Amemiya and Thomas E MaCurdy. Instrumental-variable estimation of an error-components model. *Econometrica: Journal of the Econometric Society*, pages 869–880, 1986. [6](#)
- [3] Alexandre Araujo, Laurent Meunier, Rafael Pinot, and Benjamin Negrevergne. Robust neural networks using randomized adversarial training. *arXiv preprint arXiv:1903.10219*, 2019. [3](#)
- [4] Manuel Arellano and Olympia Bover. Another look at the instrumental variable estimation of error-components models. *Journal of econometrics*, 68(1):29–51, 1995. [3](#)
- [5] Michael Baiocchi, Jing Cheng, and Dylan S Small. Instrumental variable methods for causal inference. *Statistics in medicine*, 33(13):2297–2340, 2014. [2](#), [3](#), [5](#)
- [6] Abhijit Bendale and Terrance Boulton. Towards open world recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1893–1902, 2015. [2](#)
- [7] Abhijit Bendale and Terrance E Boulton. Towards open set deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1563–1572, 2016. [2](#)
- [8] Anthony A Braga, David L Weisburd, Elin J Waring, Lorraine Green Mazerolle, William Spelman, and Francis Gajewski. Problem-oriented policing in violent crime places: A randomized controlled experiment. *Criminology*, 37(3):541–580, 1999. [2](#), [5](#)
- [9] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. [1](#), [3](#)
- [10] Hao Chen, Yali Wang, Guoyou Wang, and Yu Qiao. Lstd: A low-shot transfer detector for object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. [2](#)
- [11] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusionnet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022. [5](#)
- [12] Angus Deaton. Instruments, randomization, and learning about development. *Journal of economic literature*, 48(2):424–455, 2010. [3](#)
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [6](#)
- [14] Akshay Dhamija, Manuel Gunther, Jonathan Ventura, and Terrance Boulton. The overlooked elephant of object detection: Open set. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1021–1030, 2020. [2](#), [6](#)
- [15] Vanessa Didelez and Nuala Sheehan. Mendelian randomization as an instrumental variable approach to causal inference. *Statistical methods in medical research*, 16(4):309–330, 2007. [3](#)
- [16] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. [6](#), [8](#)
- [17] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. [3](#)
- [18] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. [2](#), [4](#), [5](#)
- [19] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. [2](#), [6](#), [8](#)
- [20] Akshita Gupta, Sanath Narayan, KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Ow-detr: Open-world detection transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9235–9244, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [8](#)
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [3](#)
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#), [6](#)
- [23] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5830–5840, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [8](#)
- [24] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996. [2](#)
- [25] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *IEEE Robotics and Automation Letters*, 7(2):5453–5460, 2022. [2](#)
- [26] Jeremias Knoblauch, Hisham Husain, and Tom Diethe. Optimal continual learning has perfect memory and is np-hard. In *International Conference on Machine Learning*, pages 5327–5337. PMLR, 2020. [6](#)
- [27] Branislav Kveton, Manzil Zaheer, Csaba Szepesvari, Lihong Li, Mohammad Ghavamzadeh, and Craig Boutilier. Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076. PMLR, 2020. [3](#)

- [28] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 4, 5, 8
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 6
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [31] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 2
- [32] Dimity Miller, Feras Dayoub, Michael Milford, and Niko Sünderhauf. Evaluating merging strategies for sampling-based uncertainty techniques in object detection. In *2019 international conference on robotics and automation (icra)*, pages 2348–2354. IEEE, 2019. 2
- [33] Dimity Miller, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. Dropout sampling for robust object detection in open-set conditions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3243–3249. IEEE, 2018. 2, 6, 8
- [34] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018. 2
- [35] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *18th international conference on pattern recognition (ICPR'06)*, volume 3, pages 850–855. IEEE, 2006. 2, 4, 6
- [36] Judea Pearl. *Causality*. Cambridge university press, 2009. 3
- [37] Judea Pearl. The do-calculus revisited. *arXiv preprint arXiv:1210.4852*, 2012. 2, 3, 5
- [38] David N Perkins, Gavriel Salomon, et al. Transfer of learning. *International encyclopedia of education*, 2:6452–6457, 1992. 5
- [39] Mohamad Amin Pourhoseingholi, Ahmad Reza Baghestani, and Mohsen Vahedi. How to control confounding effects by statistical analysis. *Gastroenterology and hepatology from bed to bench*, 5(2):79, 2012. 2
- [40] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml: An incremental task-agnostic meta-learning approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13588–13597, 2020. 2
- [41] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 1, 2, 3, 6, 8
- [43] Ellen Riloff, Janyce Wiebe, and Theresa Wilson. Learning subjective nouns using extraction pattern bootstrapping. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*, pages 25–32, 2003. 2
- [44] Julian Salazar, Davis Liang, Toan Q Nguyen, and Katrin Kirchhoff. Masked language model scoring. *arXiv preprint arXiv:1910.14659*, 2019. 3
- [45] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019. 3
- [46] Connor Shorten, Taghi M Khoshgoftaar, and Borko Furht. Text data augmentation for deep learning. *Journal of big Data*, 8:1–34, 2021. 3
- [47] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 3
- [48] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021. 2, 5, 7
- [49] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine learning*, 109(2):373–440, 2020. 2
- [50] Haotao Wang, Chaowei Xiao, Jean Kossaiji, Zhiding Yu, Anima Anandkumar, and Zhangyang Wang. Augmax: Adversarial composition of random augmentations for robust training. *Advances in neural information processing systems*, 34:237–250, 2021. 3
- [51] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*, 2020. 6
- [52] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016. 5
- [53] Marco A Wiering and Martijn Van Otterlo. Reinforcement learning. *Adaptation, learning, and optimization*, 12(3):729, 2012. 2, 3
- [54] Robert C Wilson, Elizabeth Bonawitz, Vincent D Costa, and R Becket Ebitz. Balancing exploration and exploitation with information and randomization. *Current opinion in behavioral sciences*, 38:49–56, 2021. 3
- [55] Shuo Yang, Peize Sun, Yi Jiang, Xiaobo Xia, Ruiheng Zhang, Zehuan Yuan, Changhu Wang, Ping Luo, and Min Xu. Objects in semantic topology. *arXiv preprint arXiv:2110.02687*, 2021. 1, 2, 6
- [56] Liang Zheng, Yali Zhao, Shengjin Wang, Jingdong Wang, and Qi Tian. Good practice in cnn feature transfer. *arXiv preprint arXiv:1604.00133*, 2016. 2
- [57] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 2, 3, 6

- [58] Orr Zohar, Kuan-Chieh Wang, and Serena Yeung. Prob: Probabilistic objectness for open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11444–11453, 2023. [2](#)
- [59] Alex Zwanenburg, Stefan Leger, Linda Agolli, Karoline Pilz, Esther GC Troost, Christian Richter, and Steffen Löck. Assessing robustness of radiomic features by image perturbation. *Scientific reports*, 9(1):614, 2019. [3](#)