Improving Open-World Object Localization by Discovering Background

Singh, Ashish; Jones, Michael J.; Peng, Kuan-Chuan; Chatterjee, Moitreya; Cherian, Anoop; Learned-Miller, Erik

TR2025-058 May 16, 2025

Abstract

Our work addresses the problem of learning to localize objects in an open-world setting, i.e., given the bounding box information of a limited number of object classes during training, the goal is to localize all objects, belonging to both the training and unseen classes in an image, during inference. Towards this end, recent work in this area has focused on improving the characterization of objects either explicitly by proposing new objective functions (localization quality) or implicitly using object-centric auxiliary-information, such as depth information, pixel/region affinity map etc. In this work, we address this problem by incorporating background information to guide the learning of the notion of objectness. Specifically, we propose a novel frame- work to discover background regions in an image and train an object proposal network to not detect any objects in these regions. We formulate the background discovery task as that of identifying image regions that are not discriminative, i.e., those that are redundant and constitute low information content. We conduct experiments on standard benchmarks to showcase the effectiveness of our proposed approach and observe significant improvements over the previous state-of-the-art approaches for this task.

CVPR Workshop on Domain Generalization: Evolution, Breakthroughs and Future Horizon 2025

© 2025 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Mitsubishi Electric Research Laboratories, Inc. 201 Broadway, Cambridge, Massachusetts 02139

Improving Open-World Object Localization by Discovering Background

Ashish Singh* Univ. of Mass.-Amherst ashishsingh@cs.umass.edu Michael J. Jones Mitsubishi Electric Research Labs mjones@merl.com

Anoop Cherian Mitsubishi Electric Research Labs

cherian@merl.com

Moitreya Chatterjee Mitsubishi Electric Research Labs chatterjee@merl.com

Kuan-Chuan Peng Mitsubishi Electric Research Labs

kpeng@merl.com

Erik Learned-Miller Univ.of Mass.-Amherst elm@cs.umass.edu

Abstract

Our work addresses the problem of learning to localize objects in an open-world setting, i.e., given the bounding box information of a limited number of object classes during training, the goal is to localize all objects, belonging to both the training and unseen classes in an image, during inference. Towards this end, recent work in this area has focused on improving the characterization of objects either explicitly by proposing new objective functions (localization quality) or implicitly using object-centric auxiliary-information, such as depth information, pixel/region affinity map etc. In this work, we address this problem by incorporating background information to guide the learning of the notion of objectness. Specifically, we propose a novel framework to discover background regions in an image and train an object proposal network to not detect any objects in these regions. We formulate the background discovery task as that of identifying image regions that are not discriminative, i.e., those that are redundant and constitute low information content. We conduct experiments on standard benchmarks to showcase the effectiveness of our proposed approach and observe significant improvements over the previous state-of-the-art approaches for this task.

1. Introduction

Open-world object localization (OWOL) refers to the task of localizing all objects, including novel object instances in a given image by learning from a limited set of predefined object classes. Unlike the traditional closed-world object localization setting, where the task is to localize object classes only seen during training, OWOL is a more realistic and challenging problem that evaluates generalization performance of the learned model to novel object classes.

Learning to localize object instances, not seen during train-



Figure 1. An example of a training image showing a ground truth annotated object (green box), unlabeled objects that are not in the known training classes (purple boxes), and clear non-object/background regions (red boxes) that are automatically classified as non-object regions by BOWL (**ours**) and used in training.

ing time, is important for many downstream applications like autonomous agents, robotics, and self-driving cars. As a result, there has been a growing interest in developing methods for object detection/localization in an open-world setting over the last few years [14, 34]. Previous works in OWOL have primarily focused on improving objectness learning. Specifically, Kim et al. [14] proposed incorporating localization-quality as a loss function to learn a more generalizable object representation, whereas, Wang et al. [34] explicitly enriches the training data by synthesizing pseudo-annotations for unannotated regions. While the above approaches show improvement in localizing unseen object classes, their performance is still limited. Specifically, we observe that during inference, either: (1) novel objects with features like shape, color, textures, etc. not observed in known classes are not detected, or (2) non-object regions are erroneously detected as objects due to noisy pseudo-annotations during model training. These observations allude to the challenging issue of developing and learning a generalizable notion of objectness. Learning the objectness of a region is non-trivial and most assumptions developed to characterize objectness do not generalize to unseen

^{*}Ashish did most of this work while an intern at MERL.

environments. In this work, we tackle this challenge by exploring a complementary source of information, *i.e.*, non-objectness of a region. We show that identifying non-object regions and incorporating them during the learning process significantly improves the localization of unseen object classes and the generalization to novel domains on various challenging benchmarks.

In particular, we propose to improve open-world object localization by incorporating non-objectness information during the learning process as illustrated in Figure 1. We first model the non-objectness information by constructing a codebook comprised of embeddings of exemplar non-object patches, extracted from a pre-trained model. This codebook is then used to identify seed anchors representing non-object regions while training the object localization method. This provides additional self-supervision in training the localization method along with labeled information from known classes.

Our evaluations show that by incorporating our non-object model into the object localization pipeline, our proposed method, Background-aware Open-World Localizer (BOWL), significantly improves over previous methods on multiple benchmarks for OWOL. Specifically, we surpass the previous baselines by 4.0% on the AR@100 measure for detecting non-VOC classes when the training data consist of 20 PASCAL-VOC [10] classes. Furthermore, we show a domain generalization improvement of 2.0% on the AR@100 measure for open-world localization in the wild, when evaluating a model trained with the COCO dataset [17] on the ADE20K dataset [40]. Going forward, we sometimes use the term, "background", in an informal way to mean non-object regions.

Overall, in this work, we make the following contributions:

- 1. We propose BOWL, a novel, open-world object localization method based on background-awareness.
- 2. We formalize the notion of non-objectness of image regions and propose a novel, unsupervised method for discovering non-object regions with high precision.
- 3. Our extensive evaluation shows that incorporating nonobjectness information in BOWL leads to improved recall in detecting unseen objects, outperforming competing methods for this task.

2. Related Work

In this section, we discuss related prior works and delineate the differences from our proposed approach.

Open-World Object Localization addresses the problem of localizing objects in images under the open-world setting. Early methods in object localization [1, 3, 9, 18, 41] adopted a bottom-up approach to creating a relatively small set of candidate bounding boxes that covered the objects in an image. These methods utilized hand-crafted heuristics based on raw image information like color, edge, and pixel-contrast to characterize objectness [31, 41]. In contrast, learning based methods [15, 16, 19, 22, 33] adopt a top-down approach, where the model is trained end-to-end with objectness defined by a set of

training categories in the dataset. These learning-based methods show considerable improvement over heuristics-based approaches in terms of precision and recall when localizing object classes from the training set. However, they struggle to localize novel object categories. In particular, these methods incorrectly assume any unannotated region of a training image to be a non-object region and thus suppress detection of novel objects at inference time. To resolve this issue, Kim et al. [14] proposed to modify Faster-RCNN [22] by replacing classification heads with class-agnostic objectness heads which are trained using base-class annotations, thus avoiding any incorrect assumptions about background. Saito et al. [25] proposed to instead crop labeled object regions in training images and paste them into images not containing any unlabeled objects to avoid the faulty assumption that unlabeled regions of training images do not contain any objects. While the above methods focus on mitigating the background bias, another family of methods have focused on improving object localization recall by pseudo-annotating novel objects. In particular, Wang et al. [34] learns a model to predict pairwise affinity to segment an image. Prediction from this model is then used to generate new box annotations, which is then used to train a detection model along with the original ground truth base class annotations as supervision.

Our work proposes a complementary approach, where instead of completely suppressing background or pseudo-labeling novel objects, we identify regions in training images which we can categorize as non-objects with high precision. Our experiments show that using non-object supervision for training the detector is better than completely disregarding it and avoids errors introduced by noisy pseudo-annotations.

Unsupervised Background Estimation: Most works for estimating background in images fall under the paradigm of either: (i) saliency detection [5], or (ii) figure-ground separation [23, 37-39], where the primary objective is to detect the salient foreground regions. Early attempts to solve the foreground/background separation task focused on thresholding or binary clustering on the color, depth, edge or other hand-crafted features [13, 23, 36, 37]. Recent methods in unsupervised object localization, a new formulation similar to figure-ground segmentation [28], in particular, utilize self-supervised models [6, 20] to estimate salient objects. Methods like LOST [26] and TokenCut [35] used deep features from pre-trained vision transformers (ViTs) [6]. LOST defined background regions as the patches that were most correlated with the whole image, while TokenCut investigated the effectiveness of applying spectral clustering [32] to self-supervised ViT features. Both of the above method make strong assumptions in defining objectness. As a result, they don't generalize well to images from different domains. In comparison, FOUND [27] proposed discovering the background as the primary step, with the objective of discovering foreground objects. Specifically, they first identify seed patch token with least attention value. Then background is defined as all the patches with similar representation as the

seed patch. They further improve the initial coarse background mask by refining it using a lightweight segmentation head.

While our method shares commonality with the above methods with respect to the use of self-supervised models for extraction of patch features, we differ from the above approaches with respect to our primary objective. In particular, in these prior methods, the main goal is to discover objects in the foreground region, and background estimation is a by-product. Thus they focus more on improving the precision of object discovery. Our goal on the other hand, is to improve novel unseen object localization using background as a supervisory signal. Thus for our problem statement, we want to categorize an image region as background/non-object with high confidence. In this regard, the focus of our unsupervised background estimation method is to achieve high precision with moderate recall, which is different from the above approaches.

3. Method

Our goal is to improve OWOL by incorporating non-objectness cues. Specifically, our proposed method BOWL is comprised of two main steps: (i) **Modeling non-object information:** where we construct a codebook containing embeddings of exemplar non-object patches, extracted from a pre-trained model, and (ii) **Background-aware training:** Where we train an open-world localization model with additional self-supervision using non-object seed anchors discovered using our computed non-object codebook. Fig. 3 shows an overview of our method.

3.1. Motivation

For the object localization task, we are provided with a training dataset D consisting of bounding box annotations $\{b_i\}$ from a pre-defined list S of known base classes. The goal for the object localization task (also commonly known as object proposal generation) is to generate bounding boxes on unseen images, localizing objects from the known base classes. As such, the primary objective adopted to train these models can be divided into two parts: (i) Classification objective \mathcal{L}_{cls} : which trains the model to categorize a region as 'object' or 'background,' and (ii) Bounding box regression objective \mathcal{L}_{reg} : which helps the model learn to accurately estimate the coordinates of a bounding box. Formally, given a training image I, the training loss can be formulated as:

$$\mathcal{L}(I) = \frac{1}{|A|} \sum_{i \in A} \mathcal{L}_{cls}(c_i, c_i^*) + \frac{1}{|A_K|} \sum_{i \in A_K} \mathcal{L}_{reg}(b_i, b_i^*), \quad (1)$$

where $\{c_i, b_i\}$ are the predicted label and bounding box coordinates, $\{c_i^*, b_i^*\}$ are the corresponding ground truths, A is the set of all candidate anchor boxes, A_K is the set of candidate anchor boxes with ground truth label $c_i^* = 1$. Here, $c_i^* \in C = \{0,1\}$, where $c_i^* = 1$ only when the anchor *i* can be associated with an annotated object bounding box from the known classes S (based on pre-defined intersection-over-union thresholds) and $c_i^* = 0$ otherwise. Thus, every candidate anchor box A_i not associated with an annotated object bounding box is automatically labeled as a background/non-object region. This design choice made sense for traditional methods [22], where the model was trained under the closed-world assumption, *i.e.*, at the testing time, their generalization is evaluated by detecting the objects from the known base classes. However, under the open-world setting, where the task is to localize every object in the image, which can belong to an unknown novel class $u \in U$, the training objective Eq. 1 can lead to poor test time performance. Specifically, under the training loss in Eq. 1, an unannotated object will be classified as 'background,' resulting in incorrect categorization of unknown novel object regions as 'non-objects.'

To avoid the above issue of false negative detections, Kim *et al.* [14] proposed to replace \mathcal{L}_{cls} in Eq. 1 with the objectness score prediction loss \mathcal{L}_{obj} , yielding:

$$\mathcal{L}_{OLN}(I) = \frac{1}{|A_K|} \sum_{i \in A_K} \mathcal{L}_{reg}(b_i, b_i^*) + \frac{1}{|A_K|} \sum_{i \in A_K} \mathcal{L}_{obj}(o_i, o_i^*), \quad (2)$$

where o_i and o_i^* are the predicted and ground truth objectness score with respect to an anchor box A_i . Here objectness score can refer to any metric [24, 30] that can measure bounding box localization quality with respect to ground-truth bounding boxes. Thus, in their formulation, Kim *et al.* [14] removes any dependency of predicting 'background' regions. In other words, 'background' regions are simply ignored during training.

While the above loss modification does improve performance over traditional methods with respect to localizing unknown novel objects, it underutilizes the supervision from non-object candidate anchor boxes. Specifically, Kim *et al.* [14] selects the candidate anchor boxes set A_K , by randomly sampling $|A_K|$ anchors with Intersection over Union (IoU) larger than 0.3 with the matched ground-truth boxes. While this ensures that an anchor box localized at an unknown novel object region is not incorrectly labeled as non-object, it reduces the diversity of anchor boxes in the candidate set, resulting in a reduction of the range of loss values. This issue can be potentially resolved if we knew apriori which regions correspond to non-objects. To this end, we propose the following modification to Eq. 2:

$$\mathcal{L}_{BOWL}(I) = \frac{1}{|A_K|} \sum_{i \in A_K} \mathcal{L}_{reg}(b_i, b_i^*) + \frac{1}{|A_K| + |A_B|} \sum_{i \in A_K \cup A_B} \mathcal{L}_{obj}(o_i, o_i^*),$$
(3)

where A_B refers to the set of non-object anchor boxes. Concretely, we propose to improve the diversity of sampled anchor boxes for objectness score prediction loss by sampling anchor boxes corresponding to non-object regions. Unlike previous methods, we first identify highly likely non-object regions in a given image, to ensure that we do not incorrectly label unknown novel object anchors as non-objects.

3.2. Modeling non-object patches in images

Given an arbitrary training set of unlabeled images, our goal is to identify image regions that are very unlikely to contain objects. Our model of non-object regions will consist of a representative set (called an exemplar set) of embeddings of non-object patches that cover a large percentage of all non-object patches in the training images. Note that our non-object model does not need to include every possible non-object region. Our goal is to learn a representative set of regions that with high probability are all non-object regions. When we use our non-object model to train the object localization network as described in Section 3.1 there may still be regions which are ignored because they are neither high-probability non-object regions nor labeled object regions.

To find non-object regions from an unlabeled set of training images, we observe that patches of training images that occur frequently are almost always non-object regions. For example, patches of sky, grass, concrete, brick, and white walls are all examples of non-object regions that occur frequently in many different images. This idea is consistent with the idea introduced in Sing *et al.* [29] that object patches are *representative* (meaning they occur often enough in the visual world) and *discriminative* (meaning they are different enough from the rest of the visual world). Since we are trying to model non-object patches, we are looking for patches that are representative but not discriminative. By clustering all patches of an unlabeled training set of images and taking a representative patch from the largest clusters, we will have our exemplar set of non-object regions.

Given a set of unlabeled images $\mathcal{I} = \{I_1, ..., I_N\}$ where $I_i \in \mathcal{R}^{W \times H \times 3}$, we first compute patches by dividing each image into M square patches of size $S \times S$. Given the set of image patches $\mathcal{P} = \{P_1, ..., P_{N \times M}\}$ where $P_i \in \mathcal{R}^{S \times S \times 3}$, we compute features of each patch using a self-supervised pretrained neural network (we use DINO-ViT[6] in our experiments). Each patch P_i is represented by a d-dimensional feature embedding ϕ_{P_i} computed using the pretrained model f, i.e. $f(P_i) = \phi_{P_i}$. The choice of a self-supervised pretrained model ensures that the feature representations are not biased towards any specific object category, allowing for a more general representation of patches.

Given the set of patch features, we construct an exemplar set E of patch features using a greedy nearest neighbor method. In particular, for each patch $P_i \in \mathcal{P}$, we compute:

$$s_{\max}(P_i, E) = \max_{e \in E} \frac{\phi_{P_i} \cdot \phi_e}{\|\phi_{P_i}\|_2 \|\phi_e\|_2}$$
(1)

A patch P_i is added to E if:

$$s_{\max}(P_i, E) < \lambda, \tag{2}$$

where λ is a distance threshold controlling the diversity of the exemplar set. Figure 2 illustrates the exemplar selection algorithm. This approach avoids the computational overhead of traditional clustering methods while still ensuring that the exemplar set captures the diverse range of non-object regions present in the dataset. As described in Figure 2, a count of the number of times each exemplar was the most similar exemplar to a training patch embedding is kept. These counts give the size of each exemplar cluster. As explained in Section 3.2, the largest clusters correspond to non-object patches with high probability. Our method simply uses the N clusters with the largest counts as the non-object exemplars in our background model.

It is important to note that the non-object model is intended to be learned only once and to be general enough to then be used in conjunction with learning an open-world localization network on any object detection dataset.

3.3. Background-Aware Open-World Localization

We incorporate the constructed non-object exemplar set into the training framework of an OWOL model. This integration allows us to explicitly leverage non-object regions during training, leading to more robust object localization.

Network Architecture and Base Training: We adopt the object localization network architecture which consists of two primary prediction heads:

- A regression head that predicts distances from the feature location to four sides of the ground-truth box G = (l, r, t, b).
- An objectness head estimating localization quality $c \in [0,1]$ by measuring the alignment between predicted and ground-truth boxes.

In the conventional training setup, the objectness head is trained only using foreground proposals that satisfy:

$$IoU(A_i, B_{qt}) > 0.3, \tag{3}$$

where A_i is a candidate anchor box and B_{gt} is the set of ground-truth boxes, respectively. The objectness score is supervised using an L1 regression loss:

$$\mathcal{L}_{obj} = \|1 - \operatorname{Sigmoid}(c_{A_i})\|_1, \tag{4}$$

where c_{A_i} is the centerness score of A_i . The value $c_{A_i} = 0$ when A_i has no overlap with the ground-truth box and $c_{A_i} = 1$ when there is perfect overlap.

Training with Non-object Supervision: We observe that the conventional training approach underutilizes non-object regions for supervision. To address this limitation, we augment the training process as shown in Figure 3 by incorporating explicit supervision from our non-object exemplar set E. Given the set of exemplar embeddings representing non-object regions, we categorize a region in an image (defined by anchor boxes from the detector) as a non-object region if it is similar to a member of the exemplar set, *i.e.*, the maximum similarity of the image patch embedding to all exemplar patch embeddings is above a fixed threshold. Cosine similarity between patch embeddings is used to measure similarity. We compute the negative anchor boxes at different scales defined by the detector pipeline. Specifically, for each anchor box A_i , we:

1. Extract its feature embedding ϕ_{A_i} using the same pre-trained



Image Patches

Figure 2. Exemplar selection works by first splitting the $t+1^{st}$ image into fixed-size patches and computing a feature embedding (using DINO, represented by the function, f) for each patch (represented by the grid of light blue cubes). Each patch is compared to the existing exemplars from the first t images, E(t), using the similarity function S. E(0) is initialed to the empty set. Patch embeddings that have below threshold similarity to all exemplars in E(t) (represented by dark blue cubes) are then added to the exemplar set to yield E(t+1). This repeats until all training images are processed. In addition, each exemplar in E(t) maintains a count of the number of times it was the most similar exemplar to a training patch embedding to keep track of each exemplar's cluster size.



Figure 3. Overview of training process with BOWL. First the input image on the left is passed through the feature pyramid network (FPN) backbone. Next, anchor boxes are generated which cover the image with regions of different positions, aspect ratios and scales. The anchor boxes are compared with the ground truth known object boxes to generate a set of labeled object boxes (shown in green). With BOWL, a DINO embedding for each image patch within an anchor box is also compared to the non-object exemplar set using cosine similarity, S, and embeddings that have similarity above a certain threshold to an exemplar are used as non-object regions for training.

model used in constructing E.

2. Compute its maximum similarity to the exemplar set:

$$s_{\max}(A_i, E) = \max_{e \in E} \frac{\phi_{P_i} \cdot \phi_e}{\|\phi_{A_i}\|_2 \|\phi_e\|_2}.$$
 (5)

3. Label it as a non-object region if:

s

$$\max(A_i, E) > \gamma, \tag{6}$$

where γ is a threshold parameter.

For anchor boxes identified as non-object regions, we provide negative supervision to the objectness head by setting their target objectness score to zero. The hypothesis is that the objectness head learns to disregard these regions as the loss will be high relative to ground-truth bounding box. The objectness loss can thus now be written as:

$$\mathcal{L}_{obj}^{new} = \begin{cases} \mathcal{L}_{obj} & \text{, if } \text{IoU}(A_i, B_{gt}) > 0.3\\ 1 & \text{, if } s_{\min}(A_i, E) > \gamma \end{cases}$$
(7)

 \mathcal{L}_{obj}^{new} ensures that the model learns to distinguish between object and non-object regions more effectively, leading to improved localization performance in open-world settings.

4. Experiment

4.1. Experimental Details

Datasets. We evaluate BOWL on two tasks to measure open-set generalization performance: (1) Cross-category generalization: where for a fixed data distribution, the model is evaluated on its performance on object categories not seen during model training, and (2) Cross-dataset generalization: where the model is evaluated on a dataset different from the training dataset and contains unseen object categories. For crosscategory evaluation, we use different subsets of the MS-COCO [17] dataset for training and testing and consider two settings. In the first setting, we use the subset of MS-COCO containing the 20 object classes from the PASCAL-VOC [10] for training and utilize the remaining 60 classes of MS-COCO for evaluation. We refer to this setting as VOC to NON-VOC generalization task. For the second setting, we utilize LVIS [11] benchmark, which provides annotations for 1203 classes in a long-tail distribution. We use the original 80 MS-COCO classes for training and test on the remaining 1123 classes. We refer to this setting as LVIS COCO to LVIS NON COCO generalization task

For cross-dataset evaluation, we train the model on the MS-COCO [17] dataset containing a mix of object and scene-centric images and evaluate on the ADE20K [40] dataset (a dataset of scene-centric images). We evaluate our method in two settings: (i) training with COCO subset containing 20 VOC classes and (ii) training with full COCO training set. In both settings, we use ADE20K validation set consisting of instance-level annotations of 3148 object categories.

Evaluation Metrics. Following prior work [14, 34], we calculate Average Recall (AR@k) across a range of IoU thresholds, spanning 0.5 to 0.95, and set a default detection budget k = 100. AR values are presented as percentages. We denote AR_A as the AR score for all classes, covering both base and novel categories, while AR_N represents the AR score specifically for novel classes. When measuring AR_N, the detections linked to the base classes are omitted from the budget k. This approach is consistently applied when calculating per-class AR, as well as AR values for objects of different scales: small (AR^s), medium (AR^m), and large (AR^l).

Implementation Details. For extracting non-object exemplar patches, we use the pre-trained DINO [6] ViT to extract patch-wise features. Specifically, we extract the last layer (Layer 11) 'key' features [2] using the 16×16 patch version of ViT S/16 with stride=8. Thus each image is represented as a set of 16×16 overlapping patches, with an overlap of 8 pixels. We select $\lambda = 0.2$ for exemplar selection based on held out validation set. From the extracted set of exemplars, we choose N = 1000 for training BOWL. An ablation on the effect of

Table 1. Cross-category generalization evaluation (VOC \rightarrow COCO).

Method	AR_N	AR^s_N	AR_N^m	AR^l_N
FRCNN (oracle)	52.6	37.1	60.0	73.1
FRCNN (cls-agn)	27.3	10.8	30.2	55.8
OLN [14]	33.2	18.7	39.3	58.6
GGN [34]	31.5	11.8	37.4	63.8
BOWL (ours)	37.5	20.7	46.9	58.6

Table 2. Cross-category generalization evaluation (LVIS COCO \rightarrow LVIS NON-COCO).

Method	AR	$_{\rm N}$ AR s_N	AR_N^m	AR^l_N
FRCNN (cls-	-agn) 21.0) 14.9	32.7	36.2
OLN [14] 27.4	4 17.9	44.7	53.1
GGN [34] 22.5	5 15.7	35.5	38.4
BOWL (ou	rs) 28.	3 18.3	46.5	53.9

different values for N is included in the supplemental material.

For open-world localization, we follow the same architecture as OLN [14], which uses the Faster RCNN [22] architecture with the ResNet-50 backbone [12] pretrained on ImageNet [8]. We implement BOWL using the MMDetection framework [7] and use the SGD optimizer with an initial learning rate of 0.01. We train BOWL for 16 epochs. BOWL is trained on 8 NVIDIA 1080Ti GPUs with a batch size of 2 images per device. For non-object anchor box selection, we set γ adaptively using Otsu's method [21].

4.2. Cross-category generalization

We show the results for cross-category evaluation in Table 1 for VOC to COCO classes domain generalization and Table 2 for COCO to NON-COCO domain generalization using annotations from LVIS [11] benchmark. We report Average Recall score for novel classes (AR_N@100) as our primary objective for cross-category evaluation is to assess the model's ability to localize instances from unseen novel categories.

From both results, we see that BOWL outperforms other baselines. Specifically, we observe an increase of 4.3% AR_N@100 compared to OLN for VOC to COCO generalization task. Furthermore, BOWL also outperforms OLN on COCO to LVIS NON-COCO generalization task by 0.9% AR_N@100. Considering both results, we can conclude that non-object supervision does improve performance for open-world localization task. While background supervision is also used in training the class-agnostic Faster-RCNN baseline, it significantly under performs in comparison to all the other methods. Thus improved results shown by our method can be attributed to accurate selection of non-object regions during model training using our non-object exemplar set.



Figure 4. Qualitative results of (a) GGN [33], (b) OLN [14], and (c) BOWL on MS-COCO validation images. The shown results are all predictions with objectness score higher than 0.75 generated from models trained on VOC categories. From the results we can see that while both OLN and GGN are able to localize unseen novel objects, there are significant false positive and false negative predictions. Specifically, because of noisy pseudo-annotations, GGN incorrectly predicts non-object regions as objects with high objectness score (false predictions on floor and bed in first row while wall and in second row). OLN on the other hand is not able to predict objects with shapes and scales not present in training data due to only using object supervision (false negative prediction of bed in first row and laptop in second row). BOWL mitigates the above issues by utilizing non-object supervision leading to better localization of objects. We provide more results in supplementary.

Table 3. Cross-dataset generalization evaluation (VOC \rightarrow ADE20K).

Method	AR_A	AR^s_A	AR^m_A	AR_A^l
FRCNN (cls-agn)	22.6	15.5	23.7	26.5
OLN [14]	29.2	19.7	30.7	34.4
GGN [34]	27.0	16.9	27.5	33.6
BOWL (ours)	30.2	20.7	32.5	34.8

Table 4. Cross-dataset generalization evaluation (COCO \rightarrow ADE20K).

Method	AR_A	AR^s_A	AR^m_A	AR^l_A
FRCNN (cls-agn)	25.9	20.5	28.5	27.4
OLN [14]	32.9	25.1	35.9	35.6
GGN [34]	29.8	18.9	29.1	38.2
BOWL (ours)	34.1	25.3	37.5	37.0

4.3. Open-set localization in the wild

A key requirement of any good open-world model is its ability to generalize in out-of-distribution (OOD) settings. Specifically, for the case of open-world object localization, it is important to quantify OOD generalization performance to make sure that the model is not overfitting to a particular dataset. To this end, we evaluate our method for the cross-dataset generalization task, where we train our model on MS-COCO and test it on the ADE20K dataset. For evaluation, we report Average Recall score for all categories (AR_A@100). Specifically we report AR score for all categories, since unlike the cross-category setting, similar classes across two datasets can show different properties like change in prototypical shape, scale, size *etc*.

We show our results for VOC to ADE20K and COCO to ADE20K in Table 3 and 4, respectively. Both the results show a trend, similar to cross-category evaluation tasks, where BOWL outperforms other baseline methods. Our method improves over OLN by 1% AR_A@100 when the model is trained on VOC classes and 1.2% AR_N@100 when the full COCO dataset is used for training. Improvement in both the settings suggests that our non-object supervision not only improves open-world localization but is also generalizable across data distribution shift.

4.4. Ablation Study

Comparison of different approaches for non-object supervision. We evaluate our approach for the identification of non-object regions with other methods with respect to improving localization performance. Specifically, we consider

Table 5. Cross-category generalization evaluation (VOC \rightarrow NON-VOC) of different sources for non-object supervision.

Method	AR_N	AR^s_N	AR_N^m	AR_N^l
Gradient Energy	36.2	20.9	43.1	57.4
DINO	36.8	20.3	45.4	57.7
BOWL (ours)	37.5	20.7	46.9	58.6

gradient energy [4] and image self-correlation map using DINO [6] to identify non-object regions and use the same for training the localization model. Gradient energy has been previously used in figure-ground segmentation methods [23, 37–39], where background is defined as any region in an image with below threshold gradient energy. Similarly, recent works in unsupervised object localization [26, 28] utilize self-correlation maps generated using patch features from DINO to identify foreground regions, by considering patches with high correlation to majority patches in an image as background regions.

For both methods, we first generate a binary map specifying foreground and background regions, which is then used to select anchor boxes when training the model. We categorize an anchor box as non-object, if it has at least a 90% overlap with background regions.

We compare our approach with the above methods for the VOC to COCO cross-category task. As shown in Table 5, all three methods perform better than OLN. This further confirms our initial hypothesis that accurate non-object supervision should lead to improvement in open-world object localization. Amongst the three approaches, our method performs better than the other two methods. We surmise that this is because unlike the other two approaches, our method considers dataset-level statistics to model non-object information. This makes our method more precise in identifying non-object regions. Furthermore, this result also confirms that it is more important to have higher precision in identifying non-object regions for improving open-world object localization.

Influence of the number of base classes for training. One of the advantages of open-world models is that we don't need to annotate every single object instance as these models should be able to generalize to unseen novel objects. Thus it is important to consider the effect of number of base classes used for training the model. To this end, we split the COCO dataset based on the supercategories to create six different splits (Table 6) and then trained BOWL and OLN on these splits. For evaluation, we test both methods on ADE20K validation set. As shown in Figure 5, both BOWL and OLN achieve better $AR_A@100$ scores as the number base classes in training set increases. While for the extreme case of a single base class ('Person') both methods show similar results, we see more improvement in localization performance for BOWL when number of classes is on the lower side. Thus BOWL is more efficient when considering data annotation cost.



Table 6. Base class selection for evaluating the effect of the number of base classes.

Superclasses	#Training Classes	#Training Images
Person	1	64115
(+) Vehicle	9	74152
(+) Outdoor, Animal	24	92169
(+) Accessory, Sports	39	93939
(+) Kitchen, Food	56	107036
(+) Furniture, Electronic, Appliance, Indoor	80	117266

5. Discussion and Conclusion

In this work, we introduce the idea of learning a model of nonobject regions in a self-supervised way and use it to improve the accuracy of open-world object localization. We have shown that this idea improves results on standard datasets compared to the state-of-the-art. BOWL is the first method that uses a generic model of non-object regions to improve open-world object localization. Our ablation experiments confirm that our exemplar-based non-object model achieves better results than simpler background models, while also showing that our improvement over the OLN open-world detector is independent of the number of base classes in the training set. For future work, we think further gains can be realized by combining the idea of background modeling with the creation of pseudo-labels to increase the number of labeled object for training.

Limitations: While incorporating a background model into open-world localization training improves accuracy, it does require the extra computational expense of learning a background model, although this only needs to be done once. Also, BOWL is slightly worse, in terms of performance, on large objects than some competing methods. This is likely due to competing methods' use of pseudo annotations which add more large object training examples.

References

- Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. Measuring the objectness of image windows. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2189–2202, 2012. 2
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. 6
- [3] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 328–335, 2014. 2
- [4] Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 609–617. 2023. 8
- [5] Ali Borji, Ming-Ming Cheng, Qibin Hou, Huaizu Jiang, and Jia Li. Salient object detection: A survey. *Computational visual media*, 5:117–150, 2019. 2
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings* of the IEEE/CVF international conference on computer vision, pages 9650–9660, 2021. 2, 4, 6, 8
- [7] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. arXiv preprint arXiv:1906.07155, 2019. 6
- [8] 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
- [9] Ian Endres and Derek Hoiem. Category independent object proposals. In Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part V 11, pages 575–588. Springer, 2010. 2
- [10] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 2, 6
- [11] 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. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [13] Michael J Jones and James M Rehg. Statistical color models with application to skin detection. *International journal of computer* vision, 46:81–96, 2002. 2
- [14] 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* (*RA-L*), 2022. 1, 2, 3, 6, 7
- [15] Weicheng Kuo, Bharath Hariharan, and Jitendra Malik. Deepbox: Learning objectness with convolutional networks. In *Proceedings*

of the IEEE international conference on computer vision, pages 2479–2487, 2015. 2

- [16] Hongyang Li, Yu Liu, Wanli Ouyang, and Xiaogang Wang. Zoom out-and-in network with map attention decision for region proposal and object detection. *International Journal of Computer Vision*, 127:225–238, 2019. 2
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part* V 13, pages 740–755. Springer, 2014. 2, 6
- [18] Santiago Manen, Matthieu Guillaumin, and Luc Van Gool. Prime object proposals with randomized prim's algorithm. In Proceedings of the IEEE international conference on computer vision, pages 2536–2543, 2013. 2
- [19] Pedro O O Pinheiro, Ronan Collobert, and Piotr Dollár. Learning to segment object candidates. *Advances in neural information* processing systems, 28, 2015. 2
- [20] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193, 2023. 2
- [21] Nobuyuki Otsu et al. A threshold selection method from gray-level histograms. Automatica, 11(285-296):23–27, 1975. 6
- [22] Shaoqing Ren. Faster r-cnn: Towards real-time object detection with region proposal networks. arXiv preprint arXiv:1506.01497, 2015. 2, 3, 6
- [23] Xiaofeng Ren, Charless C Fowlkes, and Jitendra Malik. Figure/ground assignment in natural images. In Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006. Proceedings, Part II 9, pages 614–627. Springer, 2006. 2, 8
- [24] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 658–666, 2019. 3
- [25] Kuniaki Saito, Ping Hu, Trevor Darrell, and Kate Saenko. Learning to detect every thing in an open world. In *European Conference on Computer Vision*, pages 268–284. Springer, 2022. 2
- [26] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. In *BMVC 2021-32nd British Machine Vision Conference*, 2021. 2, 8
- [27] Oriane Siméoni, Chloé Sekkat, Gilles Puy, Antonín Vobecký, Éloi Zablocki, and Patrick Pérez. Unsupervised object localization: Observing the background to discover objects. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 3176–3186, 2023. 2
- [28] Oriane Siméoni, Éloi Zablocki, Spyros Gidaris, Gilles Puy, and Patrick Pérez. Unsupervised object localization in the era of self-supervised vits: A survey. *International Journal of Computer Vision*, pages 1–28, 2024. 2, 8
- [29] Saurabh Singh, Abhinav Gupta, and Alexei A Efros. Unsupervised discovery of mid-level discriminative patches. In *Computer*

Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part II 12, pages 73–86. Springer, 2012. 4

- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: A simple and strong anchor-free object detector. *IEEE transactions* on pattern analysis and machine intelligence, 44(4):1922–1933, 2020. 3
- [31] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104:154–171, 2013. 2
- [32] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17:395–416, 2007. 2
- [33] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 2965–2974, 2019. 2, 7
- [34] Weiyao Wang, Matt Feiszli, Heng Wang, Jitendra Malik, and Du Tran. Open-world instance segmentation: Exploiting pseudo ground truth from learned pairwise affinity. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 4422–4432, 2022. 1, 2, 6, 7
- [35] Yangtao Wang, Xi Shen, Yuan Yuan, Yuming Du, Maomao Li, Shell Xu Hu, James L Crowley, and Dominique Vaufreydaz. Tokencut: Segmenting objects in images and videos with selfsupervised transformer and normalized cut. *IEEE transactions* on pattern analysis and machine intelligence, 2023. 2
- [36] Swee-Seong Wong and Wee Kheng Leow. Color segmentation and figure-ground segregation of natural images. In *Proceedings* 2000 International Conference on Image Processing (Cat. No. 00CH37101), pages 120–123. IEEE, 2000. 2
- [37] Zhaozheng Yin and Robert T Collins. Shape constrained figure-ground segmentation and tracking. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 731–738. IEEE, 2009. 2, 8
- [38] SU Yu and Jianbo Shi. Object-specific figure-ground segregation. In 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings., pages II–39. IEEE, 2003.
- [39] Stella X Yu, Tai Sing Lee, and Takeo Kanade. A hierarchical markov random field model for figure-ground segregation. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 118–133. Springer, 2001. 2, 8
- [40] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321, 2019. 2, 6
- [41] C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pages 391–405. Springer, 2014. 2