



O1O: Grouping of Known Classes to Identify Unknown Objects as Odd-One-Out

Misra Yavuz  and Fatma Güney 

Department of Computer Engineering, Koç University
KUIS AI Center
{myavuz21, fgüney}@ku.edu.tr

Abstract. Object detection methods trained on a fixed set of known classes struggle to detect objects of unknown classes in the open-world setting. Current fixes involve adding approximate supervision with pseudo-labels corresponding to candidate locations of objects, typically obtained in a class-agnostic manner. While previous approaches mainly rely on the appearance of objects, we find that geometric cues improve unknown recall. Although additional supervision from pseudo-labels helps to detect unknown objects, it also introduces confusion for known classes. We observed a notable decline in the model’s performance for detecting known objects in the presence of noisy pseudo-labels. Drawing inspiration from studies on human cognition, we propose to group known classes into superclasses. By identifying similarities between classes within a superclass, we can identify unknown classes through an odd-one-out scoring mechanism. Our experiments on open-world detection benchmarks demonstrate significant improvements in unknown recall, consistently across all tasks. Crucially, we achieve this without compromising known performance, thanks to better partitioning of the feature space with superclasses. Project page: <https://kuis-ai.github.io/O1O>.

Keywords: Open-World Object Detection · Geometric Proposals · Grouping of Classes

1 Introduction

Object detection methods in open-world settings are required to recognize objects from both known and unknown classes [1,19]. This task is particularly challenging due to having only a static dataset with a limited set of classes to learn from. Without any effort to recognize unknown classes, the standard object detection approach partitions the representation space between the known classes. Since unseen classes are not represented in the learned space, this approach generalizes poorly to the open-world setting. To localize objects of unseen classes, open-world detectors typically resort to additional information.

In object detection, both the localization and classification are supervised using ground truth annotations for known classes. In the case of unknown classes, however, there is no ground truth data available for training, mimicking the real-world setting. The lack of data has led to the mining of unknown objects from

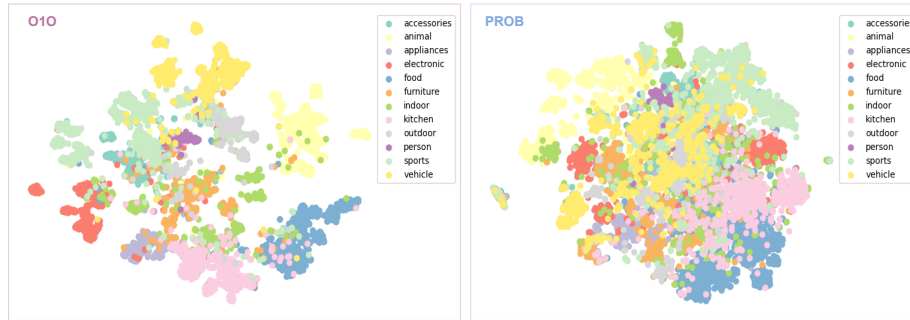


Fig. 1: **Superclass Prior vs. Single Objectness Distribution.** This figure compares the projections of object queries trained with superclass prior in our approach O1O (**left**) vs. fitting a single objectness distribution to matched queries as in PROB [44] (**right**) for S-OWOD Task 4. O1O shapes the representation space by encouraging queries of similar classes to group together, allowing us to identify odd-one-out queries as unknown objects.

highly activated areas in the background [11] or by using proposals of object localization networks [8]. This pseudo-labeling strategy generates approximate ground truth for unknown objects, enabling the model to search for objects beyond known classes. However, pseudo-labels may not necessarily correspond to object regions, and directly using noisy targets without any safeguarding degrade the performance of known classes as we show in our experiments.

A related line of work called class-agnostic object detection studies object localization without classification, i.e. separating foreground from background [27,40,20,34,28,18]. Recent work [18] improves recall by incorporating geometric cues such as predicted depth and surface normal maps into region proposal networks (RPN) [33]. However, the potential of geometric cues remains unexplored in open-world object detection (OWOD) where the goal is not only to improve unknown performance but also to maintain known performance while learning novel categories incrementally. As a first step, we trained a state-of-the-art OWOD model, PROB [44], with geometric pseudo-labels. PROB [44] learns a class-agnostic distribution to represent generic objectness. This essentially pushes all matched object queries towards an objectness center, forcing them to be similar to each other (Fig. 1). This approach degrades known performance due to noisy pseudo-labels that may not correspond to an object.

Learning a single distribution to represent any object is too simple to capture the diverse characteristics of object groups and is expected to perform poorly as the space of known concepts gets larger (Fig. 1). The other extreme is learning class prototypes by fitting a distribution to each known class [6], which is challenging because of the varying frequencies and specifics of each class, and creates more confusion for the model. To learn better objectness priors, we draw inspiration from human cognitive science. Studies in human perception suggest that humans use common attributes such as material, shape, color, and func-

tional properties to recognize and describe objects [16]. Inspired by this idea, we propose a middle ground between these two extremes: grouping similar object classes into superclasses. We learn a discriminative classifier for each known superclass, such as *person*, *vehicle*, or *furniture*. By partitioning the query embedding space into superclasses, we recalibrate the confidence scores for known classes with a conditional probability formulation. Our experiments demonstrate that recalibrated classification scores prevent degradation in the performance of known classes when using pseudo-labels for unknown classes, and achieve the best balance between known and unknown performance.

2 Related Work

2.1 Open-World Object Detection

The problem of detecting unknown objects has been tackled early on as an open-set recognition problem where the goal is to identify concepts that are semantically disjoint from the training set labels [31,12,30,5,32,13,41]. Simultaneously, methods pioneered by Joseph et al. [19] extended this formulation to open-world object detection (OWOD). In OWOD, the aim is to detect unknown objects and incrementally learn novel categories without compromising the performance of previously known classes. This setting aligns more closely with real-life requirements in ever-changing environments. Existing OWOD methods follow different approaches to separate objects from the background, whether known or unknown. Some methods extract pseudo-labels, i.e. candidate bounding boxes for unknown objects as auxiliary supervision [19,38], while others directly predict an objectness score [11,44,26]. Generally, OWOD methods can be categorized based on their pseudo-labeling approach and modeling of knowns.

RPN-based methods: Earlier approaches build on two-stage architectures in the R-CNN family [10,9,33]. Commonly, these methods [19,38] utilize the region proposal network (RPN) proposals to model unknowns. The classification head is extended to account for unknown objects. Building on [19], Yu et al. [39] learn class-specific prototypes. Wu et al. [37] attach a branch to learn objectness by predicting the IoU score of a proposal as an indicator of localization quality. In contrast to RPN proposals, unsupervised methods can be used to alleviate the bias of known classes on pseudo-labels. Wang et al. [36] first randomly generate bounding box proposals and then propose a matching score to select the ones that are most likely foreground unknown objects. In an autoencoding paradigm, Fang et al. [8] learn separate distributions for foreground and background to identify objects better, and use selective search [35] for pseudo-labels.

DETR-based methods: More recent work builds on the transformer-based architecture DETR [43]. Absent an RPN module, these methods explore various pseudo-label sources. Gupta et al. [11] aim to capture unknowns with a scoring function based on the attention activation of intermediate layers. They use two classification heads: one to separate knowns from unknowns and another to distinguish foreground from background. Ma et al. [26] combine selective search [35]

with attention activations [11] to generate pseudo-labels. They decouple localization and classification in a cascade fashion to minimize the known class bias.

Zohar et al. [44] take a different approach by omitting pseudo-labels and instead learning a distribution of objectness through joint optimization of parameters for a multivariate Gaussian on matched object queries from DETR. He et al. [15] uses a similar modeling of objectness but based on the early decoder layers to mitigate the influence of class-specific features. Additionally, they generate pseudo-labels using the Segment Anything Model (SAM) [21], which significantly contributes to their performance on unknown objects. While it is intriguing to utilize foundational models for OWOD, the definition of unknown becomes blurry for SAM, which has been trained on a large collection of internet images. Our method also builds on the variations of DETR. Differently, we explore the potential of geometric cues to detect unknowns.

2.2 Class-Agnostic Object Detection

In class-agnostic object detection, some methods utilize the rich semantics of natural language to ease localization [27,40]. Saito et al. [34] augment images by synthetically pasting foreground objects to alleviate the bias from unlabelled foreground objects in existing datasets. OLN [20] estimates the objectness of proposals based on localization quality supervised by ground truth. GOOD [18] argues that geometric cues carry more relevant information about objectness that generalizes across categories. We explore the potential of GOOD [18] as a source of structured objectness priors for open-world object detection.

2.3 Superclass Supervision

The idea of a semantic grouping of classes has previously been used in the standard object detection task to improve closed-set performance. Specifically, SGNNet [23] trains a superclass classification branch in an end-to-end fashion with a convolutional detection model. Additionally, long-tail classification methods [42,7] explored the superclass grouping to enhance recognition of tail classes by benefiting from less imbalance in superclass distributions. Differently, we encourage objects sharing similar properties to cluster within a superclass by classifying queries into superclasses. Through re-calibration of confidence scores and a novel unknown scoring function, we identify unknown objects that do not share these properties as odd-one-out (OIO). We are the first to explore the potential of superclasses for unknown object detection.

3 Methodology

We introduce Odd-One-Out (OIO), an open-world object detection method to locate and classify objects by identifying unknowns, as illustrated in Fig. 2. We first provide background on the transformer-based architectures for object detection including the extensions that we adopt in Section 3.1. Next, we introduce

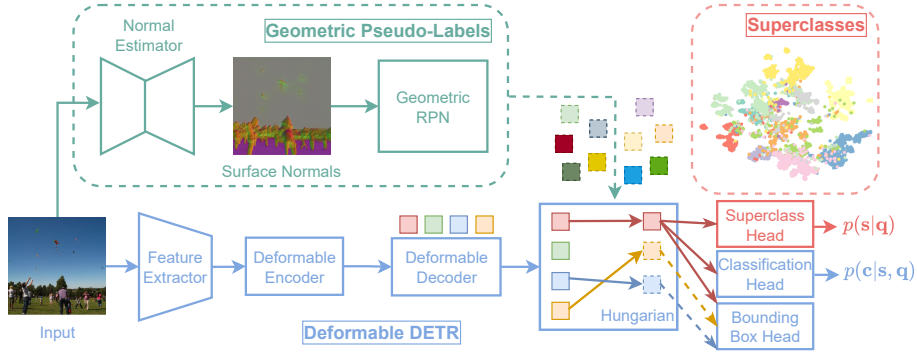


Fig. 2: **Overview of Proposed O1O.** Building on Deformable DETR [43] (blue), we first add supervision for unknowns with geometric pseudo-labels. Following GOOD [18], we extract pseudo-labels (dashed) from a Region Proposal Network (RPN) trained on surface normal maps (green). This allows us to localize unknown objects based on geometric cues in a class-agnostic manner. Noisy pseudo-labels tend to hurt the known performance of a model. To mitigate this issue, we propose to group queries into superclasses with a superclass head (red). By incorporating the learned prior from superclasses into the scoring function, we achieve the best balance between known and unknown performance.

geometric cues to obtain pseudo-labels for unknowns in Section 3.2. For known performance, we propose to group known classes into superclasses and formulate classification as a conditional probability distribution in Section 3.3. Lastly, in Section 3.4, we explain our incremental learning strategy for open-world tasks.

3.1 Background on DETR

DEtection TRansformer (DETR) [2] operates by reasoning about the relationships between objects and their interaction with image features using a transformer as shown in blue in Fig. 2. Formally, it initializes N_{query} learned positional encoding vectors. In the transformer encoder, image features interact with each other and produce weighted attention maps. In the transformer decoder, queries first interact with each other and then with the feature maps from the encoder. This results in an embedding $\mathbf{q}_i \in \mathbb{R}^d$ for each object query $i \in N_{query}$, incorporating information about other queries and the image context. Each object query is then independently decoded into box coordinates and class labels using a feedforward network (FFN) followed by a linear layer for output generation. During training, ground truth objects are matched to object queries via bipartite matching before computing the classification and bounding box losses. Queries unmatched with ground truth objects are responsible for predicting a special class label, typically denoted as $K+1$, which in closed-set methods represents the ‘no-object’ case, while in open-set methods, it is used for the unknown class. Note that unmatched queries do not receive any supervision for localization.

Extensions: Deformable DETR [43] extends DETR by incorporating a more efficient attention module known as deformable attention. This attention mechanism focuses on a small set of key sampling points around a 2D reference point, enhancing the model’s ability to attend to relevant image features. Leveraging this efficiency, Deformable DETR integrates multi-scale image features, particularly beneficial for detecting small objects. Deformable DETR significantly improves the convergence capabilities of DETR. Another variant DAB-DETR [25] improves the cross-attention operation by employing dynamically updated 4D anchor boxes consisting of a reference point (x, y) and a reference anchor size (w, h) . The more recent DN-DETR [22] further improves the convergence with denoising techniques during training. DN-DETR identifies bipartite matching as the bottleneck for slow convergence and mitigates it by injecting noisy ground truth boxes into the transformer decoder and then training the model to reconstruct the original boxes. This approach effectively halves the number of training epochs required to match the performance of Deformable DETR.

While most recent OWOD methods [11,44] build on Deformable DETR [43], we incorporate improvements from both DN-DETR [22] and DAB-DETR [25]. Our primary objective is to leverage enhanced convergence properties for rapid experimentation while maintaining the overall performance. Note that these architectural modifications improve efficiency but do not contribute to unknown performance. For fairness, we report the performance of the SOTA unknown object detection method, PROB [44] using the same architecture (See Section 4.5.)

3.2 Pseudo-Labels for Unknown Objects

In DETR-based methods for unknown object detection, the class index $K + 1$ represents the unknown class in addition to the K known classes. Since unmatched queries only receive semantic supervision, the model’s ability to locate objects beyond the known set is limited. The main challenge is obtaining labels for localizing unknown objects. Previous methods employ pseudo-labeling strategies to obtain an auxiliary signal for training to localize all objects.

Geometric Cues: Recent work in class-agnostic object localization, proposes to use geometric cues for open-world object detection [18]. While RGB-based methods suffer from over-fitting to semantic features of known classes, geometric cues can be used to mine novel object classes as pseudo-labels. Specifically, GOOD [18] trains a region proposal network (RPN) using geometric cues such as depth and surface normal maps as input. By focusing on object shapes, relative spatial locations (depth), and directional changes (normal vectors), GOOD can retrieve more accurate pseudo-labels compared to other class-agnostic detectors.

In this work, we use pseudo-labels obtained from GOOD’s RPN module to train for unknown object detection. For fairness, we adhere closely to OWOD benchmark settings regarding data partitioning into known and unknown classes. For training the RPN, we ensure that there are no objects of unknown classes according to the benchmark setting considered. This cannot be ensured while using sources like SAM [21] that is trained on a large collection of internet images.

3.3 Shaping Queries Through Superclass Supervision

Following the standard DETR-based approach, our model takes an image as input and outputs bounding box locations and class scores corresponding to objects on the image. Let $\mathbf{c} \in \{0, 1\}^{K+1}$ be a random variable representing the one-hot encoded class of a bounding box over a predefined set of K known classes and the unknown label $K + 1$. Based on the embedding of query vectors $\{\mathbf{q}_i\}_{i=1}^{N_{query}}$, we re-design the classification head.

We classify regions by further decomposing the probability of a class with superclasses. A superclass is a group of classes that share some semantic properties. For example; cats, dogs, and birds as *animals*, and cars, buses, and motorcycles as *vehicles*. In our approach, we propose to learn common properties between known classes by grouping them into superclasses. That way, we can better judge similarities between known classes, i.e. recalibrating the confidence of what is known, and then use it to accurately identify what is unknown.

Superclass Prior: In DETR, the per-class classification head outputs the probability of each query \mathbf{q} belonging to a class \mathbf{c} as $p(\mathbf{c}|\mathbf{q})$. In addition to that, we train a superclass head to group known classes into superclasses. Formally, the superclass head outputs a probability of belonging to a superclass \mathbf{s} as $p(\mathbf{s}|\mathbf{q})$ over the set of superclasses S . For each known class, we recalibrate its confidence score by multiplying it with the probability of the superclass that it belongs to:

$$p'(\mathbf{c}|\mathbf{q}) = \sum_{\mathbf{s} \in S} p(\mathbf{c} | [\mathbf{c} \in \mathbf{s}], \mathbf{q}) \cdot p(\mathbf{s}|\mathbf{q}) \quad (1)$$

where $[\cdot]$ denotes Iverson bracket. Since known classes are fixed, we can pre-define the set of superclasses based on the known classes (see Supplementary). In practice, this reduces conditioning on the superclass \mathbf{s} in the first term to a lookup operation to find the corresponding superclass for a class. Intuitively, we shape the representation space of known classes by encouraging similar properties for queries matched to classes within a superclass.

Unknown Probability: For unknown probability $p(u|\mathbf{q})$, we apply the odd-one-out approach and define its probability as the complement event of being known. Considering each class probability to be independent, we sum the recalibrated probabilities of belonging to known classes, and subtract it from 1:

$$p(u|\mathbf{q}) = 1 - \sum_{\mathbf{c} \in K} p'(\mathbf{c}|\mathbf{q}) \quad (2)$$

Essentially, the decoupled conditional probability helps the model to update its belief for a query about belonging to a class by considering the similarities with its superclass. If the model is not confident enough to label a query as known, then (2) becomes higher, and the query is predicted as unknown. In our experiments, we show that learning a structured prior with superclasses recovers the known performance by reducing the confusion caused by noisy supervision from pseudo-labels for unknown classes.

Training: We first train the RPN on only the known set of the current task and store the generated proposals. Then, we train our model on a weighted combination of localization, per-class classification, and superclass losses:

$$\mathcal{L} = \underbrace{\mathcal{L}_{\text{bbox}} + \mathcal{L}_{\text{giou}}}_{\text{localization}} + \underbrace{\mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{sup}}}_{\text{classification}} \quad (3)$$

We use the same loss function used in DETR-variants and only add the classification loss \mathcal{L}_{sup} (shown in blue) to train the superclass head. We use focal loss [24] for superclass classification to address the imbalance issue in the distribution of superclasses. We apply the classification losses only on queries matched to known targets. In other words, we only apply the localization losses for the queries matched to pseudo-labels. For the classification of unknowns, we use the knowledge of the model on known classes to reason about what is unknown.

Building on extended Deformable DETR with denoising [22], we apply a denoising loss for each component, including the superclass loss. For noisy bounding box queries, the denoising bounding box loss reconstructs original coordinates, while denoising classification losses predict the original target labels.

Inference: In inference, we use the decomposition in (1) to obtain classification confidence scores of known classes. Since one class only belongs to a specific superclass, and this matching is given in the label set, it serves as a prior rather than privileged information. We use only the current known class label set to access superclass information, with unseen objects considered part of an additional unknown superclass. For the unknown score, we use (2) and set a threshold on a training subset to ensure 95% of known instances are classified correctly.

3.4 Incremental Learning

Apart from unknown detection, OWOD formulation requires incrementally learning novel classes while still maintaining the performance on previously learned concepts. However, training a model continually on multiple tasks typically results in catastrophic forgetting [29]. To prevent such a phenomenon, the common practice is to store a small set of exemplars for fine-tuning the model on previous tasks. Most OWOD approaches select the exemplar set randomly. On the other hand, an active selection strategy can be applied for a more effective replay. For example, PROB [44] uses their objectness approximation to actively select images with high and low-scoring instances, representing easy and hard cases.

We adopt a similar active exemplar selection strategy using our class conditional probability (1). Specifically, we calculate the class-wise conditional probabilities for all queries on the training set of the previous task. Then, for each class, we select the images containing the highest 25 and the lowest 25 scoring instances and save them as exemplars to be used in the next tasks. If the number of selected images turns out to be higher than in previous work, we randomly drop out the extra number of images to be comparable.

4 Experiments

4.1 Benchmarks

We report results on two OWOD benchmarks. M-OWOD, proposed by [19], is composed of both MS-COCO and PASCAL images. 80 categories of the MS-COCO dataset are split into four non-overlapping tasks $\{T_1, \dots, T_4\}$ with 20 classes being introduced at each. In each task T_t , only the annotations of classes that have been introduced until T_t are used for training, and the rest are removed. The test set contains instances from all 80 classes, and the known vs. unknown performance is calculated according to the known-unknown split. S-OWOD, proposed by [11], only consists of MS-COCO images. Classes are split according to the supercategories they belong to. For example, all animal classes are included in Task 1, while all sports classes are introduced in Task 3. Denoting the differences, these benchmarks are also known as Mixed-superclass and Separated-superclass, respectively. For more details, please refer to [19] and [11].

4.2 Evaluation Metrics

Following the common practice in the OWOD literature, we report the mean average precision (mAP) for known performance. Since not all objects are annotated in any given dataset, unknown precision is ill-defined. A model might detect an object in a location where a real object exists, but if it does not belong to the predefined categories in the dataset, it might have been excluded from the ground truth annotations. Therefore, recall (U-Recall) is used as the main metric to evaluate the ability to retrieve unknown objects.

4.3 Implementation Details

We build our method on the extended version of Deformable DETR [43] with denoising and 4D anchor boxes [25,22]. We use the same backbone as the prior work, i.e. ResNet-50 FPN [14] pre-trained on ImageNet [4] in a self-supervised manner [3]. We set the number of queries to 100 and the query embedding dimension to 256. We use predicted surface normal maps to train RPN after comparing different sources, including RGB images and predicted depth maps (see Supplementary). To select pseudo-boxes, we threshold RPN’s confidence scores at 0.5 and apply batched NMS on ground truth and pseudo annotations with an IoU threshold of 0.5 to remove duplicates.

We train our model on 4 GPUs, with a batch size of 6, and $2e-4$ as the starting learning rate. Since Deformable DETR with the extensions converges fast, we drop the learning rate by 10 at epoch 18. We train for 30 epochs on M-OWOD Task 1 and 20 epochs on S-OWOD Task 1. Since the dataset of S-OWOD Task 1 is five times larger, our model converges earlier. For incremental tasks, we train for 10 epochs with a learning rate of $2e-5$. We perform fine-tuning from exemplar replay for 20 epochs, with a learning rate of $2e-4$ for the first half and $2e-5$ for the second half. Overall, our architectural improvements cut down training time approximately by half compared to previous DETR-based approaches.

4.4 Open-World Object Detection Performance

In Table 1, we present O1O’s performance in OWOD, comparing it to previous approaches across different tasks on two benchmarks. O1O consistently outperforms all methods across all tasks on both benchmarks except for USD [15] on S-OWOD in terms of unknown performance (U-Recall), while maintaining a competitive or the highest Known AP. Note that there’s a trade-off between known and unknown performance, as each method can make a fixed number of predictions in total. While increasing unknown recall by predicting more unknowns is possible, it would inevitably degrade known performance. Therefore, striking a balance between these two metrics is crucial.

S-OWOD vs. M-OWOD: While the two benchmarks differ in their partitioning of superclasses, our method demonstrates the ability to generalize to both, as indicated by the results. Due to the superclass separation between tasks in S-OWOD, learning becomes easier, resulting in better-known performance for all methods compared to M-OWOD. The larger dataset size of Task 1 on S-OWOD also contributes to higher known detection performances on this benchmark.

Comparison: The closest competitors in Table 1 are from DETR-based approaches, including OW-DETR [11], PROB [44], CAT [26], and USD [15]. We encountered a bug in the evaluation of M-OWOD due to duplicated filenames in the test set. After resolving this issue¹, we report the improved results of PROB as PROB[†]. PROB does not rely on any pseudo-labels but instead learns an objectness distribution. This allows them to achieve the best performance in terms of known AP on several tasks. However, in terms of unknown recall, PROB falls behind other approaches that use pseudo-labels such as MEPU [8] and USD [15]. Our method also uses pseudo-labels and outperforms all these methods except for USD [15] on S-OWOD in terms of unknown performance while being among the top three in terms of known performance.

On the Source of Pseudo-Labels: To assess the effectiveness of geometric pseudo-labels, we compare our method to previous approaches that also utilize pseudo-labels from alternative sources. Specifically, we compare against MEPU [8] which uses pseudo-labels from Object Localization Network (OLN [20]). OLN, like GOOD [18], trains the Region Proposal Network in a class-agnostic manner but relies solely on appearance cues. By leveraging geometric cues as a better source for pseudo-labels, as proposed in GOOD [18], our method outperforms MEPU across all benchmarks. Additionally, we compare our method to USD [15], which utilizes pseudo-labels from SAM [21] and achieves an impressive unknown recall. Despite the advantage of USD in terms of pseudo-label source due to the internet-scale training data of SAM [21], our method, leveraging geometric pseudo-labels, outperforms USD on M-OWOD and approaches its performance on S-OWOD. For additional comparison of pseudo-label sources, we

¹ This issue might apply to other DETR-based methods. However, since the code or checkpoints are not available for them, except for OW-DETR, which is outperformed by PROB, we could not check or re-evaluate them.

Task IDs	Task 1		Task 2				Task 3				Task 4		
	U-Recall	mAP (†)	U-Recall	mAP (†)			U-Recall	mAP (†)			mAP (†)		
Methods	(†)	Current known	(†)	Previously known	Current known	Both	(†)	Previously known	Current known	Both	Previously known	Current known	Both
ORE - ebui [19]	4.9	56.0	2.9	52.7	26.0	39.4	3.9	38.2	12.7	29.7	29.6	12.4	25.3
UC-OWOD [38]	2.4	50.7	3.4	33.1	30.5	31.8	8.7	28.8	16.3	24.6	25.6	15.9	23.2
OCPL [39]	8.26	56.6	7.65	50.6	27.5	39.1	11.9	38.7	14.7	30.7	30.7	14.4	26.7
2B- OCD [37]	12.1	56.4	9.4	51.6	25.3	38.5	11.6	37.2	13.2	29.2	30.0	13.3	25.8
OW-DETR [11]	7.5	59.2	6.2	53.6	33.5	42.9	5.7	38.3	15.8	30.8	31.4	17.1	27.8
RandBox [36]	10.6	61.8	6.3	-	-	45.3	7.8	-	-	39.4	-	-	35.4
PROB [44]	19.4	59.5	17.4	55.7	32.2	44.0	19.6	43.0	22.2	36.0	35.7	18.9	31.5
CAT [26]	23.7	60.0	19.1	55.5	32.7	44.1	24.4	42.8	18.7	34.8	34.4	16.6	29.9
MEPU-FS [8]	31.6	60.2	30.9	57.3	33.3	44.8	30.1	42.6	21.0	35.4	34.8	19.1	30.9
MEPU-SS [8]	30.3	60.0	30.6	57.0	33.1	44.5	30.0	42.2	20.5	35.0	34.3	18.9	30.4
USD - asf [15]	21.6	59.9	19.7	56.6	32.5	44.6	23.5	43.5	21.9	36.3	35.4	18.9	31.3
USD [15]	36.1	58.4	34.7	54.3	31.4	42.7	33.3	41.5	20.5	34.5	33.4	16.6	29.2
PROB [†] [44]	28.3	66.4	26.4	62.6	39.2	50.9	29.3	49.6	33.5	44.2	44.0	26.5	39.7
O1O [†] (Ours)	49.3	65.1	50.3	61.0	45.0	53.0	49.5	50.0	36.5	45.5	46.2	31.0	42.4
ORE - ebui [19]	1.5	61.4	3.9	56.5	26.1	40.6	3.6	38.7	23.7	33.7	33.6	26.3	31.8
OW-DETR [11]	5.7	71.5	6.2	62.8	27.5	43.8	6.9	45.2	24.9	38.5	38.2	28.1	33.1
PROB [44]	17.6	73.4	22.3	66.3	36.0	50.4	24.8	47.8	30.4	42.0	42.6	31.7	39.9
CAT [26]	24.0	74.2	23.0	67.6	35.5	50.7	24.6	51.2	32.6	45.0	45.4	35.1	42.8
MEPU-FS [8]	37.9	74.3	35.8	68.0	41.9	54.3	35.7	50.2	38.3	46.2	43.7	33.7	41.2
MEPU-SS [8]	33.3	74.2	34.2	67.5	41.0	53.6	33.6	50.0	37.5	45.8	43.2	33.5	40.8
USD - asf [15]	19.2	72.9	22.4	64.9	38.9	51.2	25.4	50.1	34.7	45.0	43.4	33.6	40.9
USD [15]	51.1	69.8	52.1	60.9	33.8	46.7	49.9	45.8	30.2	40.6	39.0	28.7	36.4
O1O (Ours)	49.8	72.6	51.1	65.3	44.9	54.6	48.1	49.5	41.5	46.8	47.3	42.0	45.9

Table 1: **Comparison on the OWOD Benchmarks.** We compare our method O1O to the state-of-the-art on M-OWOD (**top**) and S-OWOD (**bottom**) across 4 different tasks on each. We report recall for unknown classes (U-Recall) and mAP@0.5 for known classes, separately for previously, currently, and all known objects. In each column, we highlight the **first**, **second** and **third** best results. O1O outperforms all existing OWOD methods across all tasks in terms of U-Recall except for USD [15] on S-OWOD. Note that USD relies on SAM proposals which significantly contributes to its performance as can be seen from inferior results without SAM (USD - asf). O1O is also consistently among the top-performing methods in terms of known mAP. Since all 80 classes are known in Task 4, U-Recall is not computed. PROB[†] indicate our evaluation after fixing the duplicate issue on the test set of M-OWOD.

trained our model using RPN proposals extracted from RGB images, predicted depth or surface normal maps, in addition to high activation areas and SAM proposals. Please see Supplementary for the details of this ablation.

Incremental Learning: The active selection strategy, proposed by PROB [44] results in the top performance for both previously and the current known in Task 1 on the M-OWOD benchmark. We adopt a similar strategy using our proposed class conditional probability (1) to select the exemplars. As more classes become known in Task 2, our method starts to catch up, eventually outperforming PROB in Task 3 and 4 when most/all classes become known. This is because our superclass representation benefits more from the variety that comes with newly introduced classes in each task by improving its class conditional probabilities, and consequently enhancing our selection of exemplars for the next tasks. On

S-OWOD, other methods such as MEPU [8] and CAT [26] also perform competitively, benefiting from the availability of more data to learn from in Task 1. We observe a similar trend in our model’s performance in M-OWOD. As more classes become known through each task, our model improves its representation of superclasses, leading to increases in its ranking in terms of known performance.

4.5 Ablation Study

In Table 3, we ablate each component incrementally to show the effect of each on known and unknown performance. We use Task 1 of the M-OWOD benchmark for the component ablation study.

Architecture: Our baseline is Deformable DETR. We first augment the model with a $K+1$ unknown index and supervise each unmatched query to predict this label. This greedy strategy of predicting everything else as unknown achieves good U-Recall but leads to some loss in known AP. Next, we evaluate the impact of architectural improvements to Deformable DETR. While these changes significantly enhance convergence and slightly improve known AP performance, they do not yield substantial gains in unknown performance.

Geometric Pseudo-Labels: Unknown performance improves significantly when we add geometric pseudo-labels. For context, the previous state-of-the-art method USD reports 36.1 for unknown recall [15]. This improvement largely stems from SAM proposals, as evidenced by the lower recall without SAM (USD vs. USD - asf). Geometric pseudo-labels enable a notable performance jump to 50.7, despite using the same limited training set as the detector, unlike SAM.

Superclasses: The impressive improvement in unknown performance comes at the expense of notable degradation in known AP ($61.4 \rightarrow 55.0$), contradicting the intended goal of OWOD to preserve known performance. Our superclass approach addresses this by recalibrating the classification scores, resulting in a good balance between known and unknown performance. The visualization of the top-10 predictions by our model vs. PROB in Fig. 3 clearly shows the benefit of this recalibration. We also test our method by varying superclass groups, demonstrating its generalization to different groupings (see Supplementary).

PROB with Benefits: To further clarify our contribution with superclasses, we apply the same changes to PROB at the bottom part of Table 3. Compared to the greedy Deformable DETR baseline, PROB’s objectness probability helps maintain a high known performance, by ranking background regions lower in the confidence scores. Applying the same architectural improvements to PROB does not affect its performance. Incorporating geometric pseudo-labels improves the unknown performance of PROB by almost doubling the unknown recall ($27.8 \rightarrow 50.1$) but also negatively impacts the known performance (-7.1 in mAP), similar to what we observed with our method. PROB’s approach, which involves learning a single distribution to represent objectness, fails to mitigate the degradation in known performance. This underscores the importance of shaping the representation space with superclasses (Fig. 1).

Metrics	U-Recall (\uparrow)	mAP (\uparrow)
$1 - \max(p(\mathbf{s} \mathbf{q}))$	50.7	72.2
$1 - \max(p'(\mathbf{c} \mathbf{q}))$	52.8	68.2
$1 - \text{sum}(p'(\mathbf{c} \mathbf{q}))$	52.5	71.5
$1 - \text{sum}(p'(\mathbf{c} \mathbf{q})) + \text{thr.}$	49.8	72.6

Table 2: **Inference Score Ablation.**

We perform an ablation study on the inference score used for unknowns on S-OWOD Task 1. The first row corresponds to MSP [17] applied to superclass probabilities, and the second row is MSP applied to recalibrated class probabilities p' (1). The third row is an extended version of MSP, which uses the union of recalibrated class probabilities p' . As shown in the last row, we use the third with a threshold that is experimentally set on the training set.

Components	U-Recall (\uparrow)	mAP (\uparrow)
Deformable DETR	32.6	59.2
+ Arch. Improvements	31.6	61.4
+ Geo. Pseudo-Labels	50.7	55.0
+ Superclasses	49.3	65.1
PROB	28.3	66.4
+ Arch. Improvements	27.8	66.1
+ Geo. Pseudo-Labels	50.1	59.0

Table 3: **Component Analysis.** We perform an ablation study on Task 1 of M-OWOD by integrating our contributions one by one to the baseline, Deformable DETR, starting with architectural improvements, then geometric pseudo-labels, and finally superclasses. We also apply the same changes to PROB at the bottom part by using their objectness head instead of superclasses.

Unknown Scoring Function: We define our unknown probability based on the known class confidence scores (2). In Table 2, we report the performance of alternative scoring functions on S-OWOD, Task 1. For a confident known prediction, the maximum softmax probability (MSP) [17] becomes high, and the inverse becomes low, so we can obtain the unknown probability by simply subtracting it from 1. We apply MSP on superclass probabilities alone in the first row, and on recalibrated class probabilities (1) in the second row. In the third row, we extend the max in MSP to the union of recalibrated class probabilities. This indicates the combined belief of the model about a prediction to belong to a known class. Depending on the type of scoring function, there is a trade-off between known and unknown performances. We found the third approach to have the best balance between U-Recall and known AP. Furthermore, we calculate a threshold on a subset of the training set such that 95% of the known instances would be classified correctly. This thresholding helps us control the trade-off between known and unknown slightly better, without violating any principles.

5 Conclusion and Future Work

In this work, we build on DETR-based object detection models by better utilizing unmatched queries for unknown object detection. We first integrate pseudo-labels from a class-agnostic region proposal network based on geometrical cues, i.e. surface normal maps. Integrating pseudo-labels increases unknown recall significantly. However, it requires better modeling of known classes to handle the increased confusion caused by noisy pseudo-labels. To address this, we introduce




Fig. 3: **Qualitative Comparison of Top-10 Predictions.** We visualize top-10 predictions for PROB [44] (**left**) and our model O1O (**right**). We color known predictions in green and unknown in blue. Top-10 predictions of PROB are almost all unknown and duplicated over large regions, whereas our model can predict known classes like person, bird, and motorbike with high confidence and then focuses on unknown with lower confidence scores, indicating a better calibration of class confidence scores.

a prior on the queries using superclass groupings in addition to classification scores. Superclass supervision helps recalibrate the scores of known classes, effectively reducing confusion due to pseudo-labels. Our approach results in improved known performance with the help of superclasses while being among the top-performing methods in terms of unknown recall. We use a one-level hierarchy of superclasses, but also show that O1O generalizes across different superclass groupings. Our method relies on geometric cues for localizing unknowns which poses a challenge when these cues cannot be reliably predicted. Additionally, balancing noise reduction with accurate localization remains a critical challenge.

Acknowledgments. We thank João F. Henriques and Nazir Nayal for their valuable feedback. This project is co-funded by KUIS AI, Royal Society Newton Fund Advanced Fellowship (NAF\R\2202237), and the European Union (ERC, ENSURE, 101116486). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

1. Bendale, A., Boulton, T.E.: Towards open world recognition. In: CVPR (2015)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: CVPR (2009)
5. Dhamija, A.R., Günther, M., Ventura, J., Boulton, T.E.: The overlooked elephant of object detection: Open set (2020)
6. Du, X., Gozum, G., Ming, Y., Li, Y.: SIREN: Shaping representations for detecting out-of-distribution objects. In: NeurIPS (2022)
7. Du, Y., Shen, J., Zhen, X., Snoek, C.G.: Superdisco: Super-class discovery improves visual recognition for the long-tail. In: CVPR (2023)
8. Fang, R., Pang, G., Zhou, L., Bai, X., Zheng, J.: Unsupervised recognition of unknown objects for open-world object detection (2023)
9. Girshick, R.: Fast r-cnn. In: ICCV (2015)
10. Girshick, R.B., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2013)
11. Gupta, A., Narayan, S., Joseph, K., Khan, S., Khan, F.S., Shah, M.: OW-DETR: open-world detection transformer. In: CVPR (2022)
12. Hall, D., Dayoub, F., Skinner, J., Zhang, H., Miller, D., Corke, P., Carneiro, G., Angelova, A., Sünderhauf, N.: Probabilistic object detection: Definition and evaluation (2018)
13. Han, J., Ren, Y., Ding, J., Pan, X., Yan, K., Xia, G.: Expanding low-density latent regions for open-set object detection. In: CVPR (2022)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2015)
15. He, Y., Chen, W., Tan, Y., Wang, S.: USD: unknown sensitive detector empowered by decoupled objectness and segment anything model **2306.02275** (2023)
16. Hebart, M.N., Zheng, C.Y., Pereira, F., Baker, C.I.: Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature human behaviour* **4**(11), 1173–1185 (2020)
17. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. In: ICLR (2017)
18. Huang, H., Geiger, A., Zhang, D.: GOOD: Exploring geometric cues for detecting objects in an open world. In: ICLR (2023)
19. Joseph, K.J., Khan, S., Khan, F.S., Balasubramanian, V.N.: Towards open world object detection. In: CVPR (2021)
20. Kim, D., Lin, T.Y., Angelova, A., Kweon, I.S., Kuo, W.: Learning open-world object proposals without learning to classify **PP**, 1–1 (2021)
21. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: ICCV (2023)
22. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: DN-DETR: Accelerate detr training by introducing query denoising. In: CVPR (2022)
23. Li, K., Wang, N.Y., Yang, Y., Wang, G.: Sgnet: A super-class guided network for image classification and object detection. In: CRV (2021)

24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017)
25. Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: Dab-detr: Dynamic anchor boxes are better queries for detr. In: ICLR (2021)
26. Ma, S., Wang, Y., Fan, J., Yu Wei, Y., Li, T.H., Liu, H., Lv, F.: CAT: localization and identification cascade detection transformer for open-world object detection. In: CVPR (2023)
27. Maaz, M., Rasheed, H.A., Khan, S., Khan, F.S., Anwer, R.M., Yang, M.: Class-agnostic object detection with multi-modal transformer. In: ECCV (2021)
28. Maaz, M., Rasheed, H.A., Khan, S.H., Khan, F.S., Anwer, R.M., Yang, M.H.: Multi-modal transformers excel at class-agnostic object detection **2111.11430** (2021)
29. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: Psychology of learning and motivation, vol. 24, pp. 109–165. Elsevier (1989)
30. Miller, D., Dayoub, F., Milford, M., Sünderhauf, N.: Evaluating merging strategies for sampling-based uncertainty techniques in object detection (2018)
31. Miller, D., Nicholson, L., Dayoub, F., Sünderhauf, N.: Dropout sampling for robust object detection in open-set conditions (2017)
32. Miller, D., Sünderhauf, N., Milford, M., Dayoub, F.: Uncertainty for identifying open-set errors in visual object detection **7**, 215–222 (2021)
33. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. IEEE TPAMI **39**, 1137–1149 (2015)
34. Saito, K., Hu, P., Darrell, T., Saenko, K.: Learning to detect every thing in an open world. In: ECCV (2022)
35. Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., Smeulders, A.W.M.: Selective search for object recognition. IJCV **104**, 154 – 171 (2013)
36. Wang, Y., Yue, Z., Hua, X.S., Zhang, H.: Random boxes are open-world object detectors. In: ICCV (2023)
37. Wu, Y., Zhao, X., Ma, Y., Wang, D., Liu, X.: Two-branch objectness-centric open world detection. Proceedings of the 3rd International Workshop on Human-Centric Multimedia Analysis (2022)
38. Wu, Z., Lu, Y., Chen, X., Wu, Z., Kang, L., Yu, J.: UC-OWOD: unknown-classified open world object detection. In: ECCV (2022)
39. Yu, J., Ma, L., Li, Z., Peng, Y., Xie, S.: Open-world object detection via discriminative class prototype learning. In: ICIP (2022)
40. Zhao, S., Zhang, Z., Schuster, S., Zhao, L., B.G, V.K., Stathopoulos, A., Chandraker, M., Metaxas, D.N.: Exploiting unlabeled data with vision and language models for object detection. In: ECCV (2022)
41. Zheng, J., Li, W., Hong, J., Petersson, L., Barnes, N.: Towards open-set object detection and discovery. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2022)
42. Zhou, Y., Hu, Q., Wang, Y.: Deep super-class learning for long-tail distributed image classification. PR **80**, 118–128 (2018)
43. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable transformers for end-to-end object detection. In: ICLR (2021)
44. Zohar, O., Wang, K.C., Yeung, S.: PROB: probabilistic objectness for open world object detection. In: CVPR (2023)