

## Exploring Transformers for Open-world Instance Segmentation

Jiannan Wu<sup>1</sup>, Yi Jiang<sup>2</sup>, Bin Yan<sup>3</sup>, Huchuan Lu<sup>3</sup>, Zehuan Yuan<sup>2</sup>, Ping Luo<sup>1,4</sup>

<sup>1</sup>The University of Hong Kong <sup>2</sup>ByteDance

<sup>3</sup>Dalian University of Technology <sup>4</sup>Shanghai AI Laboratory

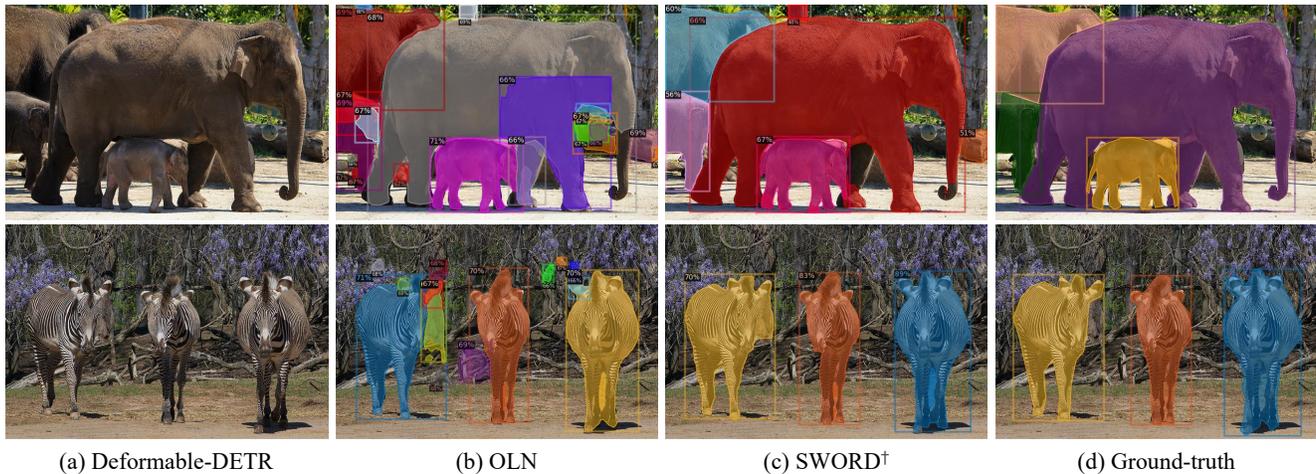


Figure 1: **Comparison of SWORD<sup>†</sup> with different methods.** The images are from the validation set of COCO [31]. All the models are trained on the PASCAL-VOC [12] classes of COCO dataset, where *elephants* and *zebras* are not seen during training. (a) Deformable-DETR [55] fails to segment the objects not labeled in the training set. (b) OLN [26] could localize the *novel* objects, however, it also produce many false positives. (c) Our proposed SWORD<sup>†</sup> can predict the correct and accurate masks for unseen categories.

### Abstract

*Open-world instance segmentation is a rising task, which aims to segment all objects in the image by learning from a limited number of base-category objects. This task is challenging, as the number of unseen categories could be hundreds of times larger than that of seen categories. Recently, the DETR-like models have been extensively studied in the closed world while stay unexplored in the open world. In this paper, we utilize the Transformer for open-world instance segmentation and present SWORD. Firstly, we introduce to attach the stop-gradient operation before classification head and further add IoU heads for discovering novel objects. We demonstrate that a simple stop-gradient operation not only prevents the novel objects from being suppressed as background, but also allows the network to enjoy the merit of heuristic label assignment. Secondly, we propose a novel contrastive learning framework to enlarge the representations between objects and back-*

*ground. Specifically, we maintain a universal object queue to obtain the object center, and dynamically select positive and negative samples from the object queries for contrastive learning. While the previous works only focus on pursuing average recall and neglect average precision, we show the prominence of SWORD by giving consideration to both criteria. Our models achieve state-of-the-art performance in various open-world cross-category and cross-dataset generalizations. Particularly, in VOC to non-VOC setup, our method sets new state-of-the-art results of 40.0% on  $AR_{100}^b$  and 34.9% on  $AR_{100}^m$ . For COCO to UVO generalization, SWORD significantly outperforms the previous best open-world model by 5.9% on  $AP^m$  and 8.1% on  $AR_{100}^m$ .*

### 1. Introduction

The standard instance segmentation models [21, 42, 11] are developed to segment the objects from a predefined taxonomy, which is not often reflective of the diversity of ob-

Table 1: The open-world generalization setups, which are established by the recent advanced approaches [26, 38, 46]. The values in the bracket indicate the class numbers.

Dataset	Train	Evaluate	Image	Mask
<b>Cross-category Generalization</b>				
COCO	VOC(20)	non-VOC(60)	95k	493k
LVIS	COCO(80)	non-COCO(1123)	100k	455k
<b>Cross-dataset Generalization</b>				
UVO	COCO(80)	non-COCO(-)	118k	860k
Objects365	COCO(80)	non-COCO(285)	118k	860k

ject classes encountered in the real world. Recently, class-agnostic **open-world instance segmentation** introduced by the advanced approaches [26, 38, 46] has gained increasing attention in the community. It requires the models to segment all objects of arbitrary categories in the image while only base-category objects can be seen during training. This task is highly challenging, as the number of unseen categories can be orders of magnitude larger than the number of seen categories. As shown in the second row of Table 1, for COCO to LVIS setup, there are 1123 non-COCO classes for out-of-domain evaluation while only objects belonging to 80 COCO classes are annotated in the training set. Besides, a critical challenge in the open-world scenario is that the *novel* objects and background co-exist in the un-annotated regions. Consequently, the closed-world instance segmentation models fail to recognize unseen objects (Figure 1a) as they equally treat the *novel* objects and background as negative samples during training.

Recently, DETR-like [3, 55] models based on Transformers [43] have exhibited superior performance in standard object detection and instance segmentation tasks. However, the study of these Transformer-based models in the field of open-world instance segmentation is still a blank page to the community, as the previous works have exclusively relied on the Mask-RCNN [21] architecture. In this work, we aim to fill in the gap by delving deeply into the recent advanced Deformable-DETR [55].

An inspiring open-world method OLN [26] proposes a classification-free network and estimates the scores of regions purely by localization quality (e.g., IoU score). In this manner, *novel* objects would not be penalized as background due to the absence of classification learning. And the localization quality score is proven to be a better objectness cue for discovering *novel* objects. Inspired by this spirit, a straightforward solution for Transformer-based models in open-world instance segmentation is to replace the classification head with IoU heads. However, this could lead to **two negative effects**: (i) The Transformer-based models are optimized with set prediction loss [3], where the classification score is indispensable for label assign-

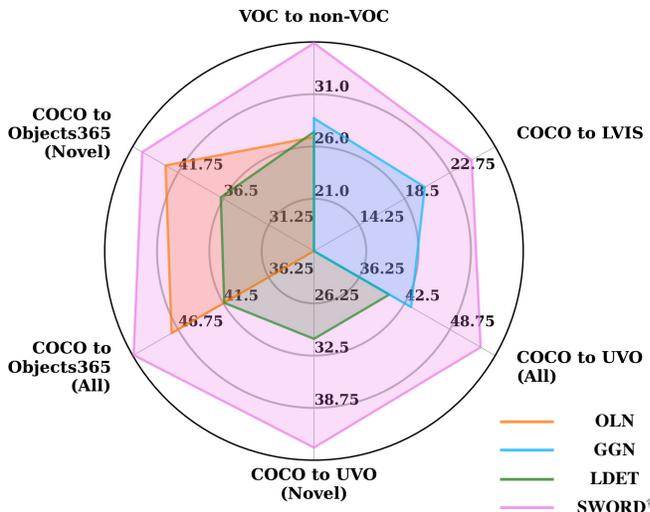


Figure 2: SWORD<sup>†</sup> achieves the state-of-the-art performance on various settings compared with other open-world methods. The results are reported based on AR<sub>100</sub><sup>m</sup> by default. The metric of Objects365 [39] is AR<sub>100</sub><sup>b</sup> since it does not provide mask annotations.

ment. Therefore, simply removing the classification head could be harmful to the process. (ii) The network would inherit the limitation of OLN that lacks the discriminative ability to differentiate the objects and background. This is because OLN is only trained with positive samples and thus fails to perceive the background. As a result, it would produce numerous false positives (Figure 1b) and result in fairly low average precision (AP). For example, in COCO to UVO (all) generalization, AR<sub>100</sub><sup>m</sup> of OLN [26] increases from 36.7% to 42.1% while AP<sup>m</sup> significantly drops from 20.7% to 14.0% when compared with Mask-RCNN [21].

In this work, we cut off the above obstacles and propose SWORD, unsealing the secrets of Transformer-based models for open-world instance segmentation. We first introduce to attach a `stop-grad` operation before the classification head and further add the IoU heads for predicting object scores. This not only prevents the *novel* object from being suppressed as background so as to improve the recall ability of network, but also allows the DETR-like models to preserve the classification head for heuristic label assignment. Using `stop-grad` alone, however, would inevitably reduce the network’s discrimination. Therefore, we then design a novel contrastive learning framework for learning the discriminative representations between objects and background. The core idea is to ensure similar representations among objects while enlarging the distinction between the objects and background in the feature space. Specifically, we maintain a universal object queue to store the annotated object embeddings during training. The averaged feature of the queue, i.e., object center, captures the common characteristics of objects and plays as the role of

query in contrastive learning. The positive and hard negative samples are dynamically selected from query embeddings according to the matching cost [44, 34] with ground-truth. Contrastive learning is the key to reducing false positives for network and greatly improves average precision.

To this end, SWORD not only reveals the strong ability in recalling *novel* objects, but also achieves high average precision. We further develop a model SWORD<sup>†</sup>, which has exactly the same architecture of Deformable-DETR [55] but is trained on the combination of annotations and pseudo ground-truths generated by SWORD. SWORD<sup>†</sup> shows clear performance gains compared with its counterpart under the open-world setups. To summarize, the contributions of this work are:

- A simple yet effective framework SWORD is presented, which is the first study of Transformer-based model for open-world instance segmentation.
- We introduce the `stop-grad` operation to kill two birds with one stone: preventing the side-effect of classification learning to discover *novel* objects in open-world setups, and enabling the heuristic label assignment for DETR-like models.
- We design a novel contrastive learning method to learn the discriminative representations between objects and background, which is essential for achieving high average precision for the network.
- Extensive experiments demonstrate that our models achieve state-of-the-art performance on several benchmarks including COCO [31], LVIS [16], UVO [47] and Objects365 [39], as shown in Figure 2.

## 2. Related Work

**Open-world Instance Segmentation.** Towards building more practical applications in the real world, the open-world-related problems [2, 6, 17, 6, 35, 27, 40, 53, 29, 52] have raised great attention recently. Kim et al. [26] firstly establish the protocol of open-world instance segmentation. Literally, open-world models should not only segment all the previously seen objects, but also localize the unseen objects during inference. There are several works [26, 38, 46, 51, 23, 45, 24] attempting to solve the problem from various aspects. We refer the readers to Appendix A for a comprehensive review of these works.

**Discussion.** To dig out the problems of current models in open-world scenario, we visualize the predicted results of a closed-world model Deformable-DETR [55] and an open-world model OLN [26] in Figure 3. (i) For Deformable-DETR, we notice that the *novel* objects have low activations in the feature map, which we term as *feature degradation*. This is because *novel* objects are treated as background during training. As shown in the middle picture of



Figure 3: **Visualization results of (a) closed-world model Deformable-DETR and (b) open-world model OLN.** For each example, we show the input images, feature maps and predicted results from left to right. Note that *elephant* is the unseen category in the training set.

Figure 3a, the elephant is the unseen category and its feature can hardly be distinguished from its surroundings. (ii) OLN [26] has the generalized ability to localize *novel* objects. Nevertheless, it suffers from the problem of producing many false positives, *e.g.*, parts of the man’s body in Figure 3b. OLN [26] is trained with positive samples and can not perceive backdrop. It would assign high scores for all the object proposals. To conclude, there are two critical issues for the open-world instance segmentation: *preventing the feature degradation of novel objects* and *learning the discrimination between objects and background*.

**Contrastive Learning.** Self-supervised learning could be divided into three groups: contrastive learning, self-distillation [4, 54] and masked image modeling [19, 8, 41]. Among which, contrastive learning [20, 7, 15, 50, 48, 25, 1, 25, 18] has been dominant for a long time. The core idea of it lies in that the positive samples are attracted while the negative samples are pulled away in the feature space to learn the discriminative representations. MoCo [20] maintains a memory queue to store a large number of negative pairs and enables the momentum update of the memory encoder to guarantee the queue feature consistency. MoCo v3 [10] discards the memory queue and conducts self-supervised training on visual Transformers. SimSiam [9] develops the extremely simple siamese network without any negative sample, and points out a stop-gradient operation plays an essential role in preventing mode collapse. In this work, we absorb the ideas from contrastive learning to learn the distinct representations of objects and background.

## 3. Method

### 3.1. Overview

Open-world instance segmentation problem introduced by recent advanced works [26, 38, 46] can be formulated as follows. Given an instance segmentation dataset (*e.g.*, COCO with 80 classes), we have the object annotations on the *base* category set  $\mathcal{C}_{\text{base}}$  (*e.g.*,  $\mathcal{C}_{\text{base}}$  contains 20 PASCAL-

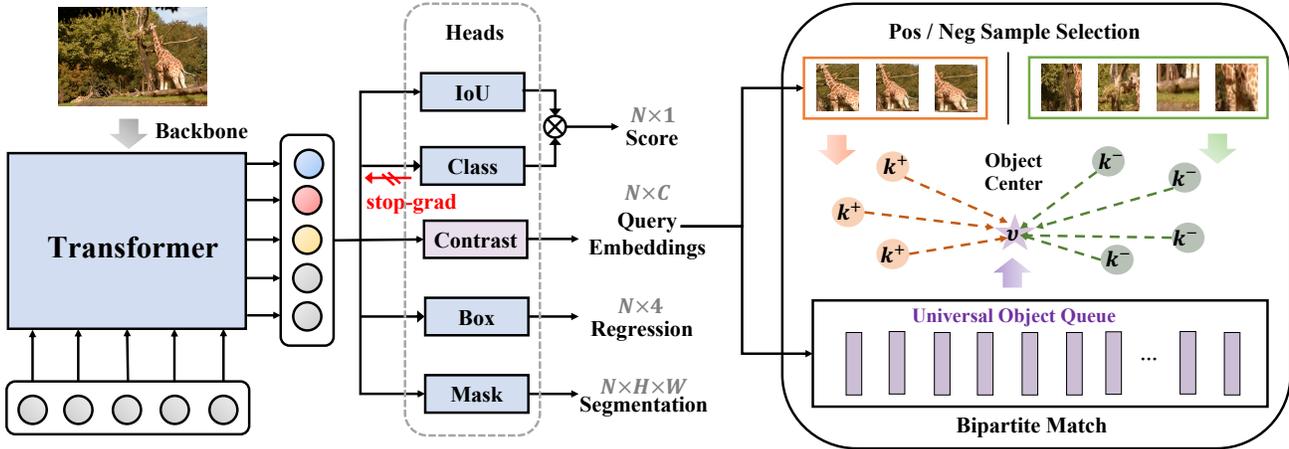


Figure 4: **The overall framework of SWORD.** We first attach a `stop-grad` operation before the classification head and add the IoU head to help discover *novel* objects. Further, a contrastive head is added on top of the Transformer decoder to predict the query embeddings for contrastive learning. It is essential for the network to learn distinct representations between objects and background. During inference, the classification scores and IoU scores are fused to produce the final scores.

VOC classes). Notably, there are also a large number of *novel* objects co-appearing in the images while remaining un-annotated. The models are trained with the base-category annotations to provide a set of class-agnostic proposals  $\mathcal{P} = \{s_i, b_i, m_i\}_{i=1}^p$  to localize *all* objects in the image, where  $s_i \in \mathbb{R}$  indicates the proposal score,  $b_i \in \mathbb{R}^4$  denotes the bounding box coordinates and  $m_i \in \mathbb{R}^{H \times W}$  is the segmentation mask for the  $i$ -th prediction. The generalization of models is evaluated by segmenting the *novel* objects from the unseen category set  $\mathcal{C}_{\text{novel}}$  (e.g., 60 non-PASCAL-VOC classes) in a class-agnostic fashion.

The overall framework of our proposed SWORD is illustrated in Figure 4. Our network is based on the closed-world model Deformable-DETR [55], and we explain how to convert it into an open-world instance segmentation model in Section 3.2. We then propose to utilize contrastive learning to generate distinct representations for objects and background in Section 3.3. Additionally, we introduce an extension of our model, denoted as SWORD<sup>†</sup>, through pseudo labeling based self-training in Section 3.5.

### 3.2. Open-world Transformer

Our network is built upon the Deformable-DETR [55] due to its simple architecture. First, we add the mask head on top of the Transformer to generate the instance masks by performing dynamic convolution [42, 11, 49]. Then, the model is transformed into an open-world model for learning class-agnostic mask proposals with following designs.

**IoU Heads.** Inspired by the philosophy of OLN [26] that localization quality is a better objectness cue than the classification score in the open-world setting, we add the extra two IoU heads on top of the Transformer decoder to predict

Table 2: The illustration of the key role of `stop-grad` operation. ‘class’ and ‘IoU’ in the first column mean classification head and IoU head, respectively.

Variants	Network	Heuristic
	Generalization	Label Assignment
class only	✗	✓
IoU only	✓	✗
class + IoU	✗	✓
class + IoU + stop-grad	✓	✓

the box IoU score  $c_b$  and mask IoU score  $c_m$ , respectively. The IoU scores are helpful to discover the *novel* objects.

**Stop-gradient Operation.** The object-or-not learning of classification can hurt the generalization of network, while DETR-like models need to preserve the classification head for heuristic label assignment (e.g., Hungarian matching [3]). *This leads to a conflict situation.* To address this issue, we propose a simple yet effective solution by introducing a `stop-grad` operation [36, 9] before the classification head. This operation prevents the gradient passing from the classification head to the network, which avoids suppressing all un-annotated regions as background. On the other hand, it can be seamlessly applied on the advanced DETR-like detectors [3, 32, 28] to facilitate the heuristic label assignment. As illustrated in Table 2, the `stop-grad` operation can confer all the desired properties upon the DETR-like models, thus enabling the networks to possess the open-world capacity in a manner akin to the OLN paradigm.

### 3.3. Contrastive Learning

The use of `stop-grad` operation would also disable the discriminative ability of network. To address the issue,

in this subsection, we propose a contrastive learning [20, 15, 9] framework to learn the distinct representations between objects and background.

**Universal Object Queue.** As shown in Figure 4, we further add the contrastive head on top of the Transformer decoder to learn the query embeddings. The parameters of the contrastive head are also copied to a momentum branch using the exponential moving average (EMA) method:

$$\theta'_c \leftarrow \alpha \theta'_c + (1 - \alpha) \theta_c \quad (1)$$

where  $\theta_c$  and  $\theta'_c$  denote the parameters of the regularly and EMA updated contrastive head, respectively.  $\alpha$  is the momentum rate. We use a universal object queue  $\mathcal{Q}$  to store the object embeddings, where each element is the projected feature from the momentum contrastive head. Given an image, we select those queries best matched to the ground-truth through bipartite matching and store their projected embeddings into  $\mathcal{Q}$  using the first-in-first-out strategy. We then average all the features in the universal object queue to get the object center  $v$ . Intuitively, the object center captures the common object characteristics and stays stable in the feature space.

**Positive and Negative Samples.** It is well known that the positive and negative samples play the essential role in contrastive learning. Here, we dynamically select the positive and negative pairs according to an optimal transport assignment method [13, 44, 34, 14]. Specifically, given an image, we take the classification results into consideration and compute the costs between predictions and ground-truths:

$$\mathcal{C} = \lambda_{cls} \cdot \mathcal{C}_{cls} + \lambda_{L1} \cdot \mathcal{C}_{L1} + \lambda_{giou} \cdot \mathcal{C}_{giou} \quad (2)$$

where,  $\mathcal{C}_{cls}$  is focal loss [30]. The box-related losses include the  $\mathcal{L}_1$  loss and generalized IoU loss [37]. Ideally, the predictions with the least cost are those objects close to the ground-truths. To improve the quality of learned embeddings, for each ground-truth object, we first dynamically choose  $k_1$  and  $k_2$  predictions with the least cost, where  $k_2 > k_1$ . Then the  $k_1$  predictions are positive samples, and the left  $k_2 - k_1$  predictions are considered as hard negatives.

**Contrastive Loss.** In the contrastive learning, we expect the positive samples  $\mathcal{K}^+$  should be close to the object center  $v$  while the negative ones  $\mathcal{K}^-$  should be pulled away. The contrastive loss [20] is formulated as:

$$\mathcal{L}_{con} = -\log \frac{\sum_{k^+ \in \mathcal{K}^+} \exp(v \cdot k^+)}{\sum_{k^+ \in \mathcal{K}^+} \exp(v \cdot k^+) + \sum_{k^- \in \mathcal{K}^-} \exp(v \cdot k^-)} \quad (3)$$

### 3.4. Training and Inference

**Training.** The label assignment for the network optimization also relies on the matching cost as Eq. (2). The predictions with the least costs are assigned to ground-truths as

positive samples and others as negatives [14]. The overall loss function for training is:

$$\begin{aligned} \mathcal{L} = & \lambda_{cls} \cdot \mathcal{L}_{cls} + \lambda_{L1} \cdot \mathcal{L}_{L1} + \lambda_{giou} \cdot \mathcal{L}_{giou} \\ & + \lambda_{mask} \cdot \mathcal{L}_{mask} + \lambda_{dice} \cdot \mathcal{L}_{dice} \\ & + \lambda_{iou} \cdot \mathcal{L}_{iou} + \lambda_{con} \cdot \mathcal{L}_{con} \end{aligned} \quad (4)$$

where,  $\mathcal{L}_{cls}$ ,  $\mathcal{L}_{L1}$  and  $\mathcal{L}_{giou}$  are the same as the components in Eq. 2. The mask-related loss is a combination of the mask binary loss and DICE loss [33]. The IoU scores are supervised by the binary cross entropy loss  $\mathcal{L}_{iou}$ .

**Inference.** We use the geometric mean of classification scores and IoU scores as final scores, i.e.,  $s = \sqrt[3]{c_c \cdot c_b \cdot c_m}$ . And the top-100 predictions are left for evaluation.

### 3.5. Extension: Pseudo Ground-truth Training

The advanced approach GGN [46] proves that using the pseudo ground-truth for training can significantly boost the performance of Mask-RCNN [21] in the open-world setup. Therefore, we also adopt the pseudo labeling method and develop an extension model, SWORD<sup>†</sup>. Following the existing practice of GGN [46], we use SWORD as teacher model to generate the pseudo boxes/masks. After filtering out those predictions having high box overlap with ground-truth, we add the remaining top- $k$  predictions to the ground-truth annotations. Finally, the standard Deformable-DETR is trained under the supervision of augmented annotations. Please see more details in Appendix B.

## 4. Experiments

In this section, we first thoroughly evaluate the performance of proposed models in two challenging open-world settings: cross-category and cross-dataset generalizations. Then we conduct extensive ablation studies to discuss the key designs and analyze the crucial issues in Sec. 4.4.

### 4.1. Experiment Settings

**Datasets.** Our experiments are conducted on COCO [31], LVIS [16], UVO [47] and Objects365 [39] datasets. **COCO** is the widely used instance segmentation benchmark with 80 categories. **LVIS** shares the same images with COCO while having a more complete label system. It has a large taxonomy of 1203 categories in a long-tailed distribution. **UVO** originates from the Kinetics400 [5] dataset and all the instance masks are exhaustively annotated. **Objects365** is a large-scale object detection dataset with 365 categories where all the COCO 80 categories are included.

We target at two challenging open-world generalization setups [26, 38, 46]: **(1) Cross-category generalization.** On COCO benchmark, we follow the common practice [26, 38] to split the annotations into two non-overlapping class sets, where the PASCAL-VOC [12] 20 classes are adopted as the *base* set and the rest of 60 non-VOC classes are *novel* set.

The second benchmark, *i.e.*, LVIS, splits the 1203 classes into 80 COCO classes for training and the remaining 1123 non-COCO classes for evaluation. For cross-category generalization, the results are reported on the unseen categories.

**(2) Cross-dataset generalization.** This is to evaluate the model’s open-world generalization ability when used in the wild. COCO is used as the training source and the models are tested on new datasets, *i.e.*, UVO<sup>1</sup> and Objects365<sup>2</sup>. In this setting, we show the results on both *novel* categories and all categories (including base and novel ones).

**Evaluation Metrics.** Following previous works [26, 38, 46], we use average recall (AR@k) and average precision (AP) over multiple IoU thresholds [0.5 : 0.95] to measure the performance. The proposal number k is set as 100 by default. The superscripts ‘b’ and ‘m’ denote the boxes and masks, respectively. And  $AR_{s/m/l}$  represent AR@100 for small, medium and large size of objects. Notably, **the most concerned metric** in open-world scenario [26] is **AR@100**.

**Implementation Details.** In all setups, models are trained and evaluated in a class-agnostic way. We use ResNet-50 [22] as backbone by default and the Transformer network has 6 encoders and 6 decoders with the hidden dimension of 256. We use 2000 object queries when the training source is VOC(COCO), otherwise the query number is set as 1000. The size of the universal object queue is set as 4096 and the EMA rate  $\alpha$  is 0.999. The values of  $k_1$  and  $k_2$  for contrastive learning are set as 10 and 100, respectively. In all our experiments, we also train the Deformable-DETR using the same setting as ours (*i.e.*, object query number, training epochs) for fair comparisons. Please see Appendix C for more implementation details.

## 4.2. Cross-category Generalization

**VOC to non-VOC.** In Table 3, we compare our methods with other state-of-the-art methods in VOC to non-VOC setup. It shows that SWORD yields the significant 7.9% gain on  $AR_{100}^b$  and 7.5% on  $AR_{100}^m$  compared with the Deformable-DETR baseline. And our model outperforms all the previous single model. The performance could be further boosted by exploiting the pseudo ground-truth of SWORD. SWORD<sup>†</sup> achieves state-of-the-art performance in all metrics, *e.g.*, 40.0% on  $AR_{100}^b$  and 34.9% on  $AR_{100}^m$ .

**COCO to LVIS.** Table 4 summarizes the performance of different methods in the COCO to LVIS setup. Compared to the Deformable-DETR baseline, SWORD shows an obvious performance gain, with +4.1% on  $AR_{100}^b$  and +4.0%

<sup>1</sup>The downsampled dense split of v1.0 contains two classes: “objects” for COCO categories and “other” for non-COCO categories. The NOVEL metrics are measured on the “other” categories. The previously released v0.5 does not distinguish the object categories and all the objects are annotated as “objects”. We report the results of ALL metrics based on this version following the previous works [38, 46]

<sup>2</sup>Objects365 only has box annotations, so we report the results regarding the box metrics.

Table 3: State-of-the-art performance in VOC to non-VOC setup.

Method	AP <sup>b</sup>	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
Mask-RCNN [21]	1.6	10.2	23.5	0.9	7.9	17.7
OLN [26]	3.7	18.0	33.5	-	16.9	-
LDET [38]	5.0	18.2	30.8	4.3	16.3	27.4
GGN [46]	5.8	17.3	31.6	4.9	16.1	28.7
GGN + OLN [46]	3.4	17.1	37.2	3.2	16.4	33.7
GOOD [23]	-	-	39.3	-	-	-
Deformable-DETR [55]	2.5	12.2	27.4	2.2	10.2	22.7
SWORD (Ours)	5.8	17.8	35.3	4.8	15.7	30.2
SWORD <sup>†</sup> (Ours)	<b>6.2</b>	<b>22.0</b>	<b>40.0</b>	<b>5.8</b>	<b>20.2</b>	<b>34.9</b>

Table 4: State-of-the-art performance in COCO to LVIS setup.

Method	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
Mask-RCNN [21]	6.1	19.4	5.6	17.2
GGN [46]	7.6	22.4	7.2	20.4
Deformable-DETR [55]	6.3	19.4	5.5	16.4
SWORD (Ours)	8.8	23.5	8.0	20.4
SWORD <sup>†</sup> (Ours)	<b>9.8</b>	<b>28.0</b>	<b>9.0</b>	<b>23.8</b>

on  $AR_{100}^m$ . Additionally, SWORD<sup>†</sup> outperforms the previous best method GGN [46] by 5.6%  $AR_{100}^b$ , which is a relative improvement of 25%.

## 4.3. Cross-dataset Generalization

**COCO to UVO.** For COCO to UVO generalization, we evaluate models trained with 20 PASCAL-VOC classes and all 80 COCO classes. Table 5 presents a thorough comparison of results for both *novel* and *all* objects. Deformable-DETR performs considerable well in this setting, outperforming previous methods by a large margin. Our proposed model, SWORD, further improved the performance over the strong baseline for all metrics. Notably, in the partially annotated VOC(COCO) to UVO scenario, SWORD demonstrated a significant 9.6%  $AR_{100}^m$  gain for *novel* objects. This highlights the effectiveness of our model in discovering *novel* objects.

Another observation is that pseudo ground-truth training is effective in improving the average recall (AR) of models but may decrease the average precision (AP). By comparing Deformable-DETR and SWORD<sup>†</sup> (both have the same architecture), we noticed a consistent gain in AR. However, AP<sup>b</sup> for all objects dropped from 29.1% to 28.1% in the last row. The reason may attribute to that false positives in the pseudo ground-truth labels can misguide the model training.

**COCO to Objects365.** The results of COCO to Objects365 generalization are listed in Table 6. Mask-RCNN-based method LDET<sup>3</sup> improves Mask-RCNN in terms of ARs but decrease APs. SWORD significantly outperforms D-DETR baseline for all metrics and achieve SOTA performance.

<sup>3</sup>We report the results of LDET [38] using the same class-agnostic evaluation for fair comparison, whereas results in the original paper are based on the class-wise evaluation.

Table 5: Comparison of state-of-the-art performance in COCO to UVO setup. Top rows: Models are trained with 20 PASCAL-VOC classes on COCO dataset. Bottom rows: Models are trained with all 80 COCO classes on COCO dataset.

Method	Train	Novel						All					
		AP <sup>b</sup>	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>	AP <sup>b</sup>	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
Mask-RCNN [21]	VOC	5.9	11.4	16.2	2.3	7.6	11.4	20.2	25.3	30.8	15.7	20.1	24.3
LDET [38]		9.3	16.0	31.9	4.9	12.3	25.2	22.7	28.1	43.3	18.7	23.9	36.0
Deformable-DETR [55]		7.2	13.5	33.5	3.4	9.5	25.3	23.4	29.4	49.8	19.1	24.0	39.4
SWORD (Ours)		11.2	16.8	43.1	6.1	13.3	34.9	<b>24.9</b>	30.6	55.3	19.6	25.3	45.2
SWORD <sup>†</sup> (Ours)		<b>11.8</b>	<b>18.4</b>	<b>45.6</b>	<b>8.4</b>	<b>16.8</b>	<b>38.1</b>	23.4	<b>31.1</b>	<b>59.2</b>	<b>21.0</b>	<b>28.4</b>	<b>49.5</b>
Mask-RCNN [21]	COCO	11.8	16.4	30.4	7.0	13.8	25.5	25.7	30.2	43.8	20.7	25.7	36.7
LDET [38]		12.9	19.0	35.9	8.2	15.9	30.5	26.0	30.9	47.0	22.1	27.3	40.7
GGN [46]		-	-	-	-	-	-	24.0	29.8	52.2	20.3	-	43.4
Deformable-DETR [55]		14.2	20.0	45.8	9.0	16.7	37.9	29.1	35.0	60.7	24.7	30.1	50.3
SWORD (Ours)		<b>17.5</b>	22.2	48.1	<b>12.8</b>	19.4	40.6	<b>32.0</b>	<b>36.5</b>	61.2	<b>28.0</b>	32.4	51.5
SWORD <sup>†</sup> (Ours)		16.6	<b>22.7</b>	<b>50.0</b>	12.7	<b>20.9</b>	<b>42.8</b>	28.1	35.2	<b>62.0</b>	25.7	<b>32.5</b>	<b>53.0</b>

Table 6: Comparison of state-of-the-art performance in COCO to Objects365 setup.

Method	Novel						All					
	AP <sup>b</sup>	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AR <sub>s</sub> <sup>b</sup>	AR <sub>m</sub> <sup>b</sup>	AR <sub>l</sub> <sup>b</sup>	AP <sup>b</sup>	AR <sub>10</sub> <sup>b</sup>	AR <sub>100</sub> <sup>b</sup>	AR <sub>s</sub> <sup>b</sup>	AR <sub>m</sub> <sup>b</sup>	AR <sub>l</sub> <sup>b</sup>
Mask-RCNN [21]	13.0	19.3	32.8	18.2	36.4	43.5	25.1	23.9	40.3	22.7	42.8	53.4
LDET [38]	12.8	20.0	36.8	20.7	40.5	48.9	22.5	22.7	41.4	22.9	44.3	54.9
Deformable-DETR [55]	12.9	19.0	40.1	22.8	43.4	54.1	27.3	25.3	48.7	27.5	50.9	65.6
SWORD (Ours)	<b>16.6</b>	22.8	43.9	25.0	48.6	57.6	<b>29.7</b>	<b>27.3</b>	50.8	28.6	54.0	67.2
SWORD <sup>†</sup> (Ours)	16.3	<b>23.5</b>	<b>45.9</b>	<b>25.9</b>	<b>50.5</b>	<b>60.7</b>	28.7	27.2	<b>51.9</b>	<b>29.4</b>	<b>55.4</b>	<b>68.4</b>

#### 4.4. Ablation Study

In this subsection, we conduct extensive ablation studies to analyze the crucial composing of our method. The experimental results are based on the COCO (80 classes) to UVO setting and the backbone is ResNet50 otherwise specified. We report the result in terms of mask metrics.

**Analysis of Key Designs.** Table 7 presents the ablation results to study the key designs of our method. In addition to the standard Deformable-DETR, we also build an OLN version Deformable-DETR by replacing the classification head with IoU heads to make a comparison. Since COCO is fully annotated and there are hardly any *novel* objects, the co-existence of objects and background in un-annotated areas is not a significant issue. Therefore, we mainly focus on the discussion of **VOC(COCO) to UVO generalization**.

Starting from the closed-world Deformable-DETR, we first introduce the *stop-grad* operation to transform it into an open-world model. This simple yet effective operation prevents un-annotated regions from being suppressed as background. The obvious background-to-foreground transition for *novel* objects leads to the simultaneous improvement of AP and AR for *novel* objects. However, the use of *stop-grad* reduces the network’s discrimination, resulting in many false positives predictions and a performance drop in AP for *all* objects. Henceforth, the proposed contrastive learning framework is indispensable. This design significantly improves AP<sup>m</sup> for *all* objects by 6.5%,

and all ARs show steady performance improvement. The two key designs are also validated in the COCO to UVO setup, where *stop-grad* improves AR and contrastive learning greatly increases AP. These results demonstrate the effectiveness of our proposed approach.

From another perspective, our proposed SWORD also demonstrates performance advantages over the OLN-version Deformable-DETR. SWORD not only reveals consistent performance advantages on ARs, but also shows 5.1 and 11.2 points gain on the AP<sup>m</sup> for VOC to UVO and COCO to UVO setups, respectively. These results proves the superiority of our method.

**Classification Cost for Sample Selection in Contrastive Learning.** To evaluate the effect of classification cost in the contrastive learning, we set  $C_{cls} = 0$  in Eq. (2) for the ablation. From Table 8, we observe that performance drops drastically without classification cost. Such phenomenon stands with the view that classification score is crucial for two potential reasons. First, the classification cost ensures the network’s consistency in assigning labels during both contrastive learning and network training. Second, the localization cost alone will introduce those predictions closest to the ground-truths as positive samples, while classification cost helps choose more discriminative samples.

**The Values of  $k_1$  and  $k_2$  in Contrastive Learning.** To study the impact on the number of positive and negative samples, we provide the ablation results of  $k_1$  and  $k_2$  in Ta-

Table 7: Ablation on the key designs of our method. ‘D-DETR’ represents Deformable-DETR [55]. The results of an OLN version Deformable-DETR are also presented for comparison. We start from the Deformable-DETR baseline and gradually add the key components. The final model is the proposed SWORD.

Variants	Novel			All		
	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
<i>VOC(COCO) to UVO</i>						
D-DETR OLN-version	5.8	11.4	31.0	14.5	23.4	43.1
D-DETR	3.4	9.5	25.3	19.1	24.0	39.4
+ stop-grad	4.7 (+1.3)	11.6 (+2.1)	34.1 (+8.8)	13.1 (−6.0)	22.1 (−1.9)	45.1 (+5.7)
+ contrastive learning	6.1 (+1.4)	13.3 (+1.7)	34.9 (+0.8)	19.6 (+6.5)	25.3 (+3.2)	45.2 (+0.1)
<i>COCO to UVO</i>						
D-DETR OLN-version	10.0	17.5	40.9	17.8	28.0	51.2
D-DETR	9.0	16.7	37.9	24.7	30.1	50.3
+ stop-grad	10.3 (+1.3)	18.4 (+1.7)	41.6 (+3.7)	21.6 (−3.1)	30.5 (+0.4)	52.0 (+1.7)
+ contrastive learning	12.8 (+2.5)	19.4 (+1.0)	40.6 (−1.0)	28.0 (+6.4)	32.4 (+1.9)	51.5 (−0.5)

Table 8: Ablation on the classification cost for sample selection in contrastive learning.

class cost	Novel			All		
	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
✗	8.9	17.0	40.7	16.6	26.2	51.4
✓	12.8	19.4	40.6	28.0	32.4	51.5

Table 9: Ablation on the values of  $k_1$  and  $k_2$  in contrastive learning.

$k_1$	$k_2$	Novel			All		
		AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>	AP <sup>m</sup>	AR <sub>10</sub> <sup>m</sup>	AR <sub>100</sub> <sup>m</sup>
1	100	11.4	19.7	40.1	25.5	31.6	50.7
10	100	12.8	19.4	40.6	28.0	32.4	51.5
50	100	8.9	17.0	40.3	19.3	28.9	51.5
10	20	12.1	19.2	40.3	27.4	32.2	51.5
10	200	13.0	20.8	39.7	28.3	32.8	50.6

ble 9. The first three rows show that increasing the number of positive samples can result in a large number of false positives, which negatively affect the average precision (AP). By comparing line 2-4-5, it indicates that more negative samples benefit AP while hurt AR@100. This is reasonable because more negative samples can aid contrastive learning in generating more distinct representations; however, it may also incorrectly identify real objects as negative samples and suppress them as background.

#### 4.5. Visualization

In Figure 5, we visualize the score distribution of different methods in VOC to non-VOC setting. Deformable-DETR [55] can only find out the seen category objects and thus its score distribution is mainly located on the low-scoring areas. OLN [26] is trained with positive samples, making it merely produce the high-scoring proposals. Although it reveals the open-world ability to locate *novel* objects, it can not effectively discriminate the objects and

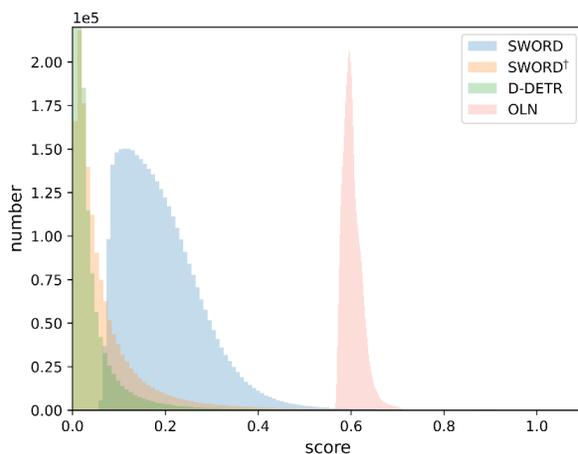


Figure 5: Comparing score distributions of proposals for different methods on COCO [31] validation set. ‘D-DETR’ represents Deformable-DETR [55]. All models are trained with 20 PASCAL-VOC classes. For fair comparisons, we select the top-100 proposals per image for all the methods.

background. The proposed SWORD not only displays the favorable open-world generalization but also provide distinct confidence scores for objects and background.

## 5. Conclusion

In this work, we present a Transformer-based framework, SWORD, for open-world instance segmentation. Specifically, we introduce the `stop-grad` operation to prevent the feature degradation of *novel* objects and propose a contrastive learning strategy to enlarge the discriminating representations between objects and background. We also develop an extension model, SWORD<sup>†</sup>, by exploiting the pseudo labels of SWORD. Extensive experiments demonstrate that the proposed models achieve state-of-the-art performance in various open-world generalization setups.

## Acknowledgements

This paper is partially supported by the National Key R&D Program of China No.2022ZD0161000 and the General Research Fund of Hong Kong No.17200622. The paper is supported in part by the National Natural Science Foundation of China under grant No.62293540, 62293542, U1903215 and the Fundamental Research Funds for the Central Universities No.DUT22ZD210.

## References

- [1] Yutong Bai, Xinlei Chen, Alexander Kirillov, Alan Yuille, and Alexander C Berg. Point-level region contrast for object detection pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16061–16070, 2022. 3
- [2] Abhijit Bendale and Terrance Boulton. Towards open world recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1893–1902, 2015. 3
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2, 4
- [4] 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. 3
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 5
- [6] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Deep metric learning for open world semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15333–15342, 2021. 3
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 3
- [8] Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. *arXiv preprint arXiv:2202.03026*, 2022. 3
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 3, 4, 5
- [10] Xinlei Chen\*, Saining Xie\*, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. 3
- [11] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 1, 4
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 1, 5
- [13] Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 303–312, 2021. 5
- [14] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. 5
- [15] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 3, 5
- [16] 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. 3, 5
- [17] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8401–8409, 2019. 3
- [18] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *Advances in Neural Information Processing Systems*, 33:5679–5690, 2020. 3
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 3
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3, 5
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 2, 5, 6, 7
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [23] Haiwen Huang, Andreas Geiger, and Dan Zhang. Good: Exploring geometric cues for detecting objects in an open world. *arXiv preprint arXiv:2212.11720*, 2022. 3, 6
- [24] Tarun Kalluri, Weiyao Wang, Heng Wang, Manmohan Chandraker, Lorenzo Torresani, and Du Tran. Open-world instance segmentation: Top-down learning with bottom-up supervision. *arXiv preprint arXiv:2303.05503*, 2023. 3

- [25] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. 3
- [26] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *IEEE Robotics and Automation Letters*, 7(2):5453–5460, 2022. 1, 2, 3, 4, 5, 6, 8
- [27] Gukyeong Kwon, Mohit Prabhushankar, Dogancan Temel, and Ghassan AlRegib. Backpropagated gradient representations for anomaly detection. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 206–226. Springer, 2020. 3
- [28] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022. 4
- [29] Chuang Lin, Peize Sun, Yi Jiang, Ping Luo, Lizhen Qu, Ghohamreza Haffari, Zehuan Yuan, and Jianfei Cai. Learning object-language alignments for open-vocabulary object detection. *arXiv preprint arXiv:2211.14843*, 2022. 3
- [30] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1, 3, 5, 8
- [32] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 4
- [33] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016. 5
- [34] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. 3, 5
- [35] Lu Qi, Jason Kuen, Yi Wang, Jiuxiang Gu, Hengshuang Zhao, Zhe Lin, Philip Torr, and Jiaya Jia. Open-world entity segmentation. *arXiv preprint arXiv:2107.14228*, 2021. 3
- [36] Limeng Qiao, Yuxuan Zhao, Zhiyuan Li, Xi Qiu, Jianan Wu, and Chi Zhang. Defrcn: Decoupled faster r-cnn for few-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8681–8690, 2021. 4
- [37] Hamid Rezaatofghi, 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. 5
- [38] Kuniaki Saito, Ping Hu, Trevor Darrell, and Kate Saenko. Learning to detect every thing in an open world. *arXiv preprint arXiv:2112.01698*, 2021. 2, 3, 5, 6, 7
- [39] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. 2, 3, 5
- [40] Jingbo Sun, Li Yang, Jiaxin Zhang, Frank Liu, Mahantesh Halappanavar, Deliang Fan, and Yu Cao. Gradient-based novelty detection boosted by self-supervised binary classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8370–8377, 2022. 3
- [41] Keyu Tian, Yi Jiang, Qishuai Diao, Chen Lin, Liwei Wang, and Zehuan Yuan. Designing bert for convolutional networks: Sparse and hierarchical masked modeling. *arXiv preprint arXiv:2301.03580*, 2023. 3
- [42] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *European conference on computer vision*, pages 282–298. Springer, 2020. 1, 4
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [44] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009. 3, 5
- [45] Cheng Wang, Guoli Wang, Qian Zhang, Peng Guo, Wenyu Liu, and Xinggang Wang. Openinst: A simple query-based method for open-world instance segmentation. *arXiv preprint arXiv:2303.15859*, 2023. 3
- [46] 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. 2, 3, 5, 6, 7
- [47] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10776–10785, 2021. 3, 5
- [48] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021. 3
- [49] Junfeng Wu, Qihao Liu, Yi Jiang, Song Bai, Alan Yuille, and Xiang Bai. In defense of online models for video instance segmentation. *arXiv preprint arXiv:2207.10661*, 2022. 4
- [50] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Peize Sun, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8392–8401, 2021. 3

- [51] Xizhe Xue, Dongdong Yu, Lingqiao Liu, Yu Liu, Satoshi Tsutsui, Ying Li, Zehuan Yuan, Ping Song, and Mike Zheng Shou. Single-stage open-world instance segmentation with cross-task consistency regularization. 2022. [3](#)
- [52] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15325–15336, 2023. [3](#)
- [53] Shuo Yang, Peize Sun, Yi Jiang, Xiaobo Xia, Ruiheng Zhang, Zehuan Yuan, Changhu Wang, Ping Luo, and Min Xu. Objects in semantic topology. *arXiv preprint arXiv:2110.02687*, 2021. [3](#)
- [54] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. [3](#)
- [55] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)