

CAT: LoCalization and IdentifiAtion Cascade Detection Transformer for Open-World Object Detection

Shuailei Ma^{1*} Yuefeng Wang^{1*} Ying Wei^{1 2†} Jiaqi Fan¹
Thomas H. Li³ Hongli Liu⁴ Fanbing Lv⁴

¹Northeast University, Shenyang, China ²Information Technology R&D Innovation Center of Peking University

³School of Electronic and Computer Engineering, Peking University Shenzhen Graduate School, Shenzhen, China

⁴Changsha Hisense Intelligent System Research Institute Co., Ltd

Abstract

Open-world object detection (OWOD), as a more general and challenging goal, requires the model trained from data on known objects to detect both known and unknown objects and incrementally learn to identify these unknown objects. The existing works which employ standard detection framework and fixed pseudo-labelling mechanism (PLM) have the following problems: (i) The inclusion of detecting unknown objects substantially reduces the model's ability to detect known ones. (ii) The PLM does not adequately utilize the priori knowledge of inputs. (iii) The fixed selection manner of PLM cannot guarantee that the model is trained in the right direction. We observe that humans subconsciously prefer to focus on all foreground objects and then identify each one in detail, rather than localize and identify a single object simultaneously, for alleviating the confusion. This motivates us to propose a novel solution called CAT: LoCalization and IdentifiAtion Cascade Detection Transformer which decouples the detection process via **the shared decoder in the cascade decoding way**. In the meanwhile, we propose the **self-adaptive pseudo-labelling mechanism** which combines the model-driven with input-driven PLM and self-adaptively generates robust pseudo-labels for unknown objects, significantly improving the ability of CAT to retrieve unknown objects. Experiments on two benchmarks, i.e., MS-COCO and PASCAL VOC, show that our model outperforms the state-of-the-art methods. The code is publicly available at <https://github.com/xiaomabuwei/CAT>.

1. Introduction

Open-world object detection (OWOD) is a more practical detection problem in computer vision, making arti-

*Equal contribution.

†Corresponding author.



Figure 1. When faced with new scenes in open world, humans subconsciously focus on all foreground objects and then identify them in detail in order to alleviate the confusion between the known and unknown objects and get a clear view. Motivated by this, our CAT utilizes the shared decoder to decouple the localization and identification process in the cascade decoding way, where the former decoding process is used for localization and the latter for identification.

cial intelligence (AI) smarter to face more difficulties in real scenes. Within the OWOD paradigm, the model's life-span is pushed by iterative learning process. At each episode, the model trained only by known objects needs to detect known objects while simultaneously localizing unknown objects and identifying them into the unknown class. Human annotators then label a few of these tagged unknown classes of interest gradually. The model given these newly-added annotations will continue to incrementally update its knowledge without retraining from scratch.

Recently, Joseph *et al.* [21] proposed an open-world object detector, ORE, based on the two-stage Faster R-CNN [38] pipeline. ORE utilized an auto-labelling step to obtain pseudo-unknowns for training model to detect unknown objects and learned an energy-based binary classifier to distinguish the unknown class from known classes. However, its success largely relied on a held-out validation set which

was leveraged to estimate the distribution of unknown objects in the energy-based classifier. Then, several methods [29, 43–45] attempted to extend ORE and achieved some success. To alleviate the problems in ORE, Gupta *et al.* [17] proposed to use the detection transformer [4, 46] for OWOD in a justifiable way and directly leveraged the framework of DDETR [46]. In addition, they proposed an attention-driven PLM which selected pseudo labels for unknown objects according to the attention scores.

For the existing works, we find the following hindering problems. (i) Owing to the inclusion of detecting unknown objects, the model’s ability to detect known objects substantially drops. To alleviate the confusion between known and unknown objects, humans prefer to dismantle the process of open-world object detection rather than parallelly localize and identify open-world objects like most standard detection models. (ii) To the best of our knowledge, in the existing OWOD PLM, models leverage the learning process for known objects to guide the generation of pseudo labels for unknown objects, without leveraging the prior conditions of the inputs (*texture, light flow, etc*). As a result, the model cannot learn knowledge beyond the data annotation. (iii) The fixed selection manner of PLM cannot guarantee that the model learns to detect unknown objects in the right direction, due to the uncertain quality of the pseudo labels. The models may be worse for detecting unknown objects.

When faced with a new scene, humans prefer focusing on all foreground objects and then analysing them in detail [6], as shown in Figure.1. Motivated by this and the aforementioned observations, we propose a novel Localization and Identification Cascade Detection Transformer. CAT comprises three dedicated components namely, **shared transformer decoder**, **cascade decoupled decoding manner** and **self-adaptive pseudo-labelling mechanism**. Via the cascade decoupled decoding manner, the shared transformer decoder decouples the localization and identification process. Therefore, the influence of the category information of the identification process on the localization process is reduced. In this case, the model can localize more foreground objects so that the model’s ability to retrieve unknown objects is improved. Meanwhile, the independent recognition process allows the model to identify with more focus, so that the influence of unknown on detecting known objects is alleviated. In this decoding way, the former decoding process is used for localization and the latter for identification. The self-adaptive PLM maintains the ability of CAT to explore the knowledge beyond the known objects and self-adaptively adjusts the pseudo-label generation according to the model training process. Our contributions can be summarized fourfold:

- We propose a novel localization and identification cascade detection transformer (CAT), which has excellent ability to retrieve unknown objects and alleviate the in-

fluence of detecting unknown objects on the detection of known ones.

- Inspired by the subconscious reactions when people face open scenes, we propose the cascade decoupled decoding way, which decouples the decoding procedure via the shared decoder.
- We introduce a novel pseudo-labelling mechanism that self-adaptively combines the model-driven and input-driven pseudo-labelling during the training process for generating robust pseudo-labels and exploring knowledge beyond known objects.
- Our extensive experiments on two popular benchmarks demonstrate the effectiveness of the proposed CAT. CAT outperforms the state-of-the-art methods for OWOD, IOD, and open-set detection. For OWOD, CAT achieves absolute gains ranging from 9.7% to 12.8% in terms of unknown recall over the SOTA method.

2. Problem Formulation

At time t , let $\mathcal{K}^t = \{1, 2, \dots, C\}$ denote the set of known object classes and $\mathcal{U}^t = \{C + 1, \dots\}$ denote the unknown classes which might be encountered at the test time. The known object categories \mathcal{K}^t are labeled in the dataset $\mathcal{D}^t = \{\mathcal{J}^t, \mathcal{L}^t\}$ where \mathcal{J}^t denotes the input images and \mathcal{L}^t denotes the corresponding labels at time t . The training image set consists of M images $\mathcal{J}^t = \{i_1, i_2, \dots, i_M\}$ and corresponding labels $\mathcal{L}^t = \{\ell_1, \ell_2, \dots, \ell_M\}$. Each $\ell_i = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N\}$ denotes a set of N object instances with their class labels $c_n \in \mathcal{K}^t$ and locations, x_n, y_n, w_n, h_n denote the bounding box center coordinates, width and height respectively. The Open-World Object Detection removes the artificial assumptions and restrictions in traditional object detection and makes object detection tasks more aligned with real life. It requires the trained model \mathcal{M}_t not only to detect the previously encountered known classes C but also to identify an unseen class instance as belonging to the unknown class. In addition, it requires the object detector to be capable of incremental update for new knowledge and this cycle continues over the detector’s lifespan. In incremental updating phase, the unknown instances identified by \mathcal{M}_t are annotated manually, and along with their corresponding training examples, update \mathcal{D}^t to \mathcal{D}^{t+1} and \mathcal{K}^t to $\mathcal{K}^{t+1} = \{1, 2, \dots, C, \dots, C + n\}$, the model adds the n new classes to known classes and updates itself to \mathcal{M}_{t+1} without retraining from scratch on the whole dataset \mathcal{D}^{t+1} .

3. Proposed method

This section elaborates the proposed CAT in details. In Sec.3.1, the overall architecture of CAT is described in detail. We propose to decouple the decoding process of the

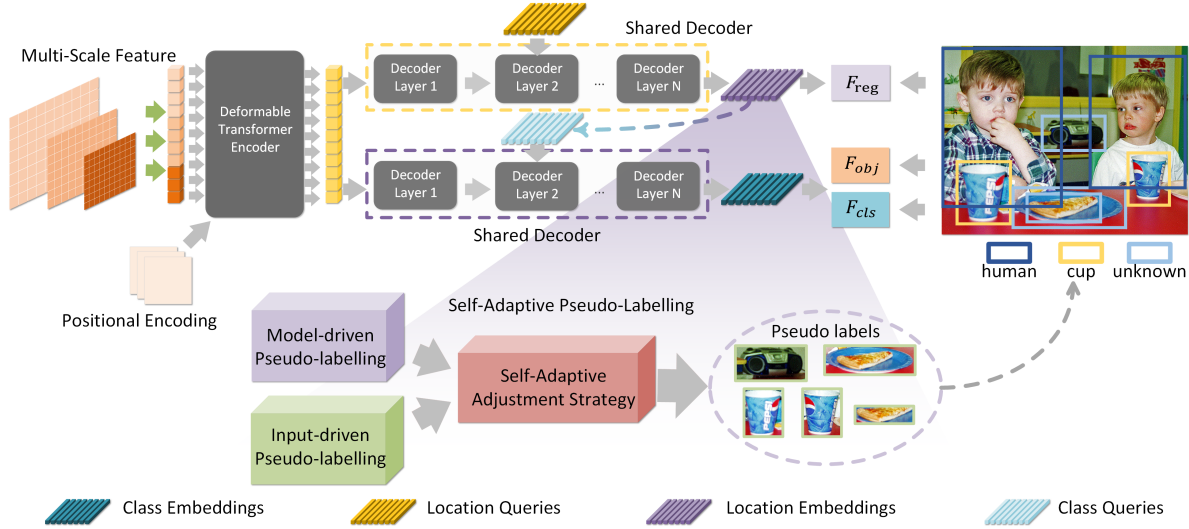


Figure 2. Overall Architecture of proposed CAT framework. The proposed CAT consists of a multi-scale feature extractor, the shared transformer decoder, the regression prediction branch, and the self-adaptive pseudo-labelling. The multi-scale feature extractor comprises the mainstream feature extraction backbone and a deformable transformer encoder, for extracting multi-scale features. The shared transformer decoder is a deformable transformer decoder and decouples the localization and identification process in the cascade decoding way. The regression prediction branch contains the bounding box regression branch F_{reg} , novelty objectness branch F_{obj} , and novelty classification branch F_{cls} . While the novelty classification and objectness branches are single-layer feed-forward networks (FFN) and the regression branch is a 3-layer FFN.

detection transformer and propose the localization and identification cascade decoupled decoding manner in Sec.3.2. A novel self-adaptive adjustment strategy for pseudo-labelling is proposed in Sec.3.3. In Sec.3.4, we illustrate the end-to-end training strategy of CAT.

3.1. Overall Architecture

As shown in Figure.2, for a given image $\mathcal{J} \in \mathbb{R}^{H \times W \times 3}$, CAT uses a hierarchical feature extraction backbone to extract multi-scale features $Z_i \in \mathbb{R}^{\frac{H}{4 \times i^2} \times \frac{W}{4 \times i^2} \times 2^i C_s}$, $i = 1, 2, 3$. The feature maps Z_i are projected from dimension C_s to dimension C_d by using 1×1 convolution and concatenated to N_s vectors with C_d dimensions after flattening out. Afterwards, along with supplement positional encoding $P_n \in \mathbb{R}^{N_s \times C_d}$, the multi-scale features are sent into the deformable transformer encoder to encode semantic features. The encoded semantic features $M \in \mathbb{R}^{N_s \times c_d}$ are acquired and sent into the shared decoder together with a set of N learnable location queries. Aided by interleaved cross-attention and self-attention modules, the shared decoder transforms the location queries $Q_{\text{location}} \in \mathbb{R}^{N \times D}$ to a set of N location query embeddings $\mathcal{E}_{\text{location}} \in \mathbb{R}^{N \times D}$. The $\mathcal{E}_{\text{location}}$ are then input to the regression branch to locate N foreground bounding boxes containing the known classes and unknown classes. Meanwhile, the $\mathcal{E}_{\text{location}}$ are used as class queries and sent into the shared decoder together with the M again. The shared decoder transforms the class queries to N class query embeddings $\mathcal{E}_{\text{class}}$ that

are corresponding to the location query embeddings. The $\mathcal{E}_{\text{class}}$ are then sent into the objectness and novelty classification branch to predict the objectness and category respectively. After selecting the unique queries that best match the known instances by a bipartite matching loss, the remaining queries are utilized to select the unknown category instances and generate pseudo labels by self-adaptive pseudo-labelling mechanism.

3.2. The Cascade Decoupled Decoding Way

Detection transformer [3, 4, 11, 25, 32, 46] leverages the object queries to detect object instances, where each object query represents an object instance. In the decoding stage, the object queries are updated to query embeddings by connecting object queries with semantic information from the encoded semantic features. The generated query embeddings couple the location and category information for both object localization and identification process simultaneously. For open-world object detection, the model requires detecting the known objects, localizing the unknown objects, and identifying them as the unknown class.

Inspired by how people react to new scenarios [6], a cascade decoupled decoding manner is proposed to decode the encoded features in a cascade way. We leverage the shared decoder to decode the encoded features twice. The first decoded embeddings are utilized to localize the foreground objects, while the second decoded embeddings are leveraged to identify the object categories and “unknown”. The

operation of localization and identification cascade decoding structure is expressed as follows:

$$\mathcal{E}_{\text{Location}} = \mathcal{F}_s(\mathcal{F}_e(\mathcal{O}(\mathcal{J}), P_n), \mathcal{Q}_{\text{Location}}, \mathcal{R}), \quad (1)$$

$$\mathcal{E}_{\text{Class}} = \mathcal{F}_s(\mathcal{F}_e(\mathcal{O}(\mathcal{J}), P_n), \mathcal{E}_{\text{Location}}, \mathcal{R}). \quad (2)$$

where $\mathcal{F}_s(\cdot)$ denotes the shared decoder. $\mathcal{F}_e(\cdot)$ is the encoder and $\mathcal{O}(\cdot)$ is the backbone. P_n stands for the positional encoding. \mathcal{R} represents the reference points and \mathcal{J} denotes the input image. In the cascade decoupled decoding phase, the location embeddings are used as class queries to generate class embeddings. Therefore, the localization process is not restricted by the category information, and the identification process can get help from the location knowledge in the cascade structure.

3.3. Self-Adaptive Pseudo-labelling Mechanism

Pseudo labels play an important role in guiding models to detect unknown object instances, determining the upper learning limitation of the model. The existing methods [17, 21] only use model-driven pseudo-labelling and do not take full advantage of the inputs' priori knowledge (light flow, textures, etc). The model-driven pseudo-labelling [17] makes the model's learning get caught up in the knowledge of known objects, for the reason that the only source of knowledge for the model is known object instances. In addition, their fixed selection manner cannot guarantee the right learning direction for unknown objects. We propose to combine model-driven with input-driven pseudo-labelling [36, 41, 47] for expanding the knowledge sources of the model. In the meanwhile, the pseudo-labels selection scheme should not be fixed, but be adapted as training and able to adjust itself when facing unexpected problems.

In this paper, inspired by [1], a novel pseudo-labelling mechanism is proposed for self-adaptively combining model-driven and input-driven pseudo-labelling according to the situation faced by the model, where the attention-driven pseudo-labelling [17] is used as the model-driven pseudo-labelling and selective search [41] is selected as the input-driven pseudo-labelling. In the self-adaptive pseudo-labelling mechanism, the model-driven pseudo-labelling generates pseudo-labels' candidate boxes P^m and the corresponding confidence s_o , and the input-driven pseudo-labelling generates pseudo-label candidate boxes P^I . The object confidence of generated pseudo labels is formulated as follows:

$$S_i = (\text{norm}(s_o))^{W_m} \cdot \left(\max_{1 \leq j \leq |P^I|} \left(\text{IOU}(P_j^I, P_i^m) \right) \right)^{W_I}, \quad (3)$$

where $\text{IOU}(\cdot)$ (Intersection-over-union [39]) is the most commonly used metric for comparing the similarity between two arbitrary shapes, i denotes the index of the

Algorithm 1 COMPUTINGADAPTIVEWEIGHTS

Input: Loss Memory: L_m ; Current Iteration: t ; Positive Momentum Amplitude: π_{pma} ; Negative Momentum Amplitude: π_{nma} ; T_{start} : Start iteration; T_b : Weight updating cycle; Loss ← Compute using Equation.9

Output: self-adaptive weights W_m^t and W_I^t

```

1: while train do
2:   if  $t \leq T_{start}$  then
3:     Initialise  $W_m^0 \leftarrow 0.8$  and  $W_I^0 \leftarrow 0.2$ 
4:     Initialise  $L_m$  using Equation.5
5:   else
6:     Update  $L_m$  using Equation.5
7:     if  $t \% T_b == 0$  then
8:       Compute  $\Delta l$  using  $L_m$  and Equation.6
9:       Compute  $\Delta w$  using  $\Delta l$  and Equation.7
10:      Update  $W_m^t$  and  $W_I^t$  using Equation.8
11:    end if
12:  end if
13: end while

```

pseudo labels. W_m and W_I are the self-adaptive weights, which are controlled by the *Measurer*, *Sensor* and *Adjuster*, as formulated below:

$$W^t = \text{Adjuster}(W^{t-1}, \text{Sensor}(\text{Measurer}(L_m))), \quad (4)$$

where L_m represents the loss memory which is stored and updated in real time during model training. The formulation is illustrated in Equation.5:

$$L_m = \text{DEQUE}(loss_{t-1}, loss_{t-2}, \dots, loss_{t-n}), \quad (5)$$

where DEQUE is the sequence function, and t is the current iteration. Considering the sensitivity of the model and the uneven quality of the data, we leverage *Measurer* to obtain the trend of the losses Δl for replacing the single loss. The formula is as follows:

$$\text{Measurer}(L_m) = \frac{\sum_{i=1}^n \alpha_i \cdot loss_{t-i}}{\sum_{j=n+1}^N \beta_j \cdot loss_{t-j}}, \quad n < N < T, \quad (6)$$

where α and β denote the weighted average weights and they are the decreasing series of equal differences (i.e. $\sum_{i=1}^n \alpha_i = \sum_{j=n+1}^N \beta_j = \frac{\alpha_i - \alpha_{i-1}}{\alpha_{i+1} - \alpha_i} = \frac{\beta_j - \beta_{j-1}}{\beta_{j+1} - \beta_j} = 1$). In the *Sensor*, the variable of the weight Δw is acquired as follows:

$$\text{Sensor}(\Delta l) = \begin{cases} \pi_{nma} \cdot \text{Sigmoid}(\Delta l - 1), & \Delta l > 1, \\ -\pi_{pma} \cdot \Delta l, & \Delta l \leq 1, \end{cases} \quad (7)$$

where π_{pma} and π_{nma} represents the positive and negative momentum amplitude (i.e. the amplitude of incremental changing), respectively. In the *Adjuster*, we use Equation.8 to update the self-adaptive weight via a incremental

way [9, 18, 21], for memory storage and enhancing the robustness.

$$\begin{cases} \mathcal{W}_m^t = \mathcal{W}_m^{t-1} + \Delta w \times \mathcal{W}_m^{t-1}, \\ \mathcal{W}_I^t = \mathcal{W}_I^{t-1} - \Delta w \times \mathcal{W}_I^{t-1}, \\ \mathcal{W}_m^t, \mathcal{W}_I^t = \text{norm}(\mathcal{W}_m^t, \mathcal{W}_I^t), \end{cases} \quad (8)$$

where $\text{norm}(\cdot)$ is the normalization operation. The update strategy for the weights during training is shown in Algorithm.1.

3.4. Training and Inference

Our CAT is trained end-to-end using the following joint loss formulation:

$$L = L_{\text{localization}} + L_{\text{identification}} + L_{\text{objectness}}, \quad (9)$$

where $L_{\text{localization}}$, $L_{\text{identification}}$ and $L_{\text{objectness}}$ denotes the loss terms for foreground localization, novelty identification and object scoring, respectively. When a set of new categories are introduced at each episode, we employ an exemplar replay based finetuning to alleviate catastrophic forgetting of learned classes and then finetune the model using a balanced set of exemplars stored for each known class. The bounding boxes and categories predictions of the known and top-k unknown objects are simultaneous used during evaluation.

4. Experiments

4.1. Datasets and Metrics

The experiments are implemented on two mainstream splits of MS-COCO [27] and Pascal VOC [14] dataset. We group the classes into a set of nonoverlapping tasks $\{T^1, \dots, T^t, \dots\}$. The class in task T^c only appears in tasks where $t \geq c$. In task T^c , classes encountered in $\{T^c : c \leq t\}$ and $\{T^c : c > t\}$ are considered as known and unknown classes, respectively.

OWOD SPLIT [21] splits the 80 classes of MS-COCO into 4 tasks and selects training set for each task from the MS-COCO and Pascal VOC training set images. Pascal VOC testing and MS-COCO validation set are used for evaluation.

MS-COCO SPLIT [17] mitigates data leakage across tasks in [21] and is more challenging. The training and testing data are selected from MS-COCO.

Metrics: Following the most commonly used evaluation metric for object detection, we use mean average precision (mAP) to evaluate the known objects. Inspired by [2, 13, 17, 21, 30], U-Recall is used as main metric for unknown objects. U-Recall measures the ability of the model to retrieve unknown object instances for OWOD problem.

4.2. Implementation Details

The multi-scale feature extractor consists of a Resnet-50 [20] pretrained on ImageNet [12] in a self-supervised [5] manner and a deformable transformer encoder whose number of layer is set to 6. For the shared decoder, we use a deformable transformer decoder and the number of layer is set to 6, too. We set the number of queries $M = 100$, the dimension of the embeddings $D = 256$ and the number of pseudo-labels $k = 5$. During inference, top-50 high scoring detections are used for evaluation for per image.

4.3. Comparison With State-of-the-art Methods

For a fair comparison, we compare CAT with ORE [21] without the energy-based unknown identifier (EBUI) that relies on held-out validation data with weak unknown object supervision and other SOTA methods [17, 43–45] to demonstrate the effectiveness of our method for OWOD problem. We present the comparison in terms of known class mAP and unknown class recall where U-Recall cannot be computed in Task 4 due to the absence of unknown test annotations, for the reason that all 80 classes are known.

OWOD SPLIT: The results compared with the state-of-the-art methods on OWOD split for OWOD problem are shown in Table.1. Benefiting from the cascade decoupled decoding manner and the self-adaptive pseudo-labelling mechanism, the ability of CAT to detect unknown objects goes substantially beyond the existing models. Compared with 2B-OCD’s [43] U-Recall of 12.1, 9.4 and 11.6 on Task 1, 2 and 3, our CAT achieves 23.7, 19.1 and 24.4 in the corresponding tasks, achieving significant absolute gains up to 12.8%. The ability to detect known objects and alleviate catastrophic forgetting of previous knowledge gains an improved performance with significant gains, achieving significant absolute gains up to 4.7% beyond OW-DETR [17]. This demonstrates the significant performance of the cascade decoding manner.

MS-COCO SPLIT: We report the results on MS-COCO split in Table.2. MS-COCO split mitigates data leakage across tasks and assign more data to each Task, while CAT receives a more significant boost compared with OWOD split. Compared with OW-DETR’s U-Recall of 5.7, 6.2 and 6.9 on Task 1, 2 and 3, our CAT achieves 24.0, 23.0 and 24.6 in the corresponding tasks, achieving significant absolute gains up to 18.3%. Furthermore, the performance on detecting known objects achieves significant absolute gains up to 9.7%. This demonstrates that our CAT has the more powerful ability to retrieve new knowledge and detect the known objects when faced with more difficult tasks.

Qualitative Results: We report qualitative results in Figure.3. We show the detection results of CAT (top row) and OW-DETR (bottom row), with **Blue** - known objects and **Yellow** - unknown objects. It is easy to see that CAT could detect more unknown objects. In the left column, OW-

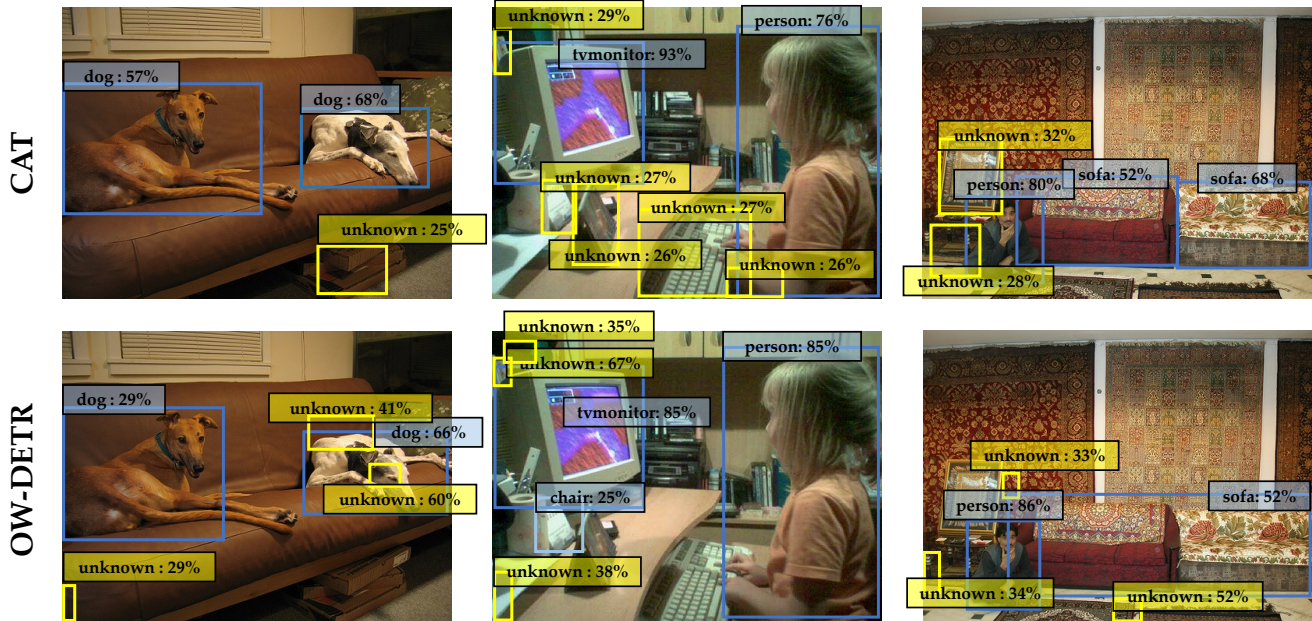


Figure 3. **Comparison of qualitative results between CAT and OW-DETR.** Detections of CAT (top row) and OW-DETR (bottom row) are displayed, with **Blue** - known and **Yellow** - unknown object detections. CAT detected more unknown objects than OW-DETR. In the left column, OW-DETR identifies the background and known objects as unknowns and the real unknown object (*carton*) as the background, and our model accurately identifies the *carton* as the unknown object. In the middle column, OW-DETR identifies the two *calendars* as the *chair* and the background, respectively, and the *keyboard* as the background, and our CAT accurately identifies them as unknown objects. The right column shows that OW-DETR not only does not detect the unknown object (*frame*) but also identifies two known objects (*sofa*) as one. Our model accurately identifies the frame as an unknown object and also accurately identifies the two *sofas*.

Table 1. **State-of-the-art comparison on OWOD split.** The comparison is shown in terms of U-Recall and known class mAP. U-Recall measures the ability of the model to retrieve unknown object instances for OWOD problem. For a fair comparison, we compare with the recently introduced methods and ORE not employing EBUI. The CAT achieves improved all metrics over the existing works across all tasks, demonstrating our model’s effectiveness for OWOD problem. U-Recall cannot be computed in Task 4 due to the absence of unknown test annotations, for the reason that all 80 classes are known.

Task IDs →	Task 1		Task 2			Task 3			Task 4				
Metrics →	Unknown Recall (↑)	Known mAP(↑) Current	Unknown Recall (↑)	Known mAP(↑)			Unknown Recall (↑)	Known mAP(↑)			Known mAP(↑)		
				Previously	Current	Both		Previously	Current	Both	Previously	Current	Both
UC-OWOD [44]	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
ORE-EBUI [21]	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
OW-DETR [17]	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
OCPL [45]	8.3	56.6	7.7	50.6	27.5	39.1	11.9	38.7	14.7	30.7	30.7	14.4	26.7
2B-OCD [43]	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
Ours: CAT	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

DETR identifies the background and known objects (*dog*) as unknowns and the real unknown object (*carton*) as the background, while our model accurately identifies the *carton* as the unknown object. As shown in the middle column, OW-DETR identifies the two *calendars* as the *chair* and the background, respectively, and the *keyboard* as the background, while our CAT accurately identifies them as unknown objects. The right column shows that OW-DETR

not only does not detect the unknown object (*frame*) but also identifies two known objects (*sofa*) as one, while our model accurately identifies the frame as an unknown object and accurately identifies the two *sofas*.

4.4. Ablation Study

We conduct abundant ablative experiments to verify the effectiveness of CAT’s components on the OWOD split.

Table 2. **State-of-the-art comparison on MS-COCO split.** The comparison is shown in terms of U-Recall and mAP. Although the MS-COCO split is more challenging, our model gets a more significant improvement on this in comparison to ORE and OW-DETR. The significant metric improvements demonstrate that our CAT has the ability to retrieve new knowledge beyond the range of closed set and would not be limited by category knowledge of existing objects.

Task IDs →	Task 1		Task 2			Task 3			Task 4				
Metrics →	Unknown	Known	Unknown	Known		Unknown	Known		Known				
	Recall	mAP(↑)	Recall	mAP(↑)		Recall	mAP(↑)		mAP(↑)				
	(↑)	Current	(↑)	Previously	Current	Both	(↑)	Previously	Current	Both	Previously	Current	Both
ORE-EBUI [21]	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 [17]	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
Ours: CAT	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

Table 3. **Component ablation experiment.** The comparison is shown in terms of known class average precision (mAP) and unknown class recall (U-Recall). **CAT-Cddw** is our model without the cascade decoupled decoding way. **CAT-Sam** is our model without the self-adaptive manner but with the prior from the selective search. We also include the performance of deformable DETR and an upper bound (D-DETR trained with ground-truth unknown class annotations) as reported by OW-DETR [17].

Task IDs →	Task 1		Task 2			Task 3			Task 4				
Metrics →	Unknown	Known	Unknown	Known		Unknown	Known		Known				
	Recall	mAP(↑)	Recall	mAP(↑)		Recall	mAP(↑)		mAP(↑)				
	(↑)	Current	(↑)	Previously	Current	Both	(↑)	Previously	Current	Both	Previously	Current	Both
Upper Bound	31.6	62.5	40.5	55.8	38.1	46.9	42.6	42.4	29.3	33.9	35.6	23.1	32.5
D-DETR [46]	-	60.3	-	54.5	34.4	44.7	-	40.0	17.7	33.3	32.5	20.0	29.4
CAT-Cddw	19.1	59.3	18.6	52.8	30.2	41.5	21.0	41.0	17.6	33.0	32.6	15.8	27.9
CAT-Sam	19.1	59.7	16.9	54.8	32.4	43.6	18.6	42.1	19.3	34.5	34.0	16.0	29.5
Final:CAT	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

Furthermore, we demonstrate the effectiveness of our model for incremental object detection and open-set detection.

Ablating Components: To study the contribution of each component, we design ablation experiments in Table 3. In comparison to the Final CAT, removing the cascade decoupled decoding manner **CAT-Cddw** reduces the performance on retrieving unknown objects and detecting known objects, achieving significant absolute gains down to 4.6, 0.5 and 3.4 points in Task 1,2,3 for U-Recall and the mAP for known objects is reduced by 0.7, 2.6, 1.8, and 2.0 in Task 1,2,3,4. The results demonstrate that the cascade decoupled decoding manner is better for the open-world object detection which contains the unknown objects and improves the ability of CAT to retrieve unknown objects and detect known objects. To ablate the self-adaptive manner component, we remove the self-adaptive manner from CAT and hold the prior from selective search. Compared with CAT, removing the self-adaptive manner **CAT-Sam** significantly reduces the performance on detecting unknown objects, achieving significant absolute gains down to 4.6, 2.2, and 5.8 points in Task 1,2,3 respectively. The results demonstrate that the self-adaptive manner could efficiently combine the input and model-driven pseudo-labelling mechanism, improving the CAT’s ability to explore unknown ob-

jects. Thus, each component has a critical role to play in open-World object detection.

Incremental Object Detection: To intuitively present our CAT’s ability for detecting object instances, we compare it to [17, 21, 35, 40] on the incremental object detection (IOD) task. We evaluate the experiments on three standard settings, where a group of classes (10, 5 and last class) are introduced incrementally to a detector trained on the remaining classes (10, 15 and 19), based on PASCAL VOC 2007 dataset [14]. As the results shown in Table 4, CAT outperforms the existing method in a great migration on all three settings, indicating the power of localization and identification cascade detection transformer for IOD.

Open-set Detection Comparison: To further demonstrate CAT’s ability to handle unknown instances in open-set data, we follow the same evaluation protocol as [17, 21, 31] and report the performance in Table 5. CAT achieves promising performance in comparison to the existing methods.

5. Relation to Prior Works

The issue of standard object detection [4, 10, 16, 19, 26, 28, 34, 37, 38, 46, 48] has been raised for several years, numerous works have investigated this problem and push the

Table 4. **Performance comparison on incremental object detection task.** Evaluation is performed on three standard settings, where a group of classes (10, 5 and last class) are introduced incrementally to a detector trained on the remaining classes (10,15 and 19). Our CAT performs favorably against existing approaches on all three settings, illustrating the power of localization identification cascade detection transformer for incremental objection detection.

10 + 10 setting	aero	cycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	bike	person	plant	sheep	sofa	train	tv	mAP
ILOD [40]	69.9	70.4	69.4	54.3	48	68.7	78.9	68.4	45.5	58.1	59.7	72.7	73.5	73.2	66.3	29.5	63.4	61.6	69.3	62.2	63.2
Faster ILOD [35]	72.8	75.7	71.2	60.5	61.7	70.4	83.3	76.6	53.1	72.3	36.7	70.9	66.8	67.6	66.1	24.7	63.1	48.1	57.1	43.6	62.1
ORE-(CC+EBUI) [21]	53.3	69.2	62.4	51.8	52.9	73.6	83.7	71.7	42.8	66.8	46.8	59.9	65.5	66.1	68.6	29.8	55.1	51.6	65.3	51.5	59.4
ORE-EBUI [21]	63.5	70.9	58.9	42.9	34.1	76.2	80.7	76.3	34.1	66.1	56.1	70.4	80.2	72.3	81.8	42.7	71.6	68.1	77	67.7	64.5
OW-DETR [17]	61.8	69.1	67.8	45.8	47.3	78.3	78.4	78.6	36.2	71.5	57.5	75.3	76.2	77.4	79.5	40.1	66.8	66.3	75.6	64.1	65.7
Ours: CAT	76.5	75.7	67.0	51.0	62.4	73.2	82.3	83.7	42.7	64.4	56.8	74.1	75.8	79.2	78.1	39.9	65.1	59.6	78.4	67.4	67.7
15 + 5 setting	aero	cycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	bike	person	plant	sheep	sofa	train	tv	mAP
ILOD [40]	70.5	79.2	68.8	59.1	53.2	75.4	79.4	78.8	46.6	59.4	59	75.8	71.8	78.6	69.6	33.7	61.5	63.1	71.7	62.2	65.8
Faster ILOD [35]	66.5	78.1	71.8	54.6	61.4	68.4	82.6	82.7	52.1	74.3	63.1	78.6	80.5	78.4	80.4	36.7	61.7	59.3	67.9	59.1	67.9
ORE-(CC+EBUI) [21]	65.1	74.6	57.9	39.5	36.7	75.1	80	73.3	37.1	69.8	48.8	69	77.5	72.8	76.5	34.4	62.6	56.5	80.3	65.7	62.6
ORE-EBUI [21]	75.4	81	67.1	51.9	55.7	77.2	85.6	81.7	46.1	76.2	55.4	76.7	86.2	78.5	82.1	32.8	63.6	54.7	77.7	64.6	68.5
OW-DETR [17]	77.1	76.5	69.2	51.3	61.3	79.8	84.2	81.0	49.7	79.6	58.1	79.0	83.1	67.8	85.4	33.2	65.1	62.0	73.9	65.0	69.4
Ours: CAT	75.3	81.0	84.4	64.5	56.6	74.4	84.1	86.6	53.0	70.1	72.4	83.4	85.5	81.6	81.0	32.0	58.6	60.7	81.6	63.5	72.2
19 + 1 setting	aero	cycle	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	bike	person	plant	sheep	sofa	train	tv	mAP
ILOD [40]	69.4	79.3	69.5	57.4	45.4	78.4	79.1	80.5	45.7	76.3	64.8	77.2	80.8	77.5	70.1	42.3	67.5	64.4	76.7	62.7	68.2
Faster ILOD [35]	64.2	74.7	73.2	55.5	53.7	70.8	82.9	82.6	51.6	79.7	58.7	78.8	81.8	75.3	77.4	43.1	73.8	61.7	69.8	61.1	68.5
ORE-(CC+EBUI) [21]	60.7	78.6	61.8	45	43.2	75.1	82.5	75.5	42.4	75.1	56.7	72.9	80.8	75.4	77.7	37.8	72.3	64.5	70.7	49.9	64.9
ORE-EBUI [21]	67.3	76.8	60	48.4	58.8	81.1	86.5	75.8	41.5	79.6	54.6	72.8	85.9	81.7	82.4	44.8	75.8	68.2	75.7	60.1	68.8
OW-DETR [17]	70.5	77.2	73.8	54.0	55.6	79.0	80.8	80.6	43.2	80.4	53.5	77.5	89.5	82.0	74.7	43.3	71.9	66.6	79.4	62.0	70.2
Ours: CAT	86.0	85.8	78.8	65.3	61.3	71.4	84.8	84.8	52.9	78.4	71.6	82.7	83.8	81.2	80.7	43.7	75.9	58.5	85.2	61.1	73.8

field to certain heights. Whereas the strong assumption that the label space of object categories to be encountered during the life-span of the model is the same as during its training results that these methods cannot meet real-world needs. The success of [7, 15, 22–24, 33, 38] demonstrates the feasibility of foreground localization based on the position and appearance of objects. Existing works [8, 17, 21, 29, 43–45] attempt to leverage the framework of standard object detection models for open-world object detection. In this paper, we propose a novel transformer [42] based framework. CAT decouples the localization and identification process and connects them in a cascade approach. In CAT, the foreground localization process is not limited by the category of known objects, whereas the process of foreground identification can use information from the localization process. Along with self-adaptive pseudo-labelling, CAT can gain information beyond the data annotation and maintain a stable learning process according to self-regulation.

6. Conclusions

We analyze the drawbacks of the parallel decoding structure for open-world object detection. Motivated by the subconscious reactions of humans when facing new scenes, we propose a novel localization and identification cascade detection transformer (CAT), which decouples the localization and identification process via the cascade decoding manner. The cascade decoding manner alleviates the influence

Table 5. **Performance comparison on open-set object detection task.** Our CAT achieves significant performance in comparison to existing works.

Evaluated on →	VOC	WRI
Standard Faster R-CNN [40]	81.8	77.1
Standard RetinaNet	79.2	73.8
Dropout Sampling [31]	78.1	71.1
ORE [21]	81.3	78.2
OW-DETR [17]	82.1	78.6
Ours: CAT	83.2	79.5

of detecting unknown objects on the detection of known objects. With the self-adaptive pseudo-labelling mechanism, CAT gains knowledge beyond the data annotations, generates pseudo labels with robustness and maintains a stable training process via self-adjustment. The extensive experiments on two popular benchmarks, *i.e.*, PASCAL VOC and MS COCO demonstrate that CAT’s performance is better than the existing methods.

Acknowledgment

This work is supported by National Nature Science Foundation of China (grant No.61871106), and the Open Project Program Foundation of the Key Laboratory of Opto-Electronics Information Processing, Chinese Academy of Sciences (OEIP-O-202002).

References

- [1] Kiam Heong Ang, Gregory Chong, and Yun Li. Pid control system analysis, design, and technology. *IEEE transactions on control systems technology*, 13(4):559–576, 2005. 4
- [2] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chelappa, and Ajay Divakaran. Zero-shot object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 384–400, 2018. 5
- [3] Josh Beal, Eric Kim, Eric Tzeng, Dong Huk Park, Andrew Zhai, and Dmitry Kislyuk. Toward transformer-based object detection. *arXiv preprint arXiv:2012.09958*, 2020. 3
- [4] 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, 3, 7
- [5] 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. 5
- [6] Marisa Carrasco. Visual attention: The past 25 years. *Vision Research*, 51(13):1484–1525, 2011. Vision Research 50th Anniversary Issue: Part 2. 2, 3
- [7] Peihao Chen, Chuang Gan, Guangyao Shen, Wenbing Huang, Runhao Zeng, and Mingkui Tan. Relation attention for temporal action localization. *IEEE Transactions on Multimedia*, 2019. 8
- [8] Peihao Chen, Dongyu Ji, Kunyang Lin, Runhao Zeng, Thomas H Li, Mingkui Tan, and Chuang Gan. Weakly-supervised multi-granularity map learning for vision-and-language navigation. *NeurIPS*, 2022. 8
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 5
- [10] Xingyu Chen, Junzhi Yu, Shihan Kong, Zhengxing Wu, and Li Wen. Joint anchor-feature refinement for real-time accurate object detection in images and videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(2):594–607, 2020. 7
- [11] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2988–2997, 2021. 3
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [13] Akshay Dhamija, Manuel Gunther, Jonathan Ventura, and Terrance Boult. The overlooked elephant of object detection: Open set. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1021–1030, 2020. 5
- [14] 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. 5, 7
- [15] Spyros Gidaris and Nikos Komodakis. Attend refine repeat: Active box proposal generation via in-out localization. *arXiv preprint arXiv:1606.04446*, 2016. 8
- [16] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 7
- [17] Akshita Gupta, Sanath Narayan, KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Ow-detr: Open-world detection transformer. In *CVPR*, 2022. 2, 4, 5, 6, 7, 8
- [18] 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. 5
- [19] 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. 7
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [21] K J Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5826–5836, 2021. 1, 4, 5, 6, 7, 8
- [22] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8420–8429, 2019. 8
- [23] 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. 8
- [24] Hongyang Li, Yu Liu, Wanli Ouyang, and Xiaogang Wang. Zoom out-and-in network with map attention decision for region proposal and object detection. *International Journal of Computer Vision*, 127(3):225–238, 2019. 8
- [25] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. 3
- [26] 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. 7
- [27] 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. 5
- [28] Yue Lu, Xingyu Chen, Zhengxing Wu, and Junzhi Yu. Decoupled metric network for single-stage few-shot object detection. *IEEE Transactions on Cybernetics*, 2022. 7

- [29] Muhammad Maaz, Hanoona Bangalath Rasheed, Salman Hameed Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Ming-Hsuan Yang. Multi-modal transformers excel at class-agnostic object detection. *arXiv*, 2021. [2](#), [8](#)
- [30] Dimity Miller, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. Dropout sampling for robust object detection in open-set conditions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3243–3249. IEEE, 2018. [5](#)
- [31] Dimity Miller, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. Dropout sampling for robust object detection in open-set conditions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3243–3249. IEEE, 2018. [7](#), [8](#)
- [32] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2917, 2021. [3](#)
- [33] Pedro O O Pinheiro, Ronan Collobert, and Piotr Dollár. Learning to segment object candidates. *Advances in neural information processing systems*, 28, 2015. [8](#)
- [34] Yanwei Pang, Tiancai Wang, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Ling Shao. Efficient featureized image pyramid network for single shot detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7336–7344, 2019. [7](#)
- [35] Can Peng, Kun Zhao, and Brian C Lovell. Faster ilod: Incremental learning for object detectors based on faster rcnn. *Pattern recognition letters*, 140:109–115, 2020. [7](#), [8](#)
- [36] Jordi Pont-Tuset, Pablo Arbeláez, Jonathan T. Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(1):128–140, 2017. [4](#)
- [37] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. [7](#)
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [1](#), [7](#), [8](#)
- [39] Adrian Rosebrock. Intersection over union (iou) for object detection. *Diambil kembali dari PYImageSearch* <https://www.pyimagesearch.com/2016/11/07/intersection-over-union-iou-for-object-detection>, 2016. [4](#)
- [40] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, pages 3400–3409, 2017. [7](#), [8](#)
- [41] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. [4](#)
- [42] 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. [8](#)
- [43] Yan Wu, Xiaowei Zhao, Yuqing Ma, Duorui Wang, and Xi-anlong Liu. Two-branch objectness-centric open world detection. In *Proceedings of the 3rd International Workshop on Human-Centric Multimedia Analysis*, pages 35–40, 2022. [2](#), [5](#), [6](#), [8](#)
- [44] Zhiheng Wu, Yue Lu, Xingyu Chen, Zhengxing Wu, Liwen Kang, and Junzhi Yu. Uc-owod: Unknown-classified open world object detection. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, pages 193–210. Springer, 2022. [2](#), [5](#), [6](#), [8](#)
- [45] Jinan Yu, Liyan Ma, Zhenglin Li, Yan Peng, and Shaorong Xie. Open-world object detection via discriminative class prototype learning. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 626–630. IEEE, 2022. [2](#), [5](#), [6](#), [8](#)
- [46] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [2](#), [3](#), [7](#)
- [47] C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *European conference on computer vision*, pages 391–405. Springer, 2014. [4](#)
- [48] Zhengxia Zou, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey. *arXiv preprint arXiv:1905.05055*, 2019. [7](#)