

Open-Vocabulary DETR with Conditional Matching

Yuhang Zang¹, Wei Li¹, Kaiyang Zhou¹, Chen Huang², Chen Change Loy¹✉
¹S-Lab, Nanyang Technological University ²Carnegie Mellon University
{zang0012, wei.l, kaiyang.zhou, cclloy}@ntu.edu.sg chen-huang@apple.com

Abstract. Open-vocabulary object detection, which is concerned with the problem of detecting novel objects guided by natural language, has gained increasing attention from the community. Ideally, we would like to extend an open-vocabulary detector such that it can produce bounding box predictions based on user inputs in form of either natural language or exemplar image. This offers great flexibility and user experience for human-computer interaction. To this end, we propose a novel open-vocabulary detector based on DETR—hence the name OV-DETR—which, once trained, *can detect any object given its class name or an exemplar image*. The biggest challenge of turning DETR into an open-vocabulary detector is that it is impossible to calculate the classification cost matrix of novel classes without access to their labeled images. To overcome this challenge, we formulate the learning objective as a binary matching one between input queries (class name or exemplar image) and the corresponding objects, which learns useful correspondence to generalize to unseen queries during testing. For training, we choose to condition the Transformer decoder on the input embeddings obtained from a pre-trained vision-language model like CLIP, in order to enable matching for both text and image queries. With extensive experiments on LVIS and COCO datasets, we demonstrate that our OV-DETR—the *first end-to-end Transformer-based open-vocabulary detector*—achieves non-trivial improvements over current state of the arts. Code is available at <https://github.com/yuhangzang/OV-DETR>.

1 Introduction

Object detection, a fundamental computer vision task aiming to localize objects with tight bounding boxes in images, has been significantly advanced in the last decade thanks to the emergence of deep learning [25,32,9,29,15]. However, most object detection algorithms are unscalable in terms of the vocabulary size, *i.e.*, they are limited to a fixed set of object categories defined in detection datasets [8,22]. For example, an object detector trained on COCO [22] can only detect 80 classes and is unable to handle new classes beyond the training ones.

A straightforward approach to detecting novel classes is to collect and add their training images to the original dataset, and then re-train or fine-tune the detection model. This is, however, both impractical and inefficient due to the large cost of data collection and model training. In the detection literature, generalization from base to novel classes has been studied as a zero-shot detection

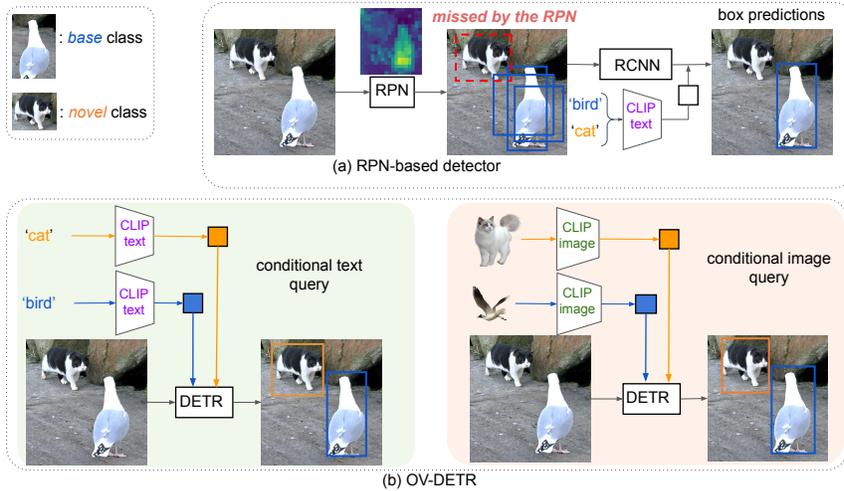


Fig. 1: **Comparison between a RPN-based detector and our Open-Vocabulary Transformer-based detector (OV-DETR) using conditional queries.** The RPN trained on closed-set object classes is easy to ignore novel classes (*e.g.*, the “cat” region receives little response). Hence the cats in this example are largely missed with few to no proposals. By contrast, our OV-DETR is trained to perform matching between a conditional query and its corresponding box, which helps to learn correspondence that can generalize to queries from unseen classes. Note we can take input queries in the form of either text (class name) or exemplar images, which offers greater flexibility for open-vocabulary object detection.

problem [1] where zero-shot learning techniques like word embedding projection [10] are widely used.

Recently, open-vocabulary detection, a new formulation that leverages large pre-trained language models, has gained increasing attention from the community [36,13]. The central idea in existing works is to align detector’s features with embedding provided by models pre-trained on large scale image-text pairs like CLIP [27] (see Fig 1 (a)). This way, we can use an aligned classifier to recognize novel classes only from their descriptive texts.

A major problem with existing open-vocabulary detectors [36,13] is that they rely on region proposals that are often not reliable to cover all novel classes in an image due to the lack of training data, see Fig. 1(a). This problem has also been identified by a recent study [17], which suggests the binary nature of the region proposal network (RPN) could easily lead to overfitting to seen classes (thus fail to generalize to novel classes).

In this paper, we propose to train end-to-end an open-vocabulary detector under the Transformer framework, aiming to enhance its novel class general-

ization without using an intermediate RPN. To this end, we propose a novel open-vocabulary detector based on DETR [2]—hence the name OV-DETR—which is trained to *detect any object given its class name or an exemplar image*. This would offer greater flexibility than conventional open-vocabulary detection from natural language only.

Despite the simplicity of end-to-end DETR training, turning it into an open-vocabulary detector is non-trivial. The biggest challenge is the inability to calculate the classification cost for novel classes without their training labels. To overcome the challenge, we re-formulate the learning objective as binary matching between input queries (class name or exemplar image) and the corresponding objects. Such a matching loss over diverse training pairs allows to learn useful correspondence that can generalize to unseen queries during testing. For training, we extend the Transformer decoder of DETR to take conditional input queries. Specifically, we condition the Transformer decoder on the query embeddings obtained from a pre-trained vision-language model CLIP [27], in order to perform conditional matching for either text or image queries. Fig. 1 shows this high-level idea, which proves better at detecting novel classes than RPN-based closed-set detectors.

We conduct comprehensive experiments on two challenging open-vocabulary object detection datasets, and show consistent improvements in performance. Concretely, our OV-DETR method achieves 17.4 mask mAP of novel classes on the open-vocabulary LVIS dataset [13] and 29.4 box mAP of novel classes on open-vocabulary COCO dataset [36], surpassing SOTA methods by 1.3 and 1.8 mAP, respectively.

2 Related Work

Open-Vocabulary Object Detection leverages the recent advances in large pre-trained language models [36,13] to incorporate the open-vocabulary information into object detectors. OVR-CNN [36] first uses BERT [6] to pre-train the Faster R-CNN detector [29] on image-caption pairs and then fine-tunes the model on downstream detection datasets. ViLD [13] adopts a distillation-based approach that aligns the image feature extractor of Mask R-CNN [15] with the image and text encoder of CLIP [27] so the CLIP can be used to synthesize the classification weights for any novel class. The prompt tuning techniques [39,38,37] for the pre-trained vision-language model have also been applied for open-vocabulary detectors, like DetPro [7]. Our approach differs from these works in that we train a Transformer-based detector end-to-end, with a novel framework of conditional matching.

Zero-Shot Object Detection is also concerned with the problem of detecting novel classes [1,20,40,28,31]. However, this setting is less practical due to the harsh constraint of limiting access to resources relevant to unseen classes [36]. A common approach to zero-shot detection is to employ word embeddings like GloVe [26] as the classifier weights [1]. Other works have found that using external resources like textual descriptions can help improve the generalization

of classifier embeddings [20,28]. Alternatively, Zhao *et al.* [31] used Generative Adversarial Network (GAN) [12] to generate feature representations of novel classes. While Zhu *et al.* [40] synthesized unseen classes using a data augmentation strategy.

Visual Grounding is another relevant research area where the problem is to ground a target object in one image using natural language input [5,3]. Different from open-vocabulary detection that aims to identify all target objects in an image, the visual grounding methods typically involve a particular single object, hence cannot be directly applied to generic object detection. There is a relevant visual grounding method though, which is called MDETR [16]. This method similarly trains DETR along with a given language model so as to link the output tokens of DETR with specific words. MDETR also adopts a conditional framework, where the visual and textual features are combined to be fed to the Transformer encoder and decoder. However, the MDETR method is not applicable to open-vocabulary detection because it is unable to calculate the cost matrix for novel classes under the classification framework. Our OV-DETR bypasses this challenge by using a conditional matching framework instead.

Object Detection with Transformers. The pioneer DETR approach [2] greatly simplifies the detection pipeline by casting detection as a set-to-set matching problem. Several follow-up methods have been developed to improve performance and training efficiency. Deformable DETR [41] features a deformable attention module, which samples sparse pixel locations for computing attention, and further mitigates the slow convergence issue with a multi-scale scheme. SMCA [11] accelerates training convergence with a location-aware co-attention mechanism. Conditional DETR [24] also addresses the slow convergence issue, but with conditional spatial queries learned from reference points and the decoder embeddings. Our work *for the first time* extends DETR to the open-vocabulary domain by casting open-vocabulary detection as a conditional matching problem, and achieves non-trivial improvements over current SOTA.

3 Open-Vocabulary DETR

Our goal is to design a simple yet effective open-vocabulary object detector that can detect objects described by arbitrary text inputs or exemplar images. We build on the success of DETR [2] that casts object detection as an end-to-end set matching problem (among closed classes), thus eliminating the need of hand-crafted components like anchor generation and non-maximum suppression. This pipeline makes it appealing to act as a suitable framework to build our end-to-end open-vocabulary object detector.

However, it is non-trivial to retrofit a standard DETR with closed-set matching to an open-vocabulary detector that requires matching against unseen classes. One intuitive approach for such open-set matching is to learn a class-agnostic module (*e.g.*, ViLD [13]) to handle all classes. This is, however, still unable to match for those open-vocabulary classes that come with no labeled images. Here we provide a new perspective on the matching task in DETR, which

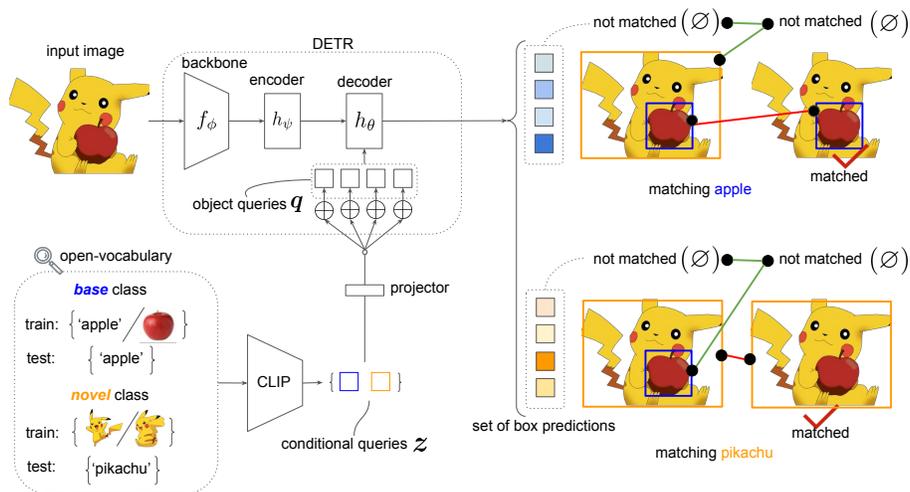


Fig. 2: **Overview of OV-DETR.** Unlike the standard DETR, our method does not separate ‘objects’ from ‘non-objects’ for a closed set of classes. Instead, OV-DETR performs open-vocabulary detection by measuring the matchability (‘matched’ vs. ‘not matched’) between some conditional inputs (text or exemplar image embeddings from CLIP) and detection results. We show such pipeline is flexible to detect open-vocabulary classes with arbitrary text or image inputs.

leads us to reformulate the fixed set-matching objective into a conditional binary matching one between conditional inputs (text or image queries) and detection outputs.

An overview of our Open-Vocabulary DETR is shown in Fig. 2. At high level, DETR first takes query embeddings (text or image) as conditional inputs obtained from a pre-trained CLIP [27] model, and then a binary matching loss is imposed against the detection result to measure their matchability. In the following, we will revisit the closed-set matching process in standard DETR in Section 3.1. We then describe how to perform conditional binary matching in our OV-DETR in Section 3.2.

3.1 Revisiting Closed-Set Matching in DETR

For input image \mathbf{x} , a standard DETR infers N object predictions $\hat{\mathbf{y}}$ where N is determined by the fixed size of object queries \mathbf{q} that serve as learnable positional encodings. One single pass of the DETR pipeline consists of two main steps: (i) set prediction, and (ii) optimal bipartite matching.

Set Prediction. Given an input image \mathbf{x} , the global context representations \mathbf{c} is first extracted by a CNN backbone f_ϕ and then a Transformer encoder h_ψ :

$$\mathbf{c} = h_\psi(f_\phi(\mathbf{x})), \quad (1)$$

where the output \mathbf{c} denotes a sequence of feature embeddings of \mathbf{q} . Taking the context feature \mathbf{c} and object queries \mathbf{q} as inputs, the Transformer decoder h_θ (with prediction heads) then produce the set prediction $\hat{\mathbf{y}} = \{\hat{\mathbf{y}}_i\}_{i=1}^N$:

$$\hat{\mathbf{y}} = h_\theta(\mathbf{c}, \mathbf{q}), \quad (2)$$

where $\hat{\mathbf{y}}$ contains both bounding box predictions $\hat{\mathbf{b}}$ and class predictions $\hat{\mathbf{p}}$ for a closed-set of training classes.

Optimal Bipartite Matching is to find the best match between the set of N predictions $\hat{\mathbf{y}}$ and the set of ground truth objects $\mathbf{y} = \{\mathbf{y}_i\}_{i=1}^M$ (including no object \emptyset). Specifically, one needs to search a permutation of N elements $\sigma \in \mathfrak{S}_N$ that has the lowest matching cost:

$$\hat{\sigma} = \operatorname{argmin}_{\sigma \in \mathfrak{S}_N} \sum_i^N \mathcal{L}_{cost}(\mathbf{y}_i, \hat{\mathbf{y}}_{\sigma(i)}), \quad (3)$$

where $\mathcal{L}_{cost}(\mathbf{y}_i, \hat{\mathbf{y}}_{\sigma(i)})$ is a pair-wise *matching cost* between ground truth \mathbf{y}_i and the prediction $\hat{\mathbf{y}}_{\sigma(i)}$ with index $\sigma(i)$. Note \mathcal{L}_{cost} is comprised of the losses for both class prediction $\mathcal{L}_{cls}(\hat{\mathbf{p}}, \mathbf{p})$ and bounding box localization $\mathcal{L}_{box}(\hat{\mathbf{b}}, \mathbf{b})$. The whole bipartite matching process produces *one-to-one* label assignments, where each prediction $\hat{\mathbf{y}}_i$ is assigned to a ground-truth annotation \mathbf{y}_j or \emptyset (no object). The optimal assignment can be efficiently found by the Hungarian algorithm [18].

Challenge. As mentioned above, the bipartite matching method cannot be directly applied to an open-vocabulary setting that contains both *base* and *novel* classes. The reason is that computing the matching cost in Eq. (3) requires access of the label information, which is unavailable for *novel* classes. We can follow previous works [13,35,7] to generate class-agnostic object proposals that may cover the *novel* classes, but we do not know the ground-truth classification labels of these proposals. As a result, the predictions for the N object queries cannot generalize to novel classes due to the lack of training labels for them. As shown in Fig. 3 (a), bipartite matching can only be performed for base classes with available training labels.

3.2 Conditional Matching for Open-Vocabulary Detection

To enable DETR to go beyond closed-set classification and perform open-vocabulary detection, we equip the Transformer decoder with conditional inputs and reformulate the learning objective as binary matching problem.

Conditional Inputs. Given an object detection dataset with standard annotations for all the training (*base*) classes, we need to convert those annotations to conditional inputs to facilitate our new training paradigm. Specifically, for each ground-truth annotation with bounding box \mathbf{b}_i and class label name $\mathbf{y}_i^{\text{class}}$, we use the CLIP model [27] to generate their corresponding image embedding $\mathbf{z}_i^{\text{image}}$ and text embedding $\mathbf{z}_i^{\text{text}}$:

$$\begin{aligned} \mathbf{z}_i^{\text{image}} &= \text{CLIP}_{\text{image}}(\mathbf{x}, \mathbf{b}_i), \\ \mathbf{z}_i^{\text{text}} &= \text{CLIP}_{\text{text}}(\mathbf{y}_i^{\text{class}}). \end{aligned} \quad (4)$$

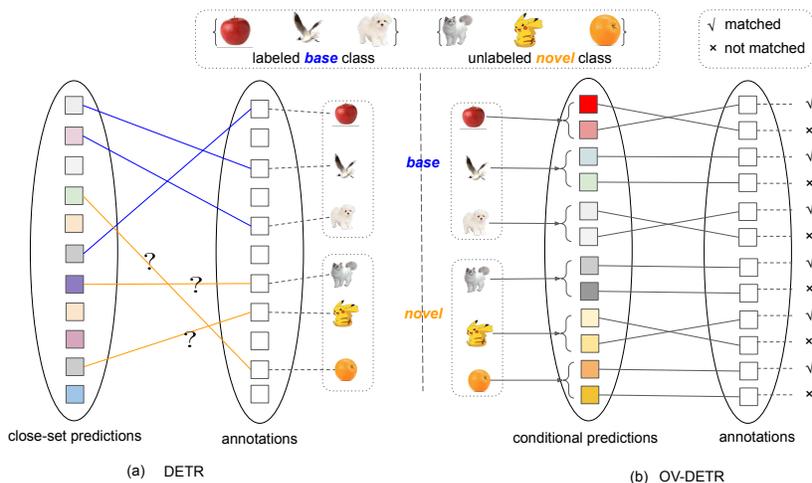


Fig. 3: **Comparing the label assignment mechanisms of DETR and our OV-DETR.** (a) In the original DETR, the set-to-set prediction is conducted via bipartite matching between predictions and closed-set annotations, in which a cost matrix in respect of the queries and categories. Due to the absence of class label annotations for *novel* classes, computing such a class-specific cost matrix is impossible. (b) On the contrary, our OV-DETR casts the open-vocabulary detection as a conditional matching process and formulate a binary matching problem that computes a class-agnostic matching cost matrix for conditional inputs.

Such image and text embeddings are already well-aligned by the CLIP model. Therefore, we can choose either of them as input queries to condition the DETR’s decoder and train to match the corresponding objects. Once training is done, we can then take arbitrary input queries during testing to perform open-vocabulary detection. To ensure equal training conditioned on image and text queries, we randomly select z_i^{text} or z_i^{image} with probability $\xi = 0.5$ as conditional inputs. Moreover, we follow previous works [13,35,7] to generate additional object proposals for *novel* classes to enrich our training data. We only extract image embeddings z_i^{image} for such novel-class proposals as conditional inputs, since their class names are unavailable to extract text embeddings. Please refer to supplementary materials for more details.

Conditional Matching. Our core training objective is to measure the matchability between the conditional input embeddings and detection results. In order to perform such conditional matching, we start with a fully-connected layer \mathbb{F}_{proj} to project the conditional input embeddings (z_i^{text} or z_i^{image}) to have the same dimension as \mathbf{q} . Then the input to the DETR decoder \mathbf{q}' is given by:

$$\mathbf{q}' = \mathbf{q} \oplus \mathbb{F}_{\text{proj}}(z_i^{\text{mod}}), \quad \text{mod} \in \{\text{text, image}\}, \quad (5)$$

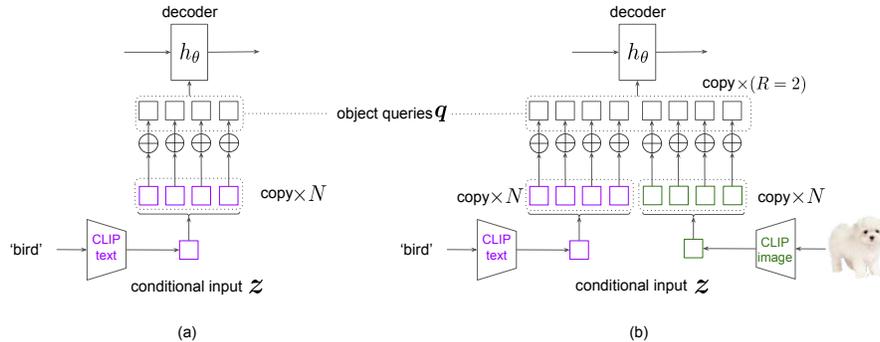


Fig. 4: DETR decoder with (a) single conditional input or (b) multiple conditional inputs in parallel.

where we use a simple addition operation \oplus to convert the *class-agnostic* object queries \mathbf{q} into *class-specific* \mathbf{q}' informed by $\mathbb{F}_{\text{proj}}(\mathbf{z}_i^{\text{mod}})$.

In practice, adding the conditional input embeddings \mathbf{z} to only one object query will lead to a very limited coverage of the target objects that may appear many times in the image. Indeed, in existing object detection datasets, there are typically multiple object instances in each image from the same or different classes. To enrich the training signal for our conditional matching, we copy the object queries \mathbf{q} for R times, and the conditional inputs ($\mathbf{z}_i^{\text{text}}$ or $\mathbf{z}_i^{\text{image}}$) for N times before performing the conditioning in Eq. (5). As a result, we obtain a total of $N \times R$ queries for matching during each forward pass, as shown in Fig. 4 (b). Experiments in the supplementary material will validate the importance of such “feature cloning” and also show how we determine N and R based on the performance-memory trade-off. Note for the final conditioning process, we further add an attention mask to ensure the independence between different query copies, as is similarly done in [4].

Given the conditioned query features \mathbf{q}' , our binary matching loss for label assignment is given as:

$$\mathcal{L}_{\text{cost}}(\mathbf{y}, \hat{\mathbf{y}}_{\sigma}) = \mathcal{L}_{\text{match}}(\mathbf{p}, \hat{\mathbf{p}}_{\sigma}) + \mathcal{L}_{\text{box}}(\mathbf{b}, \hat{\mathbf{b}}_{\sigma}), \quad (6)$$

where $\mathcal{L}_{\text{match}}(\mathbf{p}, \hat{\mathbf{p}}_{\sigma})$ denotes a new matching loss that replaces the classification loss $\mathcal{L}_{\text{cls}}(\mathbf{p}, \hat{\mathbf{p}}_{\sigma})$ in Eq. (3). Here in our case, \mathbf{p} is a 1-dimensional sigmoid probability vector that characterizes the matchability (‘matched’ vs. ‘not matched’), and $\mathcal{L}_{\text{match}}$ is simply implemented by a Focal loss [21] $\mathcal{L}_{\text{Focal}}$ between predicted $\hat{\mathbf{p}}_{\sigma}$ and ground-truth \mathbf{p} . For instance, with the ‘bird’ query as input, our matching loss should allow us to match all the bird instances in one image, while tagging instances from other classes as ‘not matched’.

3.3 Optimization

After optimizing Eq. (6), we obtain the optimized label assignments σ for different object queries. This process produces a set of detected objects with assigned box coordinates $\hat{\mathbf{b}}$ and 2-dim matching probability $\hat{\mathbf{p}}$ that we will use to compute our final loss function for modeling training. We further attach an embedding reconstruction head to the model, which learns to predict embedding \mathbf{e} to be able to reconstruct each conditional input embedding \mathbf{z}^{text} or $\mathbf{z}^{\text{image}}$:

$$\mathcal{L}_{\text{embed}}(\mathbf{e}, \mathbf{z}) = \|\mathbf{e} - \mathbf{z}^{\text{mod}}\|_1, \text{ mod} \in \{\text{text}, \text{image}\}. \quad (7)$$

Supplementary materials validate the effectiveness of $\mathcal{L}_{\text{embed}}$.

Our final loss for model training combines $\mathcal{L}_{\text{embed}}$ with bounding box losses $\mathcal{L}_{\text{match}}(\mathbf{p}, \hat{\mathbf{p}})$ and $\mathcal{L}_{\text{box}}(\mathbf{b}, \hat{\mathbf{b}})$ again:

$$\begin{aligned} \mathcal{L}_{\text{loss}}(\mathbf{y}, \hat{\mathbf{y}}) &= \mathcal{L}_{\text{match}}(\mathbf{p}, \hat{\mathbf{p}}) + \mathcal{L}_{\text{box}}(\mathbf{b}, \hat{\mathbf{b}}) + \mathcal{L}_{\text{embed}}(\mathbf{e}, \mathbf{z}) \\ &= \lambda_{L_{\text{Focal}}} \mathcal{L}_{\text{Focal}} + \lambda_{L_{L1}} \mathcal{L}_{L1} + \lambda_{L_{\text{GIoU}}} \mathcal{L}_{\text{GIoU}} + \lambda_{L_{\text{embed}}} \mathcal{L}_{\text{embed}}, \end{aligned} \quad (8)$$

where \mathcal{L}_{box} consists of the L1 loss and the generalized IoU (GIoU) [30] loss for boxes, while $\lambda_{L_{\text{Focal}}}$, $\lambda_{L_{L1}}$, $\lambda_{L_{\text{GIoU}}}$ and $\lambda_{L_{\text{embed}}}$ are the weighting parameters.

3.4 Inference

During testing, for each image, we send the text embedding \mathbf{z}^{text} of all the *base+novel* classes to the model and merge the results by selecting the top k predictions with highest prediction scores. We follow the prior work [13] to use $k = 100$ for COCO dataset and $k = 300$ for LVIS dataset. To obtain the context representation \mathbf{c} in Eq. (1), we forward the input image through the CNN backbone f_ϕ and Transformer encoder h_ψ . Note \mathbf{c} is computed only once and shared for all conditional inputs for efficiency. Then the conditioned object queries from different classes are sent to the Transformer decoder in parallel. In practice, we copy the object queries for R times as shown in Fig. 4 (b).

4 Experiments

Datasets. We evaluate our approach on two standard open-vocabulary detection benchmarks modified from LVIS [14] and COCO [22] respectively. LVIS [14] contains 100K images with 1,203 classes. The classes are divided into three groups, namely frequent, common and rare, based on the number of training images. Following ViLD [13], we treat 337 rare classes as *novel* classes and use only the frequent and common classes for training. The COCO [22] dataset is a widely-used benchmark for object detection, which consists of 80 classes. Following OVR-CNN [36], we divide the classes in COCO into 48 *base* categories and 17 *novel* categories, while removing 15 categories without a synset in the WordNet hierarchy. The training set is the same as the full COCO but only images containing at least one *base* class are used. We refer to these two benchmarks as OV-LVIS and OV-COCO hereafter.

| # | Method | AP ^m | AP ^m _{novel} | AP ^m _c | AP ^m _f |
|---|-------------|-----------------|----------------------------------|------------------------------|------------------------------|
| 1 | Mask R-CNN† | 22.5 | 0.0 | 22.6 | 32.4 |
| 2 | Def DETR | 22.4 | 0.0 | 22.4 | 32.0 |

Table 1: **Mask R-CNN and Def DETR on OV-LVIS, both trained on base classes.** †: copied from ViLD [13].

| # | P | M | AP ^m | AP ^m _{novel} | AP ^m _c | AP ^m _f |
|---|---|---|-----------------|----------------------------------|------------------------------|------------------------------|
| 1 | | | 24.2 | 9.5 | 23.2 | 31.7 |
| 2 | ✓ | | 19.9 | 6.3 | 17.4 | 28.6 |
| 3 | ✓ | ✓ | 26.6 | 17.4 | 25.0 | 32.5 |

Table 2: **Ablation study** on using object proposals (P) and our conditional binary matching mechanism (M).

Evaluation Metrics. For OV-LVIS, we report the mask mAP for rare, common and frequent classes, denoted by AP_r^m , AP_c^m and AP_f^m . The rare classes are treated as *novel* classes (AP_{novel}^m). The symbol AP^m denotes to the mAP of all the classes. For OV-COCO, we follow previous work that only reports the AP50^b metric, which means the box mAP at IoU threshold 0.5.

Extension for Instance Segmentation. For OV-LVIS, instance segmentation results are needed for the evaluation process. Although DETR [2] and its follow-ups [41,24] are developed for the object detection task, they can also be extended to the instance segmentation task. We follow DETR [2] to add an external class-agnostic segmentation head to solve the instance segmentation task. The segmentation head employs the fully convolutional network (FCN [23]) structure, which takes features extracted from the Transformer decoder as input and produces segmentation masks.

Implementation Details. Our model is based on Deformable DETR [41]. Following ViLD [13], we also use the open-source CLIP model [27] based on ViT-B/32 for extracting text and image embeddings. Please refer to our supplementary material for more training details.

4.1 Ablation Studies

We conduct ablation study on OV-LVIS to evaluate the main components in our approach.

The Architecture Difference. Previous works such as ViLD [13] are based on the RPN-based Mask R-CNN [15], while our work is based on the Transformer-based detector Deformable DETR [41]. We first study the difference of these two detectors on the open-vocabulary setting trained with *base* classes only. As shown in Table 1 row(1-2), we observe that Mask R-CNN performs a slightly better than Deformable DETR [41]. This gap is small, indicating that we have a fair starting point compared to ViLD [13].

Object Proposals. We then replace Deformable DETR’s classifier layer as text embedding provided by CLIP and trained with *base* classes only. This step is similar to the previous ViLD-text [13] method. Results is presented in Table 2 row 1. We observe that the AP_{novel}^m metric improved from 0.0 to 9.5. To further improve the AP_{novel}^m metric, we add the object proposals that may contain the

Table 3: **Main results on OV-LVIS and OV-COCO.** For OV-LVIS (w/ 886 base classes and 317 novel classes), we report mask mAP and a breakdown on novel (rare), common, and frequent classes. For OV-COCO (w/ 48 base classes and 17 novel classes), we report bounding box mAP at IoU threshold 0.5. †: zero-shot methods that do not use captions or image-text pairs. ‡: ensemble model.

| # | Method | OV-LVIS | | | | OV-COCO | | |
|---|-----------------|-----------------|----------------------------------|------------------------------|------------------------------|-------------------|------------------------------------|-----------------------------------|
| | | AP ^m | AP ^m _{novel} | AP ^m _c | AP ^m _f | AP50 ^b | AP50 ^b _{novel} | AP50 ^b _{base} |
| 1 | SB [1]† | - | - | - | - | 24.9 | 0.3 | 29.2 |
| 2 | DELO [40]† | - | - | - | - | 13.0 | 3.1 | 13.8 |
| 3 | PL [28]† | - | - | - | - | 27.9 | 4.1 | 35.9 |
| 4 | OVR-CNN [36] | - | - | - | - | 46.0 | 22.8 | 39.9 |
| 5 | ViLD-text [13] | 24.9 | 10.1 | 23.9 | 32.5 | 49.3 | 5.9 | 61.8 |
| 6 | ViLD [13] | 22.5 | 16.1 | 20.0 | 28.3 | 51.3 | 27.6 | 59.5 |
| 7 | ViLD-ens. [13]‡ | 25.5 | 16.6 | 24.6 | 30.3 | - | - | - |
| 8 | OV-DETR | 26.6 | 17.4 | 25.0 | 32.5 | 52.7 | 29.4 | 61.0 |
| | (ours vs. #6) | (+4.1) | (+1.3) | (+5.0) | (+4.2) | (+1.4) | (+1.8) | (+1.5) |

region of *novel* classes into the training stage. Because we do not know the category id of these object proposals, we observe that the label assignment of these object proposals is inaccurate and will decrease the AP^m_{novel} performance from 9.5 to 6.3.

Conditional Binary Matching. Now we replace DETR’s default close-set labeling assignment as our proposed conditional binary matching. The comparison results between Table 2 row 2-3 shows that our binary matching strategy can better leverage the knowledge from object proposals and improve the AP^m_{novel} from 9.5 to 17.4. Such a large improvement shows that the proposed conditional matching is essential when applying the DETR-series detector for the open-vocabulary setting.

4.2 Results on Open-vocabulary Benchmarks

Table 3 summarizes our results. We compare our method with SOTA open-vocabulary detection methods including: (1) OVR-CNN [36] (see Table 3 row 4). It pre-trains the detector’s projecting layer on image-caption pairs using contrastive loss and then fine-tunes on the object detection task; (2) Variants of ViLD [13] such as ViLD-text and ViLD-ensemble (see Table 3 rows 5-7). ViLD is the first study that uses CLIP embeddings [27] for open-vocabulary detection. Compared with ViLD-text, ViLD uses knowledge distillation from the CLIP visual backbone, improves AP_{novel} at the cost of hurting AP_{base}. ViLD-ens. combines the two models and shows improvements for both metrics. Such an ensemble-based method also brings extra time and memory cost.

For completeness, we also list the results of some previous zero-shot methods such as SB [1], DELO [40] and PL [28] in Table 3 rows 1-3. On *OV-LVIS* bench-

Table 4: **Generalization to other datasets.** We evaluate OV-DETR trained on LVIS when transferred to other datasets such as PASCAL VOC 2007 test set and COCO validation set by simply replacing the text embeddings. The experimental setting is the same as that of ViLD [13]. We observe that OV-DETR achieves better generalization performance than ViLD [13].

| # | Method | Pascal VOC | | COCO | | |
|---|----------------|-------------------------------|-------------------------------|-----------------|-------------------------------|-------------------------------|
| | | AP ₅₀ ^b | AP ₇₅ ^b | AP ^b | AP ₅₀ ^b | AP ₇₅ ^b |
| 1 | ViLD-text [13] | 40.5 | 31.6 | 28.8 | 43.4 | 31.4 |
| 2 | ViLD [13] | 72.2 | 56.7 | 36.6 | 55.6 | 39.8 |
| 3 | OV-DETR | 76.1 | 59.3 | 38.1 | 58.4 | 41.1 |
| | (ours vs #2) | (+3.9) | (+2.6) | (+1.5) | (+2.8) | (+1.3) |

mark, OV-DETR improves the previous SOTA ViLD by 4.1 on AP^m and 1.3 on AP_{novel}^m. Compared with ViLD, our method will not affect the performance of *base* classes when improve the *novel* classes. Even compared with the ensemble result of ViLD-ensemble, OV-DETR still boosts the performance by 1.5, 0.8, 1.0 and 2.2, respectively (%). Noted that OV-DETR only uses a single model and does not leverage any ensemble-based technique. On *OV-COCO* benchmark, OV-DETR improves the baseline and outperforms OVR-CNN [36] by a large margin, notably, the 6.6 mAP improvements on *novel* classes. Compared with ViLD [13], OV-DETR still achieves 1.4 mAP gains on all the classs and 1.8 mAP gains on *novel* classes. In summary, it is observed that that OV-DETR achieves superior performance across different datasets compared with different methods.

4.3 Generalization Ability of OV-DETR.

We follow ViLD [13] to test the generalization ability of OV-DETR by training the model on LVIS [14] dataset and evaluated on PASCAL VOC [8] and COCO [22]. We keep the same implementation details with ViLD [13]. We switch the text embeddings of the category names from the source dataset to new datasets. The text embeddings of new classes are used as conditional inputs during the inference phase. As shown in Table 4, we observe that OV-DETR achieves better transfer performance than ViLD. The experimental results show that the model trained by our conditional-based mechanism has transferability to other domains.

4.4 Qualitative Results

We visualize OV-DETR’s detection and segmentation results in Fig. 5. The results based on conditional text queries, conditional image queries, and a mixture of conditional text and image queries are shown in the top, middle and bottom row, respectively. Overall, our OV-DETR can accurately localize and precisely segment out the target objects from novel classes despite no annotations of these

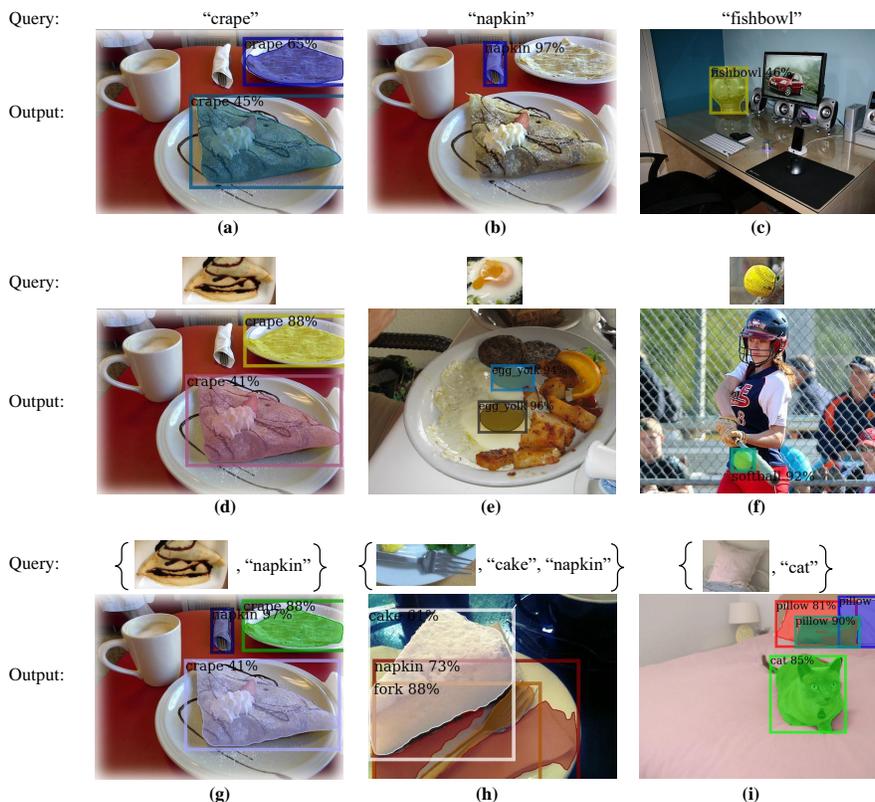


Fig. 5: **Qualitative results on LVIS.** OV-DETR can precisely detect and segment novel objects (e.g., ‘crape’, ‘fishbowl’, ‘softball’) given the conditional text query (top) or conditional image query (middle) or a mixture of them (bottom).

classes during training. It is worth noting that the conditional image queries, such as “crape” in (d) and “fork” in (h), appear drastically different from those in the target images but OV-DETR can still robustly detect them.

4.5 Inference Time Analysis

OV-DETR exhibits great potential in open-vocabulary detection but is by no means a perfect detector. The biggest limitation of OV-DETR is that the inference speed is slow when the number of classes to detect is huge like 1,203 on LVIS [13]. This problem is caused by the conditional design that requires multiple forward passes in the Transformer decoder (depending on the number of classes).

We show a detailed comparison on the inference time between Deformable DETR and OV-DETR in Table 5. Without using any tricks, the vanilla OV-

| # | Method | COCO | LVIS |
|---|------------------|------------|------------|
| 1 | Def DETR | 0.31 | 1.49 |
| 2 | Ours | 0.72 | 23.84 |
| 3 | Ours (optimized) | 0.63 | 9.57 |
| | (vs #2) | (+↓ 12.5%) | (+↓ 59.9%) |

Table 5: Comparison of the inference time (second per iteration) between Deformable DETR [41] and our OV-DETR before/after optimization on LVIS and COCO.

DETR (#2), i.e., using a single forward pass for each class, is about $2\times$ slower than Deformable DETR (#1) on COCO (w/ 80 classes) while $16\times$ slower on LVIS (w/ 1,203 classes). As discussed in Sec. 3.2 and shown in Fig. 4(b), we optimize the speed by forwarding multiple conditional queries to the Transformer decoder in parallel, which reduces the inference time by 12.5% on COCO and nearly 60% on LVIS (see #3 in Table 5). Still, there is much room for improvement.

It is worth noting that such a slow inference problem is not unique to our approach—most instance-conditional models would have the same issue [19], which is the common price to pay in exchange for better performance. The computation bottleneck of our method lies in the computation of the Transformer decoder in Eq. (2). A potential solution is to design more efficient attention modules [33,34], which we leave as future work. In human-computer interaction where users already have target object(s) in mind, *e.g.*, a missing luggage or a specific type of logo, the conditional input is fixed and low in number, thus the inference time is negligible.

5 Conclusion

Open-vocabulary detection is known to be a challenging problem due to the lack of training data for unseen classes. Recent advances in large language models have offered a new perspective for designing open-vocabulary detectors. In this work, we show how an end-to-end Transformer-based detector can be turned into an open-vocabulary detector based on conditional matching and with the help of pre-trained vision-language models. The results show that, despite having a simplified training pipeline, our open-vocabulary detector based on Transformer significantly outperforms current state of the arts that are all based on two-stage detectors. We hope our approach and the findings presented in the paper can inspire more future work on the design of efficient open-vocabulary detectors.

Acknowledgment This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s). It is also partly supported by the NTU NAP grant and Singapore MOE AcRF Tier 2 (MOE-T2EP20120-0001). This work was supported by SenseTime SenseCore AI Infrastructure-AIDC.

References

1. Bansal, A., Sikka, K., Sharma, G., Chellappa, R., Divakaran, A.: Zero-shot object detection. In: ECCV. pp. 384–400 (2018)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. pp. 213–229. Springer (2020)
3. Chen, K., Kovvuri, R., Nevatia, R.: Query-guided regression network with context policy for phrase grounding. In: ICCV. pp. 824–832 (2017)
4. Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: CVPR. pp. 1601–1610 (2021)
5. Deng, C., Wu, Q., Wu, Q., Hu, F., Lyu, F., Tan, M.: Visual grounding via accumulated attention. In: CVPR. pp. 7746–7755 (2018)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
7. Du, Y., Wei, F., Zhang, Z., Shi, M., Gao, Y., Li, G.: Learning to prompt for open-vocabulary object detection with vision-language model. In: CVPR. pp. 14084–14093 (2022)
8. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. IJCV **88**(2), 303–338 (2010)
9. Felzenszwalb, P., McAllester, D., Ramanan, D.: A discriminatively trained, multi-scale, deformable part model. In: CVPR. pp. 1–8. Ieee (2008)
10. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: Devise: A deep visual-semantic embedding model. NeurIPS (2013)
11. Gao, P., Zheng, M., Wang, X., Dai, J., Li, H.: Fast convergence of detr with spatially modulated co-attention. In: ICCV. pp. 3621–3630 (2021)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. NeurIPS **27** (2014)
13. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. In: ICLR (2022)
14. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: CVPR. pp. 5356–5364 (2019)
15. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017)
16. Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: Mdetr-modulated detection for end-to-end multi-modal understanding. In: ICCV. pp. 1780–1790 (2021)
17. Kim, D., Lin, T.Y., Angelova, A., Kweon, I.S., Kuo, W.: Learning open-world object proposals without learning to classify. RA-L (2022)
18. Kuhn, H.W.: The hungarian method for the assignment problem **2**(1-2), 83–97 (1955)
19. Li, S., Xiao, T., Li, H., Zhou, B., Yue, D., Wang, X.: Person search with natural language description. In: CVPR. pp. 1970–1979 (2017)
20. Li, Z., Yao, L., Zhang, X., Wang, X., Kanhere, S., Zhang, H.: Zero-shot object detection with textual descriptions. In: AAAI. pp. 8690–8697 (2019)
21. Lin, T.Y., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2999–3007 (2017)
22. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755. Springer (2014)

23. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR. pp. 3431–3440 (2015)
24. Meng, D., Chen, X., Fan, Z., Zeng, G., Li, H., Yuan, Y., Sun, L., Wang, J.: Conditional detr for fast training convergence. In: ICCV. pp. 3651–3660 (2021)
25. Papageorgiou, C., Poggio, T.: A trainable system for object detection. IJCV **38**(1), 15–33 (2000)
26. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP. pp. 1532–1543 (2014)
27. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021)
28. Rahman, T., Chou, S.H., Sigal, L., Carenini, G.: An improved attention for visual question answering. In: CVPR. pp. 1653–1662 (2021)
29. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. NeurIPS **28**, 91–99 (2015)
30. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: CVPR. pp. 658–666 (2019)
31. Shizhen, Z., Changxin, G., Yuanjie, S., Lerenhan, L., Changqian, Y., Zhong, J., Nong, S.: Gtnet: Generative transfer network for zero-shot object detection. In: AAAI (2020)
32. Szegedy, C., Toshev, A., Erhan, D.: Deep neural networks for object detection. NeurIPS **26** (2013)
33. Tay, Y., Bahri, D., Yang, L., Metzler, D., Juan, D.C.: Sparse sinkhorn attention. In: ICML. pp. 9438–9447. PMLR (2020)
34. Wang, S., Li, B.Z., Khabsa, M., Fang, H., Ma, H.: Linformer: Self-attention with linear complexity. arXiv preprint arXiv:2006.04768 (2020)
35. Xie, J., Zheng, S.: Zsd-yolo: Zero-shot yolo detection using vision-language knowledge distillation. arXiv preprint arXiv:2109.12066 (2021)
36. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: CVPR. pp. 14393–14402 (2021)
37. Zhang, Y., Zhou, K., Liu, Z.: Neural prompt search (2022)
38. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Conditional prompt learning for vision-language models. In: CVPR (2022)
39. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. IJCV (2022)
40. Zhu, P., Wang, H., Saligrama, V.: Don’t even look once: Synthesizing features for zero-shot detection. In: CVPR. pp. 11693–11702 (2020)
41. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR (2020)