

# Region-Aware Pretraining for Open-Vocabulary Object Detection with Vision Transformers

Dahun Kim      Anelia Angelova      Weicheng Kuo  
 Google Research, Brain Team  
 {mcahny, anelia, weicheng}@google.com

## Abstract

We present *Region-aware Open-vocabulary Vision Transformers (RO-ViT)* – a contrastive image-text pretraining recipe to bridge the gap between image-level pretraining and open-vocabulary object detection. At the pretraining phase, we propose to randomly crop and resize regions of positional embeddings instead of using the whole image positional embeddings. This better matches the use of positional embeddings at region-level in the detection finetuning phase. In addition, we replace the common softmax cross entropy loss in contrastive learning with focal loss to better learn the informative yet difficult examples. Finally, we leverage recent advances in novel object proposals to improve open-vocabulary detection finetuning. We evaluate our full model on the LVIS and COCO open-vocabulary detection benchmarks and zero-shot transfer. RO-ViT achieves a state-of-the-art 32.1  $AP_{\tau}$  on LVIS, surpassing the best existing approach by +5.8 points in addition to competitive zero-shot transfer detection. Surprisingly, RO-ViT improves the image-level representation as well and achieves the state of the art on 9 out of 12 metrics on COCO and Flickr image-text retrieval benchmarks, outperforming competitive approaches with larger models.

## 1. Introduction

The ability to detect objects in the visual world is a hallmark of computer vision and machine intelligence. It is key to enable many applications, e.g. autonomous agents adapting to new environments with many novel objects, a shopping system handling fine-grained user queries such as “fedora”, “bonnet” outside of the training vocabulary. However, modern object detectors typically rely on manual annotations of regions and class labels, which limits their vocabulary size to an order of  $10^3$  and it is prohibitively expensive to scale up further. This is orders of magnitude smaller than the objects we encounter in the visual world.

Recently, the open-vocabulary detection task (OVD) has been proposed to overcome such limitation by leveraging

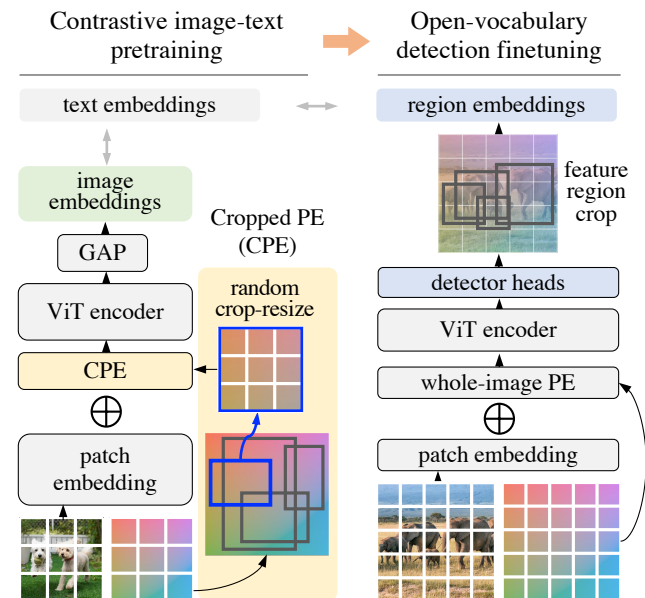


Figure 1. Existing vision-language models are designed for image-level tasks, e.g., classification and retrieval. In this paper, we aim to enhance the image-text pretraining for region-level downstream task: open-vocabulary object detection. At pretraining, we propose to randomly crop and resize regions of positional embeddings (PE) instead of using the whole image PE. This better matches the use of PE at region-level in the detection finetuning and results in better performance in detection and surprisingly also retrieval task.

abundant image-text pairs for training and ingesting the text queries provided by users at test time [54]. By treating the categories as text embeddings rather than discrete ids, open-vocabulary detectors can flexibly predict a wide range of objects unseen during the training time. Most existing works leverage image-text pre-training which contains rich semantic knowledge of open-vocabulary concepts. Knowledge distillation [13, 19], weak supervision [63], self-training [41, 57, 60], and frozen model [29] have been proposed, and CNN backbones are most commonly used. With the growing popularity of vision transformers in image understanding [12, 32, 56], multimodal [1, 2, 47], and self-

supervised tasks [5, 6, 22], it is important to understand how to build capable open-vocabulary detectors with vision transformers [12, 35].

To our best knowledge, all existing works assume pre-trained Vision-Language Models (VLM) are given, and develop adaptation or finetuning recipes to bridge the gap between image-level pretraining and object-level finetuning [13, 19, 41, 57, 60]. Since the VLMs are designed for image-level tasks *e.g.* classification, retrieval, the notion of objects/regions are not adequately utilized in the pretraining process. We believe it would be beneficial for open-vocabulary detection if we bake in locality information in the image-text pretraining.

We present RO-ViT, a simple recipe to pretrain vision transformers in a region-aware manner for open-vocabulary object detection. Standard pretraining typically uses full-image positional embeddings, which does not generalize well to detection tasks. Thus, we propose a novel positional embedding scheme called “Cropped Positional Embedding” which better matches the use of region crops in detection finetuning (see Fig. 1). In addition, we found it beneficial to replace the softmax cross entropy loss with focal loss in contrastive learning, which affords us more control to learn from harder and more informative examples. Finally, we leverage recent advances in novel object proposals [28] to improve open-vocabulary detection finetuning. The motivation is that existing approaches often miss novel objects in the object proposal stage because the proposals tend to overfit to the foreground categories.

We evaluate the approach on the standard LVIS and COCO open-vocabulary benchmarks. On LVIS, our best model achieves a state-of-the-art 32.1  $AP_r$  at the system level, surpassing the best existing approach by +5.8  $AP_r$ . Compared to the best existing ViT-based approach, ours outperforms by a healthy margin of +6.5  $AP_r$ . Our LVIS-trained model outperforms existing baselines on transfer detection to Objects365 without re-training. Although not explicitly optimized for retrieval, RO-ViT surprisingly achieves the state-of-the-art performance on 9 out of 12 metrics in image-text retrieval benchmark and outperforms strong baselines with standard positional embeddings, cross entropy loss, and larger model capacity. We conduct ablations to confirm the benefits of each proposed component. Visualization of the learnt positional embeddings shows that our approach (CPE) leads to more symmetrical and diverse patterns than the baseline. In summary:

- We propose RO-ViT to bridge the positional embeddings in image-text pretraining to open-vocabulary detection finetuning.
- We show that image-text pretraining with focal loss is more effective than existing softmax CE loss.
- We improve the open-vocabulary detection finetuning recipe with novel object proposals.

- RO-ViT achieves the state of the art on the LVIS open-vocabulary detection benchmark, and 9 out of 12 metrics on COCO and Flickr image-text retrieval benchmarks.

We hope these findings will facilitate the research community to further explore open-vocabulary detection from the perspective of image-text pretraining with potential benefits for both region-level and image-level tasks.

## 2. Related Works

### Learning open-vocabulary and zero-shot recognition.

Building general representation for open-vocabulary and zero-shot recognition is a fundamental machine learning task. DeViSE [15] and ConSE [36] are pioneering works to learn a shared image-text embedding space for zero-shot recognition by convolutional networks. As image and language often co-occur on the internet, the community has explored learning such representation from various data sources including image tags [8, 11, 27], captions/text descriptions [10, 23, 42, 45, 59], alt-texts [26, 37, 55], and image search queries [39]. Recently, contrastive learning has become a popular paradigm to learn from large image-text corpus due to its simplicity and scalability, where increasing model size and data yield consistent improvement on open-vocabulary classification benchmarks. The aforementioned works focus on image-level recognition, whereas we study how to best tailor them for region-level understanding.

### Self-supervised representation learning for localization.

Due to the challenge to scale up labeling for localization tasks, many efforts have been made to learn locality-sensitive representation in a self-supervised manner. Existing approaches roughly fall into the contrastive or generative paradigms. Contrastive approaches typically involve region or point-level contrastive learning using sliding windows [49], object proposals [24, 48], or point samples [3]. Most existing contrastive methods are CNN-based, while ViT backbones have been popular recently with generative approaches [5, 22] and has shown promise with contrastive approaches [6]. Masked image modeling is commonly used to learn the ViT features for localization tasks. Even though these self-supervised methods are well-suited for standard localization tasks, they lack the image-text correspondence necessary for open-vocabulary recognition. Thus, we study contrastive image-text representation learning for open-vocabulary detection with ViTs.

### Open-vocabulary object detection and segmentation.

Zero-shot detection was proposed to scale up detection models beyond their limited training categories. Popular approaches accomplish this by learning the alignment between region visual representation and category word embeddings [4, 9, 40, 58] or hallucinating visual features with a generative model [21, 64]. Without any visual knowledge of the novel categories, zero-shot detectors tend to struggle

with low performance. This motivates the open-vocabulary detection benchmark [54] to bridge the gap between zero-shot and fully-supervised detection.

With the rise of image-text pretraining, many works have explored adapting these pretrained models to open-vocabulary detection and segmentation [17, 19, 30, 60, 61]. For example, ViLD [19] proposes to distill the image-text knowledge into the detector. DetPro [13] improves ViLD by incorporating the idea of prompt optimization [62]. Moreover, region-text self-training has been shown effective on image caption data [60], classification data [41], or even unlabeled data [57]. Weak supervision [63] and phrase grounding [31] approaches have also been proposed. In terms of architecture, CNNs are most commonly used but vision transformers have recently been adopted as well [35, 61]. While existing works typically assume image-text pretrained models are given and focus on finetuning or adaptation strategies, our focus is to improve the upstream image-text pretraining with vision transformers.

### 3. Method

In this paper we address the problem of open-vocabulary object detection. At training time, the model has access to the detection labels of  $C_B$  base categories, but needs to detect objects from a set of  $C_N$  novel categories at test time. We follow existing works [13, 19, 29, 60] to leverage pretrained vision and language models (VLM) [39]. Unlike existing works, we explore how to best pretrain our own VLMs with vision transformers [12] for open-vocabulary object detection.

#### 3.1. Preliminaries

**Contrastive image-text pretraining.** Two-tower contrastive image-text learning involves an image encoder and a text encoder [26, 39]. The text tower is typically transformer encoder, whereas the image tower can be CNN-based or ViT as in our case. Given a set of image-text corpus, the model learns to bring each image and its corresponding text together and push other non-matching texts away in the embedding space. The most common objective is softmax cross-entropy loss. The image/text embeddings are typically obtained by taking the pre-pended class token embedding, self-attention pooling, or average pooling. We use the global average pooling followed by L2 normalization for both image and text embeddings.

**Open-vocabulary object detector.** An open-vocabulary object detector is trained with the labels of  $C_B$  base categories, but needs to detect the union of base and novel categories ( $C_B \cup C_N$ ) at test time. Most object detectors use K-way classifiers because the number of categories are the same between train and test time. To deal with additional categories at test time, the common practice is to replace the

conventional fixed-size classifier fully-connected layer with the text embeddings of base categories [19, 54]. It is important that the text embeddings come from the matching text encoder of the image encoder during pretraining, so that the open-vocabulary knowledge would be preserved. We represent the background category by a “background” phrase and the proposals not matched to any annotations in  $C_B$  are labeled as background.

During training, we compute the detection scores  $p_i$  for each region  $i$  as the cosine similarity between the RoI-Align feature (*i.e.*, region embedding) and the text embeddings of  $C_B$ , followed by a softmax. At test time, we expand the text embeddings from  $C_B$  to  $C_B \cup C_N$  plus the “background” category for open-vocabulary detection. We extract the VLM embedding of region  $i$  by RoI-Align on the output feature map of the ViT backbone, and compute the VLM region scores  $z_i$  as the cosine similarity with the  $C_B \cup C_N$  text embeddings. Similarly, the detection scores  $p_i$  are now computed with the  $C_B \cup C_N$  text embeddings. The combined open-vocabulary detection score  $s_i^{\text{OVD}}$  is obtained by geometric means [19, 29]:

$$s_i^{\text{OVD}} = \begin{cases} z_i^{(1-\alpha)} \cdot p_i^\alpha & \text{if } i \in C_B \\ z_i^{(1-\beta)} \cdot p_i^\beta & \text{if } i \in C_N \end{cases} \quad (1)$$

, where  $\alpha, \beta \in [0, 1]$  control the weights for base and novel categories. The background score comes directly from the detection score  $p_i$ , because the VLM score with “background” phrase tends to be not as reliable.

We adopt Mask R-CNN heads in our detector and use *class-agnostic* box regression and mask prediction heads following existing works [13, 19, 29, 54, 60]. Note that we use a contrastively pretrained ViT to initialize our detector backbone, and adopt the simple feature pyramid and windowed attention to handle higher-resolution image inputs as proposed by Li *et al.* [32].

#### 3.2. Region-Aware Image-Text Pretraining

Existing vision-language models are trained to match an image as a whole to a text description [26, 39]. However, the pretraining is unaware of the alignment between its region-level representations and text tokens, which is essential to downstream open-vocabulary detection. We propose a novel Cropped Positional Embeddings (CPE) to bridge this gap, and also find it beneficial to learn from hard examples with a focal loss. Our improvements introduce no extra parameters and minimal computation costs.

**Cropped Positional Embedding (CPE).** The positional embeddings are important to transformers as they provide the information of where each element in the set comes from. This information is often useful for downstream recognition and localization tasks.

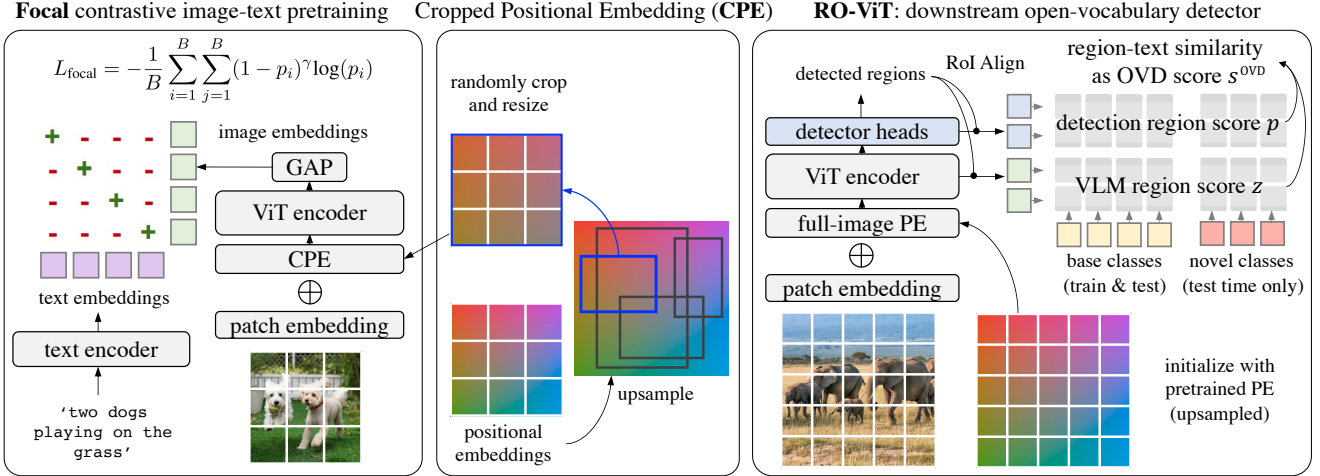


Figure 2. **RO-ViT framework overview.** Our Region-aware Open-vocabulary Vision Transformer (**RO-ViT**) attempts to bridge the gap between the image-level vision-language model (VLM) pretraining and downstream open-vocabulary detection. For the pretraining, we propose Cropped Positional Embedding (**CPE**) which randomly crops and resizes a *region* of positional embeddings instead of using the whole-image PE. In addition, we use **focal** loss instead of the common softmax cross entropy loss for contrastive learning. The pretrained ViT backbone is transferred to the downstream open-vocabulary detection by replacing the global average pooling with detector heads. RO-ViT detector takes the whole-image positional embeddings as input, and the detections are used to crop out region embeddings from the ViT backbone features. The region embeddings then match with the cached category embeddings to obtain the VLM score  $z$ , which is combined with the detection score  $p$  into the open-vocabulary detection score  $s^{\text{OVD}}$  (see Equation 1).

There is a mismatch between the way the positional embeddings are used in existing contrastive pretraining approaches and open-vocabulary detection finetuning. The pretraining approaches typically apply full-image positional embeddings [26, 39] during training, and use the same positional embeddings for downstream tasks, *e.g.*, zero-shot recognition. However, the recognition occurs at region-level for open-vocabulary detection finetuning, which requires the full-image positional embeddings to generalize to regions that they never see during the pretraining.

To bridge this gap, we propose Cropped Positional Embedding (CPE) (see center of Figure 2). First, we up-sample the positional embeddings from the image size typical for pretraining, *e.g.*, 224 to that typical for detection tasks, *e.g.*, 1024. Then we randomly crop and resize a region from the up-sampled positional embeddings and use that as the image-level positional embeddings during pretraining. The regions are uniformly sampled from the normalized coordinates as  $x_1 \sim \text{Uniform}(0, 1)$ ,  $y_1 \sim \text{Uniform}(0, 1)$ ,  $x_2 \sim \text{Uniform}(x_1, 1)$ ,  $y_2 \sim \text{Uniform}(y_1, 1)$ , while keeping the crop scale ratio in  $[0.1, 1.0]$ .

Intuitively, this causes the model to view an image not as a full image in itself, but as a region crop from some larger unknown image. This better matches the downstream use case of detection where recognition occurs at region- rather than image-level.

**Focal Loss.** We desire to have finer control over how hard examples are weighted than what the softmax cross entropy

loss can provide. Focal loss offers a natural option to tune the weights of hard examples [33]. Let  $v_i$  and  $l_i$  be the normalized image and text embeddings, and the image-to-text (I2T) contrastive losses be  $L_{\text{softmax}}$  and  $L_{\text{focal}}$  for the softmax (baseline) or focal loss (RO-ViT). We define the losses mathematically below:

$$L_{\text{softmax}} = -\frac{1}{B} \sum_{i=1}^B \log\left(\frac{\exp(v_i l_i / \tau)}{\sum_{j=1}^B \exp(v_i l_j / \tau)}\right) \quad (2)$$

$$L_{\text{focal}} = -\frac{1}{B} \sum_{i=1}^B \sum_{j=1}^B (1 - p_i)^\gamma \log(p_i) \quad (3)$$

, where  $p_i$  denotes the true class probability as below:

$$p_i = \begin{cases} \sigma(v_i l_j / \tau) & \text{if } i = j \\ 1 - \sigma(v_i l_j / \tau) & \text{if } i \neq j \end{cases} \quad (4)$$

Here  $\sigma$  denotes the sigmoid function. We adopt the simpler non alpha-balanced form of focal loss [33]. The text-to-image (T2I) contrastive losses are symmetrical with the I2T losses by simply exchanging the inner/outer summation loops. The total loss is the sum of both I2T and T2I losses.

### 3.3. Open-vocabulary Detector Finetuning

Here we present two simple techniques to improve the downstream open-vocabulary detector (Sec. 3.1). Despite the backbone features pretrained from the vast open-vocabulary data, the added detector layers (neck and heads)

method	pretrained model	detector backbone	AP <sub>r</sub>	AP
<b>ConvNet based:</b>				
DetPro-Cascade [13]	ViT-B/32	R-50	20.0	27.0
Detic-CN2 [63]	ViT-B/32	R-50	24.6	32.4
RegionCLIP [60]	R-50x4	R-50x4	22.0	32.3
ViLD-Ens [19]	ViT-B/32	R-152	18.7	26.0
ViLD-Ens [19]	ViT-L/14	EffNet-B7	21.7	29.6
ViLD-Ens [19]	EffNet-B7	EffNet-B7	26.3	29.3
VL-PLM [57]	ViT-B/32	R-50	17.2	27.0
OV-DETR [53]	ViT-B/32	R-50	17.4	26.6
Rasheed <i>et al.</i> [41]	ViT-B/32	R-50	21.1	25.9
PromptDet [14]	ViT-B/32	R-50	21.4	25.3
<b>ViT based:</b>				
OWL-ViT [35]	ViT-H/14	ViT-H/14	23.3	35.3
OWL-ViT [35]	ViT-L/14	ViT-L/14	25.6	34.7
<b>RO-ViT (ours)</b>	ViT-B/16	ViT-B/16	28.0	30.2
<b>RO-ViT (ours)</b>	ViT-L/14	ViT-L/14†	31.4	34.0
<b>RO-ViT (ours)</b>	ViT-L/16	ViT-L/16	<b>32.1</b>	34.0

Table 1. **LVIS open-vocabulary object detection (mask APs)**. RO-ViT outperforms the best existing approach by **+5.8 AP<sub>r</sub>** on novel categories. When using ViT-Base and Large RO-ViT outperforms OWL-ViT based on ViT-Large by **+3.0** and **+6.5 AP<sub>r</sub>**, respectively. All methods use the same instance-level supervision from LVIS base categories for detection training. †: We use image size 860 during detection for a fair comparison with OWL-ViT.

are newly trained with the downstream detection dataset (*e.g.*, LVIS base categories). Existing approaches often miss novel/unlabeled objects in the object proposal stage because the proposals tend to classify them as background. To remedy this, we leverage recent advances in novel object proposal method [28] and adopt the localization quality-based objectness (*i.e.*, centerness score) instead of object-or-not binary classification score. Following [28], we use a single anchor per location and combine the predicted objectness score  $o_i$  with the ensemble detection score in Equation 1 to obtain the final OVD score as:  $S_i^{\text{OVD}} = o_i^\delta \cdot s_i^{\text{OVD}}$ .

Additionally, we replace the standard classifier and mask output layer with the normalized layers [46]. This is to reduce their scale variance over categories by L2-normalizing the weights  $w$  and features  $x$  as:  $f(x; w, b, \tau) = \frac{\tau}{\|w\|_2 \|x\|_2} w^T x + b$ , where  $\tau = 20$ . Although we do *not* have rare categories at training (*i.e.*, open-vocabulary setting), we empirically found it beneficial (see Sec. 4).

## 4. Experimental Results

**Pretraining details.** Our contrastive image-text pretraining is performed from scratch. We adopt the widely used ViT-B/16 and ViT-L/16 as the image encoder. The input image size is  $224 \times 224$  which results in  $14 \times 14$  positional embeddings with patch size  $16 \times 16$ . To generate the Cropped Positional Embedding (CPE), we first interpolate the positional embeddings to size  $64 \times 64$ . We randomly crop a re-

method	pretrained model	detector backbone	novel AP	AP
<b>ConvNet based:</b>				
ViLD [19]	ViT-B/32	R-50	27.6	51.3
OV-DETR [53]	ViT-B/32	R-50	29.4	52.7
<b>w/ pseudo box labels:</b>				
XPM <i>et al.</i> [25]	R-50	R-50	27.0	41.2
RegionCLIP [60] †	R-50x4	R-50x4	39.3	55.7
PromptDet [14]	ViT-B/32	R-50	26.6	50.6
VL-PLM [57]	ViT-B/32	R-50	34.4	53.5
Rasheed <i>et al.</i> [41] ‡	ViT-B/32	R-50	36.9	51.5
<b>w/ weak supervision:</b>				
Detic-CN2 [63]	ViT-B/32	R-50	24.6	32.4
<b>ViT based:*</b>				
<b>RO-ViT (ours)</b>	ViT-B/16	ViT-B/16	30.2	41.5
<b>RO-ViT (ours)</b>	ViT-L/16	ViT-L/16	33.0	47.7

Table 2. **COCO open-vocabulary object detection (box AP50)**. RO-ViT represents the first ViT-based approach and demonstrates a very competitive novel AP without using pseudo labeling or weak supervision. †: RegionCLIP uses an off-the-shelf RPN during its pretraining. ‡: Rasheed *et al.* uses an external MViT detector [34] during pretraining. \*: The other ViT-based method [35] report their results on LVIS only.

gion with the scale ratio in  $[0.1, 1.0]$ , and the aspect ratio in  $[0.5, 2.0]$ . The region crop is resized back to the size  $14 \times 14$  (*i.e.*, CPE), and is added to the patch embeddings. We use the global average pooling at the last ViT layer to obtain the image embedding. The text encoder is a 12-layer Transformer following [39, 51], and the maximum length of the input text as 64 tokens. Both image and text embeddings are L2 normalized and trained with focal loss. We train on the same image-text dataset as Chao *et al.* [26]. Unless otherwise specified, we use a batch size of 4096 for ablation and 16384 to compare with other methods. We use the AdamW optimizer with learning rate of  $5e-4$  and a linear warmup of 10k steps, and train for 500k iterations.

**Downstream detection details.** We train RO-ViT with base categories  $C_B$  for 46.1k/11.3k iterations with image size 1024, large scale jittering [16], batch size 256/128, the SGD optimizer with weight decay  $1e-4/1e-2$ , momentum 0.9 and an initial learning rate of 0.36/0.02 for LVIS/COCO datasets. The pretrained positional embeddings are bilinearly interpolated to adjust to the size of patch embeddings of higher resolution [12]. We set the backbone learning rate lower (*e.g.*,  $0.1 \times$ ) than the rest of the model to retain the pretraining knowledge during detection finetuning. We use the score combination of  $(\alpha, \beta, \delta) = (0.65, 0.3, 3)$  in Sec. 3.3. We use CLIP [39] prompt templates and take the average text embeddings of each category.

### 4.1. Open-vocabulary Object Detection

**LVIS benchmark.** We conduct evaluation on the LVIS dataset [20] which contains a large and diverse set of 1203

method	image backbone size	MS COCO (5K test set)						Flickr30K (1K test set)					
		image-to-text			text-to-image			image-to-text			text-to-image		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R10	R@1	R@5	R@10
CLIP [39]	302M	58.4	81.5	88.1	37.8	62.4	72.2	88.0	98.7	99.4	68.7	90.6	95.2
ALIGN [26]	408M	58.6	83.0	89.7	45.6	69.8	78.6	88.6	98.7	99.7	75.7	93.8	96.8
FLAVA [44]	86M	42.7	76.8	-	38.4	67.5	-	67.7	94.0	-	65.2	89.4	-
FILIP [50]	302M	61.3	84.3	90.4	45.9	70.6	79.3	89.8	99.2	99.8	75.0	93.4	96.3
Florence [52]	637M	64.7	85.9	-	47.2	71.4	-	90.9	99.1	-	76.7	93.6	-
CoCa-Large [51]	303M	65.4	85.6	91.4	50.1	73.8	81.8	91.4	99.2	<b>99.9</b>	79.0	95.1	97.4
CoCa [51]	1B	66.3	86.2	91.8	51.2	74.2	82.0	<b>92.5</b>	<b>99.5</b>	<b>99.9</b>	80.4	95.7	<b>97.7</b>
<b>RO-ViT (ours)</b>	303M	<b>68.9</b>	<b>87.8</b>	<b>92.2</b>	<b>51.8</b>	<b>75.0</b>	<b>83.0</b>	92.1	99.4	99.7	<b>80.7</b>	<b>96.1</b>	<b>97.7</b>

Table 3. **Zero-shot image-text retrieval results on COCO and Flickr30K benchmarks.** We compare with dual-encoder methods. We achieve state-of-the-art results on COCO benchmark. We outperform CoCa-Large with the same backbone by +3.5 / +1.7 R@1, and even surpass CoCa with 3× larger backbone (ViT-Giant) by +2.6 / +0.6 R@1 on image-to-text / text-to-image retrieval. We also match or outperform the state-of-the-art methods on Flickr benchmark

object categories suitable for open-vocabulary detection. Following the existing works [13, 19, 59], we use the frequent and common categories as base categories  $C_B$  for training, and hold out the rare categories as novel categories  $C_N$  for testing. Mask  $AP_r$  is the main benchmark metric. We report the mean over three runs following [19].

As shown in Table 1, our best model achieves 32.1  $AP_r$ , which significantly outperforms best existing ViT-based approach OWL-ViT [35] by **+6.5** points. Notably, our method with smaller ViT-B/16 backbone surpasses OWL-ViT with ViT-L/14 by **+3.0**  $AP_r$ . Compared to the best existing approach (ViLD-Ens with EffNet-B7 backbone), we outperform by +5.8  $AP_r$ . We note that RO-ViT is simply trained end-to-end with the cross-entropy loss without the employment of LVIS-tailored losses such as federated loss [35, 60, 63], weak supervision [63], or self-training [41, 57, 60].

**COCO benchmark.** We also present the comparison on the COCO benchmark where the setup uses 48 base categories for training and 17 novel categories for testing [19]. The main metric is AP50 of novel categories. Due to the smaller number of training categories, we observe a tendency to overfit to these categories. Unlike most competing methods, RO-ViT does *not* use any auxiliary objectives to mitigate overfitting such as pseudo box/mask labels [14, 25, 41, 57, 60], knowledge distillation [13, 19], weak supervision [63]. Still, Table 2 shows that RO-ViT is very competitive among the other methods trained with various sources. In addition, RO-ViT represents the first ViT-based method to be reported on this benchmark, as the other method OWL-ViT [35] only benchmarks on LVIS.

## 4.2. Image-Text Retrieval

Apart from evaluating region-level representation through open-vocabulary detection, we evaluate the *image-level* representation of RO-ViT in image-text retrieval through the MS-COCO [7] and Flickr30K [38] benchmarks.

method	backbone	AP	AP <sub>50</sub>	AP <sub>75</sub>
supervised [19]	R-50	25.6	38.6	28.0
ViLD [19]	R-50	11.8	18.2	12.6
DetPro [13]	R-50	12.1	18.8	12.9
<b>RO-ViT (ours)</b>	ViT-B/16	14.0	22.3	14.9
<b>RO-ViT (ours)</b>	ViT-L/16	<b>17.1</b>	<b>26.9</b>	<b>18.5</b>

Table 4. **Transfer detection on Objects365 (Box APs).** All models are trained on the LVIS base categories and tested on Objects365 dataset, without finetuning.

We further train RO-ViT (ViT-L/16), the same model as the last row of Table 1 and 4, for 40K iterations more at a higher resolution *e.g.* 448, following the standard of existing works [26, 51].

Table 3 shows the comparison with other dual-encoder methods on zero-shot image-to-text (I2T) and text-to-image (T2I) retrieval. Surprisingly, RO-ViT outperform all published works on the MS COCO benchmark, and is on par with the state of the art [51] on Flickr. Compared to CoCa [51] with the same backbone capacity (ViT-L), RO-ViT outperforms on 11 out of 12 image-text retrieval metrics. Our model has 303M parameters and achieves the best performance overall. This shows that our pretraining method not only improves the region-level representation for open-vocabulary detection but also the global image-level representation for retrieval.

## 4.3. Transfer Object Detection

We evaluate the generalization ability of RO-ViT in zero-shot transfer detection. We use the same detector trained on the LVIS base categories (Sec. 4.1) and test on Objects365-v1 validation split [43] following the setup of ViLD [13, 19]. We replace the LVIS with Objects365 vocabulary embeddings to perform the transfer detection without finetuning.

Table 4 summarizes the box AP scores in comparison with prior works. Our best model achieves 17.1 AP, which outperforms existing works ViLD by +5.3 AP and DetPro

pretraining method	CPE	focal	AP <sub>r</sub>	AP
base			21.4	26.6
no PE			18.7	25.2
SinCos PE			21.7	26.9
feat Crop-Resize			20.6	26.6
ours	✓		23.8 (+2.4)	27.4
ours		✓	21.9 (+0.5)	27.4
ours	✓	✓	24.3 (+2.9)	27.6

(a) **Pretraining strategy:** ‘base’ has the standard learnable positional embeddings (PE). No PE, randomly crop and resize the last feature map (feat Crop-Resize), and sinusoidal PE are either slightly worse or help marginally. Our proposed Cropped PE (CPE) improves by +2.4 AP<sub>r</sub>. Focal contrastive loss brings the gain further to +2.9 AP<sub>r</sub>.

backbone	loc.obj.	norm.lyr	AP <sub>r</sub>	AP
ViT-B/16			21.4	26.6
ViT-B/16	✓		23.4 (+2.0)	27.2
ViT-B/16		✓	22.0 (+0.6)	26.8
ViT-B/16	✓	✓	23.9 (+2.5)	27.5
ViT-L/16			25.1	30.1
ViT-L/16	✓	✓	27.2 (+1.9)	30.6

(d) **Downstream detector improvement:** Combining the localization quality-based objectness from the proposal stage to the detection scoring (‘loc.obj.’) improves by  $\sim 2$  AP<sub>r</sub>. Normalized classifier and mask output layer (‘norm.lyr’) gives additional gain.

frozen backbone	AP <sub>r</sub>	AP
base	9.7	12.2
CPE	16.2 (+6.5)	17.1
CPE + focal	16.5 (+6.8)	17.1

(b) **Frozen backbone study:** Only the newly added detector layers are trained in detection finetuning, to directly evaluate the pretrained features for open-vocabulary detection. Our proposed CPE improves the baseline by a large margin of +6.5 AP<sub>r</sub>.

bblr	AP <sub>r</sub>	AP
0.0	16.5	17.1
0.001	19.7	22.9
0.01	20.4	23.0
0.1	24.3	27.6
1.0	17.9	25.1

(c) Ablating **backbone finetuning learning rate ratio** (‘bblr’) w.r.t. added detector layers.

backbone	batch	imp.(d)	CPE + focal	AP <sub>r</sub>	AP
ViT-B/16	4096			21.4	26.6
ViT-B/16	4096	✓		23.9	27.5
ViT-B/16	4096	✓	✓	26.2 (+2.3)	29.2
ViT-B/16	16384	✓		26.4	30.3
ViT-B/16	16384	✓	✓	28.0 (+1.6)	30.2
ViT-L/16	4096	✓		27.2	30.6
ViT-L/16	4096	✓	✓	29.9 (+2.7)	31.1
ViT-L/16	16384	✓	✓	32.1	34.0

(e) **Scalability w.r.t model size and contrastive batch size:** The benefit of our ‘CPE + focal’ holds with larger model and batch size, improving AP<sub>r</sub> up to +2.7 points. ‘imp.(d)’ denotes the detector improvements in sub-table (d).

Table 5. **Ablation studies.** We follow LVIS open-vocabulary detection benchmark. We train on base (‘frequent’ + ‘common’) categories, test on novel (‘rare’) categories, and report AP<sub>r</sub>. We use ViT-B/16 backbone and contrastive batch size 4096 unless otherwise noted.

by +5.0 AP. Due to the different backbone capacity (R50 vs ViT), we note that this comparison is to showcase that RO-ViT can achieve strong cross-dataset detection generalization. For transfer detection, we assume all categories are novel and set  $\alpha, \beta = (0.0, 0.65)$  in Equation 1.

#### 4.4. Ablation Study

We provide ablation studies on the LVIS open-vocabulary detection and image-text retrieval benchmarks.

**Pretraining strategy.** Table 5a ablates our proposed methods in the contrastive image-text pretraining. We start with “base” that uses the common full-image positional embeddings and cross entropy loss. We explore different PE and observe that removing the positional embeddings hurts the performance (no PE), while sinusoidal PE (SinCos PE) yields a similar performance to the baseline. We also try the random crop and resize on the last backbone feature map (“feat Crop-Resize”) to use region features during pretraining but see no improvement. In contrast, our proposed Cropped Positional Embedding (CPE) offers a clear benefit of +2.4 AP<sub>r</sub> and focal contrastive loss (focal) yields additional gains. Our method achieves a gain of +2.9 AP<sub>r</sub> in total. This shows that CPE can learn strong region-level representations, and our focal contrastive learning helps preserving the knowledge about the pretrained image-text concepts,

backbone	focal	CPE	MS COCO		Flickr30K	
			I2T	T2I	I2T	T2I
ViT-B/16			59.1	42.5	84.8	70.9
ViT-B/16	✓		60.3	44.0	85.4	71.6
ViT-B/16	✓	✓	62.2	44.2	86.5	71.8
ViT-L/16	✓	✓	67.0	49.7	89.5	77.2

Table 6. **Pretraining evaluation on zero-shot image-text retrieval (Recall@1).** We evaluate the image-level representation of our pretrained model on COCO and Flickr30k retrieval tasks. We ablate the focal contrastive loss, Cropped Positional Embedding (CPE) and backbone size.

both of which are essential for open-vocabulary detection.

**Image-text retrieval.** In Table 6, we demonstrate that our proposed focal contrastive loss and Cropped Positional Embedding (CPE) both improve the image-text retrieval. We use a batch size of 16384, image size 224, and ViT-B/16 backbone as our baseline, and report Recall@1 metrics on MS COCO and Flickr datasets. We first observe that focal loss brings clear improvements on all metrics, and adding CPE brings further improvements with a total of +3.1 / +1.7 image-to-text (I2T) and +1.7 / +0.9 text-to-image (T2I) improvements on the MS COCO / Flickr datasets. By replacing ViT-B/16 with ViT-L/16, we observe marked improvements of +4.8 / +5.5 / +3.0 / +5.4 on these metrics, showing

that our recipe is highly scalable with backbone capacity.

**Frozen backbone study.** We investigate the effectiveness of the *frozen* backbone features in open-vocabulary detection in Table 5b. This evaluation setting is more rigorous for representation learning because finetuning an entire network evaluates not only the quality of the representations but also the initialization and optimization method, as discussed in Goyal *et al.* [18]. During the detection finetuning, all ViT backbone layers are frozen, and only the added detector layers (neck and heads) are trained. We observe frozen RO-ViT backbone improves the baseline by +6.8  $AP_r$ , a much larger margin compared to the full-finetuning setting. This study shows the region-aware representation of RO-ViT is critical for downstream detection tasks.

**Backbone learning rate ratio.** As discussed in Sec. 3.1, our open-vocabulary object detector depends on the pre-trained knowledge in the backbone to recognize novel categories beyond the detector annotations. Therefore, it is important to set the backbone learning rates lower than the rest of the detector layers to prevent forgetting in the finetuning process. We present ablations on the learning rate ratio between the backbone and the detector layers in Table 5c. We found  $0.1\times$  to be the sweet spot. Higher ratios lead to forgetting and hurts the  $AP_r$ , and lower ratios hurt the ability to adapt to detection tasks and hurts AP overall.

**Downstream detector improvements.** In addition, Table 5d demonstrates our improvements in the downstream detector. Learning the objectness by localization quality, *i.e.*, centerness (“loc.obj.”) [28] and leveraging it into detection score improves  $AP_r$  by a clear gain of +1.9 points compared to using the conventional binary-classification in the proposal. This indicates the novel, unlabeled objects are indeed often missed in the proposal stage, which may limit the final open-vocabulary detection performance. Localization-quality based objectness helps remedy this. The use of normalized activation in the last layers of the classifier and mask heads additionally improves the performance.

**Scalability with respect to model size and batch size.** While increasing the model capacity and batch size have been beneficial for contrastive learning in general [26, 37, 39, 51], we study their benefits in the downstream open-vocabulary detection in Table 5e. Note the detector improvements presented in Sec. 3.3 are applied. We first observe increasing the batch size from the default 4096 to 16384 shows gains of +2.1~2.5  $AP_r$  for both ViT-B/L. Then, we notice upgrading ViT-B to ViT-L brings +3.3~3.7  $AP_r$  gain across different batch sizes. Last but not least, the gain of our proposed “CPE + focal” pretraining is consistent among all settings by improving +2.2~2.7  $AP_r$ . Putting everything together, RO-ViT best model achieves 32.1  $AP_r$ .

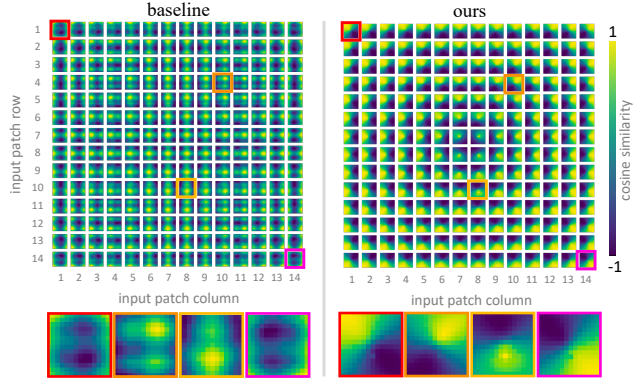


Figure 3. **Visualization of learned positional embeddings.** Tiles show the cosine similarity between the positional embedding of the patch (at the indicated row-column position) and the positional embeddings of all other patches. ViT-B/16 backbone is used.

#### 4.5. Visualization of Positional Embeddings

In Fig. 3, we visualize and compare the learned positional embeddings of RO-ViT with the baseline, based on ViT-B/16 backbone. Each tile is the cosine similarity between positional embeddings of one patch and all other patches. The brightness patterns show in both models, closer patches have more similar positional embeddings, indicating the encoded distance and locality within the image.

We observe a few differences between RO-ViT positional embeddings and the baselines. First, RO-ViT forms more distinct clusters at different patch locations, *i.e.*, ours shows symmetrical global pattern around the center patch, while the baseline has similar patterns on opposite ends of the image (*e.g.*, the pattern in top-left patch is similar to the bottom-right patch). Also, the brightness patterns of RO-ViT tend to be more concentrated and strong (*i.e.*, near 1 or -1). To summarize, the visualization indicates RO-ViT positional embeddings acquire more structure and symmetry than the baseline in the pretraining process.

## 5. Conclusion

We present RO-ViT—a contrastive image-text pretraining framework to bridge the gap between image-level pretraining and open-vocabulary detection finetuning. Our methods are simple, scalable, and easy to apply to any contrastive backbones with minimal computation overhead and no increase in parameters. Experiments show that our RO-ViT achieves the state-of-the-art on LVIS open-vocabulary detection benchmark. Moreover, RO-ViT achieves state of the art on 9 out of 12 metrics of image-text retrieval benchmark, showing that the learnt representation is not only beneficial at region-level but also highly effective on image-level. We hope this study can help the research on open-vocabulary detection from the perspective of image-text pretraining which can benefit both region-level and image-level tasks.



## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning, 2022. [1](#)
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6836–6846, 2021. [1](#)
- [3] Yutong Bai, Xinlei Chen, Alexander Kirillov, Alan Yuille, and Alexander C. Berg. Point-level region contrast for object detection pre-training. In *CVPR*, pages 16061–16070, June 2022. [2](#)
- [4] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *ECCV*, 2018. [2](#)
- [5] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. [2](#)
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, October 2021. [2](#)
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. In <https://arxiv.org/abs/1504.00325>, 2015. [6](#)
- [8] Xinlei Chen and Abhinav Gupta. Webly supervised learning of convolutional networks. In *ICCV*, 2015. [2](#)
- [9] Berkan Demirel, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis. Zero-shot object detection by hybrid region embedding. In *BMVC*, 2018. [2](#)
- [10] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *CVPR*, 2021. [2](#)
- [11] Santosh K Divvala, Ali Farhadi, and Carlos Guestrin. Learning everything about anything: Webly-supervised visual concept learning. In *CVPR*, 2014. [2](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [1](#), [2](#), [3](#), [5](#)
- [13] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *CVPR*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#)
- [14] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Promptdet: Towards open-vocabulary detection using uncurated images. In *European Conference on Computer Vision*, pages 701–717. Springer, 2022. [5](#), [6](#)
- [15] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, 2013. [2](#)
- [16] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. [5](#)
- [17] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European Conference on Computer Vision*, pages 540–557. Springer, 2022. [3](#)
- [18] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. [8](#)
- [19] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#)
- [20] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. [5](#)
- [21] Nasir Hayat, Munawar Hayat, Shafin Rahman, Salman Khan, Syed Waqas Zamir, and Fahad Shahbaz Khan. Synthesizing the unseen for zero-shot object detection. In *ACCV*, 2020. [2](#)
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, June 2022. [2](#)
- [23] Xiangteng He and Yuxin Peng. Fine-grained image classification via combining vision and language. In *CVPR*, 2017. [2](#)
- [24] Olivier J. Hénaff, Skanda Koppula, Jean-Baptiste Alayrac, Aaron van den Oord, Oriol Vinyals, and João Carreira. Efficient visual pretraining with contrastive detection. In *ICCV*, pages 10086–10096, October 2021. [2](#)
- [25] Dat Huynh, Jason Kuen, Zhe Lin, Jiuxiang Gu, and Ehsan Elhamifar. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7020–7031, 2022. [5](#), [6](#)
- [26] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [27] Armand Joulin, Laurens van der Maaten, Allan Jabri, and Nicolas Vasilache. Learning visual features from large weakly supervised data. In *ECCV*, 2016. [2](#)
- [28] 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. [2](#), [5](#), [8](#)

- [29] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. F-vm: Open-vocabulary object detection upon frozen vision and language models. *arXiv preprint arXiv:2209.15639*, 2022. 1, 3
- [30] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 3
- [31] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, 2022. 3
- [32] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *ECCV*, 2022. 1, 3
- [33] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 4
- [34] Muhammad Maaz, Hanoona Rasheed, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Ming-Hsuan Yang. Class-agnostic object detection with multi-modal transformer. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, pages 512–531. Springer, 2022. 5
- [35] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. In *ECCV*, 2022. 2, 3, 5, 6
- [36] Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. 2014. 2
- [37] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Hanxiao Liu, Adams Wei Yu, Minh-Thang Luong, Mingxing Tan, and Quoc V. Le. Combined scaling for zero-shot transfer learning. *CoRR*, abs/2111.10050, 2021. 2, 8
- [38] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, pages 2641–2649, 2015. 6
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3, 4, 5, 6, 8
- [40] Shafin Rahman, Salman Khan, and Nick Barnes. Improved visual-semantic alignment for zero-shot object detection. In *AAAI*, 2020. 2
- [41] Hanoona Rasheed, Muhammad Maaz, Muhammad Uzair Khattak, Salman Khan, and Fahad Shahbaz Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. *arXiv preprint arXiv:2207.03482*, 2022. 1, 2, 3, 5, 6
- [42] Mert Bulent Sariyildiz, Julien Perez, and Diane Larlus. Learning visual representations with caption annotations. In *ECCV*, 2020. 2
- [43] 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 *ICCV*, 2019. 6
- [44] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *CVPR*, pages 15638–15650, 2022. 6
- [45] Josiah Wang, Katja Markert, Mark Everingham, et al. Learning models for object recognition from natural language descriptions. In *BMVC*, 2009. 2
- [46] Jiaqi Wang, Wenwei Zhang, Yuhang Zang, Yuhang Cao, Jiangmiao Pang, Tao Gong, Kai Chen, Ziwei Liu, Chen Change Loy, and Dahua Lin. Seesaw loss for long-tailed instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9695–9704, 2021. 5
- [47] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. *arXiv preprint arXiv:2108.10904*, 2021. 1
- [48] Fangyun Wei, Yue Gao, Zhirong Wu, Han Hu, and Stephen Lin. Aligning pretraining for detection via object-level contrastive learning. In *NeurIPS*, 2021. 2
- [49] Tete Xiao, Colorado J Reed, Xiaolong Wang, Kurt Keutzer, and Trevor Darrell. Region similarity representation learning. In *ICCV*, pages 10539–10548, October 2021. 2
- [50] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *ICLR*, 2021. 6
- [51] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *TMLR*, 2022. 5, 6, 8
- [52] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luowei Zhou, and Pengchuan Zhang. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, November 2021. 6
- [53] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. *arXiv preprint arXiv:2203.11876*, 2022. 5
- [54] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *CVPR*, 2021. 1, 3
- [55] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *CVPR*, 2022. 2
- [56] Bowen Zhang, Zhi Tian, Quan Tang, Xiangxiang Chu, Xiaolin Wei, Chunhua Shen, and Yifan Liu. Segvit: Semantic

- segmentation with plain vision transformers. *arXiv preprint arXiv:2210.05844*, 2022. 1
- [57] Shiyu Zhao, Zhixing Zhang, Samuel Schuster, Long Zhao, Anastasis Sathopoulos, Manmohan Chandraker, Dimitris Metaxas, et al. Exploiting unlabeled data with vision and language models for object detection. In *ECCV*, 2022. 1, 2, 3, 5, 6
- [58] Ye Zheng, Ruoran Huang, Chuanqi Han, Xi Huang, and Li Cui. Background learnable cascade for zero-shot object detection. In *ACCV*, 2020. 2
- [59] Yiwu Zhong, Jing Shi, Jianwei Yang, Chenliang Xu, and Yin Li. Learning to generate scene graph from natural language supervision. In *ICCV*, 2021. 2, 6
- [60] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, and Jianfeng Gao. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. 1, 2, 3, 5, 6
- [61] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, 2022. 3
- [62] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 3
- [63] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *ECCV*, 2022. 1, 3, 5, 6
- [64] Pengkai Zhu, Hanxiao Wang, and Venkatesh Saligrama. Don't even look once: Synthesizing features for zero-shot detection. In *CVPR*, 2020. 2