



# Open Vocabulary Object Detection with Pseudo Bounding-Box Labels

Mingfei Gao<sup>✉</sup>, Chen Xing, Juan Carlos Niebles, Junnan Li, Ran Xu, Wenhao Liu, and Caiming Xiong

Salesforce Research, Palo Alto, USA  
{mingfei.gao,cxing,jniebles,junnan.li,ran.xu,wenhao.liu,cxiong}@salesforce.com

**Abstract.** Despite great progress in object detection, most existing methods work only on a limited set of object categories, due to the tremendous human effort needed for bounding-box annotations of training data. To alleviate the problem, recent open vocabulary and zero-shot detection methods attempt to detect novel object categories beyond those seen during training. They achieve this goal by training on a pre-defined base categories to induce generalization to novel objects. However, their potential is still constrained by the small set of base categories available for training. To enlarge the set of base classes, we propose a method to automatically generate pseudo bounding-box annotations of diverse objects from large-scale image-caption pairs. Our method leverages the localization ability of pre-trained vision-language models to generate pseudo bounding-box labels and then directly uses them for training object detectors. Experimental results show that our method outperforms the state-of-the-art open vocabulary detector by 8% AP on COCO novel categories, by 6.3% AP on PASCAL VOC, by 2.3% AP on Objects365 and by 2.8% AP on LVIS. *Code is available here.*

**Keywords:** Open vocabulary detection · Pseudo bounding-box labels


## 1 Introduction

Object detection [11, 19, 20, 27] is a core task in computer vision that has considerably advanced with the adoption of deep learning and continues to attract significant research effort [26, 31, 33]. Current deep object detection methods achieve astonishing performance when learning a pre-defined set of object categories that have been annotated in a large number of training images (PASCAL VOC [6], COCO [17]). Unfortunately, their success is still limited to detecting

M.Gao and C. Xing contributed equally.

---

**Supplementary Information** The online version contains supplementary material available at [https://doi.org/10.1007/978-3-031-20080-9\\_16](https://doi.org/10.1007/978-3-031-20080-9_16).



**Fig. 1.** Previous methods (left) rely on human-provided box-level annotations of predefined base classes during training and attempt to generalize to objects of novel classes during inference. Our method (right) generates pseudo bounding-box annotations from large-scale image-caption pairs by leveraging the localization ability of pre-trained vision-language (VL) models. Then, we utilize them to improve our open vocabulary object detector. Compared to the human annotations of a fixed/small set of base classes, our pseudo bounding-box label generator easily scales to a large set of diverse objects from the large-scale image-caption dataset, thus is able to achieve better detection performance on novel objects compared to previous methods

a small number of object categories (e.g., 80 categories in COCO). One reason is that most detection methods rely on supervision in the form of instance-level bounding-box annotations, hence requiring very expensive human labeling efforts to build training datasets. Furthermore, when we need to detect objects from a new category, one has to further annotate a large number of bounding-boxes in images for this new object category.

Two families of recent work have attempted to reduce the need of annotating new object categories: zero-shot object detection and open vocabulary object detection. In zero-shot detection methods [1, 22], object detection models are trained on *base* object categories with human-provided bounding box annotations to promote their generalization ability on *novel* object categories, by exploiting correlations between base and novel categories. These methods can alleviate the need for large amounts of human labeled data to some extent. Building on top of such methods, open vocabulary object detection [34] aims to improve the detection performance of novel objects with the help of image captions. However, the potential of existing zero-shot and open vocabulary methods is constrained by the small size of the base category set at training, due to the high cost of acquiring large-scale bounding-box annotations of diverse objects. As a result, it is still challenging for them to generalize well to diverse objects of novel categories in practice.

A potential avenue for improvement is to enable open vocabulary detection models to utilize a larger set of base classes of diverse objects by reducing the requirement of manual annotations. In this paper we ask: can we automatically generate bounding-box annotations for objects at scale using existing resources? Can we use these generated annotations to improve open vocabulary detection? The most recent progress on vision-language pre-training gives us some hope. Vision-language models [12, 15, 16, 21, 28] are pre-trained with large scale image-

caption pairs from the web. They show amazing zero-shot performance on image classification, as well as promising results on tasks related to text-visual region alignment, such as referring expressions, which implies strong localization ability.

Motivated by these observations, we improve open vocabulary object detection using pseudo bounding-box annotations generated from large-scale image-caption pairs, by taking advantage of the localization ability of pre-trained vision-language models. As shown in Fig. 1, we design a pseudo bounding-box label generation strategy to automatically obtain pseudo box annotations of a diverse set of objects from existing image-caption datasets. Specifically, given a pre-trained vision-language model and an image-caption pair, we compute an activation map (Grad-CAM [24]) in the image that corresponds to an object of interest mentioned in the caption. We then convert the activation map into a pseudo bounding-box label for the corresponding object category. Our open vocabulary detector is then directly trained with supervision of these pseudo labels. Our detector can also be fine-tuned with human-provided bounding boxes if they are available. Since our method for generating pseudo bounding-box labels is fully automated with no manual intervention, the size and diversity of the training data, including the number of training object categories, can be largely increased. This enables our approach to outperform existing zero-shot/open vocabulary detection methods that are trained with a limited set of base categories.

We evaluate the effectiveness of our method by comparing with the state-of-the-art (SOTA) zero-shot and open vocabulary object detectors on four widely used datasets: COCO, PASCAL VOC, Objects365 and LVIS. Experimental results show that our method outperforms the best open vocabulary detection method by 8% mAP on novel objects on COCO, when both of the methods are fine-tuned with COCO base categories. Moreover, we surprisingly find that even when not fine-tuned with COCO base categories, our method can still outperform fine-tuned SOTA baseline by 3% mAP. We also evaluate the generalization performance of our method to other datasets. Experimental results show that under this setting, our method outperforms existing approaches by 6.3%, 2.3% and 2.8% on PASCAL VOC, Objects365 and LVIS, respectively.

Our contributions are summarized as follows: (1) We propose an open vocabulary object detection method that can train detectors with pseudo bounding-box annotations generated from large-scale image-caption pairs. To the best of our knowledge, this is the first work which enables open vocabulary object detection using pseudo labels during training. (2) We introduce a pseudo label generation strategy using the existing pre-trained vision-language models. (3) With the help of pseudo labels, our method largely outperforms the SOTA methods. Moreover, when trained with only pseudo labels, our method achieves higher performance than the SOTA that rely on training with manual bounding-box annotations.

## 2 Related Work

**Object detection** aims at localizing objects in images. Traditional detection methods are supervised using human-provided bounding box annotations. Two-


stage detection methods [8, 11, 23] are one of the most popular frameworks. These methods generate object proposals in the first stage and classify these proposals to different categories in the second stage. Weakly supervised object detectors seek to relieve such heavy human annotation burden by using image-level labels such as image-level object categories [2, 29], captions [32] and object counts [7] for training. Although these approaches show promising performance, they only support objects in a fixed set of categories. Whenever one needs to detect objects from a new category, they have to collect and manually annotate instances from the new category and retrain the detector.

**Open vocabulary and zero-shot object detection** target at training an object detector with annotations on base object classes to generalize to novel object classes during inference. Most zero-shot detection methods achieve this level of generalization by aligning the visual and the text representation spaces for objects from base classes during training [1, 22, 36], and inferring novel objects during inference by exploiting correlation between base and novel objects. Recent methods encourage the visual-semantic alignment for novel objects by different strategies such as synthesizing visual representations of novel classes [35, 36] or utilizing existing object names semantically similar to their names [22]. Joseph et al. introduce OREO [13] to incrementally learn unknown objects based on contrastive clustering and energy based unknown identification. To further improve the zero-shot performance on novel object categories, Zareian et al. [34] propose open vocabulary object detection that transfers knowledge from a pre-trained vision-language model by initializing their detector with parameters of the image encoder of the vision-language model. This strategy improves the state-of-the-art by a large margin. Gu et al. [9] propose ViLD which achieves good zero-shot performance by distilling knowledge from a large-scale vision-language model (CLIP [21]). However, all these methods are trained with a small set of base categories that have bounding-box labels since acquiring bounding-box annotations of diverse objects in large-scale training data is expensive. In practice, if a novel category at inference is very different from base categories, it is still challenging for these methods to generalize well to such novel objects. In contrast, our method generates pseudo box labels for diverse objects from large-scale image-caption pairs and use them to train the detector. When human-provided box annotations are available, our framework has the flexibility to utilize them.

**Vision-language pre-training models** are trained with large-scale image-caption pairs. They have been successful not only in image-language tasks such as image retrieval, VQA and referring expression [12, 15, 16, 28], but also in pure image tasks such as zero-shot image classification [21]. Recent methods typically utilize a multi-modal module to encourage the interaction between the vision and language modalities [15, 16, 28], which may implicitly encode the word-to-region localization information inside the model. We take advantage of their localization ability and design a strategy to obtain pseudo bounding-box labels of a large and diverse set of objects from the large-scale image-caption datasets. With this strategy, we largely improve open vocabulary object detection.

### 3 Our Approach

Our framework contains two components: a pseudo bounding-box label generator and an open vocabulary object detector. Our pseudo label generator automatically generates bounding-box labels for a diverse set of objects by leveraging a pre-trained vision-language model. We then train our detector directly with the generated pseudo labels. For fair comparison with existing open vocabulary detection methods, when base object categories with human annotated bounding-boxes are available, we can also fine-tune our trained detector with such data.



**Fig. 2.** Illustration of our pseudo box label generation process. The input to the system is an image-caption pair. We use image and text encoders to extract the visual and text embeddings of the image and its corresponding caption. We then obtain multi-modal features by image-text interaction via cross-attention. We maintain objects of interest in our pre-defined object vocabulary. For each object of interest embedded in the caption (for example, *racket* in this figure), we use Grad-CAM to visualize its activation map in the image. This map indicates the contribution of the image regions to the final representation of the object word. Finally, we determine the pseudo bounding-box label of the object by selecting the object proposal that has the largest overlap with the activation

#### 3.1 Generating Pseudo Box Labels

Figure 2 illustrates the overall procedure of our pseudo label generation. Our goal is to generate pseudo bounding-box annotations for objects of interest in an image, by leveraging the implicit alignment between regions in the image and words in its corresponding caption in a pre-trained vision-language model. Before diving into our method, we first briefly introduce the general structure of the recent vision-language models.

An image  $\mathbf{I}$  and its corresponding caption,  $\mathbf{X} = \{x_1, x_2, \dots, x_{N_T}\}$ , are the inputs to the model, where  $N_T$  is the number of words in the caption (including [CLS] and [SEP]). An image encoder is used to extract image features  $\mathbf{V} \in \mathbb{R}^{N_V \times d}$  and a text encoder is utilized to get text representations  $\mathbf{T} \in \mathbb{R}^{N_T \times d}$ .  $N_V$  is the number of region representations of the image. Moreover, a multi-modal encoder with  $L$  consecutive cross-attention layers is often employed to fuse the information from both image and text encoders. In the  $l$ -th cross-attention layer, the

interaction of an object of interest  $x_t$  in the caption with the image regions is shown in Eqs. 1 and 2, where  $\mathbf{A}_t^l$  denotes the corresponding visual attention scores at the  $l$ -th cross-attention layer.  $\mathbf{h}_t^{l-1}$  indicates the hidden representations obtained from the previous  $(l - 1)$ -th cross-attention layer and  $\mathbf{h}_t^0$  is the representation of  $x_t$  from the text encoder.

$$\mathbf{A}_t^l = \text{Softmax}\left(\frac{\mathbf{h}_t^{l-1}\mathbf{V}^T}{\sqrt{d}}\right), \quad (1)$$

$$\mathbf{h}_t^l = \mathbf{A}_t^l \cdot \mathbf{V}. \quad (2)$$

From these equations, a cross-attention layer measures the relevance of the visual region representations with respect to a token in the input caption, and calculates the weighted average of all visual region representations accordingly. As a result, the visual attention scores  $\mathbf{A}_t^l$  can directly reflect how important different visual regions are to token  $x_t$ . Therefore, we visualize the activation maps based on the attention scores to locate an object in an image given its name in the caption.

We use Grad-CAM [24] as the visualization method and follow its original setting to take the final output  $s$  from the multi-modal encoder, and calculate its gradient with respect to the attention scores.  $s$  is a scalar that represents the similarity between the image and its caption. Specifically, the final activation map  $\Phi_t$  of the image given an object name  $x_t$  is calculated as

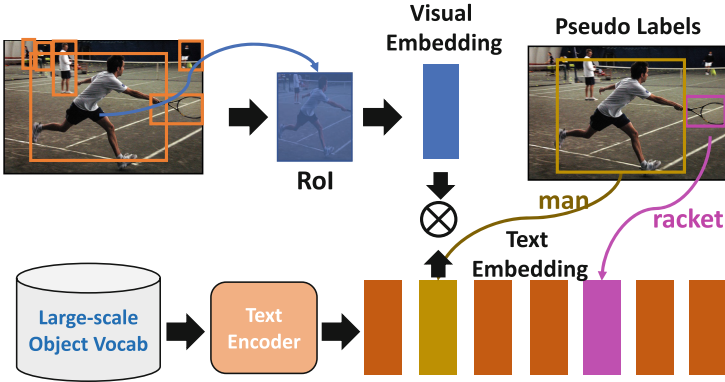
$$\Phi_t = \mathbf{A}_t^l \cdot \max\left(\frac{\partial s}{\partial \mathbf{A}_t^l}, 0\right). \quad (3)$$

If there are multiple attention heads in a cross-attention layer, we average the activation map  $\Phi_t$  from all heads as the final activation map.

After we get an activation map of an object of interest in the caption using this strategy, we draw a bounding box covering the activated region as the pseudo label of the category. We adopt existing proposal generators, e.g., [11, 30] to generate proposal candidates  $\mathbf{B} = \{b_1, b_2, \dots, b_K\}$  and select the one overlapping the most with  $\Phi_t$ :

$$\hat{b} = \arg \max_i \frac{\sum_{b_i} \Phi_t(b_i)}{\sqrt{|b_i|}}, \quad (4)$$

where  $\sum_{b_i} \Phi_t(b_i)$  indicates summation of the activation  $\Phi_t$  map within a box proposal and  $|b_i|$  indicates the proposal area. In practice, we maintain a list of objects of interest (referred as object vocabulary) during training and get pseudo bounding-box annotations for all objects in the training vocabulary (see Sect. 4.1 for details). Figure 4 shows some examples of the activation maps. As we can see, the activated regions correspond well with the relevant regions. The generated bounding boxes are of good quality. When they are directly used to train an open vocabulary object detector, the object detector significantly outperforms the current SOTA open-vocabulary/zero-shot object detectors.



**Fig. 3.** Illustration of our detector. An image is processed by a feature extractor followed by a region proposal network. Region-based features are then calculated by applying RoI pooling/RoI align over region proposals and the corresponding visual embeddings are obtained. Similarity of the visual and text embeddings of the same object are encouraged during training

### 3.2 Open Vocabulary Object Detection with Pseudo Labels

After we get pseudo bounding-box labels, we can use them to train an object detector. Since our pseudo-label generation is disentangled from detector training process, our framework can accommodate detectors with any architecture. In this work, we focus on the open vocabulary scenario where a detector aims at detecting arbitrary objects during inference.

A general open vocabulary detection system [34] is shown in Fig. 3. A feature map is extracted from an input image using a feature extractor based on which object proposals are generated. Then, region-based visual embeddings,  $\mathbf{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{N_r}\}$ , are obtained by RoI pooling/RoI align [11] followed by a fully connected layer, where  $N_r$  denotes the number of regions. In the meanwhile, text embeddings,  $\mathbf{C} = \{\mathbf{bg}, \mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_{N_c}\}$ , of object candidates from the object vocabulary are acquired by a pre-trained text encoder, where  $N_c$  is the training object vocabulary size and  $\mathbf{bg}$  indicates “background” that matches irrelevant visual regions. The goal of the open vocabulary object detector is to pull close the visual and text embeddings of the same objects and push away those of different objects. The probability of  $\mathbf{r}_i$  matches  $\mathbf{c}_j$  is calculated as

$$p(\mathbf{r}_i \text{ matches } \mathbf{c}_j) = \frac{\exp(\mathbf{r}_i \cdot \mathbf{c}_j)}{\exp(\mathbf{r}_i \cdot \mathbf{bg}) + \sum_k \exp(\mathbf{r}_i \cdot \mathbf{c}_k)}, \tag{5}$$

where text embeddings  $\mathbf{C}$  is fixed during training. The cross entropy loss is used to encourage the matching of positive pairs and discourage the negative ones.

During inference, given a group of object classes of interest, a region proposal will be matched to the object class if its text embedding has the smallest distance to the visual embedding of the region compared to all object names in the

vocabulary. This strategy is similar to other zero-shot/open vocabulary detection methods. To perform a fair comparison to prior work, we also adopt Mask-RCNN as the base of our open vocabulary detector. We set  $\mathbf{bg}=\mathbf{0}$  and include objectness classification, objectness box regression and class-agnostic box regression losses following [34].

## 4 Experiments

### 4.1 Datasets and Object Vocabulary for Training

**Training Datasets.** In our method, we generate pseudo bounding-box annotations of diverse objects from a combination of existing image-caption datasets including COCO Caption [3], Visual-Genome [14], and SBU Caption [18]. Our final dataset for pseudo label generation and detector training contains about one million images.

**Object Vocabulary.** When we generate pseudo labels for object categories from the aforementioned dataset, our default object vocabulary is constructed by the union of all the object names in COCO, PASCAL VOC, Objects365 and LVIS, resulting in 1,582 categories. We would also like to note that since our method doesn't require extra human annotation efforts, our training object vocabulary can be easily augmented.

### 4.2 Evaluation Benchmarks

**Baselines.** We compare with recent zero-shot and open vocabulary methods [1, 22, 34, 36]. Among the baselines, *Zareian et al.* [34] is the SOTA method, thus, is treated as our major baseline.

**Generalized Setting in COCO.** Most existing methods are evaluated under this setting proposed in [1]. COCO detection training set is split to base set containing 48 base/seen classes and target set including 17 novel/unseen classes. Base classes are used for training. During inference, models predict object categories from the union of base and novel classes. The performance of models is evaluated using the mean average precision over the novel classes.

Our method can be trained using the large-scale dataset with the generated pseudo labels. To perform a fair comparison with baselines, we fine-tune our detector using COCO base categories following their setup. Moreover, we also report our method's performance without fine-tuning on COCO base categories.

**Generalization Ability to Other Datasets.** We are interested in measuring the generalization ability of a model to other datasets that the model is not trained on. Therefore, we evaluate our method and the strongest baseline (both are fine-tuned using COCO base classes) on PASCAL VOC [5] test set, Objects365 v2 [25] validation set and LVIS [10] validation set<sup>1</sup>. PASCAL VOC

<sup>1</sup> We use LVIS v0.5, since the validation set of LVIS v1.0 contains images from *COCO train 2017* which our method may finetune on in some experiments.



**Table 1.** Performance on COCO dataset. Our method outperforms all the previous approaches when all models are fine-tuned on COCO base categories. When our method is not fine-tuned, it still outperforms other fine-tuned baselines

Method	Fine-tuned with Box Anno. on COCO <b>Base</b> categories	Generalized setting		
		Novel AP	Base AP	Overall AP
Bansal et al. [1]	Yes	0.3	29.2	24.9
Zhu et al. [36]	Yes	3.4	13.8	13.0
Rahman et al. [22]	Yes	4.1	35.9	27.9
Zareian et al. [34]	Yes	22.8	46.0	39.9
Our method	Yes	<b>30.8</b>	46.1	42.1
Our method	No	<b>25.8</b>	–	–

**Table 2.** Generalization performances to other datasets. Our method has better generalization performance to other datasets compared to *Zareian et al.*

Method	Fine-tuned on COCO <b>Base</b> categories	PASCAL VOC	Objects365	LVIS
Zareian [34]	Yes	52.9	4.6	5.2
Our method	Yes	<b>59.2</b>	<b>6.9</b>	<b>8.0</b>
Our method	No	44.4	<b>5.1</b>	<b>6.5</b>

is a widely used dataset by traditional object detection methods which contains 20 object categories. Objects365 and LVIS are datasets include 365 and 1,203 object categories, respectively, which makes them very challenging. When evaluating on each of these datasets (PASCAL VOC, Objects365 and LVIS), visual regions will be matched to one of the object categories (including background) of each dataset during inference.

**Evaluation Metric.** Following prior work [34], we use the standard metric in object detection tasks, i.e., mean average precision over classes, and set the IoU threshold to 0.5.

### 4.3 Implementation Details

In our main experiment, we use the ALBEF model pre-trained with 14M data<sup>2</sup> as our vision-language model for pseudo label generation. We follow the default setting of ALBEF, unless otherwise noted. The cross-attention layer used for Grad-CAM visualization is set to  $l = 8$  in Eq. 3. We conduct our main experiments using ALBEF because of its good performance in object grounding when image captions are present. Note that other pre-trained vision-language models

<sup>2</sup> <https://github.com/salesforce/ALBEF> (BSD-3-Clause License).

**Table 3.** Effect of proposal quality. All models are fine-tuned on COCO base classes. Better proposals lead to better detection performance

Method	Proposal generator	Generalized COCO novel	PASCAL VOC	Objects365	LVIS
Zareian et al. [34]	–	22.8	52.9	4.6	5.2
Our method	Selective Search	<b>28.5</b>	<b>53.0</b>	<b>5.5</b>	<b>5.9</b>
	Mask-RCNN	<b>30.8</b>	<b>59.2</b>	<b>6.9</b>	<b>8.0</b>

**Table 4.** Effect of different vision-language models. All models are fine-tuned on COCO base classes. Better VL model leads to better detection performance

Method	VL model	Generalized COCO novel	PASCAL VOC	Objects365	LVIS
Zareian et al. [34]	–	22.8	52.9	4.6	5.2
Our method	LXMERT [28]	<b>27.0</b>	<b>56.5</b>	<b>5.5</b>	<b>6.4</b>
	ALBEF [15]	<b>30.8</b>	<b>59.2</b>	<b>6.9</b>	<b>8.0</b>

can also fit our framework without major modifications or adding additional constraints on detector training. We conduct an ablation study to show our method’s performance when another pre-trained vision-language model (LXMERT [28]) is employed in Sect. 4.5. As a default option, we use a off-the-shelf Mask-RCNN with ResNet-50 trained on COCO 2017 train set as our proposal generator. To ensure there is no labels of novel categories leaking to our model, we have excluded the novel categories when training the proposal generator. We also show our results with an unsupervised proposal generator, selective search [30], in Sect. 4.5.

We use Mask-RCNN with ResNet-50 as the base of our open vocabulary detector and keep following the default settings here<sup>3</sup>. We utilize the pre-trained CLIP (ViT-B/32) text encoder to extract text embeddings of objects in our vocabulary and use the text prompts provided in [9] to ensemble the text representation. We use a batch size of 64 with learning rate of 0.02 when training the open vocabulary detector using our large-scale dataset with pseudo bounding-box labels, and a batch size of 8 with base learning rate of 0.0005 when optionally fine-tuning on COCO base classes. Models are optimized using SGD. The weight decay is set to 0.0001. The maximum iteration number is 150,000 and the learning rate is updated by a decreasing factor of 0.1 at 60,000 and 120,000 iterations.

#### 4.4 Experimental Results

As shown in Table 1, when fine-tuned using COCO base categories same as our baselines, our method outperforms our strongest baseline (*Zareian et al.*) largely

<sup>3</sup> <https://github.com/alirezazareian/ovr-cnn> (MIT License).

**Table 5.** Performance of our method when using vocabularies of different sizes for pseudo label generation.  $\mathbb{V}^-$  and  $\mathbb{V}$  contain 65 and 1.5k+ categories, respectively

Methods	Vocabulary	Generalized setting		
		Novel AP (17)	Base AP (48)	Overall AP (65)
Zareian et al. [34]	–	22.8	46.0	39.9
Our method	$\mathbb{V}^-$ (65)	<b>29.7</b>	44.3	40.4
	$\mathbb{V}$ (1.5k+)	<b>30.8</b>	46.1	42.1

**Table 6.** Performance of our method when trained with pseudo labels generated from different amounts of data. All models are fine-tuned using COCO base categories. Our performance improves when trained with pseudo labels of more data

Methods	Data of pseudo label generation	Generalized setting		
		Novel AP (17)	Base AP (48)	Overall AP (65)
Zareian et al. [34]	–	22.8	46.0	39.9
Our method	COCO Cap	<b>29.1</b>	44.4	40.4
	COCO Cap, VG, SBU	<b>30.8</b>	46.1	42.1

by 8%. When not fine-tuned using COCO base categories and only trained with generated pseudo labels, our method achieves 25.8% AP on the novel categories, which still significantly outperforms the SOTA method (*Zareian et al.*) by 3%.


Generalization ability to a wide range of datasets is also important for an open vocabulary object detector, since it makes a detector directly usable as an out-of-the-box method in the wild. Table 2 shows the generalization performance of detectors to different datasets, where both our method and our baseline are not trained using these datasets. Since objects365 and LVIS have a large set of diverse object categories, evaluation results on these datasets would be more representative to demonstrate the generalization ability. Results show that our method achieves better performance than *Zareian et al.* on all three datasets when both of the methods are fine-tuned with COCO base categories. Our method improves the results of our baseline by 2.3% in Objects365 and 2.8% on LVIS. Besides, our fine-tuned method beats the SOTA largely by 6.3% on PASCAL VOC. When not fine-tuned with COCO base categories, the performance of our method still outperforms *Zareian et al.* (fine-tuned with COCO base categories) on Objects365 and LVIS. When not fine-tuned, our method’s underperforms our fine-tuned baseline on PASCAL VOC. It is very likely because of that there is a large semantic overlap between the COCO base categories and PASCAL VOC object categories. Therefore, fine-tuning on COCO base set helps the model’s transfer ability to PASCAL VOC.

#### 4.5 Ablation Study

**How Does the Quality of Bounding-Box Proposals Affect Performance?** Our pseudo label generator combines object proposals and the acti-

**Table 7.** Performance of our method with different text encoders

Method	Text encoder	Generalized setting		
		Novel AP	Base AP	Overall AP
Zareian et al. [34]	Bert	22.8	46.0	39.9
Our method	Bert	<b>28.8</b>	45.1	40.9
	CLIP	<b>30.8</b>	46.1	42.1




**Fig. 4.** Visualization of some activation maps. Colorful blocks indicate values of Grad-CAM activation map in the corresponding regions. We zero out blocks with values smaller than half of the max value in the map so the main focus is highlighted. Black boxes indicate object proposals and red boxes indicate the final selected pseudo bounding-box labels (Color figure online)

vation map to select boxes. Generally, the better the proposals are, the more accurate our pseudo bounding-box annotations would be. The default proposal generator in our main experiments is a Mask-RCNN trained with COCO detection category excluding the novel categories. To analyze the effect of proposal quality, we also run experiments with an unsupervised proposal generator, Selective Search [30] and summarize results in Table 3. The results show that our method with Selective Search outperforms *Zareian et al.* with a clear margin. This demonstrates the effectiveness of our method even with an unsupervised proposal generator.

### What Is the Effect of Different Pre-trained Vision-Language Models?

Besides the utilization of ALBEF in our main experiment, we also experiment with LXMERT [28], which is an earlier vision-language model that also fuses information from both the vision and language modalities. Specifically, we generate pseudo labels based on the activation map in the last layer of the (text-to-vision) cross-attention module. Results are shown in Table 4. It shows that



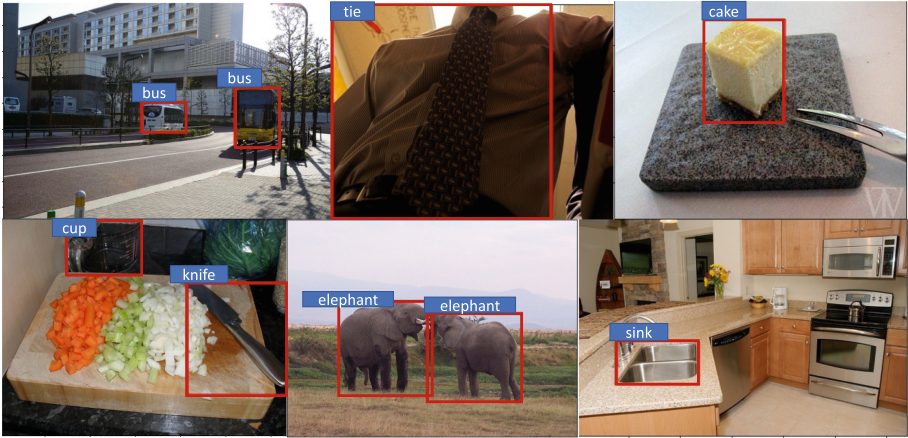
**Fig. 5.** Visualization of our generated pseudo bounding-box annotations on COCO. The red boxes indicate successful cases and the yellow one denote failure case. Our pseudo label generator can generate objects (slippers, pot and pie) that are not covered by COCO’s category list. The generator cannot capture an object if it is not shown in the caption (e.g. the car in the last column) (Color figure online)

with LXMERT, our method’s performance is slightly worse compared with our method using ALBEF. This may due to the fact that LXMERT employs less image-caption data for training. While compared with *Zareian et al.*, our method with LXMERT still performs significantly better.

**What Is the Effect of Our Object Vocabulary Size?** We utilize an object vocabulary containing over 1.5k object categories ( $\mathbb{V}$ ) by default when generating our pseudo labels. The vocabulary size is much larger than the object vocabulary of any dataset we evaluated on, such as base and novel object categories in COCO ( $\mathbb{V}^-$ ). A natural question is, for performance of COCO novel categories, for example, would it be better if we just use  $\mathbb{V}^-$  for pseudo-box generation and detector training? To answer the question, we conduct experiments using these two scales of vocabularies. Experimental results in Table 5 show that  $\mathbb{V}$  leads to better performance than  $\mathbb{V}^-$  on COCO novel categories. The results suggest that adding additional object categories outside the COCO categories during pre-training will benefit our model performance. Besides, using the larger vocabulary improves model’s generalization ability to datasets that include a large set of object categories, i.e., Objects365 and LVIS. We observe that using  $\mathbb{V}$  improves results of  $\mathbb{V}^-$  by 2.3% in Objects365 and by 2.7% in LVIS.

**Does More Data for Pseudo Label Generation Help?** Our pseudo labels are generated from image-caption pairs. Intuitively, the more data is used for pseudo label generation, the larger amount of diverse objects will be utilized for training our open vocabulary detector. As a result, our model performance should be improved. We show the performance of our method using pseudo labels with different amounts of image-caption pairs in Table 6. Results show that our method can benefit from a larger dataset. Our performance is improved by  $\sim 2\%$  on the target set when using more data. Moreover, our method still outperforms the baseline significantly even when pre-trained with *COCO Cap* only.

**What Is the Effect of Different Text Encoders?** Besides our default choice of text encoder (CLIP), we implement our method with another encoder, i.e.,



**Fig. 6.** Some example results of our open vocabulary detector. The shown categories are from novel categories in COCO (Color figure online)

Bert (base) [4] which is a widely used language model that is trained using text data only. Our main baseline *Zareian et al.* uses Bert as their text encoder as well. The comparison results are shown in Table 7. The results suggest that with Bert encoder, our method’s performance is slightly worse on COCO target set compared with our method using CLIP text encoder. This may due to the fact that CLIP text encoder is trained using image-caption pairs which results in better generalization performance for image-related tasks. While our method with Bert encoder still outperforms *Zareian et al.* by 6% AP on novel categories. It indicates that the performance improvement of our method doesn’t mainly come from a better text encoder, but from training with large-scale pseudo bounding-box labels.

**Qualitative Visualization.** We visualize examples of our generated pseudo bounding boxes in Fig. 4 and Fig. 5. As we can see, the generated pseudo labels show promising performance (see red boxes) in grounding objects and are able to cover categories, e.g., pot, slippers and pie (Fig. 5), that are not in the original object list of COCO’s ground-truth annotations. We also observe that some background objects will be missed when it is not mentioned in the caption (see the yellow box in the last column of Fig. 5). Nevertheless, instances of the same object category may show up in other image-caption pairs in the large-scale dataset and our method could generate pseudo box labels for those cases. Therefore, our detector can still recognize this object after training. We also show some examples of our detection results in Fig. 6. They are all from COCO novel categories which are not covered by COCO base annotations. The results demonstrate promising generalization ability of our method to novel objects.

## 5 Closing Remarks

We propose a novel framework that trains an open vocabulary object detector with pseudo bounding-box labels generated from large-scale image-caption pairs. We introduce a pseudo label generator that leverages the localization ability of pre-trained vision-language models to generate pseudo bounding-box annotations for diverse objects embedded in image captions. The generated pseudo labels can be used to improve open vocabulary object detection. Experimental results show that our method outperforms the state-of-the-art zero-shot and open vocabulary object detection methods by a large margin.

**Potential Negative Societal Impact.** Our method generates pseudo bounding box labels to alleviate human labeling efforts. Since our pseudo label generator mines annotations of objects from the input captions without human intervention, our pseudo labels might be biased because of the bias embedded in the language descriptions. Manually filtering out the biased object names in the vocabulary could be an effective solution.

## References

1. Bansal, A., Sikka, K., Sharma, G., Chellappa, R., Divakaran, A.: Zero-shot object detection. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11205, pp. 397–414. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01246-5\\_24](https://doi.org/10.1007/978-3-030-01246-5_24)
2. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: CVPR, pp. 2846–2854 (2016)
3. Chen, X., et al.: Microsoft COCO captions: data collection and evaluation server. arXiv preprint [arXiv:1504.00325](https://arxiv.org/abs/1504.00325) (2015)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
5. Everingham, M.: The pascal visual object classes challenge, (voc2007) results (2007). <http://pascallin.ecs.soton.ac.uk/challenges/VOC/voc2007/index.html>
6. Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: a retrospective. *Int. J. Comput. Vis.* **111**(1), 98–136 (2015)
7. Gao, M., Li, A., Yu, R., Morariu, V.I., Davis, L.S.: C-WSL: count-guided weakly supervised localization. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11205, pp. 155–171. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-01246-5\\_10](https://doi.org/10.1007/978-3-030-01246-5_10)
8. Girshick, R.: Fast R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1440–1448 (2015)
9. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Zero-shot detection via vision and language knowledge distillation. arXiv preprint [arXiv:2104.13921](https://arxiv.org/abs/2104.13921) (2021)
10. Gupta, A., Dollár, P., Girshick, R.: LVIS: a dataset for large vocabulary instance segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
11. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2961–2969 (2017)

12. Jia, C., et al.: Scaling up visual and vision-language representation learning with noisy text supervision. arXiv preprint [arXiv:2102.05918](https://arxiv.org/abs/2102.05918) (2021)
13. Joseph, K., Khan, S., Khan, F.S., Balasubramanian, V.N.: Towards open world object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5830–5840 (2021)
14. Krishna, R., et al.: Visual genome: connecting language and vision using crowd-sourced dense image annotations. *Int. J. Comput. Vis.* **123**(1), 32–73 (2017)
15. Li, J., Selvaraju, R.R., Gotmare, A.D., Joty, S., Xiong, C., Hoi, S.: Align before fuse: vision and language representation learning with momentum distillation. In: *NeurIPS* (2021)
16. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: VisualBERT: a simple and performant baseline for vision and language. arXiv preprint [arXiv:1908.03557](https://arxiv.org/abs/1908.03557) (2019)
17. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-10602-1\\_48](https://doi.org/10.1007/978-3-319-10602-1_48)
18. Ordonez, V., Kulkarni, G., Berg, T.: Im2Text: describing images using 1 million captioned photographs. *Adv. Neural. Inf. Process. Syst.* **24**, 1143–1151 (2011)
19. Papageorgiou, C., Poggio, T.: A trainable system for object detection. *Int. J. Comput. Vis.* **38**(1), 15–33 (2000)
20. Papageorgiou, C.P., Oren, M., Poggio, T.: A general framework for object detection. In: Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271), pp. 555–562. IEEE (1998)
21. Radford, A., et al.: Learning transferable visual models from natural language supervision. arXiv preprint [arXiv:2103.00020](https://arxiv.org/abs/2103.00020) (2021)
22. Rahman, S., Khan, S., Barnes, N.: Improved visual-semantic alignment for zero-shot object detection. In: *AAAI*, pp. 11932–11939 (2020)
23. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. *Adv. Neural. Inf. Process. Syst.* **28**, 91–99 (2015)
24. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 618–626 (2017)
25. Shao, S., et al.: Objects365: a large-scale, high-quality dataset for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8430–8439 (2019)
26. Sun, P., et al.: Sparse R-CNN: end-to-end object detection with learnable proposals. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14454–14463 (2021)
27. Szegedy, C., Toshev, A., Erhan, D.: Deep neural networks for object detection (2013)
28. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. arXiv preprint [arXiv:1908.07490](https://arxiv.org/abs/1908.07490) (2019)
29. Tang, P., et al.: PCL: proposal cluster learning for weakly supervised object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**(1), 176–191 (2018)
30. Uijlings, J.R., Van De Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. *IJCV* **104**(2), 154–171 (2013)
31. Xie, E., et al.: DetCo: unsupervised contrastive learning for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8392–8401 (2021)



32. Ye, K., Zhang, M., Kovashka, A., Li, W., Qin, D., Berent, J.: Cap2Det: learning to amplify weak caption supervision for object detection. In: ICCV, pp. 9686–9695 (2019)
33. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3D object detection and tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 11784–11793 (2021)
34. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: CVPR, pp. 14393–14402 (2021)
35. Zhu, P., Wang, H., Saligrama, V.: Zero shot detection. *IEEE Trans. Circuits Syst. Video Technol.* **30**(4), 998–1010 (2019)
36. Zhu, P., Wang, H., Saligrama, V.: Don't even look once: synthesizing features for zero-shot detection. In: CVPR, pp. 11693–11702 (2020)