

Aligning Bag of Regions for Open-Vocabulary Object Detection

Size Wu¹ Wenwei Zhang¹ Sheng Jin^{2,3} Wentao Liu^{3,4} Chen Change Loy^{1*}
¹S-Lab, Nanyang Technological University ²The University of Hong Kong
³SenseTime Research and Tetras.AI ⁴Shanghai AI Laboratory
 {size001, wenwei001, ccloy}@ntu.edu.sg {jinsheng, liuwentao}@sensetime.com

Abstract

Pre-trained vision-language models (VLMs) learn to align vision and language representations on large-scale datasets, where each image-text pair usually contains a bag of semantic concepts. However, existing open-vocabulary object detectors only align region embeddings individually with the corresponding features extracted from the VLMs. Such a design leaves the compositional structure of semantic concepts in a scene under-exploited, although the structure may be implicitly learned by the VLMs. In this work, we propose to align the embedding of bag of regions beyond individual regions. The proposed method groups contextually interrelated regions as a bag. The embeddings of regions in a bag are treated as embeddings of words in a sentence, and they are sent to the text encoder of a VLM to obtain the bag-of-regions embedding, which is learned to be aligned to the corresponding features extracted by a frozen VLM. Applied to the commonly used Faster R-CNN, our approach surpasses the previous best results by 4.6 box AP₅₀ and 2.8 mask AP on novel categories of open-vocabulary COCO and LVIS benchmarks, respectively. Code and models are available at <https://github.com/wusize/ovdet>.

1. Introduction

A traditional object detector can only recognize categories learned in the training phase, restricting its application in the real world with a nearly unbounded concept pool. Open-vocabulary object detection (OVD), a task to detect objects whose categories are absent in training, has drawn increasing research attention in recent years.

A typical solution to OVD, known as the distillation-based approach, is to distill the knowledge of rich and unseen categories from pre-trained vision-language models (VLMs) [21, 38]. In particular, VLMs learn aligned image and text representations on large-scale image-text pairs (Fig. 1(a)). Such general knowledge is beneficial for OVD. To extract the knowledge, most distillation-based ap-

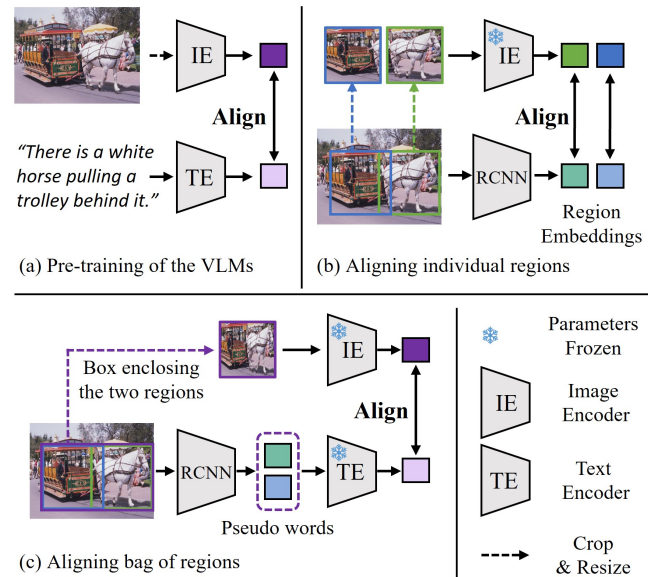


Figure 1. (a) Typical vision-language models (VLMs) learn to align representations of images and captions with rich compositional structure. (b) Existing distillation-based object detectors align each *individual* region embedding to features extracted by the frozen image encoder of VLMs. (c) Instead, the proposed method aligns the embedding of *bag of regions*. The region embeddings in a bag are projected to the word embedding space (dubbed as pseudo words), formed as a sentence, and then sent to the text encoder to obtain the *bag-of-regions* embedding, which is aligned to the corresponding image feature extracted by the frozen VLMs.

proaches [10, 15, 52] align each *individual* region embedding to the corresponding features extracted from the VLM (Fig. 1(b)) with some carefully designed strategies.

We believe VLMs have implicitly learned the inherent compositional structure of multiple semantic concepts (e.g., co-existence of stuff and things [1, 23]) from a colossal amount of image-text pairs. A recent study, MaskCLIP [57], leverages such a notion for zero-shot segmentation. Existing distillation-based OVD approaches, however, have yet to fully exploit the compositional structures encapsulated in

*Corresponding author.

VLMs. Beyond the distillation of individual region embedding, we propose to align the embedding of *BAG of RegiONs*, dubbed as BARON. Explicitly learning the co-existence of visual concepts encourages the model to understand the scene beyond just recognizing isolated individual objects.

BARON is easy to implement. As shown in Fig. 1(c), BARON first samples contextually interrelated regions to form a ‘bag’. Since the region proposal network (RPN) is proven to cover potential novel objects [15, 58], we explore a neighborhood sampling strategy that samples boxes around region proposals to help model the co-occurrence of a bag of visual concepts. Next, BARON obtains the bag-of-regions embeddings by projecting the regional features into the word embedding space and encoding these pseudo words with the text encoder (TE) of a frozen VLM [38]. By projecting region features to pseudo words, BARON naturally allows TE to effectively represent the co-occurring semantic concepts and understand the whole scene. To retain the spatial information of the region boxes, BARON projects the box shape and box center position into embeddings and add to the pseudo words before feeding them to TE.

To train BARON, we align the bag-of-regions embeddings to the teacher embeddings, which are obtained by feeding the image crops that enclose the bag of regions to the image encoder (IE) of the VLM. We adopt a contrastive learning approach [46] to learn the pseudo words and the bag-of-regions embeddings. Consistent with the VLMs’ pre-training (e.g., CLIP [38]), the contrastive loss pulls close corresponding student (the detector) and teacher (IE) embedding pairs and pushes away non-corresponding pairs.

We conduct extensive experiments on two challenging benchmarks, OV-COCO and OV-LVIS. The proposed method consistently outperforms existing state-of-the-art methods [10, 52, 58] in different settings. Combined with Faster R-CNN, BARON achieves a 34.0 (4.6 increase) box AP₅₀ of novel categories on OV-COCO and 22.6 (2.8 increase) mask mAP of novel categories on OV-LVIS. It is noteworthy that BARON can also distill knowledge from caption supervision – it achieves 32.7 box AP₅₀ of novel categories on OV-COCO, outperforming previous approaches that use COCO caption [13, 53, 56, 58].

2. Related Work

Vision-Language Pretraining and Its Applications. Vision-language pre-training aims to learn aligned image and text representations [12, 20–22, 38] on large-scale image-text pairs. There are many studies [22, 25, 26, 33] that pre-train vision-language models (VLMs) to improve the performance of downstream recognition and generation tasks. There are also studies that learn aligned vision-language representation so that the images can be classified with arbitrary texts [12, 35]. Recent attempts [21, 38, 54] push forward this direction by conducting contrastive learning in VLMs

on billion-scale image-text pairs. These models show impressive zero-shot performance when they are transferred to image classification tasks.

Inspired by the success of VLMs [38], some works try to exploit the alignment of vision-language representations for dense prediction tasks, e.g. segmentation [24, 40, 57] and detection [10, 15, 27, 44]. In particular, MaskCLIP [57] shows that the image encoder in VLMs [38] captures the stuff and things in a complex scene, where the pixel embeddings of each concept are naturally aligned with the corresponding text representations, although the original CLIP [38] model does not explicitly learn this target. This implies that VLMs, after trained on a massive amount of image-text pairs, have implicitly learned the compositional structure of multiple semantic concepts, which naturally exist in image-text pairs. This motivates us to explore the representation alignment between bag of regions and bag of words, different from previous works [15, 24] that focus on aligning the representation of individual pixels, regions, or words in VLMs.

Open-Vocabulary Object Detection. Traditional object detectors [3, 5, 32, 42, 59] are limited to pre-defined object categories. To detect objects of unseen categories, zero-shot object detection (ZSD) [2, 8, 17, 39, 55] is proposed to align individual region embeddings with the text embeddings of categories through different strategies. Recent attempts further explore open-vocabulary object detection (OVD) [53], a more general form of ZSD that leverages weak supervisions like visual grounding data [27], image captions [13, 53, 56], and image labels [58]. Large-scale pre-trained VLMs [38] are also exploited for their remarkable zero-shot recognition ability. These VLMs can generate conditional queries [52] or serve as a good teacher for knowledge distillation [15]. Specifically, distillation-based approaches [10, 15] extract embeddings on pre-computed region proposals and *individually* align them to the corresponding features obtained from the VLMs. To our best knowledge, BARON is the first attempt to lift the learning from *individual* regions to the *bag of regions* for OVD.

3. Method

Our method makes the first attempt to align embedding of *bag of regions* beyond *individual* regions for OVD. We call our method BARON. In this work, we instantiate the idea of BARON based on the commonly used Faster-RCNN [42] and modify it for OVD (Sec. 3.1). We design a simple strategy to form the bag of regions from the region proposals (Sec. 3.2). The bag of regions is treated as a bag of words to obtain the bag-of-regions embedding (Sec. 3.3), which are then aligned with the corresponding features from VLMs (Sec. 3.4). BARON is general that it can align the bag-of-regions embeddings to not only image representations but also text representations (Sec. 3.5).

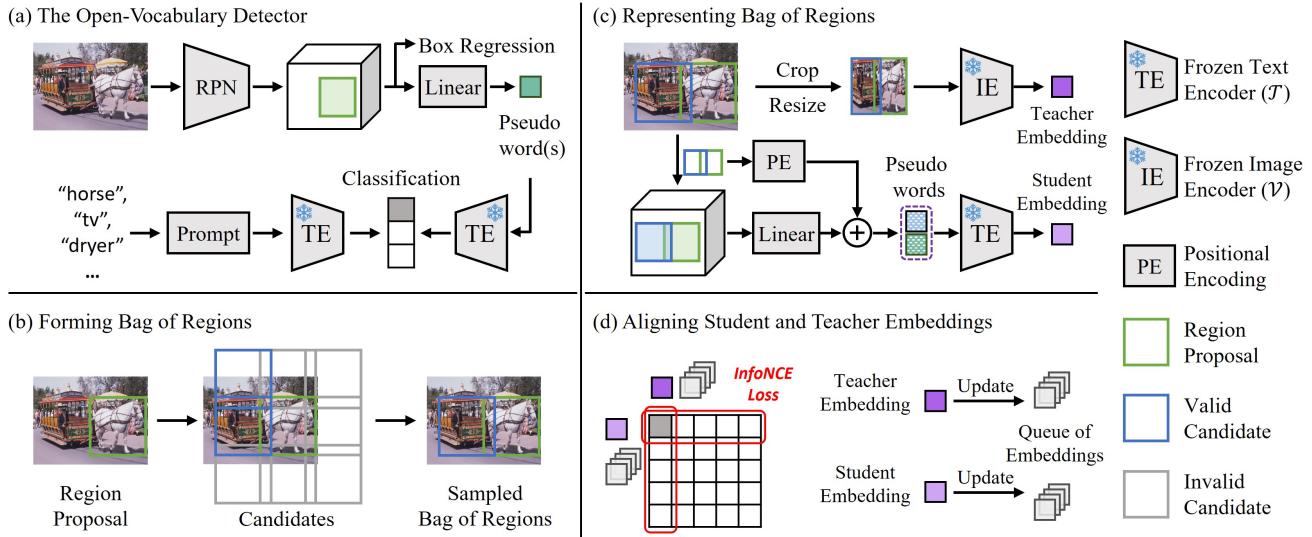


Figure 2. Overview of BARON. (a) BARON is based on a Faster R-CNN whose classifier is replaced by a linear layer to map region features into pseudo words. (b) BARON takes region proposals and its surrounding boxes to form bags of regions. (c) BARON obtains student and teacher embeddings for the bag of regions from the pre-trained VLMs. (d) BARON learns the alignment by the InfoNCE loss and maintains queues of embeddings to provide sufficient negative pairs for InfoNCE loss.

3.1. Preliminaries

In this paper, we instantiate the idea upon Faster R-CNN [42] for simplicity. The idea can also be used for other architectures [3, 30] applicable for OVD. To enable Faster R-CNN to detect objects from arbitrary vocabularies, we replace the original classifier with a linear layer that projects the region features into the word embedding space (dubbed as pseudo words) (Fig. 2(a)). In practice, the linear layer maps a region feature to multiple pseudo words to represent the rich semantic information of each object, similar to those category names consisting of multiple words (e.g., horse-driven trolley). Finally, we feed these pseudo words into the text encoder and then calculate the similarity with the category embeddings to obtain final classification results.

As shown in Fig. 2(a), given C object categories, we obtain the embedding f_c for the c -th category by feeding the category names with a prompt template, e.g., ‘a photo of {category} in the scene’ to the text encoder \mathcal{T} . For a region and its pseudo words w , the probability of the region to be classified as the c -th category is

$$p_c = \frac{\exp(\tau \cdot \langle \mathcal{T}(w), f_c \rangle)}{\sum_{i=0}^{C-1} \exp(\tau \cdot \langle \mathcal{T}(w), f_i \rangle)}, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the cosine similarity and τ is the temperature to re-scale the value.

During training, only the boxes of base categories are annotated and the learning on base categories follows the convention of Faster R-CNN with regression and classification losses [42]. To learn to detect novel categories that do not

have box annotations in training, previous distillation-based approaches [10, 15] *individually* align region embeddings (e.g., $\mathcal{T}(w)$) to the corresponding features obtained from the VLMs. To further exploit the power of VLMs that capture the compositional structure of multiple concepts, we lift the learning from *individual* regions to the *bag of regions*.

3.2. Forming Bag of Regions

Our framework is inspired by existing OVD approaches that distill knowledge from the image encoder of VLMs. Specifically, we choose the image encoder of the VLMs as the teacher and expect it to teach the detector. But different from existing approaches, we wish the detector to learn the co-existence of multiple concepts, especially the potential existence of novel objects. To effectively and efficiently learn such knowledge from the VLMs, we consider the following two properties of regions inside a bag: 1) the regions need to be close to each other because an image crop enclosing distanced regions will include a larger proportion of redundant image contents, distracting the image encoder from representing the bag of regions; 2) the regions should have similar sizes, as an imbalanced size ratio among regions will make the image representations dominated by the largest region.

According to these two requirements, we adopt a simple neighborhood sampling strategy to form the bag of regions based on the region proposals predicted by the region proposal network (RPN) [42]. Concretely, for each region proposal, we take its surrounding eight boxes (neighbors) as the candidates, which are spatially close to each other as shown in Fig. 2(b). We also allow these candidates to

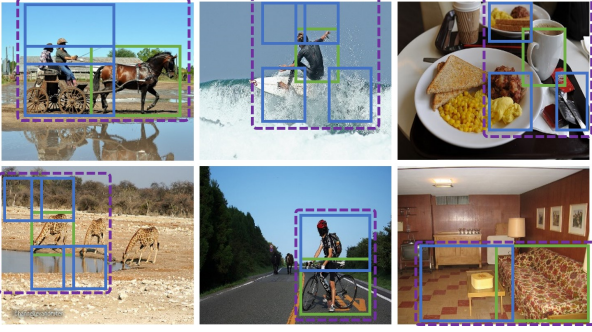


Figure 3. Visualization of sampled bags of regions. Green boxes denote the region proposals and blue boxes are sampled neighbors (candidates). Areas exceeding image boundary are cropped out.

overlap slightly at a specific Intersection over Foreground (IOF) to improve the continuity of regional representations. To balance the sizes of regions inside the bag, we simply let the eight candidate boxes have the same shape as the region proposal. In practice, candidates that exceed the boundary of the image over a certain proportion, *e.g.*, more than $\frac{2}{3}$ of their area out of the image, will be discarded. The remaining candidates are independently sampled, and the sampled boxes, together with the region proposal, form a bag of regions.

We sample G groups for each region proposal to obtain rich bag-of-regions representations. At each time of sampling, the probability of a candidate being sampled is adjusted to prevent the box enclosing the grouped regions from having an extreme aspect ratio. Assuming a base probability p_b and a scaling factor α , the probability to sample the left and right candidates is $p = p_b \times \min\{(\frac{H}{W})^\alpha, 1\}$ and the probability to sample upper and bottom candidates is $p = p_b \times \min\{(\frac{W}{H})^\alpha, 1\}$, where H and W are the height and width of the region proposal box. α is 3.0 by default.

We show some sampling results of bags of regions on COCO dataset [31] in Fig 3. The grouped regions in a bag can cover objects (in blue) co-occurring with the object in the region proposal (in green). The context of the co-occurring objects leads BARON to get a scene-level understanding of the regions, *e.g.* ‘a horse pulling the carriage’ and ‘a man surfing on the wave’. Novel object categories occasionally appear in the bag of regions, *e.g.* the ‘carriage’ in the top left image of Fig 3 and the ‘cup’ in the top right. In the following Sec 3.3 and Sec 3.4, these potential novel object categories would be learned in the context of the bag of regions.

3.3. Representing Bag of Regions

With the sampled bags of regions, BARON obtains the bag-of-regions embeddings from both the student (*i.e.*, the open-vocabulary object detector) and the teacher (*i.e.*, VLMs). We denote the j -th region in the i -th group as b_j^i and the pseudo words after the projection layer as w_j^i . For the pre-trained VLM, we use \mathcal{T} to denote the text encoder and \mathcal{V} to denote the image encoder.

Student Bag-of-Regions Embedding. Because the region features are projected to word embedding space and learned to be aligned with the text embedding of category, a straightforward way to obtain the embedding of a bag of regions is to concatenate the pseudo words and feed them to the text encoder \mathcal{T} . However, the spatial information of the regions will be lost in such a process, including relative box center positions and relative box shapes. The center position and shape indicate the spatial relationship among the regions in a bag, which is essential to induce a sentence-like interpretation for the bag of regions. And they are also encoded in the teacher (image encoder of VLMs) through positional embeddings in the input [38]. Therefore, BARON encodes the spatial information into positional embeddings p_j^i that have the same dimension with w_j^i , following the practices of Transformers [9, 47]. The positional embeddings are added to the pseudo words before concatenation. Assuming the group size is N^i , this representation can be formulated as $f_t^i = \mathcal{T}(w_0^i + p_0^i, w_1^i + p_1^i, \dots, w_{N^i-1}^i + p_{N^i-1}^i)$.

Teacher Bag-of-Regions Embedding. The image embedding of the grouped regions can be obtained by feeding the image crop that encloses the regions to the image encoder \mathcal{V} . The image crop may contain redundant contents that are outside the grouped regions; we mask out them in the attention layers of \mathcal{V} . The image feature can be formulated as $f_v^i = \mathcal{V}(b_0^i, b_1^i, \dots, b_{N^i-1}^i)$.

3.4. Aligning Bag of Regions

BARON aligns the bag-of-regions embeddings from student and teacher to make the student learn to encode the co-existence of multiple regions, which potentially contain multiple concepts. We adopt the contrastive learning approach used in vision-language pre-training [38]. Specifically, given G bags of regions, the alignment InfoNCE loss [36] between bag-of-regions embeddings is calculated as

$$\mathcal{L}_{\text{bag}} = -\frac{1}{2} \sum_{k=0}^{G-1} (\log(p_{t,v}^k) + \log(p_{v,t}^k)). \quad (2)$$

The $p_{t,v}^k$ and $p_{v,t}^k$ are calculated as

$$p_{t,v}^k = \frac{\exp(\tau' \cdot \langle f_t^k, f_v^k \rangle)}{\sum_{l=0}^{G-1} \exp(\tau' \cdot \langle f_t^k, f_v^l \rangle)} \quad (3)$$

$$p_{v,t}^k = \frac{\exp(\tau' \cdot \langle f_v^k, f_t^k \rangle)}{\sum_{l=0}^{G-1} \exp(\tau' \cdot \langle f_v^k, f_t^l \rangle)}, \quad (4)$$

respectively, where τ' is the temperature to re-scale the cosine similarity. The loss pulls positive pairs $\{f_t^k, f_v^k\}$ close to each other and pushes away negative pairs $\{f_t^k, f_v^l\} (k \neq l)$.

In practice, the number of groups G is small for a single image. We maintain two queues to save the image and text embeddings in previous iterations during training to provide sufficient negative pairs [18].

Table 1. Comparison with state-of-the-art methods on OV-COCO benchmark. We separately compare our approach with methods distilling knowledge from CLIP and approaches using COCO caption. † means using proposals produced by MAVL [34].

Method	Supervision	Backbone	Detector	AP ₅₀ ^{novel}	AP ₅₀ ^{base}	AP ₅₀
ViLD [15]	CLIP	ResNet50-FPN	FasterRCNN	27.6	59.5	51.2
OV-DETR [52]	CLIP	ResNet50	DeformableDETR	29.4	61.0	52.7
BARON (Ours)	CLIP	ResNet50-FPN	FasterRCNN	34.0	60.4	53.5
OVR-CNN [53]	Caption	ResNet50-C4	FasterRCNN	22.8	46.0	39.9
RegionCLIP [56]	Caption	ResNet50-C4	FasterRCNN	26.8	54.8	47.5
Detic [58]	Caption	ResNet50-C4	FasterRCNN	27.8	51.1	45.0
PB-OVD [13]	Caption	ResNet50-C4	FasterRCNN	30.8	46.1	42.1
VLDet [28]	Caption	ResNet50-C4	FasterRCNN	32.0	50.6	45.8
BARON (Ours)	Caption	ResNet50-C4	FasterRCNN	33.1	54.8	49.1
Rasheed <i>et al.</i> [41]†	CLIP + Caption	ResNet50-C4	FasterRCNN	36.6	54.0	49.4
BARON (Ours)†	CLIP + Caption	ResNet50-C4	FasterRCNN	42.7	54.9	51.7




Figure 4. The caption-version BARON. We align the embeddings of the bag of regions to the caption embeddings.

Aligning Individual Regions. The alignment between individual regions’ student and teacher embeddings is complementary to that of a bag of regions. Therefore, we also adopt the individual-level distillation in our implementation. For computational efficiency, we obtain the teacher embeddings from the feature map of the image encoder’s last attention layer by RoiAlign [19] instead of repeatedly passing image crops to the image encoder. Similarly, the student embeddings are extracted from the text encoder’s last attention layer by averaging pseudo-word embeddings of the same region. We apply the InfoNCE loss and keep queues of embeddings to calculate the individual-level loss $\mathcal{L}_{\text{individual}}$.

3.5. Caption Supervision

It is noteworthy that BARON can accept caption supervision too. As shown in Fig. 4, the core idea is to replace the image embedding f_v with embedding obtained by feeding the image captions to the text encoder \mathcal{T} . To obtain the region groups, we randomly sample some region proposals generated by the RPN. As an image can have multiple captions, we follow the practice in UniCL [50] to apply a soft cross entropy loss that simultaneously aligns the student bag-of-regions embedding to multiple caption embeddings. The alignment of individual regions is discarded since we cannot get the correspondence between pseudo words of regions and the actual words in a caption without grounding. In this way, BARON learns to align the bag-of-regions embedding

of the image caption, which also describes the existence of multiple concepts of the image.

4. Experiments

Datasets. We evaluate our method on the two popular object detection datasets, *i.e.*, COCO [31] and LVIS [16]. For the COCO dataset, we follow OV-RCNN [53] to split the object categories into 48 base categories and 17 novel categories. For the LVIS dataset, we follow ViLD [15] to split the 337 rare categories into novel categories and the rest common and frequent categories into base categories. For brevity, we denote the open-vocabulary benchmarks based on COCO and LVIS as OV-COCO and OV-LVIS.

Evaluation Metrics. We evaluate the detection performance on both base and novel categories for completeness. For OV-COCO, we follow OV-RCNN [53] to report the box AP at IoU threshold 0.5, noted as AP₅₀. For OV-LVIS, we report both the mask and box AP averaged on IoUs from 0.5 to 0.95, noted as mAP. The AP₅₀ of novel categories (AP₅₀^{novel}) and mAP of rare categories (AP_r) are the main metrics that evaluate the open-vocabulary detection performance on OV-COCO and OV-LVIS, respectively.

Implementation Details. We build BARON on Faster R-CNN [42] with ResNet50-FPN [29]. For a fair comparison with existing methods, we initialize the backbone network with weights pre-trained by SOCO [48] and apply synchronized Batch Normalization (SyncBN) [37] following DetPro [10]. We choose the 2× schedule (180, 000 iterations) for the main experiments on COCO and LVIS [6, 49]. For the pre-trained VLM, we choose the CLIP [38] model based on ViT-B/32 [9]. For the adaptation to caption supervision, we base our method on Faster R-CNN with ResNet50-C4 backbone [42] and adopt the 1× schedule. For the prompt of category names, we use the hand-crafted prompts in ViLD [15] for all our experiments by default. We use learned prompts only when comparing with DetPro [10].

Table 2. Comparison with state-of-the-art methods on OV-LVIS. * denotes the re-implemented ViLD [15] reported in DetPro [10].

Method	Ensemble	Learned Prompt	Object Detection				Instance segmentation			
			AP _r	AP _c	AP _f	AP	AP _r	AP _c	AP _f	AP
ViLD [15]	-	-	16.3	21.2	31.6	24.4	16.1	20.0	28.3	22.5
OV-DETR [52]	-	-	-	-	-	-	17.4	25.0	32.5	26.6
BARON (Ours)	-	-	17.3	25.6	31.0	26.3	18.0	24.4	28.9	25.1
ViLD [15]	✓	-	16.7	26.5	34.2	27.8	16.6	24.6	30.3	25.5
ViLD* [15]	✓	-	17.4	27.5	31.9	27.5	16.8	25.6	28.5	25.2
BARON (Ours)	✓	-	20.1	28.4	32.2	28.4	19.2	26.8	29.4	26.5
DetPro [10]	✓	✓	20.8	27.8	32.4	28.4	19.8	25.6	28.9	25.9
BARON (Ours)	✓	✓	23.2	29.3	32.5	29.5	22.6	27.6	29.8	27.6

Table 3. Comparison of the transfer ability of the model trained on OV-LVIS. * denotes the re-implemented ViLD [15] reported in DetPro [10]. ‡ denotes that we use hand-crafted prompts for a fair comparison with ViLD.

Method	Pascal VOC		COCO						Objects365					
	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l	AP	AP ₅₀	AP ₇₅	AP _s	AP _m	AP _l
Supervised [10]	78.5	49.0	46.5	67.6	50.9	27.1	67.6	77.7	25.6	38.6	28.0	16.0	28.1	36.7
ViLD* [15]	73.9	57.9	34.1	52.3	36.5	21.6	38.9	46.1	11.5	17.8	12.3	4.2	11.1	17.8
BARON (Ours) [‡]	74.5	57.9	36.3	56.1	39.3	25.4	39.5	48.2	13.2	20.0	14.0	4.8	12.7	20.1
DetPro [10]	74.6	57.9	34.9	53.8	37.4	22.5	39.6	46.3	12.1	18.8	12.9	4.5	11.5	18.6
BARON (Ours)	76.0	58.2	36.2	55.7	39.1	24.8	40.2	47.3	13.6	21.0	14.5	5.0	13.1	20.7

4.1. Benchmark Results

OV-COCO. We report the comparison with previous methods in Table 1. BARON surpasses previous state of the arts with either pre-trained VLMs or COCO captions [7], indicating its effectiveness and flexibility. Note that OV-DETR is based on the Deformable DETR [59], which is stronger than the Faster R-CNN [42] with higher performance on base categories. But BARON still outperforms OV-DETR by 4.6 AP₅₀ on novel categories. When using caption supervision, BARON even outperforms PB-OVD [13] that uses sophisticated pseudo-labeling. Combining CLIP image features, COCO captions, and MAVL [34] proposals, BARON outperforms Rasheed *et al.* [41] by a large margin.

OV-LVIS. We compare BARON with other methods on the OV-LVIS benchmark in Table 2. Because ViLD [15] is trained with large-scale jittering [14] and a prohibitive 32× schedule, DetPro [10] re-implemented it with backbone weights pre-trained by SOCO [48] and a regular 2× schedule. DetPro also proposes learned prompts for the category’s names. Besides, an ensembling strategy for classification scores is adopted in ViLD and DetPro. For a fair comparison, we respectively implement our method on OV-LVIS with and without these tricks. BARON achieves the best performance in all the scenarios.

Transfer to Other Datasets. We transfer the open-vocabulary detector trained on OV-LVIS to three other datasets, including Pascal VOC 2007 [11] test set, COCO [31] validation set and Objects365 [43] v2 validation set. We compare our results with DetPro [10] and the ViLD [15] implemented in DetPro. The comparison is fair

since all the models are based on the same object detector and training schedule. As shown in Table 3, our approach exhibits better generalization ability on all three datasets.

4.2. Ablation Study

In this section, we ablate the effectiveness of components in BARON on OV-COCO benchmark.

Effectiveness of Aligning Bag of Regions. We start from a baseline that only uses the individual-level loss $\mathcal{L}_{\text{individual}}$. As shown in Table 4(#1), the individual-level baseline achieves 25.7 mAP₅₀ on novel categories. We then replace $\mathcal{L}_{\text{individual}}$ with the loss for bag of regions \mathcal{L}_{bag} (Table 4(#2)). Without considering the spatial information of region boxes, the performance on novel categories is compatible with aligning individual regions. When adding the positional embedding of the region boxes’ spatial information, the performance on novel categories (Table 4(#3)) dramatically increases by 7.1 mAP₅₀. This means that spatial information is essential to effectively exploit the compositional structure of co-occurring visual concepts in a bag of regions. Finally, in Table 4(#4), we find that the individual-level loss is complementary to the bag-of-regions alignment, which brings 1.2 mAP₅₀ performance gain on novel categories.

Sampling Strategies. We explore two baselines to sample bags of regions to support the rationale of our neighborhood sampling strategy. The first is to equally split an image into grids (dubbed as grid sampling) like the pre-training stage in OVR-CNN [53] such that the fixed grids form a bag of regions. And the second is to randomly sample region proposals to form a bag of regions (dubbed as random sampling). For a fair comparison, we keep the number of

Table 4. Effectiveness of main components of BARON

#	$\mathcal{L}_{\text{individual}}$	\mathcal{L}_{bag}	PE	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
1	✓	-	-	25.7	59.6	50.6
2	-	✓	-	25.7	59.4	50.5
3	-	✓	✓	32.8	60.1	53.0
4	✓	✓	✓	34.0	60.4	53.5

Table 7. Ablation study on the sampling probability p_b

p_b	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
0.1	33.7	59.8	52.9
0.3	34.0	60.4	53.5
0.5	33.2	60.0	53.0

sampled regions in each sampling strategy roughly the same. Concretely, we split the images into $3 \times 3 = 9$ grids for the grid sampling strategy and sample 9 region proposals for the random sampling strategy. For these two strategies, we take 4 permutations of the regions to obtain richer bag-of-regions embeddings so that there would be 36 regions for each image. For the neighborhood sampling, we introduce a reduced version of the strategy that restricts the number of region proposals per image to 12 and takes 1 bag per proposal. This is because we record that the average number of regions in a bag is 3 with $p_b = 0.3$ and $\alpha = 3.0$, meaning the average sampled regions for each image in the neighborhood sampling strategy is close to $3 \times 12 = 36$.

The grid sampling strategy, whose fixed regions may either contain too many objects or only small parts of an object, achieves 25.4 mAP_{50} on the novel categories as shown in Table 5(#1). For the random sampling strategy, regions in a bag have different sizes and shapes and the distance between the regions can be large, which hinders the image encoder from representing the bag of regions exactly. It achieves 27.3 mAP_{50} as shown in Table 5(#2). While our neighborhood sampling strategy utilizes the same amount of regions, we record 32.2 mAP on the novel categories in Table 5(#3). Compared to these two baselines, our neighborhood sampling strategy captures potential objects in the vicinity of region proposals and ensures the teacher embedding exactly represents the bag of regions by sampling the neighboring boxes of region proposals. Note that our result in Table 5(#4) is achieved with 3 groups per region proposal and no limitation on the number of region proposals. More details on how we develop the sampling strategy are in the appendix.

Box Overlap between Regions in a Bag. We let the sampled regions in a group overlap at a certain IOF. In Table 6, we show how the overlap between boxes affects the performance. The scalar smaller than 0.0 in Table 6(#1) means that we keep an interval between the sampled boxes. The best performance (34.0 mAP_{50}) comes with an overlap of 0.1. Table 6(#1 – 3) indicate that the regions in a group need

Table 5. Exploring sampling strategies to obtain bag of regions

#	Sampling strategy	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
1	Grid	25.4	58.0	49.5
2	Random	27.3	53.3	46.5
3	Ours (reduced)	32.2	58.3	51.5
4	Ours	34.0	60.4	53.5

Table 8. Number of sampled bags per region proposal

#bags	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
1	32.6	58.5	51.7
3	34.0	60.4	53.5
5	33.2	60.0	53.0

Table 6. Overlap (IOF) between regions

#	Overlap	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
1	-0.1	32.5	59.7	52.6
2	0.0	33.6	60.2	53.2
3	0.1	34.0	60.4	53.5
4	0.2	33.8	59.8	53.0
5	0.3	33.7	60.0	53.1

Table 9. Number of pseudo words

#	#words	$\text{AP}_{50}^{\text{novel}}$	$\text{AP}_{50}^{\text{base}}$	AP_{50}
1	2	31.6	59.5	52.2
2	4	33.1	60.1	53.0
3	6	34.0	60.4	53.5
4	8	33.5	59.9	53.0

to have a continuity of semantics. Table 6(#3 – 5) indicate that the regions also need to cover diverse image contents.

Sampling Probability. We study the effect of the probability p_b to sample the candidate boxes, which affects the number of regions in a bag. The details are in Sec. 3.2. Note that the p_b will be adjusted by the aspect ratio scaled with the scale factor α . We fix α as 3.0 and sample 3 bags for each region proposal. We observe that the best result on novel categories (34.0 mAP_{50}) is achieved with $p_b = 0.3$ in Table 7.

Number of Bags Per Proposal. We show the effect of the number of sampled bags (#bags) for each region proposal in Table 8. For details, please refer to Sec 3.2. We fix α as 3.0 and the sampling probability p_b as 0.3. We observe that the best result on novel categories (34.0 mAP_{50}) is achieved with three bags of regions for each region proposal in Table 8.

Number of Pseudo Words Per Region. As an object category often needs many words to reach a precise description, we study the number of pseudo words (#words) predicted for each region of interest. Table 9(#1 – 3) show that the performance on novel categories increases with the number of pseudo words, indicating that stacking more pseudo words to a certain extent can strengthen the detector’s ability to distinguish object categories. However, the results in Table 9(#1 – 3) show that further increasing the number of pseudo words does not bring performance gain, and redundant words can even harm the performance.

4.3. Further Analysis

We show qualitative results in this section to further analyze the effectiveness of our method.

Co-occurrence of Objects. We first illustrate how a pre-trained VLM [38] captures the co-occurrence of objects by comparing the image-text similarity in Fig. 6. We use the CLIP-ViT-B/32 model pre-trained on more than 400 million image-text pairs to obtain the image and text embeddings. For each image, we incrementally add the object categories that appear to the text description. We observe that the similarity score between the image and text embeddings increases




Figure 5. Qualitative comparisons of BARON and the individual-level baseline in Table 4(#1). **Top:** Red boxes are for the novel categories and blue for the base categories. **Bottom:** Feature map’s responses to the queried novel object categories. From (a) to (c), the queried novel categories are ‘skateboard’, ‘airplane’ and ‘dog’, respectively. BARON detects objects of novel categories that are missed by the baseline.

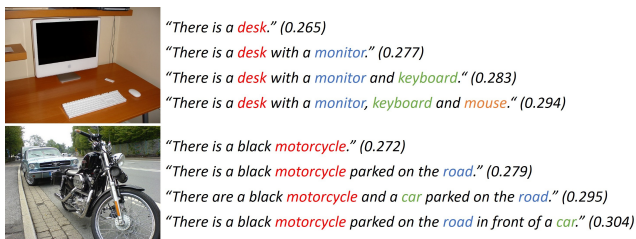


Figure 6. For each image, we incrementally add the object categories that appear in it to the text description. The similarity score between the image and text embeddings increases as the text description becomes more complete and precise.

as the text description includes more concepts as shown in Fig. 6. The examples reveal that the large-scale VLMs could capture the co-occurrence of multiple concepts in an image, although they are not explicitly trained to do so. We believe this is because each image-text pair naturally contains multiple concepts and the VLMs could implicitly learn the underlying connections when training on massive-scale image-text pairs. We also find a simple relationship between objects can also be captured by the VLMs, *e.g.*, the similarity increases when we add the relation ‘in front of’ to the text description in the second example of Fig 6.

Visualization. We further visualize the predictions of detectors learned through BARON and the individual-level baseline in Fig. 5. The images are from COCO’s validation set. We visualize the feature map’s response to the novel categories using the Grad-CAM++ [4]. We find that the model learned through BARON generates responses at locations of novel categories while the individual-level baseline induces weaker, incomplete or diffused responses. We also notice that semantically related objects can respond to the queried object category. For example, in Fig 5(c) the location of a flock of sheep chased by a dog responds to the queried object ‘dog’. Note the ‘dog’ and the ‘flock of sheep’ are not in

neighboring regions and the ‘dog’ is much smaller in size. Even though we form bags of regions by neighboring regions of equal size in training, the model has the ability to capture the relationship between objects with imbalanced sizes or large distance. This phenomenon resembles the generalization ability of human language, where learned concepts can be applied to describe or recognize new things.

5. Discussion and Conclusion

This paper goes beyond the learning of individual regions to bag of regions in OVD, exploring the ability of large-scale VLMs to represent the compositional structure of multiple concepts that naturally exists in image-text pairs. We develop a neighborhood sampling strategy to group contextually related regions into a bag and adopt the contrastive learning approach to align the bag-of-regions representations of the detector and pre-trained VLMs, which achieves new state-of-the-art performance on multiple OVD benchmarks.

The compositional structure explored in this paper is mainly about the co-occurrence of objects, and behaves like bag of words [45, 51]. The more complex compositional structure in the language is still under-explored and whether modern pre-trained vision-language models capture such structure still remains an open problem for the community. We look forward to further unveiling the behavior of the VLMs, and more importantly endowing the VLMs with human-like compositional representation to move to more generalized intelligence.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s). The work is also supported by Singapore MOE AcRF Tier 2 (MOE-T2EP20120-0001) and NTU NAP Grant. We also thank Lumin Xu and Yuanhan Zhang for valuable discussions.

References

- [1] Edward H. Adelson. On seeing stuff: the perception of materials by humans and machines. In Bernice E. Rogowitz and Thrasylvoulos N. Pappas, editors, *Human Vision and Electronic Imaging VI*, SPIE Proceedings, pages 1–12, 2001. **1**
- [2] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *Eur. Conf. Comput. Vis.*, 2018. **2**
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Eur. Conf. Comput. Vis.*, 2020. **2, 3**
- [4] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks. In *IEEE Winter App. Comput. Vis.*, 2018. **8**
- [5] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. Hybrid task cascade for instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. **2**
- [6] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. **5**
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. **6**
- [8] Berkan Demirel, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis. Zero-shot object detection by hybrid region embedding. In *Brit. Mach. Vis. Conf.*, 2018. **2**
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. Learn. Represent.*, 2021. **4, 5**
- [10] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. **1, 2, 3, 5, 6**
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. **6**
- [12] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Adv. Neural Inform. Process. Syst.*, 2013. **2**
- [13] Mingfei Gao, Chen Xing, Juan Carlos Niebles, Junnan Li, Ran Xu, Wenhao Liu, and Caiming Xiong. Open vocabulary object detection with pseudo bounding-box labels. *Eur. Conf. Comput. Vis.*, 2022. **2, 5, 6**
- [14] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. **6**
- [15] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *Int. Conf. Learn. Represent.*, 2021. **1, 2, 3, 5, 6**
- [16] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5356–5364, 2019. **5**
- [17] Nasir Hayat, Munawar Hayat, Shafin Rahman, Salman H. Khan, Syed Waqas Zamir, and Fahad Shahbaz Khan. Synthesizing the unseen for zero-shot object detection. In *Asian Conf. Comput. Vis.*, 2020. **2**
- [18] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. **4**
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Int. Conf. Comput. Vis.*, 2017. **5**
- [20] Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. In *Adv. Neural Inform. Process. Syst.*, 2014. **2**
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Int. Conf. Mach. Learn.*, 2021. **1, 2**
- [22] Wonjae Kim, Bokyung Son, and Ildoo Kim. ViLT: Vision-and-language transformer without convolution or region supervision. In *Int. Conf. Mach. Learn.*, 2021. **2**
- [23] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9404–9413, 2019. **1**
- [24] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *Int. Conf. Learn. Represent.*, 2022. **2**
- [25] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Int. Conf. Mach. Learn.*, 2022. **2**
- [26] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *Adv. Neural Inform. Process. Syst.*, 2021. **2**
- [27] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. **2**

- [28] Chuang Lin, Peize Sun, Yi Jiang, Ping Luo, Lizhen Qu, Gholamreza Haffari, Zehuan Yuan, and Jianfei Cai. Learning object-language alignments for open-vocabulary object detection. *Int. Conf. Learn. Represent.*, 2023. 5
- [29] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *Int. Conf. Comput. Vis.*, 2017. 5
- [30] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Int. Conf. Comput. Vis.*, 2017. 3
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Eur. Conf. Comput. Vis.*, 2014. 4, 5, 6
- [32] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul W. Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. *Int. J. Comput. Vis.*, 2020. 2
- [33] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. ViLbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Adv. Neural Inform. Process. Syst.*, 2019. 2
- [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 *Eur. Conf. Comput. Vis.*, pages 512–531, 2022. 5, 6
- [35] Mohammad Norouzi, Tomáš Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. In *Int. Conf. Learn. Represent.*, 2014. 2
- [36] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4
- [37] Chao Peng, Tete Xiao, Zeming Li, Yuning Jiang, Xiangyu Zhang, Kai Jia, Gang Yu, and Jian Sun. MegDet: A large mini-batch object detector. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2018. 5
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn.*, 2021. 1, 2, 4, 5, 7
- [39] Shafin Rahman, Salman H. Khan, and Nick Barnes. Transductive learning for zero-shot object detection. In *Int. Conf. Comput. Vis.*, 2019. 2
- [40] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. DenseCLIP: Language-guided dense prediction with context-aware prompting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2
- [41] Hanoona Abdul Rasheed, Muhammad Maaz, Muhammad Uzair Khattak, Salman Khan, and Fahad Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. In *Adv. Neural Inform. Process. Syst.*, 2022. 5, 6
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, 2015. 2, 3, 5, 6
- [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 *Int. Conf. Comput. Vis.*, 2019. 6
- [44] Hengcan Shi, Munawar Hayat, Yicheng Wu, and Jianfei Cai. ProposalCLIP: Unsupervised open-category object proposal generation via exploiting clip cues. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2
- [45] Ajinkya Tejanekar, Maziar Sanjabi, Bichen Wu, Saining Xie, Madian Khabza, Hamed Pirsiavash, and Hamed Firooz. A fistful of words: Learning transferable visual models from bag-of-words supervision. *CoRR*, abs/2112.13884, 2021. 8
- [46] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Eur. Conf. Comput. Vis.*, 2020. 2
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inform. Process. Syst.*, 2017. 4
- [48] Fangyun Wei, Yue Gao, Zhirong Wu, Han Hu, and Stephen Lin. Aligning pretraining for detection via object-level contrastive learning. *Adv. Neural Inform. Process. Syst.*, 2021. 5, 6
- [49] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 5
- [50] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 5
- [51] Mert Yükekşönül, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bags-of-words, and what to do about it? *CoRR*, abs/2210.01936, 2022. 8
- [52] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary DETR with conditional matching. *Eur. Conf. Comput. Vis.*, 2022. 1, 2, 5, 6
- [53] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 2, 5, 6
- [54] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2
- [55] Ye Zheng, Ruoran Huang, Chuanqi Han, Xi Huang, and Li Cui. Background learnable cascade for zero-shot object detection. In *Asian Conf. Comput. Vis.*, 2020. 2
- [56] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. RegionCLIP: Region-based language-image pretraining. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2, 5
- [57] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *Eur. Conf. Comput. Vis.*, 2022. 1, 2

- [58] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *Eur. Conf. Comput. Vis.*, 2022. [2](#), [5](#)
- [59] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable transformers for end-to-end object detection. *Int. Conf. Learn. Represent.*, 2021. [2](#), [6](#)