

# Generative Region-Language Pretraining for Open-Ended Object Detection

Chuang Lin<sup>1</sup>   Yi Jiang<sup>2\*</sup>   Lizhen Qu<sup>1</sup>   Zehuan Yuan<sup>2</sup>   Jianfei Cai<sup>1\*</sup>  
<sup>1</sup> Monash University   <sup>2</sup> ByteDance Inc.

## Abstract

In recent research, significant attention has been devoted to the open-vocabulary object detection task, aiming to generalize beyond the limited number of classes labeled during training and detect objects described by arbitrary category names at inference. Compared with conventional object detection, open vocabulary object detection largely extends the object detection categories. However, it relies on calculating the similarity between image regions and a set of arbitrary category names with a pretrained vision-and-language model. This implies that, despite its open-set nature, the task still needs the predefined object categories during the inference stage. This raises the question: What if we do not have exact knowledge of object categories during inference? In this paper, we call such a new setting as generative open-ended object detection, which is a more general and practical problem. To address it, we formulate object detection as a generative problem and propose a simple framework named *GenerateU*, which can detect dense objects and generate their names in a free-form way. Particularly, we employ Deformable DETR as a region proposal generator with a language model translating visual regions to object names. To assess the free-form object detection task, we introduce an evaluation method designed to quantitatively measure the performance of generative outcomes. Extensive experiments demonstrate strong zero-shot detection performance of our *GenerateU*. For example, on the LVIS dataset, our *GenerateU* achieves comparable results to the open-vocabulary object detection method GLIP, even though the category names are not seen by *GenerateU* during inference. Code is available at: <https://github.com/FoundationVision/GenerateU>.

## 1. Introduction

Object detection, a crucial task in computer vision, has made significant progress in recent years, particularly with the advent of deep learning [6, 8, 18, 41, 46]. The objective of this task is to locate objects in images with tight bounding boxes and recognize their respective categories. However,

\*Corresponding Authors

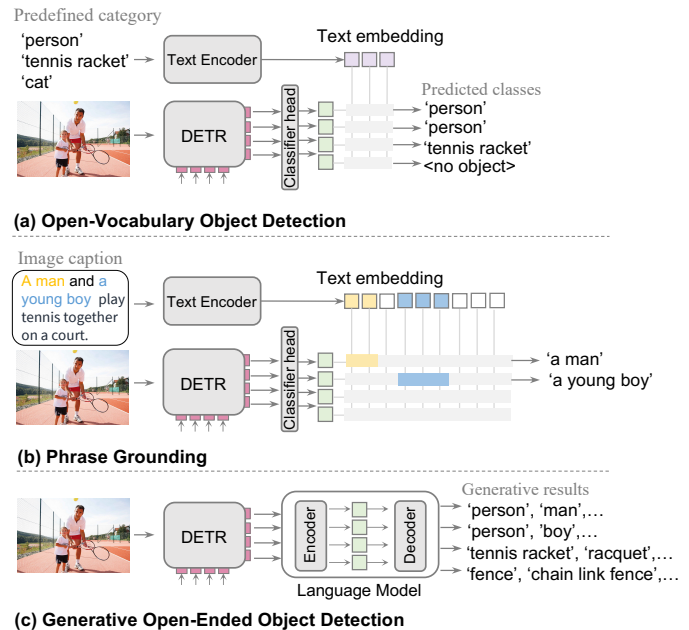


Figure 1. Comparing generative open-ended object detection with other open-set object detection tasks. Open-vocabulary object detection and phrase grounding typically require predefined categories or phrases in text prompts to align with image regions. In contrast, our introduced generative open-ended object detection is a more general and practical setting where categorical information is not explicitly defined. Such a setting is especially meaningful for scenarios where users lack precise knowledge of object categories during inference.

most existing object detection algorithms are restricted to a predetermined set of object categories defined in detection datasets [9, 14, 17, 26, 42]. For example, an object detector trained on COCO [9] can well identify 80 pre-defined classes but encounters difficulty in recognizing new classes that have not been seen or defined during training.

Open-vocabulary learning has been introduced to remedy this problem and considerable efforts have been made [7, 22, 28, 29, 35, 55, 65]. Open-vocabulary object detection addresses the closed-set limitation typically via

weak supervision from language, *i.e.*, image-text pairs, or large pretrained Vision-Language Models (VLMs), such as CLIP [39]. The text embeddings from the text encoder of CLIP can well align with the visual regions for the novel categories or the noun phrases in image captions, as shown in Fig. 1. However, despite open-set nature, open-vocabulary object detection still requires predefined object categories during the inference stage, while in many practical scenarios, we might not have exact knowledge of test images. Even if we have prior knowledge of test images, manually defining object categories might introduce language ambiguities (*e.g.*, similar object names like “person”, “a man”, “young boy” in Fig. 1) or be not comprehensive enough (*e.g.* missing “chain link fence” in Fig. 1), ultimately restricting flexibility. If we do not have prior knowledge of test images, to cover all possible objects that may appear in the images, the common solution to design a large and comprehensive label set that would be complex and time-consuming. The above discussions raise the research question: *Can we do open-world dense object detection that does not require predefined object categories during inference?* We call this problem *open-ended object detection*.

Therefore, in this paper, we formulate open-ended object detection as a generative problem and propose a simple architecture termed GenerateU to address it. Our GenerateU contains two major components: a visual object detector to localize image regions and a language model to translate visual regions to object names. This model is trained end-to-end, optimizing both components simultaneously. To broaden the knowledge of GenerateU, we train it using a small set of human-annotated object-language paired data and scale up the vocabulary size with massive image-text pairs of data. Furthermore, due to the limitation in image captions whose nouns often fall short of accurately representing all objects in the images, we employ a pseudo-labeling method to supplement missing objects in images.

Our major contributions are summarized as follows.

- We introduce a new and more practical object detection problem: **open-ended object detection**, and formulate it as a generative problem. The reformulation avoids manual effort in predefining object categories during both training and inference, aiming at a more general and flexible architecture.
- We develop a novel end-to-end learning framework, termed GenerateU, which directly generates object names in a free-form manner. GenerateU is trained with a small set of human-annotated object-language paired data and massive image-text paired data.
- Compared with open-vocabulary object detection models, GenerateU achieves comparable results on zero-shot LVIS, even though it does not see object categories during inference.

## 2. Related Work

### 2.1. Open-Vocabulary Object Detection

Learning visual models from language supervision becomes popular in image recognition tasks [19, 20, 39, 40, 58, 59, 62, 64]. Recently, many works [23, 23–25, 34, 43–45, 48, 49, 52–54, 56, 57] have extended the success in open-vocabulary image recognition to object detection, *i.e.*, open-vocabulary object detection (OVOD). As the pioneering work of OVOD, OVRCNN [61] successfully transfers the vision and language model trained on image-text pairs to the detection framework. To leverage the capability of CLIP [39], ViLD [16] and RegionCLIP [63] turn to learn the visual region feature from classification-oriented models by knowledge distillation. OV-DETR [60] proposes a novel open-vocabulary detector built upon DETR [8] by formulating the classification as a binary matching between the input queries and the referring objects. Different from distilling frozen vision encoder, OWL [35] proposes a two-step framework for contrastive image-text pre-training and end-to-end detection fine-tuning, using a standard Vision Transformer [13]. Detic [65] proposes to extend the detector vocabulary with large-scale image classification data by applying supervision to the proposal with the largest spatial size. To reduce the annotation cost for open-vocabulary object detection, VLDet [29] proposes a novel framework to directly learn from image-text paired data by formulating object-language alignment as a set matching problem.

Open vocabulary object detection aims at generalization to novel categories that were not seen during the training phase. Existing OVOD works typically rely on calculating the similarity between image regions and a set of arbitrary category names with a pretrained vision-language model. It remains necessary to predefine the categories during inference. In contrast, this work delves deeper into a new challenge - how to effectively handle scenarios where the exact categories remain unknown during the inference process.

### 2.2. Multimodal Large Language Model

Inspired by the remarkable capability of LLM [1, 4, 47, 50], multimodal foundation models [2, 15, 27, 30, 37, 38] leverage LLM as the core intelligence and are trained with extensive image-text pairs, resulting in enhanced performance across various vision-and-language tasks such as VQA and image captioning. BLIP-2 [27] proposes a generic and efficient pre-training strategy that bootstraps vision-language pre-training from off-the-shelf frozen pre-trained image encoders and frozen large language models. By inserting adapters into LLaMA’s transformer, LLaMA-Adapter [15] introduces only 1.2M learnable parameters, and turns a LLaMA into an instruction-following model. LLaVA [30] represents a novel end-to-end trained large multimodal model that combines a vision encoder and Vicuna [36]

for general-purpose visual and language understanding. It achieves impressive chat capabilities, emulating the essence of the multimodal GPT-4.

Based on this amazing progress, subsequent studies further explore enhancing the object localization capabilities of multimodal models. Kosmos-2 [37] unlocks the grounding capability, enabling the model to provide visual answers such as bounding boxes, thereby supporting a wider array of vision-language tasks, including referring expression comprehension. UNINEXT[55] reformulates diverse instance perception tasks into a unified object discovery and retrieval paradigm and can flexibly perceive different types of objects by simply changing the input prompts. DetGPT [38] locates the objects of interest based on the language instructions and the LLM enables users to interact with the system, allowing for a higher level of interactivity like visual reasoning. In contrast to these works that focus on understanding natural language expressions referring to specific objects in an image, *our work aims to identify and localize all objects of interest in an image, associating each with a class name.*

### 2.3. Dense Captioning

Dense captioning aims to predict a set of descriptions across various image regions. DenseCap [21] develops a fully convolutional localization network to extract image regions, which are then described using an RNN language model. In dense captioning, it treats each region as an individual image and tries to generate captions encompassing multiple objects and their attribute, such as “cat riding a skateboard” and “bear is wearing a red hat”. To evaluate the generated captions, dense captioning employs the METEOR [21] metric, originally designed for machine translation, which assesses the generation results from the sentence perspective. In comparison, our proposed generative open-ended object detection focuses on describing individual objects as category names in a zero-shot setting.

Note that recently, CapDet [32] and GRiT [51] attempted to integrate open-vocabulary object detection and dense captioning by introducing an additional language model as a captioning head. However, for open-world object detection, they either still require a predefined category space during inference, or are limited to the close-set object detection problem.

## 3. Method

Our goal is to build an open-world object detector capable of localizing all objects in an image and providing their corresponding category names in a free-form way. Figure 2 gives an overview of our proposed open-ended object detection model, GenerateU, which comprises two major components: an open-world object detector and a language model. We describe the open-world object detector in Sec. 3.1, which consists of an image encoder and a detection head,

and the language model, which can be directly borrowed from an existing Multimodal Large Language Model (including a V-to-L adaptor and a language model, Sec. 3.2) or retrained (Sec. 3.3). An additional pseudo-labeling method is described in Sec. 3.4 for enriching the label diversity.

### 3.1. Open-World Object Detection

The first crucial step is to extract accurate object regions from images, for which we develop an open-world object detector. Although we are not constrained to any specific object detection model, we choose Deformable DETR [66] for the following reasons. First, it alleviates heuristics when compared to other popular one or two-stage frameworks, and provides a versatile query-to-instance pipeline. Besides, Deformable DETR [66] demonstrates faster convergence speed and enhanced accuracy compared to the original DETR [8]. Deformable DETR or the original one employs Hungarian matching to learn a mapping from predicted queries to ground truth objects, followed by training the matched object query to regress to its corresponding ground truth with a combination of classification loss and bounding box regression loss. In the context of open-set problems, we do not rely on the object class information. Instead, we adopt an open-world detection approach, *i.e.* class-agnostic object detector, wherein matched queries are only classified into foreground or background. Consequently, the detection process involves three losses: binary cross-entropy, generalized IoU, and L1 regression loss.

### 3.2. Transferring from a Frozen Multimodal LLM

Once object candidates are localized, a natural idea is to leverage a pretrained multimodal large language model (MLLM) to reduce the training cost and offer strong zero-shot language abilities for generating object names. Recently developed MLLMs [2, 15, 27, 30, 37, 38] typically consist of three components: an image encoder to extract visual features; an adaptor network to bridge the modality gap and reduce the number of image embeddings, converting visual feature into token representations; a language model to generate predictions for vision-centric tasks. As shown in Figure 2 Top, we integrate a class-agnostic Deformable DETR, with a frozen image encoder to extract object queries, into an MLLM. We keep the MLLM frozen, and the only trainable module is the detection head for learning object representations. Besides, images are preprocessed to a resolution of  $224 \times 224$  pixels during training in KOSMOS-2 [37]. In object detection, it is common to employ a larger input size for multi-scale feature extraction. To handle changes in image size, we utilize linear interpolation to enlarge the original position embedding.

Despite its simplicity, our empirical study shows that such a direct open-end object detection based on pretrained MLLM results in poor performance large due to the do-

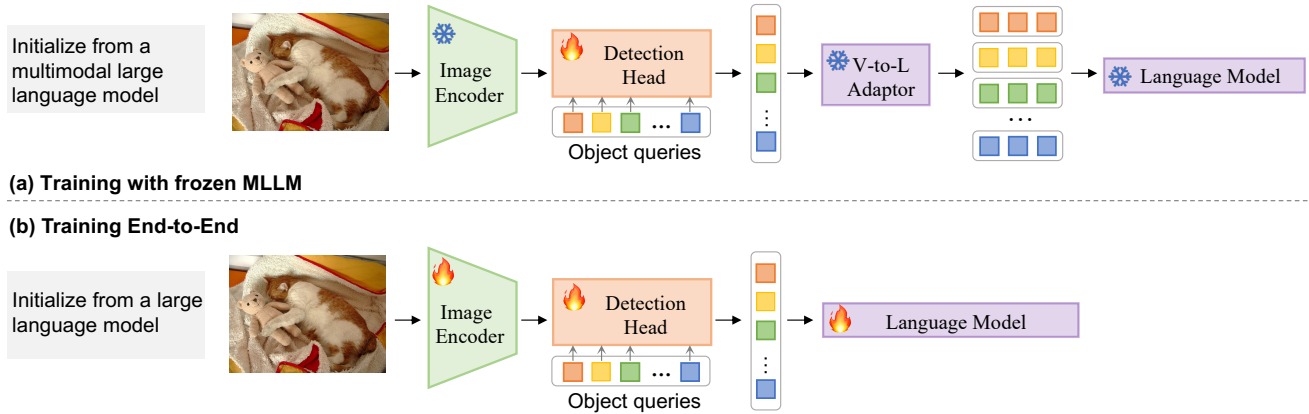


Figure 2. Overview of our proposed open-ended object detection model, GenerateU, which comprises two major components: an object detector and a language model. We compare two training strategies: (*Top*) We incorporate the class-agnostic DETR (with frozen image encoder) into a pre-trained and frozen Multimodal Large Language Model (including Adaptor and Language Model), to facilitate a smooth transfer of knowledge from the language domain to object detection; (*Bottom*) We activate the image encoder and language model as trainable components, taking an end-to-end approach to seamlessly integrate region-level understanding into the language model.

main shift. MLLMs were trained on large-scale image-text paired data, where they generate text descriptions given each whole image as a condition, without a fine-grained understanding of individual objects.

### 3.3. Generative Region-Language Pretraining

Thus, instead of employing a frozen MLLM, we propose GenerateU to directly link the open-world object detector with a language model, and activate both the image encoder and the language model as trainable components, as shown in Figure 2. Particularly, we employ an encoder-decoder-based language model, where visual representation serves as input for the encoder, and the associated text serves as the generation target for the language decoder. Here we only use the object queries that have been successfully matched with ground truth through Hungarian matching, ignoring the unmatched queries in the generating process. Specifically, each matched object query is first treated as an individual region, fed to a projection function to align its dimensions with the language model. The encoder comprises blocks with self-attention layers and feed-forward networks. The decoder, similar in structure, adds a cross-attention mechanism for interfacing with the encoder. The decoder employs an autoregressive self-attention mechanism, limiting the model’s attention to past outputs, and its final output is processed through a dense layer with a softmax function. We apply the language modeling loss to train the model.

**Exploring on the region-word alignment loss.** Benefiting from the pretrained vision-language model, we incorporate a region-word alignment loss during training to aid in learning to distinguish region features. In detail, for an object query, we treat its corresponding word as a positive sample

and consider other words in the same minibatch as negative samples. By calculating the similarity between the region features from class-agnostic Deformable DETR and text features from the fixed CLIP text encoder [39], the model is optimized by a binary cross-entropy loss. This region-word alignment loss is jointly trained with the object detection losses and the language modeling loss.

### 3.4. Enrich Label Diversity

Considerable efforts have been invested in collecting detection data that is both semantically rich and voluminous. However, human annotation proves to be both costly and resource-restrictive. Previous works, such as Kosmos-2 [37] and GLIP [28], leverage image-text pairs to expand the vocabulary scale. They use pretrained grounding models to generate bounding boxes aligned with noun phrases from captions. These boxes, in turn, serve as pseudo detection labels for training the model. However, given that captions may not comprehensively describe every object in an image, the count of nouns in the caption is typically much lower than the actual number of objects present in the corresponding image. Consequently, the quantity of pseudo-labels generated is constrained. To address this, we use GenerateU pretrained on the available labels to generate pseudo labels as a supplement for missing objects in images. Furthermore, we observed that beam search in language models naturally generates synonyms, thereby providing diverse object labels. Figure 3 provides pseudo-label examples, illustrating the generation of boxes covering nearly all objects in the image, along with a diverse set of rich vocabulary labels.

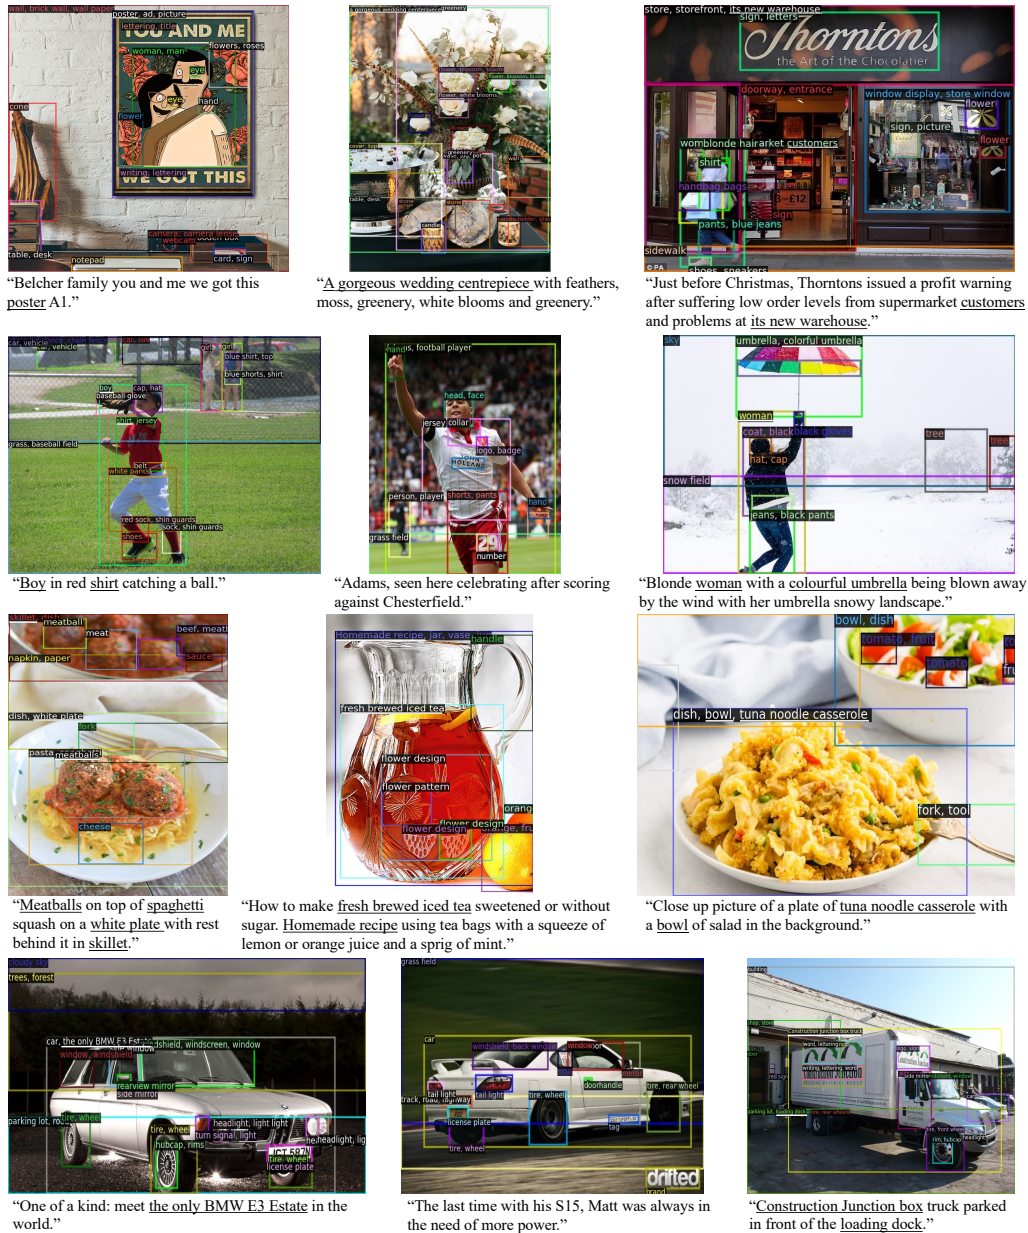


Figure 3. Selected pseudo-label examples highlight the generation of bounding boxes covering nearly all objects in the images. A varied set of descriptive labels showcases the model’s capability in producing diverse and linguistically rich vocabulary. A white underline in a figure indicates that the pseudo label comes from the noun phrase (black underline) in the caption.

## 4. Experiments

### 4.1. Datasets

Our model is trained with hybrid supervision from two types of data: detection data and image-text data. For detection data, we use Visual Genome (VG) [26] which contains 77,398 images for training. Different from previous works [32, 51], we do not use the region captioning annotations of VG, since they may contain descriptions that

encompass multiple objects or even describe actions. Our focus is directed towards the object detection task, and thus we leverage the object annotations of VG which annotate the same bounding box regions with concise and specific object names, covering a wide range of objects. For pre-training, we use about 5 million image-text pairs GRIT [37] from COYO700M [5] with pseudo grounding labels generated by KOSMOS-2 [37] using GLIP [28].

Method	Backbone	Pre-Train Data	Open-ended	Open-set	LVIS			
					APr	APc	APf	AP
MDETR [22]	RN101	LVIS	-	-	20.9	24.9	24.3	24.2
MaskRCNN [18]	RN101	LVIS	-	-	26.3	34.0	33.9	33.3
Deformable DETR [66]	Swin-T	LVIS	-	-	24.2	36.0	38.2	36.0
GLIP (C) [28]	Swin-T	O365, GoldG	-	✓	17.7	19.5	<b>31.0</b>	24.9
GLIP [28]	Swin-T	O365, GoldG, CAP4M	-	✓	<b>20.8</b>	21.4	<b>31.0</b>	26.0
GenerateU	Swin-T	VG	✓	✓	17.4	22.4	29.6	25.4
GenerateU	Swin-T	VG, GRIT5M	✓	✓	20.0	<b>24.9</b>	29.8	<b>26.8</b>
GenerateU	Swin-L	VG, GRIT5M	✓	✓	<b>22.3</b>	<b>25.2</b>	<b>31.4</b>	<b>27.9</b>

Table 1. Zero-shot domain transfer to LVIS. Our method achieves comparable performance to prior models in the close-ended setting, without requiring access to category names during inference.

## 4.2. Evaluation protocol and metrics

Addressing an open-ended problem presents challenges in evaluating performance. For instance, when annotating a laptop, the human-annotated ground truth category may be “laptop”, “computer”, “PC”, “Microcomputer”, *etc.* To assess the effectiveness of the open-end generative results, we propose two settings: (a) We calculate the similarity score between the generated and annotated category names with a fixed pretrained text encoder. Notice that this text encoder is not utilized by the generative detection model; rather, it is only used for quantitative performance measurement. Any text encoder capable of computing similarity between two sentences can be adopted for this purpose. In this study, we select the popular CLIP [39] text encoder and BERT-large [12] model. In this way, we can map the generative results to predefined categories in a dataset for evaluation. (b) We use METEOR [3] to assess the generated text quality, which is the same as the evaluation in dense captioning [21]. METEOR is widely used in NLP for evaluating machine translations. Following [21], we measure the mean Average Precision (AP) for both localization, using Intersection over Union (IoU) thresholds of .3, .4, .5, .6, and .7, and language accuracy, applying METEOR score thresholds at 0, .05, .1, .15, .2, and .25.

Following GLIP [28], we evaluate our GenerateU primarily on LVIS [17] which contains 1203 categories and report APs for rare categories ( $AP_r$ ), common categories ( $AP_c$ ), frequency categories ( $AP_f$ ) and all categories ( $mAP$ ), respectively. Specifically, following GLIP [28] and MDETR [22], we evaluate on the 5k minival subset and report the zero-shot fixed AP [11] for a fair comparison.

## 4.3. Implementation Details

We attempt two different backbones, Swin-Tiny and Swin-Large [31] as the visual encoder. The Deformable DETR architecture follows [66] with 6 encoder layers and 6 decoder layers. The number of object queries  $N$  is set to 300.

We adopt FlanT5-base [10] as the language model and initialize the weights from it. The optimizer is AdamW [33] with a weight decay of 0.05. The learning rate is set to  $2 \times 10^{-4}$  for the parameters of the visual backbone and the detection head, and  $3 \times 10^{-4}$  for the language model. For the frozen MLLM approach in Sec 3.2, we explore the state-of-the-art model KOSMOS-2 [37] which was trained with with the multimodal corpora. We fix all parameters and set the learning rate to  $2 \times 10^{-4}$  for the detection head. Our model is trained on 16 A100 GPUs.

## 4.4. Generative Open-Ended Detection Results

We assess the model’s capacity to identify diverse objects in a zero-shot setting on LVIS, presenting the results in Table 1, where rows 1-3 include methods utilizing LVIS as pre-training data, representing fully supervised models. Remarkably, we observe that even with only VG as training data, our model demonstrates commendable performance on zero-shot LVIS. This result suggests that predefined class names might not be necessary for open-world object detection during inference, especially when models are trained with extensive visual concepts. Furthermore, introducing additional training data from cost-effective image-text pairs consistently improves performance. For instance, when GenerateU trained with additional image-text pair data (GRIT [37]), we observe a substantial performance gain of 3.6 on  $AP_r$  and 1.4 on  $AP$ . In summary, the semantic richness derived from image-text data significantly enhances the model’s ability to recognize rare objects. Notably, our model achieves comparable performance to GLIP under the same training data size but without requiring access to category names during inference. Moreover, we enhanced our model by integrating a larger backbone, *i.e.* Swin-Large, confirming that our method can scale effectively with increased model complexity.

**Transfer to other datasets.** Our generative open-ended object detector can transfer to any dataset without any mod-

Methods	Image Encoder	Detection Head	Language Model	LVIS			
				APr	APc	APf	AP
Init from MLLM	Freeze	Train	Freeze*	5.3	10.2	19.7	14.3
	Train	Train	Freeze*	10.9	17.2	28.1	21.9
Init from LLM	Train	Train	Freeze	11.2	20.1	27.4	22.9
	Train	Train	Train	20.0	24.9	29.8	26.8

Table 2. Comparison of our GenerateU with two different training strategies on a pretrained LLM on zero-shot LVIS. The results highlight generative object detection benefits from end-to-end training. Here, \* refers to freezing both V-to-L adaptor and language model.

Methods	Pre-Train Data	COCO	Objects365
Supervised	-	46.5	25.6
GenerateU	VG	33.0	10.1
GenerateU	VG, GRIT	33.6	10.5

Table 3. We evaluate our GenerateU on COCO and Objects365 validation set in a zero-shot setting. GenerateU can transfer to various downstream datasets without requiring any modifications.

Metrics	APr	APc	APf	AP
CLIP [39]	20.0	24.9	29.8	26.8
BERT [12]	12.5	19.9	25.1	21.7
METEOR [3]	-	-	-	9.2

Table 4. Effect of different evaluation metrics on zero-shot LVIS.

ifications. In contrast to open vocabulary object detection methods, we do not need to switch the classifier to the specific category text embeddings. We evaluate the transferability of the proposed GenerateU on COCO [9] and Objects365 [42] validation sets, as shown in Table 3. The results demonstrate the effectiveness of our method and underscore the importance of scaling up the training data.

#### 4.5. Ablation Study

##### Comparison with transferring from a Frozen MLLM.

In Table 2, we evaluate the approach discussed in Sec 3.2, *i.e.* transferring a frozen MLLM to the generative object detection task. Following the state-of-the-art MLLM model KOSMOS-2 [37], we maintain its parameters fixed and only train the detection head. However, we observe that such direct transfer of the MLLM to object detection leads to poor performance in zero-shot LVIS. We attribute this result to the pretrained image encoder, which is trained on images of size  $224 \times 224$ , which is insufficient for the object detection task. Besides, the MLLM originally supervised by image-level annotations such as image captions, lacks understanding in both object localization and fine-grained object categories. We also experiment with making the image encoder trainable while keeping the language model frozen. This




Figure 4. Effect of beam size. Beam search plays a crucial role in generating rare object names, effectively addressing the long-tail problem.

adjustment significantly improves the results. We further fine-tune the entire model, initializing the Large Language Model (LLM) from FlanT5-base [10], which achieves the best performance. The above analysis underscores the importance of end-to-end training in the open-ended object detection task.

**Beam Search.** Towards a general and practical object detection system, we believe that beam search plays a crucial role in generating outputs that are closely aligned with human intuition. This is because it explores a broader array of possibilities during the decoding phase. Beam search naturally contains the multi-level object names and their synonyms, enhancing the system’s ability to recognize objects with varied linguistic expressions. This contributes to a more adaptable and nuanced understanding of the objects being detected. As shown in Figure 4, beam search largely improves the rare category performance, proving its importance for the long tail problem, and emphasizing the system’s effectiveness in handling less frequent object categories. In our experiments, we set the default beam size to 3.

**Evaluation Metrics.** To evaluate the results of generative object detection, we use different metrics: METEOR, evaluating the quality of text, and similarity scores, calculated by a fixed text encoder. Note that the off-the-shelf text encoder is not used for the proposed method, but just for the




Figure 5. Qualitative prediction results from GenerateU and ground truth on LVIS [17]. GenerateU produces complete and precise predictions, showcasing its ability to go beyond fixed vocabulary constraints.

Pre-Train Data	Region-word Alignment	APr	APc	APf	AP
VG	-	16.0	19.1	26.7	22.5
	✓	17.4 <sub>↑1.4</sub>	22.4 <sub>↑3.9</sub>	29.6 <sub>↑2.9</sub>	25.4 <sub>↑2.9</sub>
VG, GRIT	-	19.2	23.5	28.9	25.7
	✓	20.0 <sub>↑0.8</sub>	24.9 <sub>↑1.4</sub>	29.8 <sub>↑0.9</sub>	26.8 <sub>↑0.9</sub>

Table 5. Effect of region-word alignment loss.

evaluation. Here we compare two offline models: CLIP text encoder [39] and BERT [12], as presented in Table 4. We find that the CLIP text encoder is particularly adept at handling noun phrases, while the BERT encoder, predominantly trained on sentence corpora, may exhibit different strengths in processing textual information.

**Region-word Alignment.** We conduct ablation experiments to evaluate the effectiveness of the region-word alignment loss in our method, as presented in Table 5. We observe that training with only the detection loss and the language modeling loss yields promising results. Adding the region-word alignment loss provides additional supervision and further improves the performance. This further demonstrates the effectiveness of the proposed generative object detection method.

## 4.6. Qualitative Results

Here we provide qualitative visualizations in showcasing the superior object detection capabilities of our model. As shown in Figure 5, our model can identify a more diverse range of objects compared to ground truth annotations. These visualizations not only affirm the reliability of our detection framework but also demonstrate its ability to detect a wider variety of objects, providing valuable insights into the model’s strengths.

## 5. Conclusion

This paper introduces GenerateU, a novel approach for generative open-ended object detection. By reformulating object detection as a generative task, the model eliminates the need for predefined categories during inference, offering a more practical solution. Our method achieves comparable results with open-vocabulary models on zero-shot LVIS, even though categories are not seen in our GenerateU. Overall, this work contributes to the fast-evolving area of object detection towards to open-world scenarios. Future works include investigating the effect of training data scales and exploring methods beyond naive pseudo-labeling.



## References

- [1] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023. **2**
- [2] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. **2, 3**
- [3] Satyanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. **6, 7**
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. **2**
- [5] Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>, 2022. **5**
- [6] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6154–6162, 2018. **1**
- [7] Zhaowei Cai, Gukyeong Kwon, Avinash Ravichandran, Erhan Bas, Zhuowen Tu, Rahul Bhotika, and Stefano Soatto. X-detr: A versatile architecture for instance-wise vision-language tasks. *arXiv preprint arXiv:2204.05626*, 2022. **1**
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. **1, 2, 3**
- [9] 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. **1, 7**
- [10] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. **6, 7**
- [11] Achal Dave, Piotr Dollár, Deva Ramanan, Alexander Kirillov, and Ross Girshick. Evaluating large-vocabulary object detectors: The devil is in the details. *arXiv preprint arXiv:2102.01066*, 2021. **6**
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **6, 7, 8**
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **2**
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. **1**
- [15] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xianguyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. **2, 3**
- [16] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. **2**
- [17] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. **1, 6, 8**
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017. **1, 6**
- [19] Dat Huynh, Jason Kuen, Zhe Lin, Jiuxiang Gu, and Ehsan Elhamifar. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. *arXiv preprint arXiv:2111.12698*, 2021. **2**
- [20] 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 *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. **2**
- [21] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Denscap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4565–4574, 2016. **3, 6**
- [22] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. **1, 6**
- [23] Prannay Kaul, Weidi Xie, and Andrew Zisserman. Multi-modal classifiers for open-vocabulary object detection. *arXiv preprint arXiv:2306.05493*, 2023. **2**
- [24] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Contrastive feature masking open-vocabulary vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15602–15612, 2023.
- [25] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11144–11154, 2023. **2**

- [26] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. [1](#), [5](#)
- [27] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. [2](#), [3](#)
- [28] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. *arXiv preprint arXiv:2112.03857*, 2021. [1](#), [4](#), [5](#), [6](#)
- [29] 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. *arXiv preprint arXiv:2211.14843*, 2022. [1](#), [2](#)
- [30] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. [2](#), [3](#)
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. [6](#)
- [32] Yanxin Long, Youpeng Wen, Jianhua Han, Hang Xu, Pengzhen Ren, Wei Zhang, Shen Zhao, and Xiaodan Liang. Capdet: Unifying dense captioning and open-world detection pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15233–15243, 2023. [3](#), [5](#)
- [33] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [34] Chuofan Ma, Yi Jiang, Xin Wen, Zehuan Yuan, and Xiaojuan Qi. Codet: Co-occurrence guided region-word alignment for open-vocabulary object detection. *arXiv preprint arXiv:2310.16667*, 2023. [2](#)
- [35] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with vision transformers. *arXiv preprint arXiv:2205.06230*, 2022. [1](#), [2](#)
- [36] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023. [2](#)
- [37] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [38] Renjie Pi, Jiahui Gao, Shizhe Diao, Rui Pan, Hanze Dong, Jipeng Zhang, Lewei Yao, Jianhua Han, Hang Xu, and Lingpeng Kong Tong Zhang. Detgpt: Detect what you need via reasoning. *arXiv preprint arXiv:2305.14167*, 2023. [2](#), [3](#)
- [39] 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 *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [2](#), [4](#), [6](#), [7](#), [8](#)
- [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. *arXiv preprint arXiv:2112.01518*, 2021. [2](#)
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015. [1](#)
- [42] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. [1](#), [7](#)
- [43] Cheng Shi and Sibe Yang. Edadet: Open-vocabulary object detection using early dense alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15724–15734, 2023. [2](#)
- [44] Hengcan Shi, Munawar Hayat, and Jianfei Cai. Open-vocabulary object detection via scene graph discovery. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 4012–4021, 2023.
- [45] Hengcan Shi, Munawar Hayat, and Jianfei Cai. Unified open-vocabulary dense visual prediction. *arXiv preprint arXiv:2307.08238*, 2023. [2](#)
- [46] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, 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*, pages 14454–14463, 2021. [1](#)
- [47] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. [2](#)
- [48] Luting Wang, Yi Liu, Penghui Du, Zihan Ding, Yue Liao, Qiaosong Qi, Bialong Chen, and Si Liu. Object-aware distillation pyramid for open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11186–11196, 2023. [2](#)
- [49] Tao Wang. Learning to detect and segment for open vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7051–7060, 2023. [2](#)
- [50] BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*, 2022. [2](#)
- [51] Jialian Wu, Jianfeng Wang, Zhengyuan Yang, Zhe Gan, Zicheng Liu, Junsong Yuan, and Lijuan Wang. Grit: A generative region-to-text transformer for object understanding. *arXiv preprint arXiv:2212.00280*, 2022. [3](#), [5](#)

- [52] Jianzong Wu, Xiangtai Li, Henghui Ding, Xia Li, Guangliang Cheng, Yunhai Tong, and Chen Change Loy. Betrayed by captions: Joint caption grounding and generation for open vocabulary instance segmentation. *arXiv preprint arXiv:2301.00805*, 2023. [2](#)
- [53] Size Wu, Wenwei Zhang, Sheng Jin, Wentao Liu, and Chen Change Loy. Aligning bag of regions for open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15254–15264, 2023.
- [54] Xiaoshi Wu, Feng Zhu, Rui Zhao, and Hongsheng Li. Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7031–7040, 2023. [2](#)
- [55] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15325–15336, 2023. [1](#), [3](#)
- [56] Lewei Yao, Jianhua Han, Youpeng Wen, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, Chunjing Xu, and Hang Xu. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. *Advances in Neural Information Processing Systems*, 35:9125–9138, 2022. [2](#)
- [57] Lewei Yao, Jianhua Han, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, and Hang Xu. Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23497–23506, 2023. [2](#)
- [58] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. [2](#)
- [59] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [2](#)
- [60] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. In *European Conference on Computer Vision*, pages 106–122. Springer, 2022. [2](#)
- [61] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. [2](#)
- [62] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. *arXiv preprint arXiv:2111.07991*, 2021. [2](#)
- [63] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. [2](#)
- [64] Chong Zhou, Chen Change Loy, and Bo Dai. Denseclip: Extract free dense labels from clip. *arXiv preprint arXiv:2112.01071*, 2021. [2](#)
- [65] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *arXiv preprint arXiv:2201.02605*, 2022. [1](#), [2](#)
- [66] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [3](#), [6](#)