

# Boosting Segment Anything Model Towards Open-Vocabulary Learning

Xumeng Han<sup>1\*</sup>, Longhui Wei<sup>2†</sup>, Xuehui Yu<sup>1</sup>, Zhiyang Dou<sup>1</sup>, Xin He<sup>2</sup>,  
Kuiran Wang<sup>1</sup>, Yingfei Sun<sup>1</sup>, Zhenjun Han<sup>1†</sup>, Qi Tian<sup>2</sup>

<sup>1</sup>University of Chinese Academy of Sciences

<sup>2</sup>Huawei Inc.

{hanxumeng19, yuxuehui17, douzhiyang23, wangkuiran19}@mailsucas.ac.cn,  
{weilh2568, whut.hexin}@gmail.com, {yfsun, hanzhj}@ucas.ac.cn, tian.qi1@huawei.com

## Abstract

The recent Segment Anything Model (SAM) has emerged as a new paradigmatic vision foundation model, showcasing potent zero-shot generalization and flexible prompting. Despite SAM finding applications and adaptations in various domains, its primary limitation lies in the inability to grasp object semantics. In this paper, we present **Sambor** to seamlessly integrate SAM with the open-vocabulary object detector in an end-to-end framework. While retaining all the remarkable capabilities inherent to SAM, we boost it to detect arbitrary objects from human inputs like category names or reference expressions. Building upon the SAM image encoder, we introduce a novel SideFormer module designed to acquire SAM features adept at perceiving objects and inject comprehensive semantic information for recognition. In addition, we devise an Open-set RPN that leverages SAM proposals to assist in finding potential objects. Consequently, Sambor enables the open-vocabulary detector to equally focus on generalizing both localization and classification sub-tasks. Our approach demonstrates superior zero-shot performance across benchmarks, including COCO and LVIS, proving highly competitive against previous state-of-the-art methods. We aspire for this work to serve as a meaningful endeavor in endowing SAM to recognize diverse object categories and advancing open-vocabulary learning with the support of vision foundation models.

## 1 Introduction

Vision foundation models (Radford et al. 2021; He et al. 2022; Oquab et al. 2024) serve as robust backbones that excel across a diverse spectrum of vision tasks. The recent Segment Anything Model (SAM) (Kirillov et al. 2023) has garnered widespread attention within the community as a foundational visual model for general image segmentation. Trained with billion-scale mask labels, it demonstrates impressive zero-shot segmentation performance, seamlessly applied across a variety of applications through simple prompting (Ren et al. 2024; Wang et al. 2023). While exhibiting outstanding performance, SAM is constrained to scenarios with class-agnostic applications, necessitating a

\*This work was done during the internship in Huawei Inc.

†Corresponding Authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

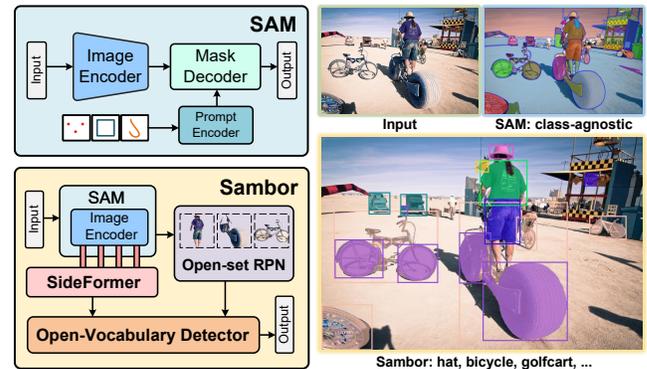


Figure 1: We develop an end-to-end open-vocabulary object detector called **Sambor**, building upon the vision foundation model SAM. Sambor enables SAM to recognize arbitrary object categories, bridging semantic gaps. It also leverages SAM’s generalization and interactive capabilities to enhance zero-shot performance and extend versatility.

more profound exploration to enhance its semantic understanding. In this work, we boost SAM towards open-vocabulary learning to detect objects of arbitrary categories, a paradigm commonly referred to as open-vocabulary object detection (Zareian et al. 2021; Wu et al. 2024).

Recent works (Gu et al. 2021; Zhong et al. 2022; Minderer et al. 2022; Kuo et al. 2023; Li et al. 2022a; Yao et al. 2022; Liu et al. 2023) commonly follow two lines to achieve open-vocabulary object detection. One seeks to expand the cognitive categories by transferring knowledge from pre-trained vision-language (VL) models (Radford et al. 2021; Jia et al. 2021). The other unifies the formulation of object detection and phrase grounding tasks, extending the available data scope from object detection to a diverse range of image-text pairs (Sharma et al. 2018; Vicente et al. 2016; Thomee et al. 2016). The utilization of large-scale data encourages the model to learn the feature alignment between object regions and language phrases. Previous efforts have been built upon conventional object detection models, introducing the VL paradigm to extend the object detection classifier from the close-set to the open-set domain. With the emergence of the powerful vision foundation model SAM,

its remarkable zero-shot localization capabilities and the flexibility for interactive prompting raise the prospect of elevating open-vocabulary object detection to new heights.

In this paper, we introduce **Sambor** (which stands for **SAM BOosteR**), as illustrated in Fig. 1. Sambor is an end-to-end open-vocabulary object detector that seamlessly integrates all functionalities from the SAM, inheriting its powerful zero-shot generalization and flexible prompting. We establish a ladder side transformer adapter, named SideFormer, on the frozen SAM image encoder. It incorporates an extractor designed to integrate features from the image encoder, drawing object perception capabilities from SAM. Subsequently, we devise an injector to augment the original features by introducing additional semantic information to assist category recognition. In pursuit of the injected features imbued with rich semantics, CLIP (Radford et al. 2021), another vision foundation model, naturally emerges as our excellent choice, given it achieves outstanding zero-shot classification through VL alignment. As a result, infusing CLIP visual features not only enhances semantic understanding but also narrows the gap with the text feature domain, providing convenience for the open-vocabulary object classifier (Gu et al. 2021; Kuo et al. 2023).

Moreover, SAM exhibits the capacity to generate high-quality class-agnostic proposals. To fully leverage this advantage and further augment the zero-shot localization ability, we develop the Open-set RPN. Specifically, Sambor adopts the two-stage object detection architecture (Ren et al. 2015), decoupling the detector into a first stage dedicated to generating high-quality proposals to ensure sufficient recall and a second stage focused on open-vocabulary classification. Open-set RPN complements the vanilla RPN (Ren et al. 2015) by introducing proposals oriented towards open-set scenarios. Thanks to the flexibility afforded by SAM, these additional proposals can be obtained through various prompts or the automatic mask generation pipeline.

We follow the GLIP (Li et al. 2022a) protocol and conduct experiments to comprehensively evaluate the effectiveness of Sambor in open-vocabulary object detection. Benefiting from the effective designs, Sambor demonstrates superior open-vocabulary detection performance on COCO (Lin et al. 2014) and LVIS (Gupta, Dollar, and Girshick 2019) benchmarks. It implies that we endow SAM with the capability to recognize arbitrary objects, boosting it to be more versatile. From another perspective, incorporating SAM into an end-to-end framework provides greater versatility compared to current state-of-the-art open-vocabulary detectors (Li et al. 2022a; Liu et al. 2023; Yao et al. 2022, 2023). For instance, it allows for seamless conversion of object detection results into instance segmentation or facilitates human-machine interaction through prompts. These capabilities are previously either unavailable or required the cascading of multiple models (Ren et al. 2024), introducing additional complexity and operational challenges. Given these encouraging results, we aspire to endow the vision foundation model SAM with recognition capabilities to address a broader spectrum of applications and offer a potential way for the development of open-vocabulary object detection.

## 2 Related Work

**General Object Detection**, a crucial computer vision task, consists of two sub-problems: finding the object (localization) and naming it (classification). Convolution-based detectors are typically divided into two-stage (Ren et al. 2015; He et al. 2017; Cai and Vasconcelos 2018) or single-stage (Redmon et al. 2016; Lin et al. 2017; Zhang et al. 2020), based on hand-crafted anchors or reference points. Recent transformer-based methods (Carion et al. 2020; Zhu et al. 2021) try to formulate object detection as a set prediction problem. These methods are constrained to predefined categories, whereas our approach aims to find and recognize objects of arbitrary category in an open domain.

**Open-Vocabulary Object Detection** (Zareian et al. 2021; Wu et al. 2024) has emerged as a new trend for modern object detection, which aims to detect objects of unbounded concepts using a more universal and practical paradigm (Zhou et al. 2022; Minderer et al. 2022; Zang et al. 2022; Xu et al. 2023; Cheng et al. 2024). Some of these methods (Gu et al. 2021; Zhong et al. 2022; Du et al. 2022; Kuo et al. 2023; Ma et al. 2023; Wu et al. 2023b,a) leverage vision-language models (VLMs) (Radford et al. 2021) to align information between regions and words through a multi-stage pre-training strategy. GLIP (Li et al. 2022a) pioneers an alternative approach that transforms the detection data into a grounding format and introduces a fusion module for simultaneous learning of vision-language alignment and object localization. Compared to previous methods, we utilize SAM to facilitate the finding of potential objects in the open domain, mitigating the issue where VLMs primarily focus on handling the intricate alignment between regions and texts, thus lacking in localization generalization.

**Segment Anything Model** (Kirillov et al. 2023) is an innovative image segmentation model trained on the dataset comprising over 1 billion masks, designed to robustly segment any object guided by diverse prompts. Influenced by its development, the community is focusing on developing more versatile and functional segmentation models (Zou et al. 2023a; Zhang et al. 2023b; Zou et al. 2023b; Li et al. 2023). In contrast, our method focuses on detecting objects of arbitrary categories and learning rich semantics from detection and grounding data, making it more suitable for open-world scenarios. We utilize SAM’s prompt functionality to obtain masks directly from predicted bounding boxes, thereby eliminating the need for training a mask head and overcoming limitations posed by segmentation data volume.

## 3 Preliminaries

### 3.1 Open-Vocabulary Object Detection

Given an image  $\mathbf{I} \in \mathbb{R}^{3 \times h \times w}$ , object detection involves solving the two sub-problems of (1) localization: find all objects with their location, represented as a box  $\mathbf{b}_j$  and (2) classification: assign a class label  $c_j \in \mathcal{C}^{\text{test}}$  to the  $j$ -th object. Here  $\mathcal{C}^{\text{test}}$  is the class vocabulary provided by the user at test time. Traditional object detection considers  $\mathcal{C}^{\text{test}} = \mathcal{C}^{\text{train}}$ , where  $\mathcal{C}^{\text{train}}$  denotes the vocabulary of detection dataset used during training. Open-vocabulary object detection allows  $\mathcal{C}^{\text{test}} \neq \mathcal{C}^{\text{train}}$ . Taking GLIP (Li et al. 2022a)

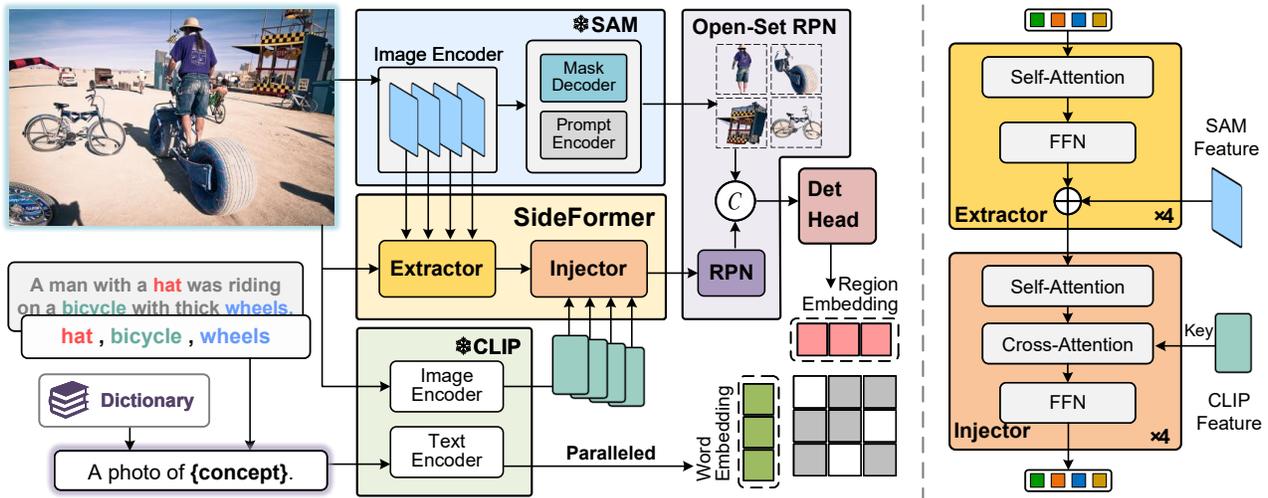


Figure 2: **Overall architecture of Sambor.** (Left) We adopt the SAM image encoder as the backbone and construct a SideFormer module to extract features and inject CLIP visual information for enhancing semantic understanding. Sambor is built upon a two-stage detector, with the first stage designed as an Open-set RPN that enhances the vanilla RPN using open-set proposals generated by SAM. The second stage is equipped with a CLIP language branch for parallel concept encoding, thereby endowing the detector with open-vocabulary classification. (Right) The specific implementations of the extractor and injector.

as an example, it reformulates detection as a grounding task, aligning each visual region with corresponding class content in text prompts. Following CLIP (Radford et al. 2021), the model takes image-text pairs as input, extracting features from both modalities through dedicated encoders. By replacing the linear classification layer with a region-word matching dot product layer, it converts from a traditional detector to an open-vocabulary detector. However, previous open-vocabulary detection approaches primarily focus on improving classification across arbitrary categories, while their localization capability is still heavily reliant on bounding box annotations provided in detection datasets (Yao et al. 2023). In this paper, we compensate for the previously overlooked ability to find potential objects in the open domain, allowing the detector to equally focus on enhancing zero-shot generalization for object localization and classification.

### 3.2 Vision Foundation Model

**SAM** (Kirillov et al. 2023) is composed of three modules: (1) Image encoder: a robust ViT-based (Dosovitskiy et al. 2020) backbone excels at extracting features from high-resolution images. (2) Prompt encoder: encoding the interactive positional information from the input points, boxes, or masks. (3) Mask decoder: a lightweight transformer-based (Vaswani et al. 2017) decoder efficiently translates the image and prompt embeddings into masks.

**CLIP** (Radford et al. 2021) leverages web-scale image-text pairs crawled from the Internet and simply aligning image features with text features via contrastive learning, delivering impressive results in zero-shot image classification. In this paper, we choose the CNN-based CLIP model (e.g., RN50×64 (He et al. 2016)) over the ViT-based (Dosovitskiy et al. 2020) one due to its superior compatibility with high-resolution image inputs (Kuo et al. 2023).

## 4 Methodology

An overview framework of Sambor is illustrated in Fig. 2. Our approach fully integrates SAM to leverage its exceptional zero-shot localization capabilities and the flexibility for interactive prompting, thus enhancing the performance and functionality of Sambor in open-world scenarios. Specifically, we employ the SAM image encoder as the backbone and freeze the parameters during training. Due to the absence of a semantic prior in SAM features, the performance in category-aware tasks falls short of optimal. To address this issue, we introduce a ladder-side transformer adapter named SideFormer (Sec. 4.1), conceptually similar to Sung, Cho, and Bansal (2022); Chen et al. (2022). It is designed to extract features from the backbone and inject comprehensive semantic information from CLIP into them.

Subsequently, we construct an open-vocabulary object detector on the backbone, which is based on the two-stage ViT-Det (Li et al. 2022b) with Cascade R-CNN (Cai and Vasconcelos 2018). For the two-stage detector, the first stage involves the RPN (Ren et al. 2015) for generating object proposals. We extend the RPN into an Open-set RPN (Sec. 4.2) by incorporating more zero-shot generalized proposals from SAM. In the second stage, we equip it with a text encoder for open-vocabulary classification (Sec. 4.3). Therefore, we utilize these two stages to foster zero-shot generalization in both localization and classification for Sambor.

### 4.1 SideFormer

We design an extractor for acquiring SAM features adept at perceiving objects and an injector for absorbing knowledge from CLIP to enhance semantic understanding.

**SAM Extractor.** We adopt the patch embedding layer with a structure identical to that in SAM (without parameter sharing), initially projecting the input image into visual tokens.

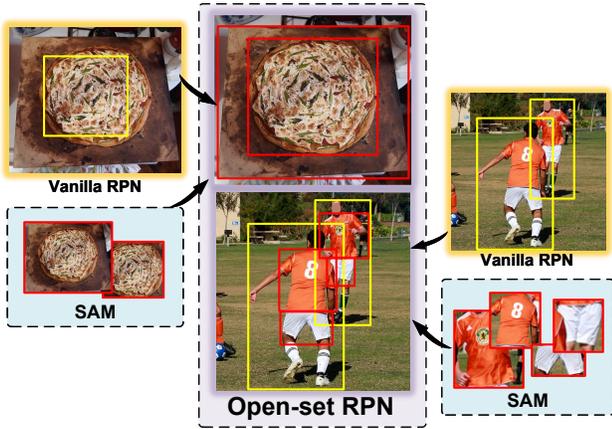


Figure 3: **An illustration of Open-set RPN.** We demonstrate two examples where SAM proposals effectively complement the vanilla RPN: **(Top-Left)** precise determination of object edge positions, and **(Bottom-Right)** clear capture of specific parts of an object, *e.g.*, a person’s clothing.

The visual tokens are first summed with those from SAM and then incrementally fused with the deeper SAM features via a series of transformer layers (Vaswani et al. 2017). Following the ViTDet (Li et al. 2022b) architecture, SAM divides the ViT (Dosovitskiy et al. 2020) into four blocks, applying global attention at the last layer in each block. In view of this, we design four corresponding transformer layers to extract features from each global attention output.

For each transformer layer, the input visual tokens  $\mathcal{F}_{\text{side}} \in \mathbb{R}^{\frac{hw}{16^2} \times D}$ , where  $D$  is the feature dimension, are encoded through a self-attention layer and a feed-forward network (FFN). Subsequently, we use a direct summation to accomplish the fusion with the extracted SAM features  $\mathcal{F}_{\text{sam}} \in \mathbb{R}^{\frac{hw}{16^2} \times D}$ . This process can be formulated as:

$$\hat{\mathcal{F}}_{\text{side}} = \mathcal{F}_{\text{side}} + \text{Attn}(\text{norm}(\mathcal{F}_{\text{side}})), \quad (1)$$

$$\mathcal{F}'_{\text{side}} = \mathcal{F}_{\text{sam}} + \gamma \cdot (\hat{\mathcal{F}}_{\text{side}} + \text{FFN}(\text{norm}(\hat{\mathcal{F}}_{\text{side}}))), \quad (2)$$

where  $\text{norm}(\cdot)$  is LayerNorm (Ba, Kiros, and Hinton 2016). We apply a learnable vector  $\gamma \in \mathbb{R}^D$  to modulate the transformer output, which is initialized with  $\mathbf{0}$ . This strategy guarantees that the feature distribution of SideFormer commences from SAM without undergoing drastic alterations.

**CLIP Injector.** After integrating features extracted from SAM, we imbue them with semantically rich information from the CLIP (Radford et al. 2021) visual encoder. The injector is also equipped with four transformer layers, supplemented by cross-attention modules for feature interaction.

We treat the input feature  $\mathcal{F}'_{\text{side}}$  as the query, and the CLIP feature  $\mathcal{F}_{\text{clip}}$  as the key and value. Initiate a self-attention encoding on the query first. Subsequently, we employ cross-attention on queries  $\mathcal{F}'_{\text{side}}$  and  $\mathcal{F}_{\text{clip}}$ , facilitating the assimilation of knowledge from CLIP. Finally, an FFN is appended, constituting the entirety of this transformer layer. Eq. 3 details the cross-attention module in the injector, while omitting the self-attention and FFN for brevity.

$$\mathcal{F}^*_{\text{side}} = \mathcal{F}'_{\text{side}} + \gamma \cdot \text{Attn}(\text{norm}(\mathcal{F}'_{\text{side}}), \text{norm}(\mathcal{F}_{\text{clip}})), \quad (3)$$

where the learnable parameter  $\gamma$  is introduced to likewise balance the injected CLIP features with the original inputs.

## 4.2 Open-Set RPN

The primary aim of object detection is to thoroughly find potential objects, which implies the necessity for the RPN in the first stage to achieve a sufficiently high recall. Hence, for the open-vocabulary object detector, possessing a robust RPN tailored for open-set domains is imperative.

To achieve this, we designed an Open-set RPN that integrates two proposal sources: the vanilla RPN (Ren et al. 2015) and SAM (Kirillov et al. 2023). We expect that the Open-set RPN will facilitate proposals from both sources to become complementary, aiming to find potential objects as comprehensively as possible. The vanilla RPN learns from detection data and is adept at handling scenarios involving common objects. However, since the RPN trained on closed-set data needs to expand its generalization, we incorporate proposals from SAM as a valuable supplement. We use the SAM head (consisting of the prompt encoder and the mask decoder) to generate masks and extract bounding boxes of these masks as object proposals. Given SAM’s robust zero-shot localization capability, its proposals can better cover areas that RPN has overlooked. Fig. 3 primarily illustrates two possible scenarios to show how SAM proposals serve as supplements. One is the enhanced precision in delineating the object contours, where SAM outperforms RPN in identifying the edge positions of overlapping objects or those with irregular shapes. The other scenario shows that SAM clearly captures parts of the whole object, *e.g.*, a person’s shirt and pants, areas where RPN tends to fall short on detail. Nevertheless, more than relying solely on SAM is required; a trainable RPN is crucial as it reinforces handling common objects and compensates for deficiencies with small objects (Kirillov et al. 2023) of SAM. Please refer to Sec. 5.3 and 5.4 for more detailed analysis and demonstrates.

Specifically, we leverage the automatic mask generation pipeline (Kirillov et al. 2023) with a point grid as prompts to predict mask proposals. Adjusting point density allows for control over the number of proposals, which in this paper defaults to a  $32 \times 32$  point grid as a balance between quantity and computational cost. We adopt the straightforward NMS to merge two sets of proposals. Given the domain gap in prediction scores between the two sets, it is unreasonable to apply NMS directly based on their scores. Therefore, we first merge the proposals from RPN (which have already been de-duplicated) with those from SAM. Then, we perform NMS (with the threshold set to 0.7) on SAM proposals using RPN boxes as references to filter out highly overlapping ones, thereby preserving areas not covered by RPN.

## 4.3 Open-Vocabulary Classification

The second stage of the detector focuses on transforming the proposals into a set of predicted bounding boxes  $B = \{\mathbf{b}_k\}_{k=1}^K$  ( $K$  is the number of predictions) along with the classification features  $\mathcal{F}^B \in \mathbb{R}^{K \times D}$ . To expand the classification into open vocabulary, we introduce a language branch, *i.e.*, the CLIP text encoder. We insert each concept

Method	Backbone	#Trainable Params	Pre-Train Data	COCO 2017val			
				AP <sup>box</sup>	AP <sup>box</sup> <sub>50</sub>	AP <sup>mask</sup>	AP <sup>mask</sup> <sub>50</sub>
DyHead-T <sup>†</sup> (Dai et al. 2021)	Swin-T (Liu et al. 2021)	≈100M	-	49.7	68.0	-	-
DyHead-T (Dai et al. 2021)	Swin-T (Liu et al. 2021)	≈100M	O365	43.6	-	-	-
GLIP-T (B) (Li et al. 2022a)	Swin-T (Liu et al. 2021)	232M	O365	44.9	61.5	-	-
GLIP-T (C) (Li et al. 2022a)	Swin-T (Liu et al. 2021)	232M	O365,GoldG	46.7	63.4	-	-
GLIP-T (Li et al. 2022a)	Swin-T (Liu et al. 2021)	232M	O365,GoldG,CC3M,SBU	46.6	63.1	-	-
DINO <sup>†</sup> (Zhang et al. 2023a)	Swin-T (Liu et al. 2021)	49M	-	54.4	72.9	-	-
DINO (Zhang et al. 2023a)	Swin-T (Liu et al. 2021)	49M	O365	46.2	-	-	-
G-DINO-T (Liu et al. 2023)	Swin-T (Liu et al. 2021)	172M	O365	46.7	-	-	-
G-DINO-T (Liu et al. 2023)	Swin-T (Liu et al. 2021)	172M	O365,GoldG	48.1	-	-	-
G-DINO-T (Liu et al. 2023)	Swin-T (Liu et al. 2021)	172M	O365,GoldG,Cap4M	48.4	64.4	-	-
ViTDet <sup>†</sup> (Li et al. 2022b)	ViT-B (Dosovitskiy et al. 2020)	141M	-	54.0	72.2	46.7	69.8
Sambor (Ours)	ViT-B (Dosovitskiy et al. 2020)	160M	O365	47.3	64.7	36.5	59.8
Sambor <sup>★</sup> (Ours)	ViT-B (Dosovitskiy et al. 2020)	160M	O365	<b>48.6</b>	<b>66.1</b>	<b>37.1</b>	<b>60.6</b>

Table 1: **Zero-shot transfer performance on COCO benchmark.** <sup>★</sup>denotes the application of Open-set RPN. <sup>†</sup>denotes supervised approaches.

name into the prompt template to form a complete sentence. These sentences are separately forwarded to the text encoder for obtaining sentence embeddings  $\mathcal{F}^T \in \mathbb{R}^{M \times D}$ , where  $M$  is the number of concepts sampled in each batch. We calculate the similarity matrix  $S = \mathcal{F}^B \cdot (\mathcal{F}^T)^\top \in \mathbb{R}^{K \times M}$  to construct the word-region alignment loss (Li et al. 2022a). In our design, text embeddings are solely utilized to predict word-region similarity scores without the additional cross-modal fusion adopted in (Li et al. 2022a; Liu et al. 2023).

**Unified Data Formulation.** Following (Yao et al. 2022, 2023), we employ a paralleled formulation to unify the data formats from object detection and phrase grounding.

- **Object Detection.** A concept set, based on dataset category names, designates present categories as positives and absent ones as negatives. Given that the category names in the dataset remain unchanged, extracting and storing category features before training can avoid redundant feature extraction and improve efficiency. Additionally, each category feature is the average of all its prompt templates, a policy applied during both training and inference.
- **Phrase Grounding.** We extract phrases corresponding to labeled objects from the caption to form a positive concept set. We utilize the large-scale and information-dense Bamboo (Zhang et al. 2022b) to diversify negative concepts. Given the variability of concept names in images, we dynamically select a subset of concepts as negatives for each batch, setting the total number of positive and negative concepts to 150. For efficiency, we randomly choose a prompt template to extract text features.

## 5 Experiments

### 5.1 Implementation Details

**Training Datasets.** For object detection, we use the Objects365 (Shao et al. 2019) dataset (referred to as O365),

comprising 365 categories. For phrase grounding, we use the GoldG (Kamath et al. 2021) dataset, which contains well-annotated images from sources including Flickr30K (Plummer et al. 2015), VG Caption (Krishna et al. 2017), and GQA (Hudson and Manning 2019). We deliberately exclude COCO (Lin et al. 2014) images to ensure a more equitable evaluation of zero-shot transfer performance.

**Training Details.** We pre-train our models using SAM with ViT-B (Dosovitskiy et al. 2020) as the backbone and CLIP with RN50×64 (He et al. 2016), using a batch size of 64. We select AdamW (Loshchilov and Hutter 2019) optimizer with a 0.05 weight decay, an initial learning rate  $4 \times 10^{-4}$ , and a cosine annealing learning rate decay. The default training schedule is 12 epochs. The input image size is  $1,024 \times 1,024$  with standard scale jittering (Ghiasi et al. 2021). This size is uniformly employed as the input for SAM, CLIP, and SideFormer. The max token length for each input sentence follows the CLIP default setting of 77. MMDetection (Chen et al. 2019) code-base is used.

### 5.2 Zero-Shot Transfer Performance

**COCO Benchmark** (Lin et al. 2014), comprising 80 common object categories, stands as the most widely utilized dataset for object detection. Considering that O365 covers all 80 categories and is frequently employed as pre-training data for COCO, we focus on evaluating the zero-shot transfer performance for models pre-trained with O365.

We provide a comparison between GLIP (Li et al. 2022a) and Grounding DINO (G-DINO) (Liu et al. 2023), along with their underlying detectors in Table 1. Our approach outperforms previous methods on the zero-shot transfer settings. When pre-trained on the same O365 dataset, Sambor shows +**2.4** AP and +**0.6** AP compared to GLIP and G-DINO, respectively. When employing the Open-set RPN (see more details in Sec. 5.3), Sambor demonstrates the best performance, surpassing even models that utilize larger

Method	Backbone	#Trainable Params	Pre-Train Data	MiniVal			Val v1.0	
				AP	AP <sub>r</sub> /AP <sub>c</sub> /AP <sub>f</sub>	AP	AP <sub>r</sub> /AP <sub>c</sub> /AP <sub>f</sub>	
MDETR <sup>†</sup> (Kamath et al. 2021)	RN101	186M	GoldG+,RefC	24.2	20.9/24.9/24.3	-	- / - / -	
Mask R-CNN <sup>†</sup> (Kamath et al. 2021)	RN101	69M	-	33.3	26.3/34.0/33.9	-	- / - / -	
GLIP-T (B) (Li et al. 2022a)	Swin-T	232M	O365	17.8	13.5/12.8/22.2	11.3	4.2/7.6/18.6	
GLIP-T (C) (Li et al. 2022a)	Swin-T	232M	O365,GoldG	24.9	17.7/19.5/31.0	16.5	7.5/11.6/26.1	
GLIP-T (Li et al. 2022a)	Swin-T	232M	O365,GoldG,Cap4M	26.0	20.8/21.4/31.0	17.2	10.1/12.5/25.5	
GLIPv2-T (Zhang et al. 2022a)	Swin-T	232M	O365,GoldG,Cap4M	29.0	- / - / -	-	- / - / -	
DetCLIP-T (A) (Yao et al. 2022)	Swin-T	-	O365	28.8	26.0/28.0/30.0	22.1	18.4/20.1/26.0	
DetCLIP-T (B) (Yao et al. 2022)	Swin-T	-	O365,GoldG	34.4	26.9/33.9/36.3	27.2	21.9/25.5/31.5	
DetCLIP-T (Yao et al. 2022)	Swin-T	-	O365,GoldG,YFCC1M	35.9	33.2/35.7/36.4	28.4	25.0/27.0/31.6	
DetCLIPv2-T (Yao et al. 2023)	Swin-T	-	O365	28.6	24.2/27.1/30.6	-	- / - / -	
DetCLIPv2-T (Yao et al. 2023)	Swin-T	-	O365,CC3M	31.3	29.4/31.7/31.3	-	- / - / -	
DetCLIPv2-T (Yao et al. 2023)	Swin-T	-	O365,GoldG,CC3M	38.4	<b>36.7/37.9/39.1</b>	-	- / - / -	
G-DINO-T (Liu et al. 2023)	Swin-T	172M	O365,GoldG	25.6	14.4/19.6/32.2	-	- / - / -	
G-DINO-T (Liu et al. 2023)	Swin-T	172M	O365,GoldG,Cap4M	27.4	18.1/23.3/32.7	-	- / - / -	
Sambor <sup>★</sup> (Ours)	ViT-B	160M	O365	33.1	29.6/32.0/34.7	26.3	20.9/24.4/30.9	
Sambor <sup>★</sup> (Ours)	ViT-B	160M	O365,GoldG	<b>39.6</b>	<b>34.6/39.3/40.7</b>	<b>32.8</b>	<b>30.9/31.0/35.7</b>	

Table 2: **Zero-shot object detection performance on LVIS benchmark.** AP<sub>r</sub>, AP<sub>c</sub>, and AP<sub>f</sub> indicate the AP values for rare, common, and frequent categories, respectively. <sup>★</sup>denotes the application of Open-set RPN. <sup>†</sup>denotes supervised approaches.

Method	Pre-Train Data	MiniVal			Val v1.0
		AP	AP <sub>r</sub> /AP <sub>c</sub> /AP <sub>f</sub>	AP	AP
X-Decoder (T)	COCO,C4M <sup>‡</sup>	-	- / - / -	-	9.6
OpenSeeD (T)	O365,COCO	-	- / - / -	-	19.4
OpenSeeD (L)	O365,COCO	-	- / - / -	-	21.0
Sambor <sup>★</sup> (Ours)	O365	27.6	27.3/27.5/27.7	21.7	21.7
Sambor <sup>★</sup> (Ours)	O365,GoldG	<b>35.7</b>	<b>31.4/37.1/35.3</b>	<b>29.4</b>	<b>29.4</b>

Table 3: **Zero-shot instance segmentation performance on LVIS benchmark.** AP<sub>r</sub>, AP<sub>c</sub>, and AP<sub>f</sub> indicate the AP values for rare, common, and frequent categories, respectively. <sup>‡</sup> includes Conceptual Captions (Sharma et al. 2018), SBU Captions (Vicente et al. 2016), Visual Genome (Krishna et al. 2017), and COCO Captions (Chen et al. 2015). <sup>★</sup>denotes the application of Open-set RPN.

datasets. We additionally report the zero-shot instance segmentation performance, which is effortlessly achievable by feeding the detection outputs into the SAM head. Notably, in comparison to other models, Sambor has a lower count of trainable parameters. This is attributed to the fact that, aside from SideFormer and the detection head, the parameters of the remaining components are frozen.

**LVIS Benchmark** (Gupta, Dollar, and Girshick 2019) contains 1,203 categories, including numerous rare categories that are seldom encountered in pre-training datasets. We report the *Fixed* AP (Dave et al. 2021) on both the MiniVal (Kamath et al. 2021) subset, comprising 5,000 images, and the complete validation set v1.0.

The zero-shot transfer performance on LVIS is presented in Table 2. Here, we also report the performance with and without the use of Open-set RPN, revealing an improvement of 0.4 AP when employed. Under comparable volumes of training data, Sambor outperforms competitors by a large margin. Specifically, in the scenario of training solely on O365 and evaluating on LVIS MiniVal, our model outperforms GLIP by **15.3** AP, and surpasses DetCLIP/DetCLIPv2 by **4.3/4.5** AP, respectively. The advantage is similarly pronounced when compared to G-DINO. Further, we incorporate the phrase grounding dataset GoldG to facilitate the generalization of Sambor. Instead of training from scratch, we fine-tune the O365 pre-trained model for 3 epochs while keeping other settings consistent. Evidently, Sambor exhibits enhanced zero-shot transfer performance across all categories. It surpasses even the results achieved with larger datasets compared to previous methods. Moreover, we present the zero-shot instance segmentation performance in Table 3, where the mask results are generated by prompting the detection boxes to the SAM head. Compared to X-Decoder (Zou et al. 2023a) and OpenSeeD (Zhang et al. 2023b), our Sambor exhibits superior performance, indicating its robust zero-shot generalization.

### 5.3 Ablation Studies

We conduct a series of ablation studies on Sambor, training with default settings on O365 unless specified otherwise.

**Effectiveness of SideFormer.** Table 4 validates the effectiveness of our designed SideFormer on COCO and LVIS. The Sambor baseline represents directly connecting the ViT-Det to the SAM image encoder. We first incorporate the

Strategy	COCO val			LVIS MiniVal			
	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP	AP <sub>r</sub>	AP <sub>c</sub>	AP <sub>f</sub>
Sambor baseline	39.0	54.7	42.5	27.7	21.5	26.5	29.9
+ SAM Extractor	42.4	58.7	46.2	29.0	25.2	27.5	31.0
+ CLIP Injector	<b>47.3</b>	<b>64.7</b>	<b>51.3</b>	<b>32.7</b>	<b>29.5</b>	<b>32.1</b>	<b>33.9</b>

Table 4: **Ablation studies on the components in Side-Former.** The combination of SAM Extractor and CLIP Injector shows the best performance.

Strategy	COCO val		LVIS MiniVal	
	AR@1000	AP	AR@1000	AP
Vanilla RPN	65.4	47.3	49.1	32.7
Open-set RPN <sup>#</sup>	<b>67.5 (+2.1)</b>	46.7 (-0.6)	<b>54.8 (+5.7)</b>	29.3 (-3.4)
Open-set RPN	<b>67.5 (+2.1)</b>	<b>48.6 (+1.3)</b>	<b>54.8 (+5.7)</b>	<b>33.1 (+0.4)</b>

Table 5: **Ablation studies on Open-set RPN.** Without further fine-tuning using additional region proposals from Open-set RPN (denoted as <sup>#</sup>), the improvement in proposal quality (AR@1000) brought by Open-set RPN cannot be directly translated to an increase in detection performance (AP). Fine-tuning serves as a remedy for this discrepancy, yielding superior results.

SAM Extractor to obtain multi-level features and perform fine-tuning. To further enhance the model’s recognition capability, we introduce the CLIP Injector to improve semantic representation, achieving optimal performance. Furthermore, we utilize region-word alignment for open-vocabulary classification, but there is a noticeable domain gap between CLIP text features and region features. Consequently, incorporating CLIP visual features effectively bridges this gap and provides benefits for Sambor.

**Effectiveness of Open-set RPN.** As elaborated in Sec. 4.2, the automatic mask generation pipeline of SAM allows for producing a number of high-quality open-set proposals, serving as a valuable complement to the vanilla RPN. We use a  $32 \times 32$  grid of points to generate open-set proposals, with a post-processing NMS threshold set to 0.7. Table 5 illustrates the average recall (AR@1000) for proposals, and it is evident that incorporating these open-set proposals significantly improves AR. However, the improved quality of region proposals does not manifest as superior detection performance; instead, there has been a decline. We posit that this is attributed to the detection head in the second stage not being exposed to the additional region proposals during training, hindering the ability to process them effectively.

To address this issue, we conduct a minor-scale fine-tuning adopting the Open-set RPN, *i.e.*, incorporating these open-set proposals during training. Specifically, with considerations for training efficiency, we use a  $32 \times 32$  grid of points to fine-tune for 1 epoch on approximately one-fifth of the O365 dataset. Maintaining all other hyper-parameters constant, employing a reduced learning rate of  $4 \times 10^{-5}$  contributes to the efficacy of fine-tuning. It effectively eradi-

Proposal	AR@1000	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>
<i>only</i> RPN	65.4	47.3	64.7	51.3	<b>33.5</b>	<b>53.2</b>	61.4
<i>only</i> SAM	59.4	46.8	63.3	51.1	29.8	52.2	<b>65.0</b>
Open-set RPN	<b>67.5</b>	<b>48.6</b>	<b>66.1</b>	<b>52.7</b>	<b>33.5</b>	<b>53.2</b>	64.2

Table 6: **Ablation studies on region proposal sources** for COCO 2017val. Using region proposals only from SAM exhibits a noticeable performance gap, particularly for small objects. Open-set RPN can effectively combine two sets of proposals, thereby achieving optimal performance.

cates inconsistencies in results, leading to superior performance for the Open-set RPN. Compared to the vanilla RPN, there is an improvement of 1.3 AP on COCO and 0.4 AP on LVIS. Unless specified otherwise, the Open-set RPN described elsewhere in this paper is fine-tuned.

**Region Proposal Sources.** After establishing the effectiveness of the Open-set RPN, a relevant question arises: *What would be the impact if we solely rely on object proposals from SAM?* We conduct ablation studies on the model fine-tuned with Open-set RPN. The evaluations include performance using proposals only from RPN, proposals only from the SAM head, and a combination of both sets, as shown in Table 6. Only RPN performance remains consistent compared to before fine-tuning (first row in Table 5). This confirms that the performance improvement in Open-set RPN is attributed to the supplementary proposals rather than the fine-tuning impact on RPN. To ensure an adequate quantity when relying solely on proposals from the SAM head, we increase the density of grid points to  $64 \times 64$  and set the NMS threshold to 0.95. There is a noticeable decrease in detection performance, especially for small objects. The difficulty of precisely targeting small objects with sampling points contributes to the inability to recall them, aligning with the results shown in Kirillov et al. (2023). Moreover, continuously increasing the density is impractical as it introduces unbearable computational and time consumption. Hence, the integration of the trainable RPN is crucial. The Open-set RPN effectively combines the two sets of region proposals in a complementary fashion, achieving optimal performance.

## 5.4 Analyses

**Open-Vocabulary Object Detector Built upon SAM.** The release of SAM has generated significant interest within the community, leading to the development of numerous derivative models. Instead of merely combining with SAM in a cascade manner, our Sambor seamlessly integrates SAM with an open-vocabulary detector into a unified end-to-end framework. This not only enables feature sharing but also facilitates the efficiency of interactive operations within the system. The advantages of Sambor are evident in the mutually beneficial relationship between SAM and the open-vocabulary object detector. (1) SAM provides the detector with powerful generalization features, yielding competitive zero-shot performance after semantic information supplementation. Moreover, the Open-set RPN leveraging proposals generated by SAM further enhances the object recall in

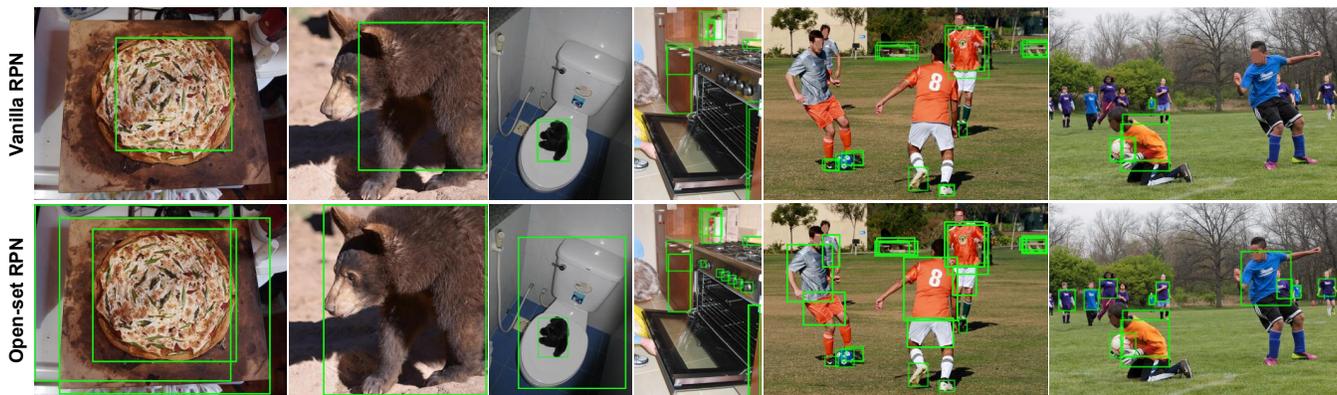


Figure 4: **Visualization comparison between Open-set RPN and the vanilla RPN.** For clarity, we only display high-quality proposals with an IoU greater than 0.7 with the ground truth boxes. In the first two examples, the vanilla RPN fails to generate proposals meeting this criterion; thus, we show the one with the highest IoU.



Figure 5: **Visualization of Sambor for open-vocabulary object detection and instance segmentation.** For better mask visual effects, we adopt HQ-SAM (Ke et al. 2023) as the mask decoder.

open-world scenarios. (2) The open-vocabulary object detector endows SAM to recognize arbitrary objects. Consequently, when utilizing functionalities from SAM, *e.g.*, interactive prompts, the detector can predict the categories while outputting segmentation results. This facilitates a more effortless and accurate acquisition of the desired targets.

**Visualization of Open-set RPN.** The comparison of object proposals between the Open-set RPN and the vanilla RPN is illustrated in Fig. 4. Thanks to incorporating more generalized proposals, the Open-set RPN effectively compensates for the recall of some objects, thereby addressing more challenging scenarios. For example, it can accurately capture the overall edges of an object and effectively leverage color information to delineate a person’s clothing, areas where the vanilla RPN performs inadequately.

**Visualization of Sambor.** We present the open-vocabulary object detection results of Sambor in Fig. 5, where the desired objects are detected by inputting category concepts. Additionally, we feed the detection boxes into the SAM head to obtain instance segmentation masks.

## 6 Conclusion

Sambor is an end-to-end open-vocabulary object detector that integrates the vision foundation model SAM. It effectively leverages the SAM features along with class-agnostic proposals to boost object detection. On the other hand,

the open-vocabulary object detector supplements SAM with the lacking classification capability, thus empowering it to extend beyond segmenting anything to recognizing arbitrary categories. Experiments demonstrate the superior open-vocabulary performance of Sambor and the contributions of the proposed modules. We aim for this effort to be an effective attempt, equipping SAM with recognition capabilities to address a wider array of application needs and offering a potential direction for open-vocabulary learning.

## Acknowledgements

This work was supported in part by the Key Deployment Program of the Chinese Academy of Sciences under Grant KGFZD145-23-18 and the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant E1XA310103.

## References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Cai, Z.; and Vasconcelos, N. 2018. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 6154–6162.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*, 213–229.

- Chen, K.; Wang, J.; Pang, J.; Cao, Y.; Xiong, Y.; Li, X.; Sun, S.; Feng, W.; Liu, Z.; Xu, J.; Zhang, Z.; Cheng, D.; Zhu, C.; Cheng, T.; Zhao, Q.; Li, B.; Lu, X.; Zhu, R.; Wu, Y.; Dai, J.; Wang, J.; Shi, J.; Ouyang, W.; Loy, C. C.; and Lin, D. 2019. MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv preprint arXiv:1906.07155*.
- Chen, X.; Fang, H.; Lin, T.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft COCO Captions: Data Collection and Evaluation Server. *arXiv preprint arXiv:1504.00325*.
- Chen, Z.; Duan, Y.; Wang, W.; He, J.; Lu, T.; Dai, J.; and Qiao, Y. 2022. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. YOLO-World: Real-Time Open-Vocabulary Object Detection. *arXiv preprint arXiv:2401.17270*.
- Dai, X.; Chen, Y.; Xiao, B.; Chen, D.; Liu, M.; Yuan, L.; and Zhang, L. 2021. Dynamic head: Unifying object detection heads with attentions. In *CVPR*, 7373–7382.
- Dave, A.; Dollár, P.; Ramanan, D.; Kirillov, A.; and Girshick, R. B. 2021. Evaluating Large-Vocabulary Object Detectors: The Devil is in the Details. *arXiv preprint arXiv:2102.01066*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, Y.; Wei, F.; Zhang, Z.; Shi, M.; Gao, Y.; and Li, G. 2022. Learning to prompt for open-vocabulary object detection with vision-language model. In *CVPR*, 14084–14093.
- Ghiasi, G.; Cui, Y.; Srinivas, A.; Qian, R.; Lin, T.; Cubuk, E. D.; Le, Q. V.; and Zoph, B. 2021. Simple Copy-Paste Is a Strong Data Augmentation Method for Instance Segmentation. In *CVPR*.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*.
- Gupta, A.; Dollar, P.; and Girshick, R. 2019. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 5356–5364.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. B. 2022. Masked Autoencoders Are Scalable Vision Learners. In *CVPR*, 15979–15988.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hudson, D. A.; and Manning, C. D. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 6700–6709.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 4904–4916.
- Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*, 1780–1790.
- Ke, L.; Ye, M.; Danelljan, M.; Liu, Y.; Tai, Y.-W.; Tang, C.-K.; and Yu, F. 2023. Segment Anything in High Quality. *arXiv preprint arXiv:2306.01567*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *arXiv preprint arXiv:2304.02643*.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123: 32–73.
- Kuo, W.; Cui, Y.; Gu, X.; Piergiovanni, A.; and Angelova, A. 2023. Open-Vocabulary Object Detection upon Frozen Vision and Language Models. In *ICLR*.
- Li, F.; Zhang, H.; Sun, P.; Zou, X.; Liu, S.; Yang, J.; Li, C.; Zhang, L.; and Gao, J. 2023. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.; Chang, K.; and Gao, J. 2022a. Grounded language-image pre-training. In *CVPR*, 10965–10975.
- Li, Y.; Mao, H.; Girshick, R.; and He, K. 2022b. Exploring plain vision transformer backbones for object detection. In *ECCV*, 280–296.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *ICCV*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 10012–10022.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Ma, C.; Jiang, Y.; Wen, X.; Yuan, Z.; and Qi, X. 2023. CoDet: Co-Occurrence Guided Region-Word Alignment for Open-Vocabulary Object Detection. *arXiv preprint arXiv:2310.16667*.
- Minderer, M.; Gritsenko, A.; Stone, A.; Neumann, M.; Weissenborn, D.; Dosovitskiy, A.; Mahendran, A.; Arnab, A.; Dehghani, M.; Shen, Z.; Wang, X.; Zhai, X.; Kipf, T.; and Houlsby, N. 2022. Simple Open-Vocabulary Object Detection with Vision Transformers. In *ECCV*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; Assran, M.; Ballas, N.; Galuba, W.; Howes, R.; Huang, P.; Li, S.; Misra, I.; Rabbat, M.; Sharma,

- V.; Synnaeve, G.; Xu, H.; Jégou, H.; Mairal, J.; Labatut, P.; Joulin, A.; and Bojanowski, P. 2024. DINOv2: Learning Robust Visual Features without Supervision. *TMLR*, 2024.
- Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, 2641–2649.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *CVPR*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; et al. 2024. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*.
- Shao, S.; Li, Z.; Zhang, T.; Peng, C.; Yu, G.; Zhang, X.; Li, J.; and Sun, J. 2019. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 8430–8439.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2556–2565.
- Sung, Y.-L.; Cho, J.; and Bansal, M. 2022. Lst: Ladder side-tuning for parameter and memory efficient transfer learning. In *NeurIPS*, 12991–13005.
- Thomee, B.; Shamma, D. A.; Friedland, G.; Elizalde, B.; Ni, K.; Poland, D.; Borth, D.; and Li, L.-J. 2016. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 59(2): 64–73.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*.
- Vicente, T. F. Y.; Hou, L.; Yu, C.-P.; Hoai, M.; and Samaras, D. 2016. Large-scale training of shadow detectors with noisily-annotated shadow examples. In *ECCV*, 816–832.
- Wang, T.; Zhang, J.; Fei, J.; Zheng, H.; Tang, Y.; Li, Z.; Gao, M.; and Zhao, S. 2023. Caption Anything: Interactive Image Description with Diverse Multimodal Controls. *arXiv preprint arXiv:2305.02677*.
- Wu, J.; Li, X.; Xu, S.; Yuan, H.; Ding, H.; Yang, Y.; Li, X.; Zhang, J.; Tong, Y.; Jiang, X.; et al. 2024. Towards open vocabulary learning: A survey. *IEEE TPAMI*.
- Wu, S.; Zhang, W.; Jin, S.; Liu, W.; and Loy, C. C. 2023a. Aligning Bag of Regions for Open-Vocabulary Object Detection. In *CVPR*.
- Wu, X.; Zhu, F.; Zhao, R.; and Li, H. 2023b. CORA: Adapting CLIP for Open-Vocabulary Detection with Region Prompting and Anchor Pre-Matching. In *CVPR*.
- Xu, Y.; Zhang, M.; Fu, C.; Chen, P.; Yang, X.; Li, K.; and Xu, C. 2023. Multi-modal Queried Object Detection in the Wild. *arXiv preprint arXiv:2305.18980*.
- Yao, L.; Han, J.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; and Xu, H. 2023. Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In *CVPR*.
- Yao, L.; Han, J.; Wen, Y.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; Xu, C.; and Xu, H. 2022. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. In *NeurIPS*, 9125–9138.
- Zang, Y.; Li, W.; Zhou, K.; Huang, C.; and Loy, C. C. 2022. Open-Vocabulary DETR with Conditional Matching. In *ECCV*.
- Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-vocabulary object detection using captions. In *CVPR*, 14393–14402.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L. M.; and Shum, H. 2023a. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. In *ICLR*.
- Zhang, H.; Li, F.; Zou, X.; Liu, S.; Li, C.; Yang, J.; and Zhang, L. 2023b. A simple framework for open-vocabulary segmentation and detection. In *ICCV*, 1020–1031.
- Zhang, H.; Zhang, P.; Hu, X.; Chen, Y.-C.; Li, L.; Dai, X.; Wang, L.; Yuan, L.; Hwang, J.-N.; and Gao, J. 2022a. Glipv2: Unifying localization and vision-language understanding. In *NeurIPS*, volume 35, 36067–36080.
- Zhang, S.; Chi, C.; Yao, Y.; Lei, Z.; and Li, S. Z. 2020. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 9759–9768.
- Zhang, Y.; Sun, Q.; Zhou, Y.; He, Z.; Yin, Z.; Wang, K.; Sheng, L.; Qiao, Y.; Shao, J.; and Liu, Z. 2022b. Bamboo: Building Mega-Scale Vision Dataset Continually with Human-Machine Synergy. *arXiv preprint arXiv:2203.07845*.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *CVPR*.
- Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022. Detecting twenty-thousand classes using image-level supervision. In *ECCV*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *ICLR*.
- Zou, X.; Dou, Z.-Y.; Yang, J.; Gan, Z.; Li, L.; Li, C.; Dai, X.; Behl, H.; Wang, J.; Yuan, L.; Peng, N.; Wang, L.; Lee, Y. J.; and Gao, J. 2023a. Generalized Decoding for Pixel, Image, and Language. In *CVPR*, 15116–15127.
- Zou, X.; Yang, J.; Zhang, H.; Li, F.; Li, L.; Gao, J.; and Lee, Y. J. 2023b. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*.