

Simple Image-Level Classification Improves Open-Vocabulary Object Detection

Ruohuan Fang¹, Guansong Pang^{2*}, Xiao Bai^{1,3*}

¹School of Computer Science and Engineering, Beihang University

²School of Computing and Information Systems, Singapore Management University

³State Key Laboratory of Software Development Environment, Jiangxi Research Institute, Beihang University

ruohuanfang@gmail.com, baixiao@buaa.edu.cn, gspang@smu.edu.sg

Abstract

Open-Vocabulary Object Detection (OVOD) aims to detect novel objects beyond a given set of base categories on which the detection model is trained. Recent OVOD methods focus on adapting the image-level pre-trained vision-language models (VLMs), such as CLIP, to a region-level object detection task via, *e.g.*, region-level knowledge distillation, regional prompt learning, or region-text pre-training, to expand the detection vocabulary. These methods have demonstrated remarkable performance in recognizing regional visual concepts, but they are weak in exploiting the VLMs' powerful global scene understanding ability learned from the billion-scale image-level text descriptions. This limits their capability in detecting hard objects of small, blurred, or occluded appearance from novel/base categories, whose detection heavily relies on contextual information. To address this, we propose a novel approach, namely Simple Image-level Classification for Context-Aware Detection Scoring (SIC-CADS), to leverage the superior global knowledge yielded from CLIP for complementing the current OVOD models from a global perspective. The core of SIC-CADS is a multi-modal multi-label recognition (MLR) module that learns the object co-occurrence-based contextual information from CLIP to recognize all possible object categories in the scene. These image-level MLR scores can then be utilized to refine the instance-level detection scores of the current OVOD models in detecting those hard objects. This is verified by extensive empirical results on two popular benchmarks, OV-LVIS and OV-COCO, which show that SIC-CADS achieves significant and consistent improvement when combined with different types of OVOD models. Further, SIC-CADS also improves the cross-dataset generalization ability on Objects365 and OpenImages. Code is available at <https://github.com/mala-lab/SIC-CADS>.

Introduction

Open-Vocabulary Object Detection (OVOD) is a challenging task that requires detecting objects of novel categories that are not present in the training data. To tackle this problem, conventional approaches focus on leveraging external image-text data as weak supervisory information to expand the detection vocabulary of the categories. In recent years, large-scale pre-trained vision-language models (VLMs), *e.g.*,

CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021), which are trained using billion-scale internet-crawled image-caption pairs (see Fig. 1(a) for an example), have been widely used to empower OVOD.

Existing VLM-based OVOD studies focus on how to adapt the image-level pre-trained CLIP to a region-level object detection task. Typically they adopt a regional concept learning method, such as Region-level Knowledge Distillation (Du et al. 2022; Bangalath et al. 2022) that aligns region embeddings to their corresponding features extracted from the image encoder of CLIP, Regional Prompt Learning (Wu et al. 2023c; Du et al. 2022; Feng et al. 2022) that learns continuous prompt representations to better pair with region-level visual embeddings, Region-Text Pre-training (Zhong et al. 2022) that explicitly aligns image regions and text tokens during vision-language pre-training, or Self-Training (Zhou et al. 2022) that generates pseudo-labels of novel objects on the image-level labeled datasets (*e.g.*, ImageNet-21k (Deng et al. 2009) and Conceptual Captions (Sharma et al. 2018)).

These methods have demonstrated remarkable performance in recognizing regional visual concepts, but they are weak in exploiting the VLMs' powerful global scene understanding ability that can capture important relations between different visual concepts (Wu et al. 2023b; Zhong et al. 2021). These relations can be the co-occurrence of different objects such as the tennis ball and the tennis racket, or their interdependence to commonly shared background environment such as the tennis ball and tennis court, as shown in Fig. 1(b). This weakness can limit their capability in detecting **hard objects** that are of small, blurred, or occluded appearance, whose detection relies heavily on contextual features of other objects in the same image. Further, the OV detectors have naturally learned the context information for base categories when aligning regional features with corresponding text embeddings, since the network can automatically capture the context features within the receptive field that are related to base objects. However, the context features related to novel objects are not learned due to the absence of novel object annotations during the training of OV detectors. This can be one of the key reasons that incurs the performance gap between base and novel objects, especially the novel hard objects.

To address this issue, this work instead aims to utilize the global scene understanding ability for OVOD. The key

*Corresponding authors: G. Pang and X. Bai

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

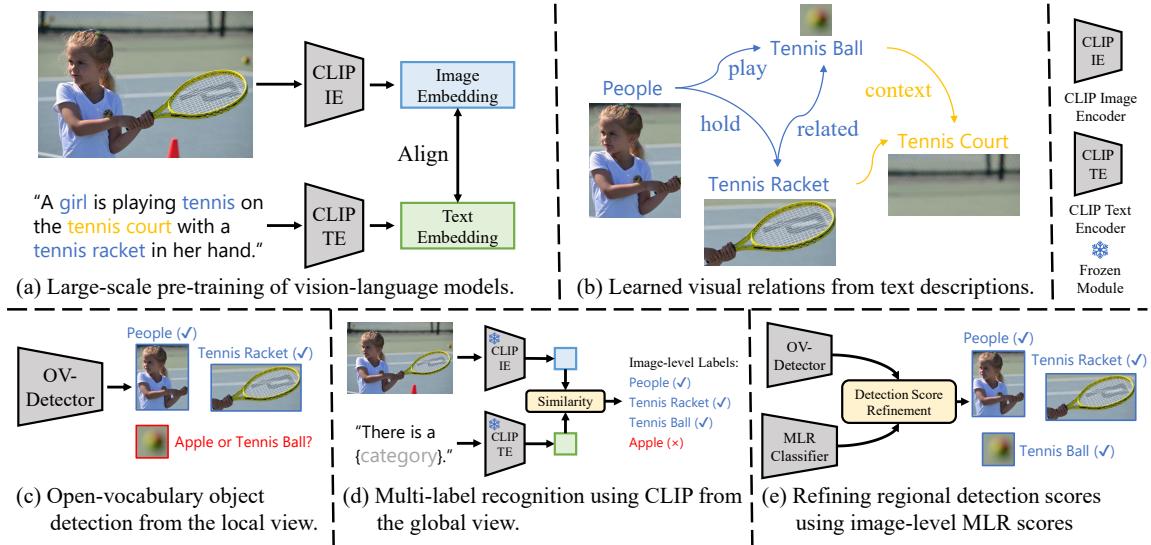


Figure 1: Motivation of our method. (a) Large VLMs like CLIP (Radford et al. 2021) can understand the image globally by learning rich knowledge from a huge amount of image-text pairs. Such knowledge can include diverse relations between visual concepts in the scene, as shown in (b). As illustrated in (c), current OVOD approaches focus on regional visual concept detection but they are weak in exploiting this global knowledge, which can fail to detect novel hard objects, such as ambiguous objects like the blurred tennis ball. (d) We instead learn an image-level multi-label recognition (MLR) module to leverage the global knowledge yielded from CLIP for recognizing those hard objects. (e) The image-level MLR scores are then utilized to refine the instance-level detection scores from a global perspective for more effective OVOD.

motivation is that the image-level embeddings extracted from CLIP’s image encoder carry global features about various objects in the entire scene, which are semantically related in the natural language descriptions. This knowledge can then provide important contextual information for detecting the aforementioned hard objects, *e.g.*, the small, blurred tennis ball in Fig. 1(c), which are otherwise difficult to detect using only regional features.

Inspired by this, we propose a novel approach that utilizes a Simple Image-level Classification module for Context-Aware Detection Scoring, termed SIC-CADS. Our image-level classification task is specified by a Multi-Label Recognition (MLR) module which learns multi-modal knowledge extracted from CLIP. The MLR module predicts image-level scores of different possible object categories that could exist in a specific scene. For example, as shown in Fig. 1(d), context information of the tennis racket and tennis court helps recognize the blurred tennis ball in such a sports-related scene. Thus, the image-level MLR scores can be used to refine the instance-level detection scores of existing OVOD models for improving their detection performance from a global perspective, as shown in Fig. 1(e).

Our main contributions are summarized as follows. (i) We propose a novel approach SIC-CADS that utilizes a MLR module to leverage VLMs’ global scene knowledge for improving OVOD performance. (ii) SIC-CADS is a simple, lightweight, and generic framework that can be easily plugged into different existing OVOD models to enhance their ability to detect hard objects. (iii) Extensive experiments on OV-LVIS, OV-COCO, and cross-dataset generalization

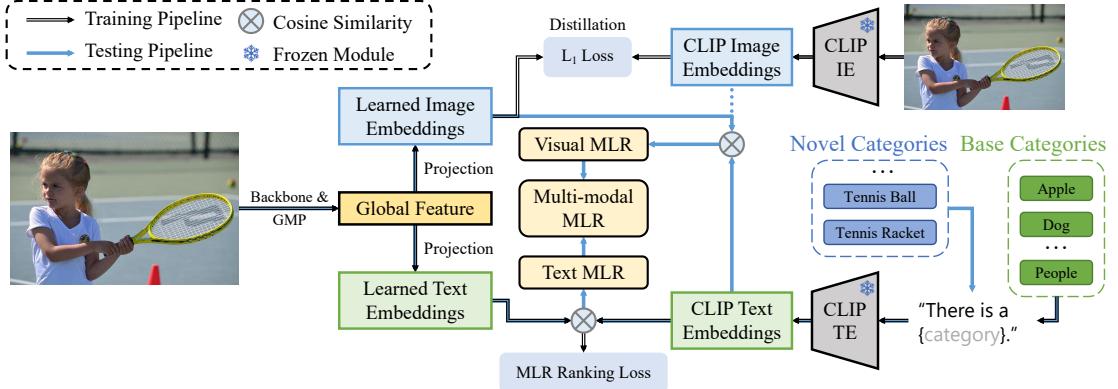
benchmarks show that SIC-CADS significantly boosts the detection performance when combined with different types of state-of-the-art (SOTA) OVOD models, achieving 1.4 - 3.9 gains of AP_r for OV-LVIS and 1.7 - 3.2 gains of AP_{novel} for OV-COCO. Besides, our method also largely improves their cross-dataset generalization ability, yielding 1.9 - 2.1 gains of mAP_{50} on Objects365 (Shao et al. 2019) and 1.5 - 3.9 gains of mAP_{50} on OpenImages (Kuznetsova et al. 2020).

Related Work

Vision-Language Models (VLMs): The task of vision-language pre-training is to learn aligned multimodal representations from image-text datasets. Early works (Joulin et al. 2016; Li et al. 2017) demonstrate that CNNs trained to predict words in image captions can learn representations competitive with ImageNet training. Recently, the large-scale VLMs (*e.g.*, CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021)) which are contrastively pre-trained on billion-scale internet-sourced image-caption datasets, have shown powerful zero-shot performance on image classification tasks. With the advent of large VLMs, they have been applied to various downstream vision tasks, including object detection (Gu et al. 2021; Du et al. 2022; Zhou et al. 2022), instance segmentation (Xu et al. 2022; Ghiasi et al. 2021), image generation (Nichol et al. 2021; Ramesh et al. 2021), and anomaly detection (Zhou et al. 2023; Wu et al. 2023a).

Open-Vocabulary Object Detection (OVOD): Traditional object detection models can only detect the base categories that are presented during training. OVOD aims to extend the vocabulary of object detectors using additional large

(a) Training and Inference Process of Our Multi-modal MLR Model



(b) Detection Score Refinement Using Image-level MLR Scores

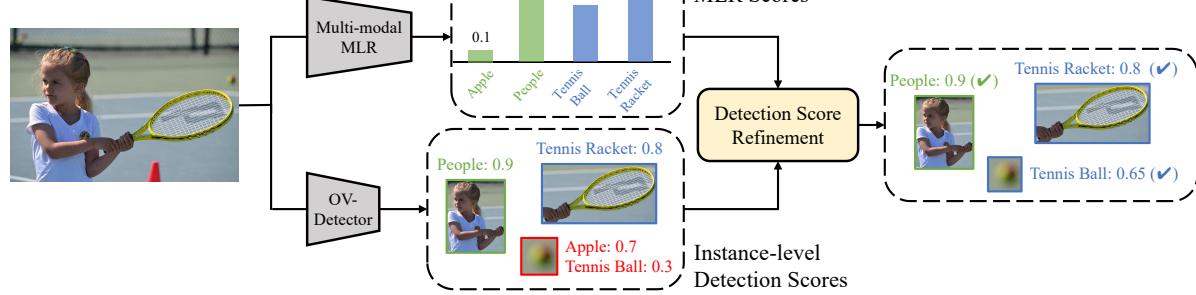


Figure 2: Overview of our approach SIC-CADS. (a) During training, our proposed MLR module learns CLIP’s global multi-modal knowledge from the text encoder (text MLR) and image encoder (visual MLR). During inference, the two branches demonstrate superior performance in recognizing objects from base and novel classes respectively. Hence, we combine the two branches to make our full model, multi-modal MLR. (b) Our MLR module can be plugged into existing OVOD models, via a simple detection score refinement process, to boost the performance in detecting hard objects from a global perspective.

image-text datasets. OVR-CNN (Zareian et al. 2021) first formulates this problem and proposes its baseline method by aligning the regional features with words in the caption. ViLD (Gu et al. 2021) addresses the problem of OVOD by distilling the regional representations using CLIP’s image encoder. Detic (Zhou et al. 2022) adopts self-training which produces pseudo-labels of novel objects on the datasets such as ImageNet-21k (Deng et al. 2009) to expand the detection vocabulary. DetPro (Du et al. 2022), PromptDet (Feng et al. 2022), and POMP (Ren et al. 2023) use prompt tuning to adapt the image-level pre-trained VLMs to the region-level object detection task. RegionCLIP (Zhong et al. 2022) learns region-level visual representations by region-text pre-training. F-VLM (Kuo et al. 2022) directly utilizes the regional features of frozen VLMs for object recognition and eliminates the need for regional knowledge distillation. Our simple MLR network adopts a similar score fusion strategy as ViLD and F-VLM. The key difference is that we recognize all objects via multi-modal MLR, while ViLD and F-VLM use an ensemble score for region-level object classification. Plus, our method is also orthogonal to ViLD and F-VLM, as shown in Tab. 1 that our SIC-CADS can be plugged into ViLD to largely improve its performance.

Zero-Shot Multi-Label Recognition (ZS-MLR): The task

of multi-label recognition is to predict the labels of all objects in one image. As an extension of Zero-Shot Learning (ZSL) (Romera-Paredes and Torr 2015; Xian, Schiele, and Akata 2017; Xian et al. 2019; Xu et al. 2020), ZS-MLR aims to recognize both seen and unseen objects in the image. The key of ZS-MLR is to align the image embeddings with the category embeddings. Previous ZS-MLR methods (Zhang, Gong, and Shah 2016; Ben-Cohen et al. 2021; Huynh and Elhamifar 2020; Narayan et al. 2021; Gupta et al. 2021) mostly use the single-modal embeddings from language models (e.g., GloVe (Pennington, Socher, and Manning 2014)) and adopt different strategies (e.g., attention mechanism (Narayan et al. 2021; Huynh and Elhamifar 2020) or generative models (Gupta et al. 2021)) to boost the performance. He *et al.* (He et al. 2022) further propose Open-Vocabulary Multi-Label Recognition (OV-MLR) task which adopts the multi-modal embeddings from CLIP for MLR. They use the transformer backbone to extract regional and global image representations and a two-stream module to transfer knowledge from CLIP. Recognize Anything Model (RAM) (Zhang et al. 2023) leverages the large-scale image-text pairs by parsing the captions for automatic image labeling. Then, they train an image-tag recognition decoder which reveals strong ZS-MLR performance.

Method

This work aims to exploit CLIP’s image-level/global knowledge to better recognize hard objects for more effective OVOD. We first train a multi-label recognition (MLR) module by transferring global multi-modal knowledge from CLIP to recognize all existing categories in the entire scene. Then during inference, our MLR module can be easily plugged into existing OVOD models, via a simple detection score refinement process, to boost the detection performance.

Preliminaries

In OVOD, typically we have an object detector trained with a detection dataset \mathcal{D}_{det} which contains the exhaustively annotated bounding-box labels for a set of base categories \mathcal{C}_B . Some external image-caption datasets may be available that can be used to expand the detection vocabulary, enabling the detector to recognize novel categories beyond the closed set of base categories. During inference, the categories in the testing set comprise both base categories \mathcal{C}_B and novel categories \mathcal{C}_N , *i.e.*, $\mathcal{C}_{\text{test}} = \mathcal{C}_B \cup \mathcal{C}_N$ and $\emptyset = \mathcal{C}_B \cap \mathcal{C}_N$. Therefore, the OVOD models are required to solve two subsequent problems: (1) the effective localization of all objects in the image, and (2) the correct classification of each object into one of the class labels in $\mathcal{C}_{\text{test}}$. Many VLM-based OVOD methods (Zhou et al. 2022; Gu et al. 2021; Bangalath et al. 2022) adopt the two-stage detection framework (*e.g.*, Mask-RCNN (He et al. 2017)) as the base detector. In the first stage, a RPN (*region proposal network*) takes an image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ as input, and produces a set of class-agnostic object proposals $\mathcal{P} \subset \mathbb{R}^4$ which denotes the coordinates for the proposal boxes. In the second stage, a RoI (*region of interest*) Align head computes the pooled representations $\mathcal{E} = \{e_r\}_{r \in \mathcal{P}} \subset \mathbb{R}^d$ for the proposals \mathcal{P} . To classify each proposal into one of the categories in $\mathcal{C}_B \cup \mathcal{C}_N$, for each category c , we obtain its text embedding t_c by feeding the category name in a prompt template, *e.g.*, ‘a photo of a {category name}’ into the text encoder \mathcal{T} of a pre-trained VLM like CLIP. The probability of a region proposal r being classified into category c is computed as:

$$p(r, c) = \frac{\exp(\cos(e_r, t_c)/\tau)}{\sum_{c' \in \mathcal{C}_{\text{test}}} \exp(\cos(e_r, t_{c'})/\tau)}, \quad (1)$$

where $\cos(\cdot, \cdot)$ denotes a cosine similarity and τ is a temperature scaling factor.

CLIP-driven Multi-modal MLR Modeling

Overall Framework: Our proposed MLR module is designed to leverage the superior scene understanding ability of CLIP for effective detection of both novel and base categories. Specifically, as shown in Fig. 2(a), given an input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, we first extract the multi-scale feature maps $\mathcal{P} = \{P_2, P_3, P_4, P_5\}$ using a ResNet50-FPN backbone, followed by a Global Max Pooling (GMP) operator and a concatenation of all feature vectors in different FPN levels to obtain a global image embedding e^{global} . The global image embedding is then utilized in two branches, including one **text MLR** branch which aligns e^{global} with the text

embeddings of different categories yielded from the text encoder of CLIP, and one **visual MLR** branch which distills the global image embedding from the image encoder of CLIP into e^{global} . During inference, the text MLR is weaker in recognizing novel categories since it is trained with only base categories, while the visual MLR can recognize both base and novel categories by distilling the zero-shot recognition ability of CLIP. Therefore, we combine the scores of both branches to achieve better recognition performance for novel and base categories, which is noted as **Multi-modal MLR**. Below we introduce them in detail.

Text MLR: This branch aims to align the image embeddings with the corresponding category text embeddings yielded from the text encoder of CLIP, so that each image can be classified based on CLIP text embeddings. We first use a linear layer $f(\cdot)$ to project the image embedding e^{global} to a new feature space, namely learned text embedding space, and obtain $e^{\text{text}} = f(e^{\text{global}}) \in \mathbb{R}^d$. We then use the commonly used prompt template ‘there is a {category}’ and feed the template filled with category names into the CLIP text encoder to obtain category-specific text embeddings t . The classification score of each category c for an image i is computed using cosine similarity:

$$s_{i,c}^{\text{text}} = \cos(e_i^{\text{text}}, t_c). \quad (2)$$

To train this branch, we adopt a powerful MLR ranking loss to increase the scores of positive categories (*i.e.*, categories that appear in the image) and decrease the scores of negative categories (*i.e.*, categories that this image does not contain). Specifically, we define the rank loss as:

$$L_{\text{rank}} = \sum_i \sum_{p \in \mathcal{N}(i), n \notin \mathcal{N}(i)} \max(1 + s_{i,n}^{\text{text}} - s_{i,p}^{\text{text}}, 0), \quad (3)$$

where $\mathcal{N}(i)$ denotes the image-level labels of image i , and $s_{i,n}$ and $s_{i,p}$ denote the classification scores of positive and negative categories w.r.t. the image, respectively. During training, text MLR learns to align the text embeddings of multi-label base categories for each image by minimizing L_{rank} . For a test image j , the classification score w.r.t. a category c is defined as $s_{j,c}^{\text{text}}$.

Visual MLR: In text MLR, since Eq. 3 includes only the base categories in the training set, the resulting learned text embeddings have weak zero-shot recognition of novel categories. To address this issue, we propose the visual MLR branch, which distills knowledge from the CLIP image encoder to achieve better zero-shot recognition performance. This is because the image encoder of CLIP has been trained on a huge amount of images using corresponding text supervision. This enables it to extract more generalized embeddings of images that include features of both base and novel categories, as well as important background environments. To distill these generalized embeddings, similar to text MLR, we use another linear layer $f(\cdot)$ to transform the global image embedding into the learned image embedding $e^{\text{image}} = f(e^{\text{global}}) \in \mathbb{R}^d$. Then, we minimize the distance between e^{image} and the embedding from the CLIP image encoder using a \mathcal{L}_1 loss:

$$L_{\text{dist}} = \sum_i \|IE(i) - e_i^{\text{image}}\|_1, \quad (4)$$

| Method | Supervision | AP _r | AP _c | AP _f | AP |
|--|--|---------------------|---------------------|---------------------|---------------------|
| BoxSup (Zhou et al. 2022) + SIC-CADS | LVIS-Base + CLIP | 16.4 20.3 (+3.9) | 31.0 31.8 (+0.8) | 35.4 35.5 (+0.1) | 30.2 31.2 (+1.0) |
| ViLD (Gu et al. 2021) + SIC-CADS | LVIS-Base + CLIP | 16.8 18.7 (+2.1) | 25.6 26.4 (+0.8) | 28.5 28.7 (+0.2) | 25.2 26.1 (+0.9) |
| RegionCLIP (Zhong et al. 2022) + SIC-CADS | LVIS-Base + CLIP + CC3M (Sharma et al. 2018) | 19.7 21.9 (+2.2) | 28.2 29.1 (+0.9) | 30.7 30.9 (+0.2) | 27.7 28.5 (+0.8) |
| OC-OVD (Bangalath et al. 2022) + SIC-CADS | LVIS-Base + CLIP + LMDet (Maaz et al. 2021) + IN21k | 21.1 22.5 (+1.4) | 25.0 25.6 (+0.6) | 29.1 29.0 (-0.1) | 25.9 26.3 (+0.4) |
| Detic (Zhou et al. 2022) + SIC-CADS | LVIS-Base + CLIP + IN21k (Deng et al. 2009) | 24.9 26.5 (+1.6) | 32.5 33.0 (+0.5) | 35.6 35.6 (+0.0) | 32.4 32.9 (+0.5) |
| POMP (Ren et al. 2023) + SIC-CADS | LVIS-Base + CLIP + IN21k (Deng et al. 2009) | 25.2 26.6 (+1.4) | 33.0 33.3 (+0.3) | 35.6 35.6 (+0.0) | 32.7 33.1 (+0.4) |

Table 1: Enabling different SOTA OVOD models on the OV-LVIS benchmark.

where $IE(i)$ denotes the embedding of image i from the CLIP image encoder. Since $IE(i)$ contains rich knowledge of novel categories and important contextual features, the learned e^{image} can well complement e^{text} in text MLR in detecting novel object categories, especially objects whose detection requires the recognition of the global image contexts. During inference, the classification score of a test image j to a category c can be computed as follows:

$$s_{j,c}^{image} = \cos(e_j^{image}, t_c). \quad (5)$$

Since there can still exist a performance gap between visual MLR (student) and CLIP IE (teacher), directly using the CLIP image embedding $IE(j)$ to replace e_j^{image} for obtaining $s_{j,c}^{image}$ in Eq. 5 often enables better performance. This variant is denoted as **visual MLR⁺**. Notably, we adopt the ViT-B/32 CLIP image encoder which takes one 224×224 resolution image as input during testing, so it incurs only minor computation and memory overhead.

Multi-modal MLR: As discussed above, text MLR and visual MLR branches complement each other in identifying base and novel categories, so we combine their classification scores to have a multi-modal MLR. Particularly, the projection function $f(\cdot)$ in each branch is trained independently. After that, we obtain the learned text embedding e_i^{text} and image embedding e_i^{image} for each image i . Inspired by (Gu et al. 2021; Kuo et al. 2022), we ensemble the two probability scores of an image j as follows:

$$p_{j,c}^{mmlr} = \begin{cases} (p_{j,c}^{text})^{\lambda_B} \cdot (p_{j,c}^{image})^{1-\lambda_B}, & c \in \mathcal{C}_B \\ (p_{j,c}^{text})^{1-\lambda_N} \cdot (p_{j,c}^{image})^{\lambda_N}, & c \in \mathcal{C}_N, \end{cases} \quad (6)$$

where $p_{j,c}^{text} = \text{sigmoid}(\tilde{s}_{j,c}^{text})$ is a sigmoid classification probability based on a normalized text embedding similarity score $\tilde{s}_{j,c}^{text} = \frac{s_{j,c}^{text} - \mu(s_{j,c}^{text})}{\sigma(s_{j,c}^{text})}$, and similarly, $p_{j,c}^{image} = \text{sigmoid}(\tilde{s}_{j,c}^{image})$ is a classification probability using the image embedding similarity score $s_{j,c}^{image}$ normalized in a similar way as $\tilde{s}_{j,c}^{text}$ ($s_{j,c}^{image}$ can be obtained from visual MLR or visual MLR⁺), and λ_B and λ_N is the hyperparameters for

controlling the combination of two scores. The zero-mean normalization is used based on the fact that the sigmoid function has a max gradient near zero, so that probability values of positive and negative categories can be better separated, which is beneficial for the detection score refinement process.

Context-aware OVOD with Image-Level Multi-modal MLR Scores

Our multi-modal MLR learns to recognize both novel and base objects from the global perspective. We show that it can be plugged into different existing OVOD models that are often focused on regional visual concepts via a post-process step to enhance their detection performance. As shown in Fig. 2(b), given a test image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, the OVOD model produces a set of instance predictions $\{(b, p_{j,c}^{ovod})_j\}$ where b_j is the bounding box coordinates and $p_j^{ovod} = \{p_{j,1}^{ovod}, p_{j,2}^{ovod}, \dots, p_{j,C}^{ovod}\}$ represents the classification probability scores of the j -th instance for all $C = |\mathcal{C}_{test}|$ categories. Then, our MLR model predicts the image-level classification scores $p_j^{mmlr} = \{p_1^{mmlr}, p_2^{mmlr}, \dots, p_C^{mmlr}\}$. Although the MLR model can not localize the objects, it provides scene context information and prior knowledge about the types of objects that may exist in the whole image. This contextual information enhances the region-level detection performance, especially in detecting the aforementioned hard objects. Therefore, we utilize the following weighted geometric mean to combine the image-level score p_j^{mmlr} and instance-level score p_j^{ovod} as follows:

$$p_j^{cads} = (p_j^{mmlr})^\gamma \odot (p_j^{ovod})^{1-\gamma}, \quad (7)$$

where p_j^{cads} denotes the context-aware detection score of the j -th instance, \odot denotes the element-wise multiplication, and γ is a hyperparameter to balance the two types of scores.

Experiments

Datasets

We evaluate our method on LVIS v1.0 (Gupta, Dollar, and Girshick 2019) and COCO (Lin et al. 2014) under the open-vocabulary settings, as defined by recent works (Zareian et al.

| Method | Supervision | AP _{novel} | AP _{base} | AP |
|---|--|---------------------|---------------------|---------------------|
| Detic (Zhou et al. 2022) + SIC-CADS | COCO-Base + CLIP + COCO Captions (Chen et al. 2015) | 27.8 31.0 (+3.2) | 51.1 52.4 (+1.3) | 45.0 46.8 (+1.7) |
| BARON (Wu et al. 2023b) + SIC-CADS | COCO-Base + CLIP + COCO Captions (Chen et al. 2015) | 35.1 36.9 (+1.8) | 55.2 56.1 (+0.9) | 49.9 51.1 (+1.2) |
| RegionCLIP (Zhong et al. 2022) + SIC-CADS | COCO-Base + CLIP + CC3M (Sharma et al. 2018) | 39.3 41.4 (+2.1) | 61.6 61.7 (+0.1) | 55.7 56.3 (+0.6) |
| OC-OVD (Bangalath et al. 2022) + SIC-CADS | CLIP + LMDet (Maaz et al. 2021) + COCO Captions (Chen et al. 2015) | 40.7 42.8 (+2.1) | 54.1 55.1 (+1.0) | 50.6 51.9 (+1.3) |
| CORA (Wu et al. 2023c) + SIC-CADS | COCO-Base + CLIP + COCO Captions (Chen et al. 2015) | 41.6 43.3 (+1.7) | 44.7 45.7 (+1.0) | 43.9 45.1 (+1.2) |

Table 2: Enabling different SOTA OVOD models on the OV-COCO benchmark.

| Method | Objects365 | OpenImages |
|-------------------|---------------------|---------------------|
| BoxSup + SIC-CADS | 26.6 28.7 (+2.1) | 46.4 50.3 (+3.9) |
| Detic + SIC-CADS | 29.3 31.2 (+1.9) | 53.2 54.7 (+1.5) |

Table 3: Results on the cross-dataset evaluation.

2021; Gu et al. 2021), with the benchmarks named as OV-COCO and OV-LVIS respectively.

OV-LVIS: LVIS is a large-vocabulary instance segmentation dataset containing 1,203 categories. The categories are divided into three groups based on their appearance frequency in the dataset: frequent, common, and rare. Following the protocol introduced by (Gu et al. 2021), we treat the frequent and common categories as base categories (noted as LVIS-Base) to train our model. It considers 337 rare categories as novel categories during testing. We report the instance segmentation mask-based average precision metrics for rare (novel), common, frequent, and all classes, denoted as AP_r, AP_c, AP_f, and AP respectively.

OV-COCO: We follow the open-vocabulary setting defined by (Zareian et al. 2021) and split the categories into 48 base categories and 17 novel categories. Only base categories in COCO train2017 (noted as COCO-Base) are used for training. We report the box mAP of novel, base, and all classes measured at IoU threshold 0.5, denoted as AP_{novel}, AP_{base}, and AP respectively.

Cross-dataset Generalization: To validate the generalization performance, we train the proposed model on the training set of LVIS and evaluate on two other datasets, Objects365 (Shao et al. 2019) and OpenImages (Kuznetsova et al. 2020). We report the AP₅₀ result evaluated at IoU threshold 0.5 for both datasets.

Implementation Details

Network Architecture: We use ResNet-50 (He et al. 2016) with FPN (Lin et al. 2017) as the default backbone network. To ensure a fair comparison with previous OVOD methods, we adopt the same VLM model, the ViT-B/32 CLIP model,

as the teacher model for our MLR module by default.

Hyperparameters: The hyperparameters λ_B , λ_N , and γ are set to 0.8, 0.8, 0.5 for OV-LVIS and 0.8, 0.5, 0.7 for OV-COCO based on the ablation results.

Training Pipeline: We adopt the offline training strategy that trains our MLR module separately from the OVOD models. We use AdamW (Loshchilov and Hutter 2017) optimizer with an initial learning rate of 0.0002 to train our MLR module. For OV-LVIS, our MLR model is trained using the image-level labels of LVIS-Base for 90,000 iterations with a batch size of 64 (48 epochs). As for OV-COCO, we train our model for 12 epochs for a fair comparison with previous OVOD models. Note that some OVOD models may leverage extra image-level labeled datasets, *e.g.*, ImageNet-21k (Deng et al. 2009), COCO Caption (Chen et al. 2015). Therefore, we also train our MLR model using the same dataset for the same iterations when plugging our method into these methods. We use images of size 480×480, augmented with random resized cropping and horizontal flipping during training.

Inference: During inference, we use the trained multi-modal MLR module to predict the image-level multi-label scores. We then combine the MLR scores with the instance predictions of the trained OVOD models to obtain the final results based on Eq. 7. Since visual MLR⁺ generally performs better than its primary version, it is used by default to compute $p_{j,c}^{image}$ in Eq. 6 during inference in our experiments, denoted as **multi-modal MLR⁺**. We show the full results of both versions in Tab. 4 and our appendix.

Main Results

To evaluate the overall performance and flexibility of our method **SIC-CADS**, we combine it with various OVOD models that use different strategies, including knowledge distillation (ViLD (Gu et al. 2021), OC-OVD (Bangalath et al. 2022)); prompt learning (POMP (Ren et al. 2023), CORA (Wu et al. 2023c)); self-training (Detic (Zhou et al. 2022)); and region-text pre-training (RegionCLIP (Zhong et al. 2022)). The results of these original OVOD models and our SIC-CADS-enabled versions on OV-LVIS, OV-COCO, and cross-dataset generalization are presented below.

OV-LVIS: The results of SIC-CADS combined with different OVOD models on OV-LVIS are shown in Tab. 1. Our

| Method | AP _r | AP _c | AP _f | R _{novel} ^{mlr} | R _{base} ^{mlr} | GFLOPs | FPS |
|--|-----------------|-----------------|-----------------|-----------------------------------|----------------------------------|--------|-----|
| Base Model (BoxSup) | 16.4 | 31.0 | 35.4 | - | - | 215 | 4.8 |
| w/ Text MLR | 17.6 (+1.2) | 31.5 (+0.5) | 35.5 (+0.1) | 9.7 | 65.0 | 233 | 4.6 |
| w/ Visual MLR | 20.1 (+3.7) | 31.1 (+0.1) | 34.7 (-0.7) | 19.6 | 32.6 | 233 | 4.6 |
| w/ Multi-modal MLR | 19.9 (+3.5) | 31.7 (+0.7) | 35.4 (+0.0) | 21.1 | 56.8 | 233 | 4.6 |
| w/ Visual MLR ⁺ | 20.6 (+4.2) | 31.3 (+0.3) | 34.6 (-0.8) | 38.1 | 31.9 | 238 | 4.5 |
| w/ Multi-modal MLR ⁺ | 20.3 (+3.9) | 31.8 (+0.8) | 35.5 (+0.1) | 34.4 | 56.5 | 238 | 4.5 |
| w/ Multi-modal MLR ⁺ (ViT-L/14 CLIP) | 21.9 (+5.5) | 32.0 (+1.0) | 35.5 (+0.1) | 37.8 | 58.5 | 348 | 3.9 |

Table 4: Results of different variants of our MLR module on OV-LVIS.

| γ | AP _r | AP _c | AP _f | AP |
|----------|-----------------|-----------------|-----------------|-------------|
| 0.3 | 19.3 | 31.6 | 35.5 | 31.0 |
| 0.4 | 20.0 | 31.8 | 35.5 | 31.2 |
| 0.5 | 20.3 | 31.8 | 35.5 | 31.3 |
| 0.6 | 20.6 | 31.6 | 35.3 | 31.1 |
| 0.7 | 20.7 | 31.1 | 34.8 | 30.7 |
| 0.8 | 21.1 | 29.8 | 33.9 | 30.8 |

Table 5: Effectiveness of hyperparameter γ .

| λ_B | λ_N | AP _r | AP _c | AP _f | AP |
|-------------|-------------|-----------------|-----------------|-----------------|------|
| 0.8 | 1.0 | 20.3 | 31.7 | 35.4 | 31.1 |
| 0.8 | 0.8 | 20.3 | 31.8 | 35.5 | 31.3 |
| 0.8 | 0.6 | 19.9 | 31.8 | 35.5 | 31.2 |
| 0.8 | 0.5 | 19.1 | 31.8 | 35.5 | 30.8 |
| 1.0 | 0.8 | 20.2 | 31.5 | 35.5 | 31.1 |
| 0.6 | 0.8 | 20.3 | 31.9 | 35.3 | 31.2 |
| 0.5 | 0.8 | 20.3 | 31.6 | 35.0 | 30.8 |

Table 6: Effectiveness of hyperparameter λ_B and λ_N .

method SIC-CADS can consistently and significantly improve all five OVOD models in detecting novel categories, achieving maximal 3.9 gains and setting a new SOTA performance of 26.6 in AP_r when plugged into POMP. Further, SIC-CADS also consistently improves the detection performance of common base categories in AP_c, having 0.5-0.9 gains across the models. It retains similar performance in detecting frequent categories in AP_f (increase/decrease in the $[-0.1, 0.2]$ range). These results demonstrate that our proposed multi-modal MLR module effectively learns important global knowledge from CLIP that is complementary to the current regional OVOD models in detecting objects of both novel and base categories. Impressively, this improvement holds for the recent best-performing OVOD models, regardless of whether they exploit external supervision.

OV-COCO: SIC-CADS shows similar enabling superiority on the OV-COCO benchmark, as reported in Tab. 2. It is remarkable that SIC-CADS consistently enhances all five SOTA OVOD models in all three metrics, AP_{novel}, AP_{base}, and AP. Particularly, it increases the AP_{novel} scores by 1.7-3.2 points and AP_{base} by 0.1-1.3 points over the five base models. Notably, RegionCLIP and CORA adopt stronger

ResNet-50 \times 4 CLIP image encoder as the backbone, and our method still improves their overall performance without using such a strong and large backbone.

Cross-dataset Generalization: In the cross-dataset generalization experiment, we train our MLR module using image-level labels of LVIS and combine it with BoxSup and Detic to evaluate the performance on Objects365 and OpenImages. As presented in Tab. 3, SIC-CADS largely increases AP₅₀ by 1.9-2.1 points on Objects365 and 1.5-3.9 points on OpenImages. These results demonstrate our excellent improvement in the generalization ability from the cross-dataset aspect.

Further Analysis of SIC-CADS

Ablation Study: Different MLR variants can be used in our SIC-CADS approach. We evaluate the use of different MLR variants on top of the baseline method BoxSup to demonstrate the importance of multi-modal MLR. In addition to the above AP evaluation metrics, We also adopt the recall rate for novel and base categories of the top-10 MLR predictions as an auxiliary metric to evaluate the performance of the MLR module (denoted as R_{novel}^{mlr} and R_{base}^{mlr} respectively). The ablation study is shown in Tab. 4. Overall, Text MLR can slightly boost both AP_r and AP_c, but it still biases toward the base categories, resulting in very limited gains for novel categories. Visual MLR and Visual MLR⁺ can significantly increase AP_r by 3.7 - 4.2 points, but AP_f drops by 0.7 - 0.8 points, showing strong zero-shot recognition but relatively weak base category recognition. Multi-modal MLR and Multi-modal MLR⁺ combine the strengths of the two branches and obtain the best recall rate in the base or novel categories, and as a result, they achieve equally excellent results for novel and base category detection, yielding 3.5 - 3.9 gains for AP_r and 0.7 - 0.8 gains for AP_c, without harming AP_f. Additionally, we also show that using the stronger VLM models, *i.e.*, ViT-L/14 CLIP, can obtain even better performances, achieving up to 5.5 gains for AP_r. We also provide the results of using large MLR models (*e.g.*, BLIP(Li et al. 2022) and RAM(Zhang et al. 2023)) in our appendix.

Computational Costs: In Tab.4, we also show the additional computational overload yielded from our plugged-in module. Generally, our model incurs small additional inference costs when plugged into current OVOD models, resulting in a slight increase in GFLOPs from 215 to 233-238, and a decrease in FPS from 4.8 to 4.6-4.5. We provide more detailed information in our appendix.

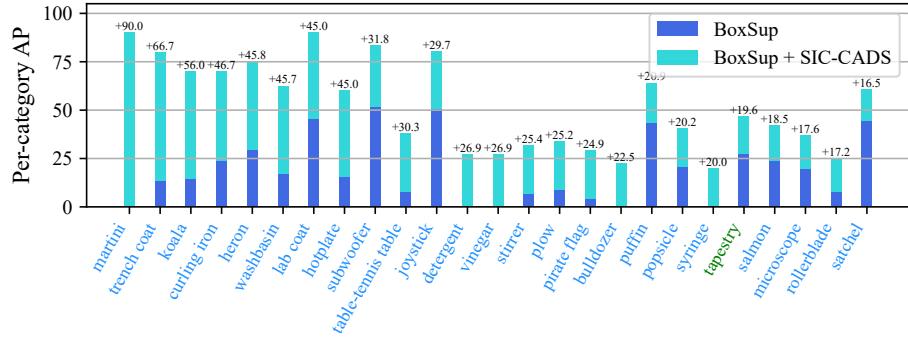


Figure 3: Class-specific AP of the top 25 categories with the largest improvement on OV-LVIS.

Analysis of Hyperparameters: Tab. 5 shows the results of varying values of γ in Eq. 7 when plugged into BoxSup on OV-LVIS. γ controls the combination of the MLR scores and the detection scores. When using a smaller γ , the MLR scores have only limited impact on the instance scores, leading to smaller gains for AP_r . However, if a large γ is used, e.g., 0.8, all detection scores will be very close to the MLR score, making all metrics drop sharply. Since $\gamma = 0.5$ achieves a good trade-off between novel and base category detection and obtains the best overall AP, we choose it as the default hyperparameter for OV-LVIS. Besides, Tab. 6 shows the results of varying values of λ_B and λ_N in Eq. 6 on LVIS. We find that our method is not sensitive when varying λ_B and λ_N within the range of [0.8, 1.0]. However, AP_f drops by 0.5 when $\lambda_B = 0.5$, and AP_r drops by 1.2 when $\lambda_N = 0.5$. So we choose $\lambda_B = 0.8$ and $\lambda_N = 0.8$ for OV-LVIS. The ablation results of γ , λ_B , and λ_N for OV-COCO are shown in our appendix.

Qualitative Analysis: To further investigate the impact of our method on different object categories, we present a qualitative analysis of the class-specific AP of the top 25 categories with the largest improvement when applying SIC-CADS to BoxSup. The results are illustrated in Fig. 3, together with detection examples for some of those categories in Fig. 4, where novel and base categories are in blue and green, respectively. Our proposed method exhibits significant improvement (maximally 90.0 gains in the class-specific AP) in detecting small-sized, blurred, or occluded objects, such as those from the martini, curling iron, washbasin, and table-tennis table categories. These categories are often ambiguous for the regional OVOD models like BoxSup, but our method can effectively detect them due to the contextual knowledge offered by our multi-modal MLR. As also shown in Fig. 4(a-b), SIC-CADS can help recognize fine-grained categories, which are otherwise recognized as coarse-grained categories, e.g., ‘table’ in (a) vs ‘table-tennis table’ in (b). Due to the learned contextual knowledge, SIC-CADS also helps correct wrong detection, e.g., ‘Soup Bowl’ in (c) is corrected as ‘Washbasin’ in (d).

Conclusion and Future Work

This paper proposes SIC-CADS, a novel approach for open-vocabulary object detection that leverages global scene under-

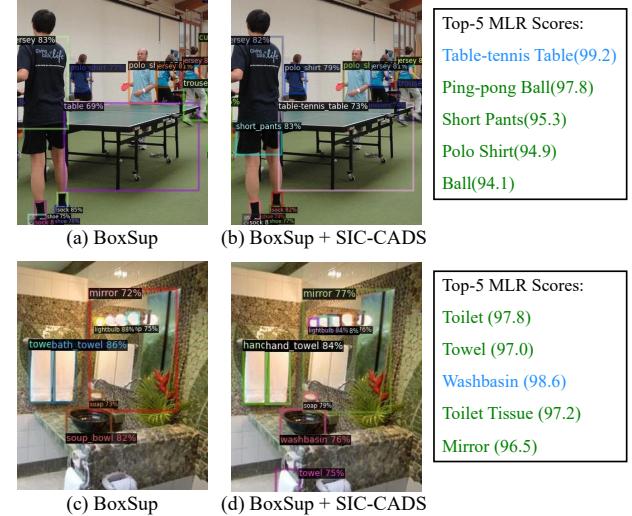


Figure 4: Visualizations results of combining SIC-CADS with BoxSup on OV-LVIS.

standing capabilities of VLMs for generalized base and novel category detection. The core in SIC-CADS is a multi-modal MLR module that enables the recognition of different types of objects based on their contextual co-occurrence relations. The resulting MLR scores help largely refine the instance-level detection scores yielded by different types of current SOTA OVOD models that are focused on regional visual concept recognition, enabling significantly improved OVOD performance. This is supported by extensive empirical results on the OV-LVIS, OV-COCO, and the cross-dataset benchmarks: Objects365 and OpenImages. Our qualitative analysis shows that one main source of the significant improvement gained by SIC-CADS is from its superior performance in detecting ambiguous novel categories, on which current OVOD models fail to work well. Despite the promising results, there are still limitations in our method. Particularly, our method may fail when context information does not match the object categories. We discuss such failure cases in our appendix and will improve the method in future work.

Acknowledgments

In this work, R. Fang and X. Bai are supported by the National Natural Science Foundation of China 62276016, 62372029. Due to space limitation, our appendix is made available in our pre-print version at <https://arxiv.org/abs/2312.10439>.

References

Bangalath, H.; Maaz, M.; Khattak, M. U.; Khan, S. H.; and Shahbaz Khan, F. 2022. Bridging the gap between object and image-level representations for open-vocabulary detection. *Advances in Neural Information Processing Systems*, 35: 33781–33794.

Ben-Cohen, A.; Zamir, N.; Ben-Baruch, E.; Friedman, I.; and Zelnik-Manor, L. 2021. Semantic diversity learning for zero-shot multi-label classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 640–650.

Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14084–14093.

Feng, C.; Zhong, Y.; Jie, Z.; Chu, X.; Ren, H.; Wei, X.; Xie, W.; and Ma, L. 2022. Promptdet: Towards open-vocabulary detection using uncurated images. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 701–717. Springer.

Ghiasi, G.; Gu, X.; Cui, Y.; and Lin, T.-Y. 2021. Open-vocabulary image segmentation. *arXiv preprint arXiv:2112.12143*.

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. Lvvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5356–5364.

Gupta, A.; Narayan, S.; Khan, S.; Khan, F. S.; Shao, L.; and van de Weijer, J. 2021. Generative multi-label zero-shot learning. *arXiv preprint arXiv:2101.11606*.

He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

He, S.; Guo, T.; Dai, T.; Qiao, R.; Ren, B.; and Xia, S.-T. 2022. Open-Vocabulary Multi-Label Classification via Multi-modal Knowledge Transfer. *arXiv preprint arXiv:2207.01887*.

Huynh, D.; and Elhamifar, E. 2020. A shared multi-attention framework for multi-label zero-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8776–8786.

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 *International Conference on Machine Learning*, 4904–4916. PMLR.

Joulin, A.; Van Der Maaten, L.; Jabri, A.; and Vasilache, N. 2016. Learning visual features from large weakly supervised data. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII* 14, 67–84. Springer.

Kuo, W.; Cui, Y.; Gu, X.; Piergiovanni, A.; and Angelova, A. 2022. F-vlm: Open-vocabulary object detection upon frozen vision and language models. *arXiv preprint arXiv:2209.15639*.

Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Malloci, M.; Kolesnikov, A.; et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7): 1956–1981.

Li, A.; Jabri, A.; Joulin, A.; and Van Der Maaten, L. 2017. Learning visual n-grams from web data. In *Proceedings of the IEEE International Conference on Computer Vision*, 4183–4192.

Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, 12888–12900. PMLR.

Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.

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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13, 740–755. Springer.

Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.

Maaz, M.; Rasheed, H. B.; Khan, S. H.; Khan, F. S.; Anwer, R. M.; and Yang, M.-H. 2021. Multi-modal transformers excel at class-agnostic object detection. *arXiv*.

Narayan, S.; Gupta, A.; Khan, S.; Khan, F. S.; Shao, L.; and Shah, M. 2021. Discriminative region-based multi-label zero-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8731–8740.

Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.

Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, 8821–8831. PMLR.

Ren, S.; Zhang, A.; Zhu, Y.; Zhang, S.; Zheng, S.; Li, M.; Smola, A.; and Sun, X. 2023. Prompt Pre-Training with Twenty-Thousand Classes for Open-Vocabulary Visual Recognition. *arXiv preprint arXiv:2304.04704*.

Romera-Paredes, B.; and Torr, P. 2015. An embarrassingly simple approach to zero-shot learning. In *International conference on machine learning*, 2152–2161. PMLR.

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 *Proceedings of the IEEE/CVF international conference on computer vision*, 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 *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2556–2565.

Wu, P.; Zhou, X.; Pang, G.; Zhou, L.; Yan, Q.; Wang, P.; and Zhang, Y. 2023a. Vadclip: Adapting vision-language models for weakly supervised video anomaly detection. *arXiv preprint arXiv:2308.11681*.

Wu, S.; Zhang, W.; Jin, S.; Liu, W.; and Loy, C. C. 2023b. Aligning Bag of Regions for Open-Vocabulary Object Detection. *arXiv preprint arXiv:2302.13996*.

Wu, X.; Zhu, F.; Zhao, R.; and Li, H. 2023c. CORA: Adapting CLIP for Open-Vocabulary Detection with Region Prompting and Anchor Pre-Matching. *arXiv preprint arXiv:2303.13076*.

Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-shot learning—the good, the bad and the ugly. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4582–4591.

Xian, Y.; Sharma, S.; Schiele, B.; and Akata, Z. 2019. f-vaegan-d2: A feature generating framework for any-shot learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10275–10284.

Xu, J.; De Mello, S.; Liu, S.; Byeon, W.; Breuel, T.; Kautz, J.; and Wang, X. 2022. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18134–18144.

Xu, W.; Xian, Y.; Wang, J.; Schiele, B.; and Akata, Z. 2020. Attribute prototype network for zero-shot learning. *Advances in Neural Information Processing Systems*, 33: 21969–21980.

Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14393–14402.

Zhang, Y.; Gong, B.; and Shah, M. 2016. Fast zero-shot image tagging. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5985–5994. IEEE.

Zhang, Y.; Huang, X.; Ma, J.; Li, Z.; Luo, Z.; Xie, Y.; Qin, Y.; Luo, T.; Li, Y.; Liu, S.; et al. 2023. Recognize Anything: A Strong Image Tagging Model. *arXiv preprint arXiv:2306.03514*.

Zhong, Y.; Shi, J.; Yang, J.; Xu, C.; and Li, Y. 2021. Learning to generate scene graph from natural language supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1823–1834.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16793–16803.

Zhou, Q.; Pang, G.; Tian, Y.; He, S.; and Chen, J. 2023. Anomalyclip: Object-agnostic prompt learning for zero-shot anomaly detection. *arXiv preprint arXiv:2310.18961*.

Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022. Detecting twenty-thousand classes using image-level supervision. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 350–368. Springer.