

# Efficient Feature Distillation for Zero-shot Annotation Object Detection

Zhuoming Liu\*, Xuefeng Hu\*, Ram Nevatia†  
 University of Southern California  
 {liuzhuom, xuefengh, nevatia}@usc.edu

## Abstract

We propose a new setting for detecting unseen objects called Zero-shot Annotation object Detection (ZAD). It expands the zero-shot object detection setting by allowing the novel objects to exist in the training images and restricts the additional information the detector uses to novel category names. Recently, to detect unseen objects, large-scale vision-language models (e.g., CLIP) are leveraged by different methods. The distillation-based methods have good overall performance but suffer from a long training schedule caused by two factors. First, existing work creates distillation regions biased to the base categories, which limits the distillation of novel category information. Second, directly using the raw feature from CLIP for distillation neglects the domain gap between the training data of CLIP and the detection datasets, which makes it difficult to learn the mapping from the image region to the vision-language feature space. To solve these problems, we propose Efficient feature distillation for Zero-shot Annotation object Detection (EZAD). Firstly, EZAD adapts the CLIP’s feature space to the target detection domain by re-normalizing CLIP; Secondly, EZAD uses CLIP to generate distillation proposals with potential novel category names to avoid the distillation being overly biased toward the base categories. Finally, EZAD takes advantage of semantic meaning for regression to further improve the model performance. As a result, EZAD outperforms the previous distillation-based methods in COCO by 4% with a much shorter training schedule and achieves a 3% improvement on the LVIS dataset. Our code is available at <https://github.com/dragonlzm/EZAD>

## 1. Introduction

Object detection is a fundamental task in computer vision. Typically, an object detector is trained on a dataset with specific categories and can’t be extended to new categories. To detect instances of categories not seen dur-

\*Equal contribution.

†Corresponding author



Figure 1. The difference between our zero-shot annotation object detection (ZAD), zero-shot object detection (ZSD), and open-vocabulary detection (OVD). Our ZAD can have novel instances in training images and novel category names, but no additional annotations are allowed.

ing training, researchers started studying Zero-Shot object Detection (ZSD) and Open-Vocabulary object Detection (OVD). In both ZSD and OVD settings, a set of “base” categories is provided with annotated training examples (bounding boxes), and the rest, called “novel” categories, are not provided with any annotated examples. In ZSD, the names of the novel categories are known at training time, but it is further posited [1] that the training data should not have any instances of the novel categories; this is very restrictive and requires training data to be screened for novel categories which defeats the purpose of the user adding categories at will. OVD does not impose this restriction but assumes that the novel category names are unavailable during training and may be added freely at inference time. Also, to detect more novel categories, the training of OVD sometimes needs additional images or captions.

We propose an alternate setting that we call Zero-shot Annotation object Detection (ZAD) task, which allows the training data to contain unannotated novel category instances and a model to mine this information at training time if novel category names are available in advance. However, additional category names may be added at inference time, as in OVD. Thus, ZAD falls in between ZSD and OVD, as Fig 1 shows. We believe ZAD is closer to a real-life setting: when no additional data or annotation is provided, a user wanting to add categories not considered at annotation time but exploiting the latent information as novel category names are added at the training time.

To recognize novel instances when only their name is available, it is natural to consider a vision-language feature space. In recent years, some large-scale vision-language

classification models trained on millions of image-text pairs, for instance, CLIP [27], ALIGN [14] have become available. Different solutions are proposed to leverage these models to enable a detector to detect novel objects. Some researchers [7, 35, 37, 38] use these models to generate pseudo-label and train their model with pseudo-label and the image-caption pair, while other trains a prompt [6, 16, 32] to turn the CLIP into a detector. However, all these methods need additional captions or images, which are unavailable in our ZAD setting. In contrast, pure distillation-based methods [9, 34] learn a mapping from image regions to CLIP feature space by distillation for detecting the novel object, which is applicable when no additional data is provided.

To know how good the features are for distillation, we first apply CLIP to classify the instances in the COCO [22] dataset. We found that the classification accuracy (ACC) is only 46% which is much lower than the ACC of the classifier in Faster R-CNN [29] (about 90%). This indicates the domain gap between the training data of CLIP and the detection dataset, making the mapping from the image region to the vision-language feature space harder to learn. In addition, since the distillation is conducted on some specific image regions, how to select such a region is an important question. ViLD [9] trains an RPN with base category annotations and applies the RPN to the training image again to generate proposals as distillation regions. As table 1 shows, these proposals are biased toward the region with base categories instances, which limits the novel information obtained by the detector and harms the distillation efficiency.

To address these problems, we propose Efficient Feature Distillation for Zero-shot Annotation object Detection (EZAD). For bridging the domain gap, we find that simply finetuning the LayerNorm layers in the CLIP with the base category instances significantly improves the ACC on both base and novel (Fig 2, 1). For the distillation regions, we expect these regions could contain novel objects so that some novel category information can be introduced into the detector. To make the best use of the only information we have in the ZAD setting, the name of novel categories, we decided to use CLIP to select these regions with the help of the novel category names. The selected regions are named CLIP Proposals, in which CLIP believes there is a novel category instance(Fig 2, 2).

After adapting the feature space and generating the CLIP Proposals, EZAD learns a mapping from the image regions to the vision-language feature space by distillation, which is achieved by minimizing the L1 loss between the features of the CLIP Proposals from the CLIP, and the one from our model (Fig 2, 3). Once the model is trained, in all potential regions given by the RPN, EZAD recognizes the novel objects by using the novel category name’s text embedding, which has a high cosine similarity score with the image fea-

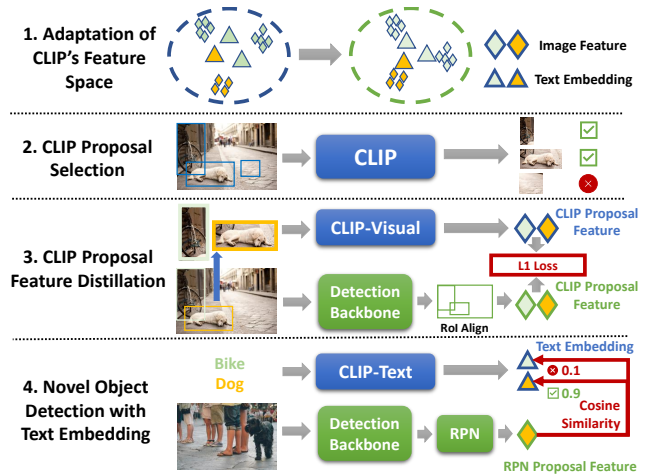


Figure 2. Overview of our method. The key contributions of EZAD are we bridge the domain gap and create better distillation regions (CLIP proposals).

Dataset	Eval On	AR@100	AR@300	AR@1000
COCO	Base	<b>56.69</b>	<b>61.45</b>	<b>64.32</b>
	Novel	34.66	43.62	51.34
LVIS	Base	<b>42.14</b>	<b>49.84</b>	<b>55.09</b>
	Novel	35.63	45.07	51.82

Table 1. The Recall of RPN on Novel and Base categories on COCO or LVIS training set. The RPN biases to the base category.

ture of the novel objects (Fig 2, 4). To further improve the model performance, we introduce a semantic-based regressor, which takes the text embedding as additional information for regression.

By only providing the name of the novel categories to EZAD, EZAD can outperform ViLD by 4% in novel categories with a much shorter training schedule in COCO. On the LVIS dataset, our method achieves better performance than ViLD on both base and novel categories. This indicates that the adapted feature and CLIP Proposals benefit both distillation quality and training efficiency.

## 2. Related Work

**The Large-scale Pretraining** already exist in Vision (MOCO [11], MOCOv2 [3], etc.) and Language (BERT [5], GPT [2], etc.) for a long time. Recently, there is a trend to use free-form supervision (raw text) to train the vision model which has evolved to large-scale vision-language pre-trained models. For instance, the CLIP [27] and ALIGN [14] are trained with a large-scale dataset with hundreds of millions of image-text pairs using contrastive learning. GLIP [19] is pre-trained at an object-level with a large-scale grounding dataset. These multi-modal models’ feature space and their knowledge are useful and can be applied to many other tasks, such as zero-shot classification and zero-shot detection. However, the training data of these models are usually noisy and there is a domain gap between

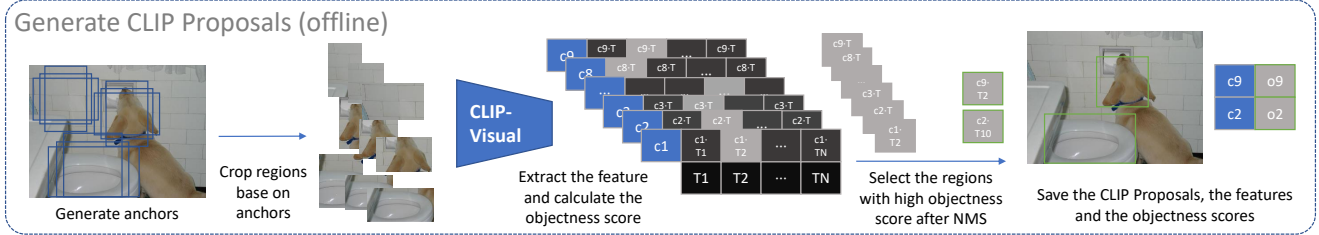


Figure 3. The pipeline of CLIP Proposals generation. The CLIP Proposals and their feature are generated before the detector training.

these data and the datasets of downstream tasks.

**Domain Adaptation** is necessary when we apply a pre-trained model to other datasets. In computer vision, the most common method for bridging the domain gap is to finetune the whole network on the new dataset or add one extra MLP layer at the end. In the natural language processing community, Prompt tuning [17, 20] surfaced as an important tool for domain adaptation in recent years. Besides, the simple renormalization method is found to be effective by Perez [26] and Lu [24]. For adapting the large-scale multi-modal model, Kim et al. [15] discuss the effectiveness of different ways in adapting the CLIP [27] to new classification datasets. As far as we know, no paper discusses bridging the domain gap between image-level and instance-level tasks.

**Zero-shot Detection object (ZSD) and Open Vocabulary object Detection (OVD)** aim to learn knowledge from the base categories and to generalize the knowledge to the novel categories, enabling the model to detect novel objects. Bansal et al. [1], Zhu et al. [39], Rahman et al. [28], and Xie et al. [33] focus on ZSD, in which novel instances cannot exist in the training image. However, in most object detection datasets, not all the instances on the image will be annotated. This makes the ZSD setting too restricted. Zhong et al. [37], Zhou et al. [38] and Zareian et al [35] pre-train their model with image-caption pair to learn the vision-language feature space for open vocabulary detection. Gao et al. [7], Long et al. [23], Wang et al. [30] Zhao et al. [36] train its model with pseudo-label. Feng et al. [6], Kuo et al. [16], and Wu et al. [32] design a prompt to make use of CLIP for detection. Ma et al. [25] and Gu et al. [9] distill visual features from CLIP, enabling object detectors to detect novel instances. Wang et al. [31] propose CondHead which can be added to different OVD models to further improve their performance. All these methods aim to solve the OVD problem, and most of them need additional training data except the distillation-based [9, 25, 31] and some prompt-based methods [16, 32]. The distillation-based and prompt-based methods can fit in with our ZAD setting, while the distillation-based method has better overall performance.

### 3. Method

**Setting definition.** Our method aims to handle the **Zero-shot Annotation object Detection (ZAD)**. In ZAD, all cate-

gories are divided into two parts: the base categories  $C_b$  and the novel categories  $C_n$ . The name of the base and expected novel category are known before model training. The training image may have novel category instances, but only the base category will be annotated. No additional data will be available for the training except the training image and the base categories annotations. The model is trained on a set of base classes and tested on both base and novel classes.

**Method Overview.** Inspired by the previous paper ViLD [9], our proposed method, **Efficient Feature Distillation for Zero-shot Annotation object Detection (EZAD)**, map the image feature to the CLIP’s multi-modal feature space. The *mapping* is learned by distilling the knowledge from CLIP in some selected *distillation regions* of each training image. The knowledge distillation is conducted by optimizing the L1 loss between the *feature from CLIP* and the feature from our model in the selected regions. After learning the mapping, given the proposal from the RPN, our model can recognize and detect novel objects by using the text embedding of the novel categories.

In Section 3.1, we describe how to adapt the *feature from CLIP* to the domain of detection datasets. In Section 3.2, we demonstrate how to select the *distillation regions* on each image. In Section 3.3, we discuss our model structure and how to train our model and learn the *mapping* with the adapted feature from the selected regions.

#### 3.1. Adapt Vision-language Model Feature

To understand how good the image feature from CLIP is, we first apply the CLIP to classify the instances in the COCO [22] dataset. We first extract the feature for each instance from the vision encoder of the CLIP. The feature of the instance  $i$  can be expressed as:  $ins_i = V(Crop(I, GT_{i(1.2x)}))$ , where  $V$  is the vision encoder of the CLIP and  $Crop(I, GT_{i(1.2x)})$  means cropping the region from Image  $I$  base on 1.2x enlarged GT bboxes  $GT_{i(1.2x)}$ . We use the 1.2x enlarged bboxes since the enlarged bboxes help CLIP yield the best classification accuracy. We present more details in the supplementary material. We generate the text embedding for each COCO category from the text encoder of the CLIP. We calculate the cosine similarity between the image feature and text embeddings and select the category with the highest cosine score as the predicted category.

We notice that when directly applying the CLIP to clas-

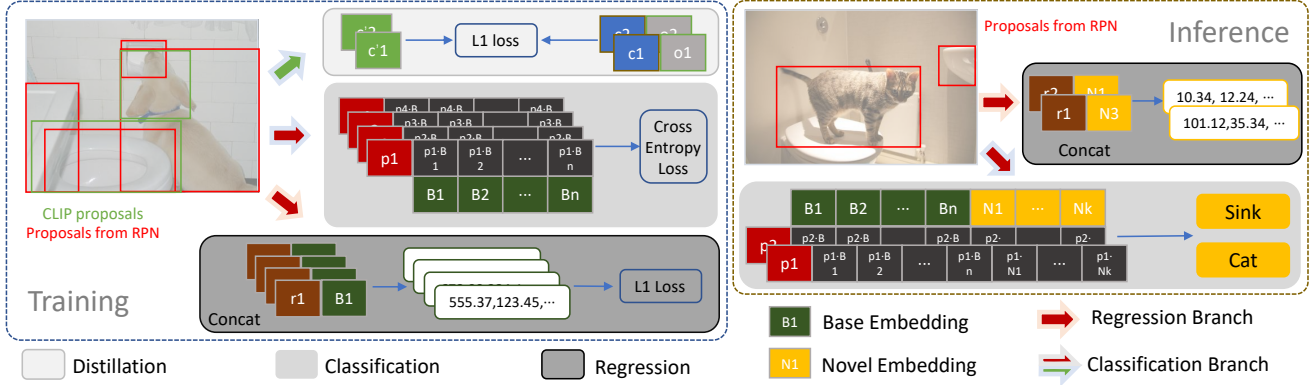


Figure 4. The training and inference of our model. The figure presents the model structures in the RoI head of the two-stage detector. We use the per CLIP Proposal distillation weight and semantic-based regressor to further improve the model performance.

sify the COCO instances, the classification accuracy (ACC) is only about 50% which is much lower than the ACC of the classifier in a well-trained detector, indicating that there is a huge distribution gap between the training data of the CLIP and detection datasets. To bridge the gap, inspired by [15], we simply fine-tune the CLIP’s layer of normalization layers by minimizing the cross-entropy loss, using all base categories instances in the detection dataset. This simple method boosts the ACC on COCO to about 80%. Also, using the adapted CLIP feature for distillation helps improve detection results.

### 3.2. Generate CLIP Proposals

To obtain useful information for the novel categories, we need to select some meaningful image regions for distillation. we expect these regions to contain some novel category objects, and introduce information of the novel categories. ViLD trains an RPN with base category annotations and uses the proposals from this RPN as the distillation regions. However, since the RPN is trained to predict base categories, which makes the proposals bias toward the areas that contain the base categories and ignore the areas that potentially have the novel category instance. Instead of using the RPN to determine where to distill, we decided to use CLIP [27]. Trained with 400 million image and text pairs collected from the internet, CLIP is trained to discriminate a large number of categories. Therefore, if the region contains a novel object, the CLIP should yield high confidence in it. We name these regions CLIP Proposals. Fig 3 demonstrates how to generate CLIP Proposal.

To select image regions as the CLIP Proposals, we first generate anchors over the image  $I$ . Then we crop the image based on the anchors and extract the feature for these anchors from CLIP’s vision encoder. We generate the text embeddings  $T_i$  for a given dictionary from the CLIP’s text encoder. In addition to the information of the novel categories, we also need knowledge of the base categories. Thus, we use all category names as the dictionary. We clas-

sify each anchor by calculating the cosine similarity score between its image feature and all text embeddings. We use the score after the softmax of the predicted category as the objectness score  $o_i$  of the anchor. We finally select the anchors with high objectness scores after the Non-Maximum Suppression (NMS) as CLIP Proposals.

All the CLIP Proposals  $C_i$  and their feature from CLIP  $c_i$  are generated offline. The  $c_i$  can be expressed like this:  $V(Crop(I, C_i))$ , where  $V$  and  $Crop$  are the same as those in  $ins_i$  definition. We also add 1.2x enlarged base categories’ GT bbox as part of the CLIP Proposals. Our experiments show that even though these CLIP Proposals are noisy, they are still meaningful regions for distillation and help our detector perform better on novel categories.

### 3.3. Model Structure

Compared with the traditional two-stage detector, our model structure has three main differences: CLIP Proposals’ feature distillation, cosine-similarity-based classifier, and semantic-based regressor. The model structure overview is shown in Fig 4.

**Proposals’ Feature Distillation.** To obtain knowledge from CLIP and map the image feature from our model into the CLIP’s feature space, we distill the knowledge from CLIP by minimizing the L1 loss between the CLIP Proposals’ feature from the CLIP’s vision encoder  $c_i$  and the one from our model  $c'_i$ . The  $c'_i$  can be expressed as  $Conv_c(Align(Bbone(I), C_i))$ , where  $Bbone(I)$  means getting the feature map by passing the image  $I$  through the backbone  $Bbone$ , and  $Align$  means doing the RoIAlign base on the CLIP Proposal  $C_i$  on the feature map.  $Conv_c$  means passing the feature after the RoIAlign through the convolution and linear layers in the classification branch.

In the CLIP Proposal generation, we know the objectness score of each proposal. For the proposal with a higher objectness score, it has a higher probability of having an object in it. Therefore, we should assign a higher weight to these proposals. We directly use the objectness score as the



weight, and the distillation loss is formulated like this:

$$L_{dist} = \frac{1}{M} \sum_{i=1}^M o_i |c_i - c'_i|_1 \quad (1)$$

where  $o_i$  is the objectness score of the CLIP Proposal  $C_i$  and  $M$  is the number of the CLIP Proposal on one image.

**Cosine-similarity-based Classifier.** By distilling the knowledge from CLIP, we are able to map the image feature of our model to CLIP feature space. Instead of using a learnable linear layer as the classifier, we use text embedding generated from CLIP’s text encoder. In the training phase, we only need the name of the base categories, which are then converted into text embedding  $B_i$ . For each proposal  $P_i$  given by the Region Proposal Network(RPN), we generate its feature  $p_i$ . The  $p_i$  can be expressed as:  $p_i = Conv_c(Align(Bbone(I), P_i))$ , where  $Conv_c$ ,  $Align$ , and  $Bbone$  are the same as those in  $c'_i$  definition. The classification loss is given by:

$$L_{cls} = \frac{1}{N} \sum_{i=1}^N L_{CE}(softmax(\mathbf{cos}_i), y_i) \quad (2)$$

where  $N$  is total number of proposals,  $y_i$  is the assigned label for the proposal  $P_i$ . The vector  $\mathbf{cos}_i$  for the proposal  $P_i$  is defined as  $[cos(p_i, B_1), \dots, cos(p_i, B_n), cos(p_i, BG)]$  in which  $n$  is number of base categories,  $cos$  is the cosine similarity score, and  $BG$  is a learnable vector for background.

At inference time, we also need to detect the novel categories. We generate the text embedding for both the base  $B_i$  and the novel  $N_i$ . The vector  $\mathbf{cos}_i$  become  $[cos(p_i, B_1), \dots, cos(p_i, B_n), cos(p_i, N_1), \dots, cos(p_i, N_k), cos(p_i, BG)]$  where  $k$  is the number of the novel categories.

**Semantic-based Regressor.** To improve the performance of the regression module, we add the semantic information of each category into consideration. We concatenate the text embedding with the proposal feature to predict the bbox. For each foreground proposal  $P_i$  given by the Region Proposal Network(RPN), we generate its feature for regression  $r_i$ . The  $r_i$  can be expressed as  $Conv_r(Align(Bbone(I), P_i))$ , where  $Align$  and  $Bbone$  are the same as those in  $c'_i$  definition, and  $Conv_r$  means passing the feature after the RoIAlign through the convolution layers and linear layers in the regression branch. The regression loss is defined as:

$$L_{reg} = \frac{1}{K} \sum_{i=1}^K L_1(Linear(Cat(r_i, B_{y_i})), a_i) \quad (3)$$

where  $K$  is the total number of the foreground proposals,  $Linear$  is the linear layer for bbox prediction,  $Cat$  mean concatenation,  $B_{y_i}$  means the text embedding of the assigned GT label  $y_i$  for the proposal  $i$ , and  $a_i$  is the GT bbox

for the proposal  $P_i$ . At inference time, since we no longer have the GT label. We concatenate the  $r_i$  with the text embedding  $B_{pred_i}$  or  $N_{pred_i}$  of the predicted category of the proposal  $P_i$ , where  $pred_i = \arg \max(\mathbf{cos}_i)$ .

Finally, our overall loss function is given by:

$$L = L_{dist} + L_{cls} + L_{reg} \quad (4)$$

## 4. Experiments

We first present our model result on COCO [22] and LVIS [10] detection benchmark in section 4.1. In section 4.2 we compare our method with ViLD to show the efficiency of our method. Finally, we conduct the ablation study with visualization analysis.

**Implementation Details.** We use the publicly available pre-trained CLIP model ViT-B/32 as the open-vocabulary classification model, with an input size of  $224 \times 224$ .

We finetune the LayerNorm layers in the CLIP with base categories instances in COCO or LVIS based on the setting and maintain all other parameters fixed. All the instances are cropped by 1.2x enlarged GT bboxes. We use an AdamW optimizer with a learning rate of 0.0001 and clip the L2 norm of the gradients when larger than 0.1. We finetune the model for 12 epochs.

For CLIP Proposal generation, we first resize the image with the image ratio maintained. We generate the anchors on each image with a stride of 32 pixels and with 5 different sizes (32, 64, 128, 256, 512), and 3 different ratios (1:1, 2:1, 1:2). We select the top 1000 anchors after NMS as CLIP Proposals on each image. If we will use the adapted CLIP features to train our detector we use the adapted CLIP to generate the CLIP Proposals. Otherwise, we use the unadapted CLIP for CLIP Proposal generation. In model training, we randomly select a fixed subset with 200 CLIP Proposals on each image for training. We provide more implementation details in the supplementary material.

### 4.1. Comparison with Current Methods

In this section, we evaluate EZAD in the COCO and LVIS detection benchmarks.

**Datasets and Evaluation Metrics.** We evaluate EZAD on COCO and LVIS(v1). For the COCO dataset, we use train2017 for training and val2017 for validation. Following [35], we divide COCO into 48 base and 17 novel categories. For the LVIS dataset, we use the training/validation images for training/evaluation. Previous work uses Frequent and Common categories as the base (866 categories), and Rare categories as the novel (337 categories). We argue that in this split, the rare category objects are so sparse (less than 0.5% annotations in the validation set) that the model’s performance on it is not representative. Therefore, we propose a new split called LVIS-Fbase, which uses the

Method	Mask	Epoch	Novel Category Knowledge Source	AP50		
				Base	Novel	Overall
ZSD-YOLO [33]	✗	50	CLIP’s feature	31.7	13.6	19.0
CORA [32]	✗	5	CLIP	35.5	<b>35.1</b>	35.4
F-VLM [16]	✓	6	CLIP	-	28.0	38.0
PBBL [7]	✓	-	ALBEF [18], CLIP, Pseudo-label	46.1	30.8	42.1
OVOS [25]	✗	36	CLIP’s feature	51.3	20.3	43.2
VTP-OVD [23]	✓	-	CLIP’s feature, Pseudo-label	31.5	29.8	46.1
OADP [30]	✓	-	CLIP’s feature, Pseudo-label	53.3	30.0	47.2
CondHead [31]	✓	107	CLIP’s feature	60.8	29.8	<b>52.7</b>
OV-DETR [34]	✗	50	CLIP’s feature	<b>61.0</b>	29.4	<b>52.7</b>
ViLD [9]	✓	107	CLIP’s feature	59.5	27.6	51.3
Ours	✓	36	CLIP’s feature, Novel Category Name	59.9	31.6	52.1

Table 2. Evaluation results on COCO benchmark. All the models are trained with the ResNet-50 backbone. Mask indicates whether the model is trained with Mask annotations. Our model achieves 4% improvement on Novel with 1/3 training time of ViLD.

Method	Epoch	AP			
		Freq	Comm	Rare	All
ViLD*	24	24.9	12.2	11.2	17.5
Ours	24	<b>30.9</b>	<b>14.3</b>	<b>12.5</b>	<b>20.5</b>
ViLD*	48	26.4	13.2	11.3	18.5
Ours	48	<b>31.9</b>	<b>15.2</b>	<b>13.1</b>	<b>21.3</b>

Table 3. Evaluation results on LVIS-Fbase benchmark. The ViLD\* is our reproduced result of ViLD with a shorter training schedule. EZAD outperforms ViLD in both 2x and 4x settings due to the efficient feature distillation.

Frequent categories as the base (405 categories), and both Common and Rare categories as the novel (common has 461 categories). On COCO, AP50 is used as the evaluation metric, while on LVIS the AP is used.

**Model.** We train a Mask R-CNN [12] model with ResNet-50 [13] FPN [21] backbone. The backbone is pre-trained on ImageNet [4]. We use SGD as the optimizer with batch size 4, learning rate 0.005, momentum 0.9, and weight decay 0.0001. We adopt linear warmup for the first 500 iterations, with a warm-up ratio is 0.001. On COCO, we train our model with 36 epochs and divide the learning rate by 10 at epoch 27 and epoch 33. On LVIS, We train our model with 48 epochs and divide the learning rate by 10 at epoch 32 and epoch 44. We train our model with multi-scale train-time augmentation. **Baselines.** We compare our method with different Zero-Shot Detection (ZSD) and Open-Vocabulary Detection (OVD) methods, which do not use additional datasets so that they can fit in the Zero-shot Annotation Detection (ZAD) setting. We mainly compare EZAD with ViLD [9] and CondHead [31], which uses distillation to obtain information on the novel categories from CLIP. They have the best overall performance on the COCO dataset. However, it uses the data augmentation of large-scale jittering [8] with an extremely long training schedule.

**Results.** Table 2 shows the results of EZAD in the COCO detection benchmark. Both the CORA [32] and F-VLM [16] are prompt-base methods. Although they need

Method	Epoch	AP50		
		Base	Novel	Overall
ViLD*	12	48.3	17.2	40.2
Ours	12	<b>55.7</b>	<b>30.4</b>	<b>49.0</b>
ViLD*	36	56.0	24.2	48.5
Ours	36	<b>59.9</b>	<b>31.6</b>	<b>52.1</b>

Table 4. Evaluation results on COCO. EZAD outperforms ViLD in both 1x and 3x settings, showing our distillation is more efficient.

less training time and have strong novel category performance, their performance in the base categories is much worse than other methods. ZSD-YOLO [33] and OVOS [25] use a one-stage detector and suffer from high false positives in the detection results, which causes low AP in novel categories. PBBL, VTP-OVD, and OADP are trained with pseudo-labels, introducing noise and harming the model performance on base categories. OV-DETR [34], CondHead [31] and ViLD [9] are pure distillation-based method. While the OV-DETR uses the transformer architecture which is heavier and has a better result than ViLD, the CondHead proposes a new detection head that can improve ViLD performance. Our EZAD achieves 59.9% and 31.6% on base and novel, respectively. Its performance is 4% better than ViLD in the novel and has a better performance in base and overall settings, with the use of 1/3 of ViLD’s training time.

Table 3 shows the results of EZAD in our LVIS-Fbase benchmark. The ViLD is trained 468 epochs on LVIS, which is too long to be fully reproduced. We reproduce the ViLD with 24 and 48 epochs. Our method uses ViLD’s model structure and distillation pipeline, and the improvement of our method comes from the distillation feature and distillation region. Therefore, reproducing the ViLD with a shorter training schedule will not affect the fairness of the comparison. With 48 epochs, our method shows a 5% improvement in Freq and 2% improvement in both Comm and Rare over the ViLD with the same training schedule. In

Method	Base				Novel				General			
	L	M	S	Avg	L	M	S	Avg	L	M	S	Avg
w/o Adaptation	69.9	70.2	46.6	66.8	90.7	<b>82.3</b>	48.5	77.7	61.3	62.2	36.9	53.9
w/ Adaptation	<b>92.5</b>	<b>89.5</b>	<b>80.3</b>	<b>87.5</b>	<b>91.1</b>	81.6	<b>66.3</b>	<b>81.9</b>	<b>84.8</b>	<b>80.7</b>	<b>68.8</b>	<b>78.3</b>

Table 5. Adapting CLIP to the detection dataset’s domain. The table presents the classification accuracy (ACC) of CLIP(w/ or w/o adaptation) when it is applied to classify the COCO dataset’s instance. The ACC is aggregated based on the size of the instances. After the adaptation, the ACC is improved by a huge margin in three different settings, especially for small objects.

CLIP Feature	Distill Region	Base	Novel	Overall
Raw	RPN Proposal	48.8	17.5	40.6
Adapted	RPN Proposal	56.9	24.6	48.5
Raw	CLIP Proposal	48.7	19.3	41.7
Adapted	CLIP Proposal	55.7	<b>30.4</b>	49.0

Table 6. Ablation study on CLIP’s feature adaptation and CLIP Proposal using COCO detection benchmark.

SB Reg	PPDW	Base	Novel	Overall
		55.5	28.2	48.3
✓		55.8	29.8	48.5
✓	✓	55.7	<b>30.4</b>	49.0

Table 7. Ablation study on Semantic-based regressor and per CLIP Proposal distillation weight using COCO detection benchmark. SB Reg means the Semantic-based regressor, and PPDW means the per CLIP Proposal distillation weight.

Proposal	IoGT	#(IoGT $\geq$ 0.8)	#(IoGT $\geq$ 0.5)
RPN	0.340	362818 (9%)	610157 (15%)
CLIP(Ours)	<b>0.365</b>	<b>563799 (14%)</b>	<b>870830 (21%)</b>

Table 8. The effective distillation region of different proposals. The table presents the Intersection between the proposal and the novel GT bboxes over novel GT bboxes (IoGT), the number of proposals that have high IoGT (IoGT $\geq$ 0.8, IoGT $\geq$ 0.5), and the percentage of these proposals in all proposals. Our CLIP proposals can cover more novel categories instances thus improving the distillation efficiency.

the CLIP proposals generation, EZAD only uses frequent and common object names to generate proposals. Therefore, EZAD does not take advantage of the category name on the Rare. We believe that the enhancement in performance for the rare and frequent categories can be attributed to eliminating the domain gap, whereas the improvement in common categories results from both the CLIP proposals and domain adaptation.

Because we have utilized the ZAD setting, there are novel instances present in our training images. This enables us to effectively leverage the information from these novel instances, thus achieving better performance than existing methods. However, the same cannot be achieved with Zero-Shot Detection (ZSD). On the other hand, we don’t need additional annotations like many OVD methods do, meaning our ZAD setting can be applied to more real-life settings.

## 4.2. Efficiency Evaluation

In this section, we compare EZAD with our reproduced ViLD to show the efficiency of our method. In Table 4, we

present our model and our reproduced ViLD on COCO with 1x and 3x training schedules. Our method is consistently better than ViLD in two different settings. Table 3 shows EZAD and our reproduced ViLD on the LVIS dataset with 2x and 4x training schedules. Our method shows substantial improvement over the ViLD with the same training schedule. These results suggest that the adapted feature space and the CLIP Proposals improve the distillation quality and efficiency. Thus, the model performance is improved.

## 4.3. Ablation Study and Visualization

In this section, we conduct ablation studies using the COCO detection benchmark. All the experiment details are the same as mentioned in section 4.1. We train our detector for 12 epochs in all experiments of this section.

**CLIP’s Feature Adaptation and CLIP Proposal.** Table 5 presents the classification result of adapting the CLIP to the COCO dataset’s domain. We evaluate the classification accuracy(ACC) on the instances in the COCO validation set. We use the splits in Section 4.1. Since the novel setting is a 17-ways classification which is easier than the other two settings, the ACC in the novel setting is much higher than the one in the other two.

Before the adaptation, the ACC in the general setting is only about 53.9%, which is much lower than the ACC of the classifier in a well-trained detector. This phenomenon indicates that there is a huge domain gap between the training data of CLIP and the detection dataset. While the objects in CLIP’s training image are clear and large and at the center of the image, the objects in the detection dataset might be small and occluded. Although we only fine-tuned the CLIP on base categories instances, we observed a huge improvement in all three settings, especially for the small objects. The average ACC reaches 87.5%, 81.9%, and 78.3% in the Base, Novel, and General settings, respectively. The ACC for the small objects improved by 33.7%, 17.8%, and 31.9% in three different settings. This indicates that the simple fine-tuning method can effectively bridge the domain gap and make the feature more discriminating.

Table 6 shows the effectiveness of the CLIP Proposals and the CLIP’s feature adaptation in the detection setting. Using the adapted CLIP’s feature for distillation can consistently improve both base and novel categories performance, no matter which kinds of distillation regions we use. Our model performance on novel benefits from using CLIP Pro-



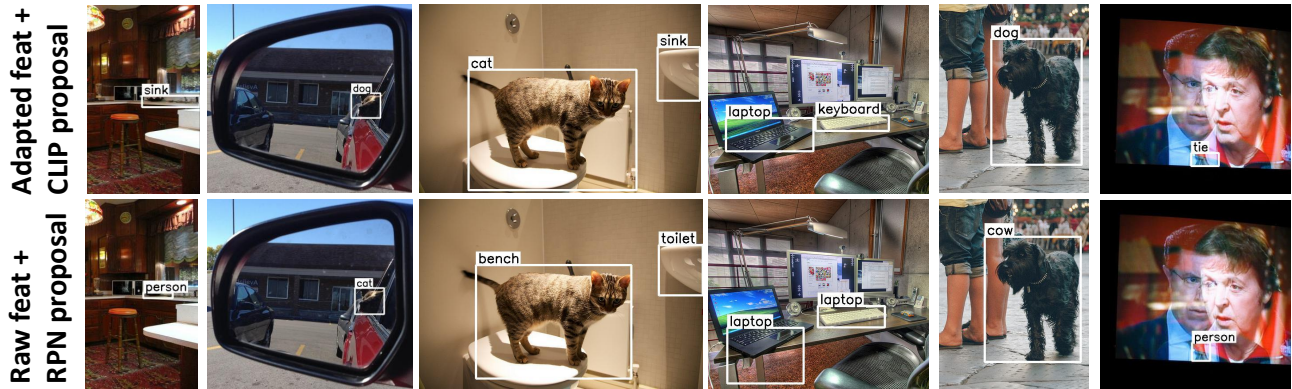


Figure 5. The visualization result on the COCO. The first row presents the results of the model trained with the adapted CLIP features from CLIP Proposals. The second row presents the results of the model trained with raw CLIP features from RPN proposals.

posals as distillation regions (30.4% vs. 24.6% with adapted CLIP features and 19.3% vs 17.5% with raw CLIP features). We believe that the adaptation of the CLIP and CLIP Proposal are complementary to each other, which makes the improvement given by CLIP Proposals with adapted CLIP (5.8%) larger than CLIP Proposals with raw CLIP (1.8%).

The performance on base categories of models trained with CLIP Proposals is slightly worse than those trained with RPN proposals since the CLIP Proposals focus more on the regions with novel categories. Our experiment results show that with a longer training schedule, the performance gap on the base categories can be eliminated while the advantage on the novel will be maintained. We will provide these results in the supplementary material.

**Semantic-based Regressor and Per CLIP Proposal Distillation Weight.** Table 7 shows the effects of semantic-base regressor and per CLIP Proposal distillation weight. All experiments use adapted CLIP’s features and CLIP Proposals as distillation regions. The semantic-base regressor helps the model perform better on both base and novel categories, showing that the semantic meaning of the categories does provide useful information to the regressor. Combining the semantic-base regressor and per CLIP Proposal distillation weight, the AP50 on novel reaches 30.4%. This indicates the reweighting of different distillation boxes further improves the distillation quality.

**Statistic and Visualization Analysis.** To compare which proposals can provide more meaningful novel categories information, we compare the effective distillation region of different proposals in an 8000 images subset of COCO training set in Table 8. We calculate the Intersection between the proposal and the novel GT bboxes over novel GT bboxes (IoGT) and present the number and the percentage of the proposals that have high IoGT in all proposals. Our CLIP proposals are 6% higher than the RPN proposal in the percentage of the high IoGT proposals, meaning that the CLIP proposal can cover more potential novel instances and improve the distillation efficiency.

Figure 5 provides some visualizations of the detected novel objects on the COCO benchmark. The first row presents the results from the model trained with the adapted CLIP features from CLIP Proposals (AFCP). The second row presents the results from the model trained with raw CLIP features from the RPN proposal (RFRP). The results show that though two models can localize the object correctly, the AFCP model has a higher classification accuracy than the RFRP model thanks to the adapted features and the more meaningful distillation regions.

## 5. Conclusion

We propose a new setting for detecting novel instances, called Zero-shot Annotation object Detection (ZAD), to move one step closer to the real-life scenario. Our method Efficient Feature Distillation for Zero-shot Annotation Detection (EZAD) in this paper. EZAD successfully bridging the domain gap between the classification and detection datasets, and selecting the meaningful distillation regions (CLIP Proposals) for obtaining knowledge from CLIP. Benefiting from these two improvements, EZAD outperforms previous distillation-based work in both COCO and LVIS with a much shorter training schedule. We believe our work provides a solid solution for applying zero-shot annotation detection in real life, and we hope our method can inspire other works in the future.

## Acknowledgement

This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA8750-19-1-1000. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation therein. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Laboratory, DARPA or the U.S. Government.



## References

- [1] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 384–400, 2018. 1, 3
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2
- [3] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [6] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Promptdet: Expand your detector vocabulary with uncurated images. *arXiv preprint arXiv:2203.16513*, 2022. 2, 3
- [7] Mingfei Gao, Chen Xing, Juan Carlos Niebles, Junnan Li, Ran Xu, Wenhao Liu, and Caiming Xiong. Towards open vocabulary object detection without human-provided bounding boxes. *arXiv preprint arXiv:2111.09452*, 2021. 2, 3, 6
- [8] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021. 6
- [9] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 2, 3, 6
- [10] Agrim Gupta, Piotr Dollár, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. 5
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 6
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2
- [15] Konwoo Kim, Michael Laskin, Igor Mordatch, and Deepak Pathak. How to adapt your large-scale vision-and-language model, 2022. 3, 4
- [16] Weicheng Kuo, Yin Cui, Xiuye Gu, AJ Piergiovanni, and Anelia Angelova. F-vlm: Open-vocabulary object detection upon frozen vision and language models. *arXiv preprint arXiv:2209.15639*, 2022. 2, 3, 6
- [17] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 3
- [18] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 6
- [19] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022. 2
- [20] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 3
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 6
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 3, 5
- [23] Yanxin Long, Jianhua Han, Runhui Huang, Xu Hang, Yi Zhu, Chunjing Xu, and Xiaodan Liang. P3 ovd: Fine-grained visual-text prompt-driven self-training for open-vocabulary object detection. *arXiv preprint arXiv:2211.00849*, 2022. 3, 6
- [24] Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*, 2021. 3
- [25] Zongyang Ma, Guan Luo, Jin Gao, Liang Li, Yuxin Chen, Shaoru Wang, Congxuan Zhang, and Weiming Hu. Open-vocabulary one-stage detection with hierarchical visual-language knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14074–14083, June 2022. 3, 6
- [26] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 3
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry,

- Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 3, 4
- [28] Shafin Rahman, Salman Khan, and Nick Barnes. Improved visual-semantic alignment for zero-shot object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11932–11939, 2020. 3
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 2
- [30] Luting Wang, Yi Liu, Penghui Du, Zihan Ding, Yue Liao, Qiaosong Qi, Biaolong Chen, and Si Liu. Object-aware distillation pyramid for open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11186–11196, 2023. 3, 6
- [31] Tao Wang. Learning to detect and segment for open vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7051–7060, 2023. 3, 6
- [32] Xiaoshi Wu, Feng Zhu, Rui Zhao, and Hongsheng Li. Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7031–7040, 2023. 2, 3, 6
- [33] Johnathan Xie and Shuai Zheng. Zero-shot object detection through vision-language embedding alignment. In *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 1–15. IEEE, 2022. 3, 6
- [34] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. In *European Conference on Computer Vision*, pages 106–122. Springer, 2022. 2, 6
- [35] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. 2, 3, 5
- [36] Shiyu Zhao, Zhixing Zhang, Samuel Schuster, Long Zhao, BG Vijay Kumar, Anastasis Stathopoulos, Manmohan Chandraker, and Dimitris N Metaxas. Exploiting unlabeled data with vision and language models for object detection. In *European Conference on Computer Vision*, pages 159–175. Springer, 2022. 3
- [37] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. 2, 3
- [38] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. 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*, pages 350–368. Springer, 2022. 2, 3
- [39] Pengkai Zhu, Hanxiao Wang, and Venkatesh Saligrama. Don’t even look once: Synthesizing features for zero-shot detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11693–11702, 2020. 3