

Point, Segment and Count: A Generalized Framework for Object Counting

Zhizhong Huang¹ Mingliang Dai¹ Yi Zhang² Junping Zhang¹ Hongming Shan^{3†}

¹ Shanghai Key Lab of Intelligent Information Processing, School of Computer Science,
Fudan University, Shanghai 200433, China

² School of Cyber Science and Engineering, Sichuan University, Chengdu 610065, China

³ Institute of Science and Technology for Brain-inspired Intelligence & MOE Frontiers Center for
Brain Science & Key Laboratory of Computational Neuroscience and Brain-Inspired Intelligence,
Fudan University, Shanghai 200433, China

{zzhuang19, mldai21, jpzhang, hmshan}@fudan.edu.cn, yzhang@scu.edu.cn

Code: <https://github.com/Hzzone/PseCo>

Abstract

Class-agnostic object counting aims to count all objects in an image with respect to example boxes or class names, a.k.a few-shot and zero-shot counting. In this paper, we propose a generalized framework for both few-shot and zero-shot object counting based on detection. Our framework combines the superior advantages of two foundation models without compromising their zero-shot capability: (i) SAM to segment all possible objects as mask proposals, and (ii) CLIP to classify proposals to obtain accurate object counts. However, this strategy meets the obstacles of efficiency overhead and the small crowded objects that cannot be localized and distinguished. To address these issues, our framework, termed PseCo, follows three steps: point, segment, and count. Specifically, we first propose a class-agnostic object localization to provide accurate but least point prompts for SAM, which consequently not only reduces computation costs but also avoids missing small objects. Furthermore, we propose a generalized object classification that leverages CLIP image/text embeddings as the classifier, following a hierarchical knowledge distillation to obtain discriminative classifications among hierarchical mask proposals. Extensive experimental results on FSC-147, COCO, and LVIS demonstrate that PseCo achieves state-of-the-art performance in both few-shot/zero-shot object counting/detection.

1. Introduction

Object counting has attracted growing research interest in recent years. It aims to estimate the specific object counts in an

† Corresponding author.

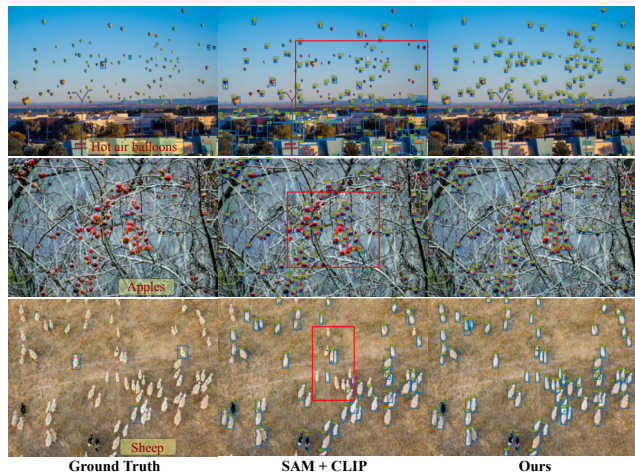


Figure 1. Sample results of vanilla SAM + CLIP and the proposed method. Given the class name (zero-shot) or example boxes (few-shot), our method can detect all objects in the image for counting.

image, especially in extremely crowded scenes that cannot be distinguished or counted one-by-one by humans. Traditional object counting methods typically focus on specific categories such as humans [4], animals, or cars. One well-known direction is crowd counting which counts all presented persons in an image. However, it requires a labor-intensive amount of training data with point annotations and is limited to the pre-defined category once the model is trained. Therefore, the recent efforts in object counting resort to class-agnostic counting, which counts arbitrary categories with a few guidance from example images [5, 10, 24, 25, 28, 31, 33] and class names [30, 34, 38]. It achieves satisfactory performance even though the categories are unseen during training, which thus reduces the burden of data starvation.

Most of class-agnostic object counting methods involve generating density maps, which can be summed to derive the object counts. Typically, they compute the similarity between visual features of input and example images to guide the object counting [5, 24, 31, 33], *i.e.*, few-shot object counting if examples are provided. In contrast, zero-shot object counting [38] only uses class names to select the best examples in the image. In summary, they mainly focus on how to improve the quality of similarity maps to produce density maps through better examples [38], transformer [24], and attention [5]. However, these density-based methods lack interpretability as density maps are hard to verify. On the other hand, detecting target objects for counting is a potential solution but box annotations are much more cumbersome to collect than points. Although the literature explores using only a few boxes and all point annotations [28], it presents inferior performance on both tasks as the model usually overfits training categories [36].

Recent years have witnessed a great breakthrough in the foundation models of computer vision, *e.g.*, Segment Anything Model (SAM) [21] for segmentation and Contrastive Language-Image Pre-Training (CLIP) [29] for the vision language model. Both of them have shown great zero-shot potential in generalizing to novel scenes. A simple solution for object counting is to employ SAM to segment everything in an image and use CLIP to classify the regions with respect to given examples/texts. However, it remains challenging to combine their advantages for object counting. Fig. 1 reveals the problems of directly applying SAM and CLIP to counting. *First*, the small objects tend to be missed by SAM, which cannot be localized by uniform grid point prompts. *Second*, it is still time-consuming if directly using CLIP to classify the cropped image regions. *Third*, object counting needs more discriminative classifications, especially for small objects, otherwise most of them cannot be distinguished from the background.

In this paper, we propose a generalized framework for object counting, termed PseCo, to address these issues in the following aspects. *First*, instead of using a predefined uniform grid point prompts for SAM to segment everything in an image, we propose a class-agnostic object localization that estimates a heatmap of all objects, from which we can infer the object coordinates of each object. Subsequently, it can provide accurate but least point prompts for SAM, segmenting small objects while reducing the computation costs. *Second*, we propose to leverage the CLIP text/image embeddings to classify the image regions, formulating a generalized framework for both zero-shot and few-shot object counting. Hence, our framework can detect and count arbitrary classes using examples/texts. *Third*, we propose a hierarchical knowledge distillation to distill the zero-shot capabilities of CLIP to our PseCo. It discriminates among the hierarchical mask proposals produced by SAM, enabling

our PseCo to distinguish the desired objects from a large number of proposals. Fig. 1 presents the results of PseCo, which effectively detects and distinguishes the target objects, including those that are very small.

The contributions are summarized as follows:

- We present PseCo, a generalized framework that leverages the advantages of SAM and CLIP for both few-shot and zero-shot object detection and counting.
- We propose a novel class-agnostic object localization, which can provide sufficient, accurate but least point prompts for SAM to segment all possible objects.
- We further propose a generalized object classification to identify target objects, with a hierarchical knowledge distillation to learn discriminative classification.
- Experimental results on various benchmarks, including FSC-147, COCO, and LVIS datasets, demonstrate the effectiveness of PseCo on object detection/counting.

2. Related Work

Class-agnostic object counting. Traditional object counting focuses on specific categories such as car [27] and human [12, 13, 18, 22, 35, 44], which can be divided into density-based and detection-based methods. Density-based methods [12, 22, 44] predict and sum over density maps to infer the counting results. The detection-based methods [13, 18] resort to object detection for counting. The former performs well in crowded scenes, and the latter provides better interpretability but needs box annotations. The class-agnostic counting [3, 5, 10, 19, 24, 25, 30, 33, 34, 38] is not limited to specific categories. Instead, they count the target objects through some exemplar bounding boxes of a new category (few-shot) [5, 10, 14, 24, 25, 33, 34], or a class name (zero-shot) [19, 38]. Most of them compute the similarity maps between visual features of images and examples to infer density maps. Although zero shot, [38] selects the best proposals with respect to class names to construct examples. C-DETR [10] detects objects only trained on point and a few box annotations. SAM-Free [34] generates mask prior from grid points and selects better point prompts from the CLIP similarity map without training. However, it is inferior, especially for small objects.

CLIP-based object detection. CLIP [29] learns well-aligned text-image embeddings, and can be applied to object detection [15, 41–43, 45, 46], segmentation [6–8], adversarial attack [40], and generation [37]. Typically, ViLD [15] distills the knowledge of CLIP to the region classifier of Mask R-CNN [17] so that CLIP can extract the classification weights from text or image for any novel class. Region-CLIP [45] retrains CLIP on the data of region-text pairs. OV-DETR [41] is transformed into a binary matching problem conditioned on CLIP embeddings. CLIP enables these methods with zero-shot capability for open-vocabulary ob-

ject detection. In contrast, the proposed PseCo is a generalized framework for few-shot/zero-shot object detection and uses a novel hierarchical knowledge distillation to encourage discriminative classifier between mask proposals.

3. The Proposed Approach

3.1. Preliminaries

SAM [21] consists of three key components: (i) an image encoder to extract image features; (ii) a prompt encoder to encode the user-input prompts; and (iii) a mask decoder to predict the segmentation masks under the control of given prompts. In particular, it utilizes points and boxes as prompts to generate high-quality hierarchical object masks, including whole, part, and subpart segmentations of each prompt. Due to massive training data, SAM exhibits promising zero-shot segmentation performance on various benchmarks. In particular, SAM receives uniform grid point prompts to segment every possible object in an image.

CLIP [29] learns well-aligned vision-language representations with contrastive loss from large-scale image-text pairs. It contains two separate encoders for each modality but maps the data into the same embedding space. In zero-shot classification, CLIP builds the image classifier using a predefined template, *e.g.*, ‘A photo of dog’ when only the class name ‘dog’ is available.

Two-stage object detection such as Faster R-CNN [32] divide the object detection into two stages. The first stage uses a region proposal network (RPN) to produce a coarse set of object boxes (proposals) and class-agnostic objectness scores. The second stage takes these proposals for classification and refines the box coordinates. In particular, open-vocabulary object detection [15, 43, 45, 46] focuses on improving zero-shot classification of the second stage, which uses CLIP embeddings as the classification weights.

SAM and CLIP have shown their potential in zero-shot segmentation and classification. In this paper, we study a challenging problem about how to synergize their advantages for object counting under a similar framework of two-stage object detection, without compromising their zero-shot capabilities when generalizing to novel scenes.

3.2. Problem Formulation and Framework

As presented in Fig. 2, given an input image, our goal is to count the target objects with respect to a set of image/text queries. Instead of predicting a density map, we formulate the object counting as object detection; that is, detect and count them all.

Inspired by two-stage object detection, we build our framework into the following steps: point, segment, and count as shown in Fig. 2 under the help of SAM and CLIP. Specifically, PseCo (i) points out all possible objects using least point coordinates, (ii) uses SAM to generate the corre-

sponding mask proposals conditioned on the point prompts, and (iii) classifies and post-processes all proposals to count target objects. In a sense, point and segment steps perform a similar role as the RPN in Faster R-CNN to provide sufficient class-agnostic object proposals for the subsequent object classification. On the other hand, the counting step can filter the undesired proposals by thresholding the scores with respect to examples/texts. At this step, we can further use non-maximum suppression (NMS) to remove duplicate proposals like object detection.

A simple baseline. Following the spirit of point, segment, and count, we design a simple baseline that leverages both SAM and CLIP. Uniform grid points, *e.g.*, 32×32 are used as prompts for SAM to segment all objects. The image regions are cropped from the input image by predicted proposals, which are then fed into CLIP for classification.

However, such a simple baseline has the following limitations. *First*, 32×32 grid points may be insufficient to enumerate all objects, especially under crowded scenes. Consequently, many small objects could be ignored, which is not suitable for object counting. Although increasing the number of points can alleviate this problem, it could inevitably waste heavy computational costs as many points are located in the background, which is impractical in real-world applications. *Second*, CLIP can only produce image-level representations, and it is computationally expensive to crop proposals during inference time. Although the vision encoder of CLIP can be distilled [15], the representations are not discriminative enough for small regions, as shown in Fig. 1.

To address the above problems, we propose a novel framework in Fig. 2, termed PseCo, for generalized object counting and detection. Specifically, PseCo only trains a point decoder to point out all objects using the least points under the setting of keypoint estimation and a classifier to classify all proposals. Both of them are built upon the image features extracted from the pre-trained image encoder of SAM, leading to negligible computational costs.

We detail their designs in the following.

3.3. Class-agnostic Object Localization

As presented in Fig. 3 (a), (b), (c), SAM ignores some small objects using uniform grid point prompts. To localize all possible objects with the least points, we propose to formulate this problem as keypoint estimation. Let $\mathbf{I} \in \mathbb{R}^{H \times W}$ be the input image, the point decoder aims to produce the class-agnostic keypoint heatmap $\widehat{\mathbf{H}} \in [0, 1]_{\frac{H}{s} \times \frac{W}{s}}$, where s is output stride. In SAM, \mathbf{I} is resized to $H = W = 1024$ along the longest side, and the point decoder shares the same architecture as the mask decoder with $s = 4$.

Although ground-truth points are available during training, the point decoder is prone to overfit the categories in the training data [36], which cannot generalize to zero-shot scenarios that the categories during testing are unseen by

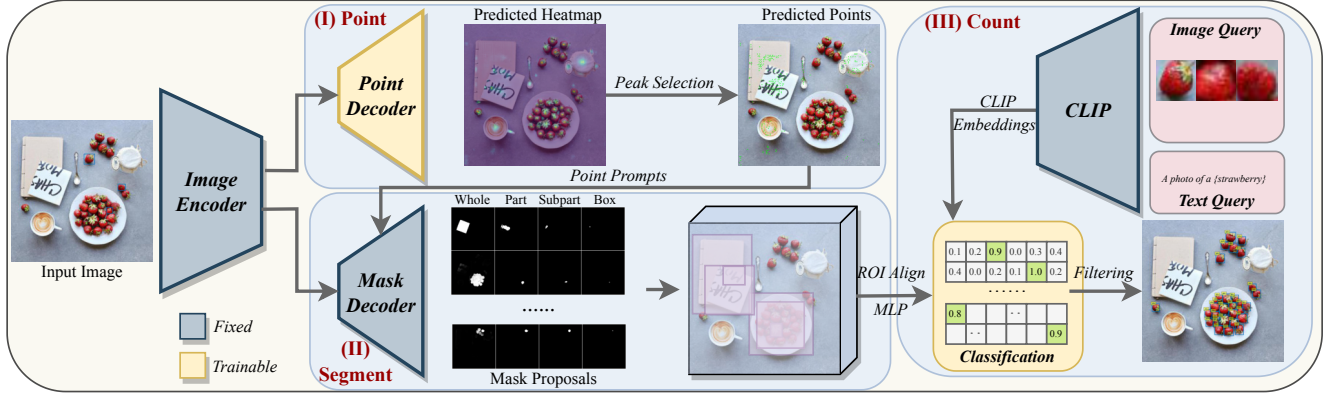


Figure 2. Illustration of the proposed PseCo, following the steps: point, segment, and count. Given an input image, the point decoder predicts the class-agnostic heatmap to point out all objects. The image encoder and mask decoder from SAM are fixed during training (the prompt encoder is omitted here) and output the mask proposals. The proposals are classified with respect to CLIP image/text embeddings.

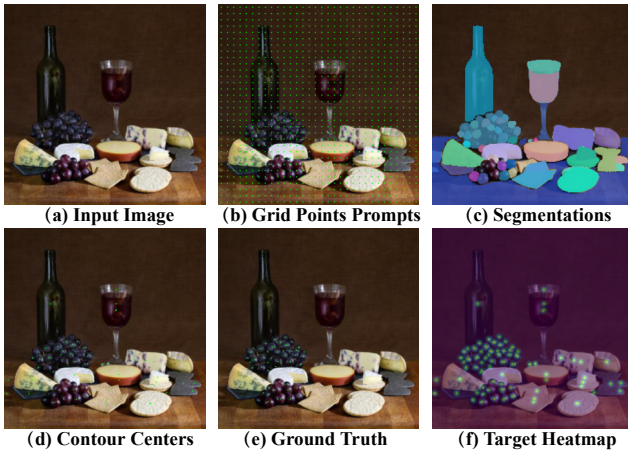


Figure 3. Sample results to generate the class-agnostic target heatmaps. Given (a) input image and (b) uniform grid point prompts, SAM predicts all (c) segmentation. We combine (d) all contour centers of segmentations to avoid bad point prompts and (e) ground-truth point annotations to produce (f) target heatmap. The resultant heatmap will be used to supervise the point decoder.

the model. In other words, the point annotations do not contain all possible objects in the image, and thus the novel categories could be misclassified as background. Benefitting from the zero-shot capabilities of SAM, we can combine the detected objects in Fig. 3 (d) and the ground-truth ones in Fig. 3 (e) to produce the final target heatmap in Fig. 3 (f) to train the point decoder. Consequently, the target map can include as many objects as possible, and retain the ones missed by SAM at the same time.

Here, we compute the contour centers of all mask proposals predicted from uniform prompts, since we find that most point prompts may not be accurately located at the center of objects. It is worthwhile to note that there may be a few duplicated points for the same object; their mask proposals

can be removed during post-processing.

We train the point decoder following [2]. We splat all estimated keypoints into a heatmap $\mathbf{H} \in [0, 1]_{\frac{H}{s} \times \frac{W}{s}}$ using a Gaussian kernel $\mathbf{H}_{xy} = \exp(-\frac{(sx-p_x)^2 + (sy-p_y)^2}{2\sigma^2})$, where $p \in \mathbb{R}^2$ is the keypoint and $\sigma = 2$ is the standard deviation according to [2]. x and y are the coordinates on the heatmap. If the Gaussians of different objects overlap, we take the element-wise maximum [47]. The point-wise mean squared error is employed for training:

$$\mathcal{L}_{\text{point}} = \|\widehat{\mathbf{H}} - \mathbf{H}\|_2^2. \quad (1)$$

It is worthwhile to note that the quantization errors caused by output stride are not taken into account. The goal of the point decoder is to provide good class-agnostic point prompts for SAM, instead of an accurate object localization. SAM can segment certain objects as long as they are pointed out.

During inference, a 3×3 max-pooling is applied to the heatmap to extract the peak key points, and the top K of them with the scores above the threshold are selected. As a result, we can detect all points from 256×256 grid with only $K = 1000$, which has the same computational costs as 32×32 grid points. In practice, there are fewer points than K as studied in our experiments.

3.4. Generalized Object Classification

Given all proposals produced in Sec. 3.3, this section aims to provide scores with respect to the image/text queries. In object counting, the image queries are cropped from input images according to the example bounding boxes. We construct the classification weights $\mathbf{W} \in \mathbb{R}^{C \times D}$ from the fixed CLIP language embeddings of class names or image embeddings of example boxes. C is the arbitrary number of queries, D is the dimension of CLIP embeddings and novel queries can be appended to the end of \mathbf{W} . The region features \mathbf{r} are extracted from the image features processed by

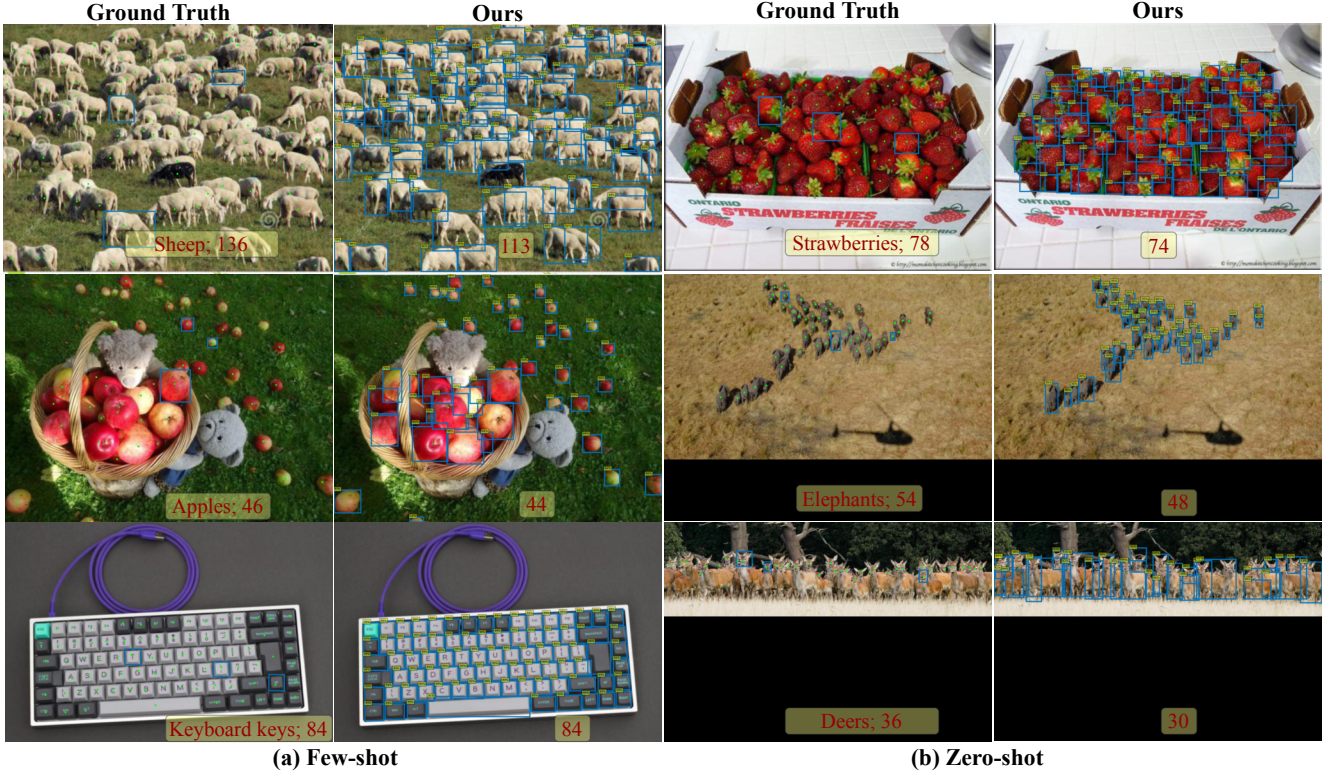


Figure 4. Qualitative results for (a) few-shot and (b) zero-shot object counting and detection. The class names, ground-truth counts, and our predicted counts are in color boxes. Zoom in for better view.

ROI align [17] and a two-layer MLP. The object detector is supervised by the annotations in the image:

$$\mathcal{L}_{\text{cls}} = \text{BCE}(\mathbf{W}r, \mathbf{c}), \quad (2)$$

where BCE is the binary cross-entropy loss following [46], and \mathbf{c} is the ground-truth labels. \mathbf{c} can be all zeros if the proposals are not matched with any ground-truth boxes.

In practice, this design yields unsatisfying results in generalizing to novel classes, as zero-shot capability of CLIP can be compromised when simply applying classification loss to the object classifier. Existing solutions include knowledge distillation [15] and enlarged vocabulary with image-caption training data [46]. They are limited to object counting, which needs more discriminative representations since most scenes of object counting are crowded with small objects.

Hierarchical knowledge distillation. We instead propose to align the region features and CLIP image embeddings of the hierarchical mask proposals from SAM. Similar to Eq. 2, for the mask proposals obtained from the same point, we build the classification weights from the CLIP image embeddings of cropped image regions. The region features are discriminated with corresponding CLIP embeddings according to their overlapping. In doing so, the image encoder can be distilled to the classifier which meanwhile becomes more

discriminative. This loss can be written as:

$$\mathcal{L}_{\text{kd}} = \frac{1}{M} \sum_{i=1}^M \text{BCE}(\mathbf{W}'r^{(i)}, \mathbf{c}'), \quad (3)$$

where M is the number of proposals of each point, $\mathbf{W}' \in \mathbb{R}^{M \times D}$ is the CLIP embeddings of image regions, and $\mathbf{c}' \in \mathbb{R}^M$ is filled 1 if the IoU between two proposals is larger than 0.5, otherwise 0. It is found that SAM usually fails to segment small objects in crowded scenes. To this end, we opt to use an additional 16×16 box around each point to improve the segmentation of small objects, and only the first mask is selected. We note that the image regions and corresponding CLIP embeddings can be prepared in advance before training, similar to [15]. The visual illustration is shown in supplementary Fig. 6.

3.5. Training and Inference

The training loss function of our framework is the combination of Eqs. (1), (2), and (3):

$$\mathcal{L} = \mathcal{L}_{\text{point}} + \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{kd}}. \quad (4)$$

During inference, non-maximum suppression is applied to all proposals to remove duplicate proposals, and the object counts are the number of detected bounding boxes with a score larger than the predefined threshold.

4. Experiments

4.1. Implementation Details

Datasets. For object counting, FSC-147 dataset [31] is used to evaluate our method. It includes 6135 images of 147 categories, where the categories do not overlap between different splits. For object detection, FSC-147 and FSCD-LVIS [16] annotated by [28] are employed. Since there are no box annotations in FSC-147 training data, we generate the pseudo-labels through the ground-truth point prompts. Although the pseudo-labels may be noisy, we find that it is sufficient to train a good classification network.

Training details. PseCo is trained for 50k iterations with a mini-batch size of 32, Adam optimizer [20] with a fixed learning rate of 10^{-4} , weight decay of 10^{-5} , $\beta_1 = 0.9$ and $\beta_2 = 0.99$. We utilize ViT-H [9] of SAM and ViT-B of CLIP. The point decoder is initialized by the mask decoder of SAM. $K = 1000$ and the threshold of heatmap is 0.05. 256 proposals and 16 pairs of each sample are randomly selected to train the classifier with \mathcal{L}_{cls} and \mathcal{L}_{kd} , making sure 25% positive proposals. No augmentation is used. The IoU threshold of NMS is 0.5.

Evaluation metrics. For object counting, we adopt two standard metrics in the literature, including Mean Average Error (MAE) and Root Mean Squared Error (RMSE). In particular, $MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$ and $RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}$, where N is the number of samples, y and \hat{y} are the ground-truth and predicted object counts. We also reported Normalized Relative Error (NAE) and Squared Relative Error (SRE) in supplementary Tab. 6. For object detection, we use AP and AP50 strictly following [28].

4.2. Experimental Results

4.2.1 Qualitative Results

Fig. 4 showcases example results on few-shot/zero-shot object counting/detection. Our PseCo has produced distinct detection results and accurate counts. PseCo can detect these small objects and discriminate well between the objects and background with the given example images/texts. However, our method is slightly not robust to the occlusion of objects. This is because SAM cannot distinguish occluded objects. Interestingly, PseCo can employ the text prompts to detect the objects accurately, even though given bad examples. For example, the example boxes of the deer in the last samples are only annotated around the head. We can address this problem by text prompts, whereas [38] selects better example boxes. We show failure cases in supplementary Fig. 7.

In addition, Fig. 5 presents the qualitative comparisons with few-shot object counting: C-DETR [28], BMNet+ [33] and FamNet [31], and zero-shot object counting: SAM-Free [34] and ZSOC [38]. We directly referred the results

from their published papers to avoid any potential bias of self-implementation. PseCo has more interpretable detection results than density maps and performs competitively at extremely crowded scenes. PseCo also presents superior and discriminative detections than C-DETR and SAM-Free.

4.2.2 Quantitative Results

We evaluate the proposed PseCo on the crowded class-agnostic counting dataset FSC-147 [39], under the setting of few-shot and zero-shot object counting. As a detection-based counting method, we also evaluate on the object detection datasets, FSC-147 [39] and FSCD-LVIS [28].

Results on few-shot object counting. In the few-shot counting scenario, each image provides three bounding box annotations of exemplar objects, which are used to count the target object in this image. Example images are cropped from input image and used to extract CLIP image embeddings as classification weights. Tab. 1 shows the quantitative comparisons with recent state-of-the-art methods, including detection-based and density-based methods.

Our PseCo achieves comparable MAE and RMSE with state-of-the-art density-based methods such as CounTR [24] and BMNet+ [33]. PseCo also shows significant improvements over the detection-based methods. FR [10], FSOD [10] and C-DETR [28] detect all proposals based on the state-of-the-art Faster R-CNN [32] or DETR [1]. Their performance is limited since there are no sufficient training data to enable the detection models for better generalization ability. SAM [21, 26] and SAM-Free [34] segment all objects and compute their similarities with examples to identify desired objects. However, they use only 32×32 grid point prompts, leading to the failure in detecting small objects. Our PseCo can address this problem with class-agnostic object localization for more accurate point prompts, and generalize well to novel categories under the help of CLIP.

Results on zero-shot object counting. Similar to the few-shot setting, we use the CLIP text embeddings as classification weights when only known class names. The results are shown in Tab. 2. We have reproduced ViLD [15] on the proposed class-agnostic localization (CAL) for better comparisons. ViLD significantly outperforms SAM [21, 34] and SAM-Free [34], validating the effectiveness of class-agnostic object localization. Replacing ViLD with the proposed classifier, the performance is further improved. We find that there exists a great gap between few-shot and zero-shot settings, which may be caused by the ambiguous class names in FSC-147 dataset, such as the go pieces labeled as ‘go game’.

Results on object detection. We evaluate the performance of object detection under both few-shot and zero-shot settings on the test set of FSC-147 and FSCD-LVIS [10]. The results are reported in Tab. 3. Our PseCo achieves almost $2 \times$ performance improvements compared to C-DETR [28],

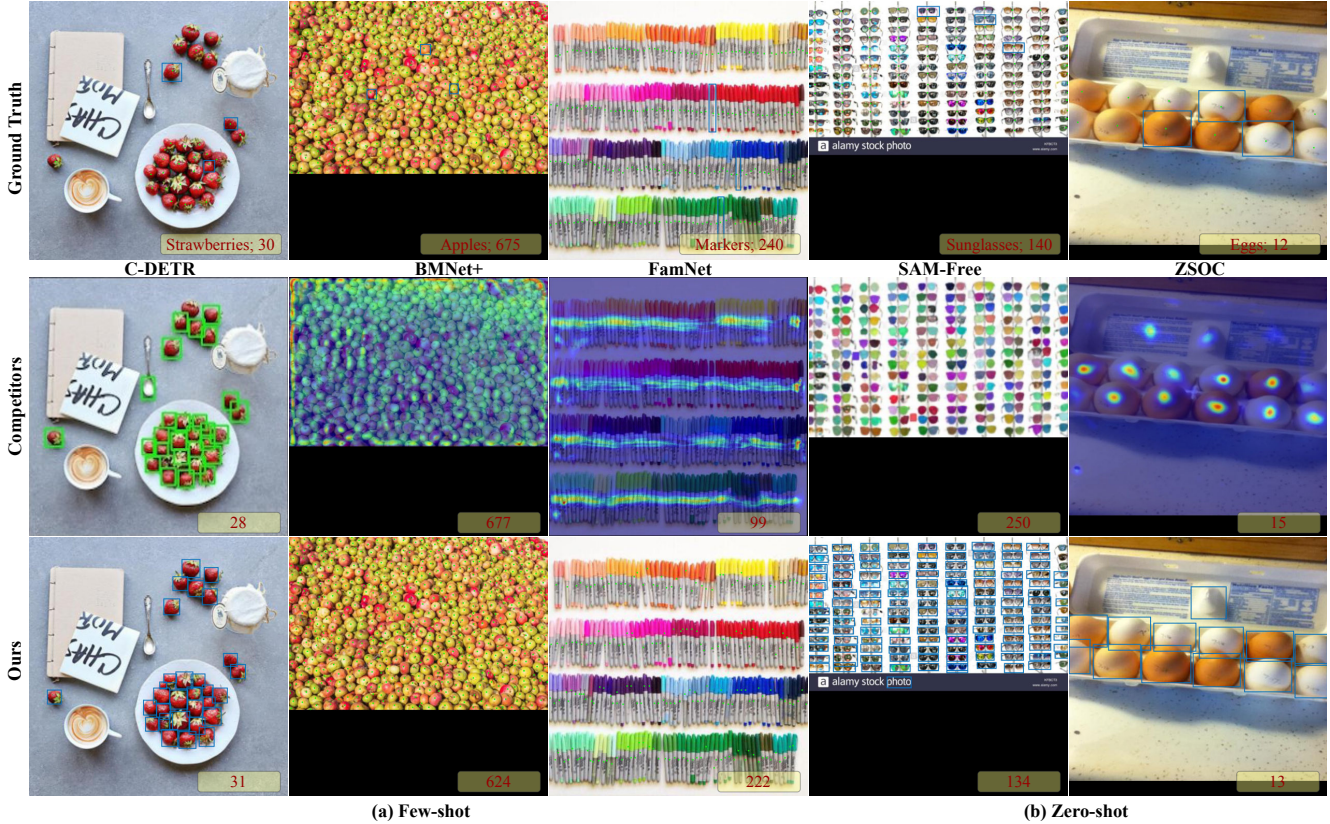


Figure 5. Qualitative comparisons for (a) few-shot (the first 3 columns) and (b) zero-shot (the last 2 columns) object counting. Only final points are placed in the second and third columns due to crowded predicted boxes. Zoom in for better view.

due to the use of SAM and the proposed components. SAM is trained on large-scale datasets so that it can generalize well to extract accurate mask proposals for unseen categories. CAL+ViLD would degenerate the performance, but is still better than baselines, demonstrating the effectiveness of class-agnostic object localization (CAL). Interestingly, PseCo behaves oppositely between two datasets under zero-shot/few-shot settings; that is, PseCo performs better with few shots than zero shots on FSC-147, versus on FSCD-LVIS. We think it happens since the example images in FSCD-LVIS are worse than text prompts.

4.3. Ablation Study

We evaluate different components in terms of few-shot detection/counting on FSC-147, and report the results of test set in Tab. 4. The baseline is detailed in Sec. 3.2.

(i) **Ablation on the localization.** We compare our proposed class-agnostic object localization with two variants: (i) grid points and (ii) point decoder trained with only ground-truth points. The same trained classification network is used for fair comparisons. The object localization from the point decoder can improve both detection and counting performance compared to grid points. However, the point decoder

	Val set		Test set	
	MAE \downarrow	RMSE \downarrow	MAE \downarrow	RMSE \downarrow
GMN [25]	29.66	89.81	26.52	124.57
MAML [11]	25.54	79.44	24.90	112.68
FamNet [31]	23.75	69.07	22.08	99.54
BMNet+ [33]	15.74	58.53	14.62	91.83
CounTR [24]	13.13	49.83	11.95	91.23
FR [10]	45.45	112.53	41.64	141.04
FSOD [10]	36.36	115.00	32.53	140.65
C-DETR [28]	-	-	16.79	123.56
SAM [21, 26]	31.20	100.83	27.97	131.24
SAM-Free [34]	-	-	19.95	132.16
Ours	15.31	68.34	13.05	112.86

Table 1. Results on few-shot object counting. The first and second parts contain the density-based and detection-based methods. We note that detecting objects for counting is much more challenging than predicting density maps. The best results are shown in bold.

could overfit the training categories if only using ground-truth annotations, which thus harms its generalization performance to unseen categories in testing data. PseCo combines all possible objects to alleviate this problem.

	Val set		Test set	
	MAE \downarrow	RMSE \downarrow	MAE \downarrow	RMSE \downarrow
RepRPN-C [30]	30.40	98.73	28.32	128.76
CounTR [24]	17.40	70.33	14.12	108.01
ZSOC [38]	26.93	88.63	22.09	115.17
SAM [21, 34]	-	-	42.48	137.50
SAM-Free [34]	-	-	24.79	137.15
CAL+ViLD [15]	28.81	111.13	17.80	130.88
Ours	23.90	100.33	16.58	129.77

Table 2. Results on zero-shot object counting.

	FSC-147		FSCD-LVIS	
	AP \uparrow	AP50 \uparrow	AP \uparrow	AP50 \uparrow
C-DETR [28]	22.66	50.57	4.92	14.49
Ours	43.53	74.64	22.37	42.56
CAL+ViLD [15]	40.56	67.21	19.67	39.33
Ours	41.14	69.03	23.93	44.54

Table 3. Results on few-shot/zero-shot object detection.

(ii) **Ablation on the SAM.** We use ViT-H as the default SAM in this paper. By comparing different versions of SAM, as expected, heavier SAM introduces more improvements. However, our method of ViT-B version still performs better than C-DETR [28] as in Tab 1.

(iii) **Ablation on \mathcal{L}_{kd} .** We remove \mathcal{L}_{kd} from PseCo and find that the performance drops. \mathcal{L}_{kd} can distill the knowledge of CLIP to the object classification network, and thus greatly improve the detection performance on unseen classes. It can further enable the object classification network to discriminate the hierarchical and small mask proposals from SAM. As a result, the counting performance, especially RMSE, is significantly improved.

(iv) **Ablation on the computation costs.** Compared to vanilla SAM in [34] that employs 32×32 grid point prompts, PseCo only selects an average of 378/388 candidate points for each image in the FSC-147 test/val sets. These points are selected from 256×256 heatmaps, 64 times more than SAM. Detailed discussion is in supplementary Sec. 6.

4.4. Results on Large-scale Datasets

Although most methods evaluate their object counting performance on FSC-147 dataset due to a large number of objects in each image, we further evaluate PseCo on two more practical, complex but sparse datasets: COCO [23] and LVIS [16] under open-vocabulary object detection. Specifically, the categories in testing data may be unseen by the detection model, and only class names are known during testing. We strictly follow the experimental settings in the state-of-the-art object detection method Detic [46]. On COCO, we report the AP50_n for novel classes and AP50 for all classes. Sim-

	Detection		Counting	
	AP \uparrow	AP50 \uparrow	MAE \downarrow	RMSE \downarrow
Baseline	39.63	67.50	21.24	129.62
(i) Grid Heatmap (only GT)	41.66	71.53	17.15	121.17
(ii) ViT-B	39.83	70.72	16.41	120.40
ViT-L	42.58	73.18	14.65	118.64
(iii) w/o \mathcal{L}_{kd}	41.69	70.59	14.57	127.95
Ours	43.53	74.64	13.05	112.86

Table 4. Ablation study of different components in PseCo.

ilarly, we report mask AP on LVIS, *i.e.*, AP_n^m and AP^m on the novel and all classes. For fair comparisons, additional trained backbone in [46] and caption data is used to train the classification network. Note that we do not report the counting performance due to their nature of sparseness.

The results in Tab. 5 show that our PseCo achieves significant performance improvements over Detic.

	COCO		LVIS	
	AP50 _n \uparrow	AP50 \uparrow	AP _n ^m \uparrow	AP ^m \uparrow
OVR-CNN [42]	22.8	46.0	-	-
ViLD [15]	27.6	51.3	16.1	22.5
Detic [46]	27.8	45.0	17.8	26.8
Ours	32.9	47.5	20.6	30.8

Table 5. Results on large-scale but sparse detection datasets. We strictly follow the settings in Detic [46] for fair comparisons. The first part may use unfair conditions, *e.g.*, training data; all results are adopted from their corresponding papers.

5. Conclusion

In this paper, we introduce PseCo, a generalized framework for few-shot/zero-shot object detection/counting. PseCo follows the spirits: point, segment, and count, which synergizes the advantages of SAM and CLIP. It employs a class-agnostic object localization to provide good point prompts for SAM. Extensive experiments validate the effectiveness of PseCo on both object detection/counting, including FSC-147, large-scale COCO, and LVIS datasets. In the future, we will explore how to achieve fine-grained object counting inspired by current great success in multi-modal LLM.

Acknowledgement This work was supported in part by STI2030-Major Projects (No. 2021ZD0200204), National Natural Science Foundation of China (Nos. 62176059 and 62101136), Shanghai Municipal Science and Technology Major Project (No. 2018SHZDZX01) and ZJLab, and Shanghai Center for Brain Science and Brain-inspired Technology.

References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [2] Bowen Cheng, Bin Xiao, Jingdong Wang, Honghui Shi, Thomas S Huang, and Lei Zhang. Higherhrnet: Scale-aware representation learning for bottom-up human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5386–5395, 2020.
- [3] Hisham Cholakkal, Guolei Sun, Fahad Shahbaz Khan, and Ling Shao. Object counting and instance segmentation with image-level supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12397–12405, 2019.
- [4] Mingliang Dai, Zhizhong Huang, Jiaqi Gao, Hongming Shan, and Junping Zhang. Cross-head supervision for crowd counting with noisy annotations. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [5] Nikola Djukic, Alan Lukezic, Vitjan Zavrtanik, and Matej Kristan. A low-shot object counting network with iterative prototype adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18872–18881, 2023.
- [6] Jiahua Dong, Yang Cong, Gan Sun, Zhen Fang, and Zhengming Ding. Where and how to transfer: Knowledge aggregation-induced transferability perception for unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(3):1664–1681, 2024.
- [7] Jiahua Dong, Yang Cong, Gan Sun, Bineng Zhong, and Xiaowei Xu. What can be transferred: Unsupervised domain adaptation for endoscopic lesions segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4022–4031, June 2020.
- [8] Jiahua Dong, Duzhen Zhang, Yang Cong, Wei Cong, Henghui Ding, and Dengxin Dai. Federated incremental semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3934–3943, June 2023.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [10] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4013–4022, 2020.
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [12] Jiaqi Gao, Zhizhong Huang, Yiming Lei, Hongming Shan, Wang JZ, and Junping Zhang. Deep rank-consistent pyramid model for enhanced crowd counting. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [13] Eran Goldman, Roei Herzig, Aviv Eisenschstat, Jacob Goldberger, and Tal Hassner. Precise detection in densely packed scenes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5227–5236, 2019.
- [14] Shenjian Gong, Shanshan Zhang, Jian Yang, Dengxin Dai, and Bernt Schiele. Class-agnostic object counting robust to intraclass diversity. In *European Conference on Computer Vision*, pages 388–403. Springer, 2022.
- [15] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference On Learning Representations*, 2021.
- [16] 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.
- [17] 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.
- [18] Meng-Ru Hsieh, Yen-Liang Lin, and Winston H Hsu. Drone-based object counting by spatially regularized regional proposal network. In *Proceedings of the IEEE international conference on computer vision*, pages 4145–4153, 2017.
- [19] Ruixiang Jiang, Lingbo Liu, and Changwen Chen. Clip-count: Towards text-guided zero-shot object counting. *arXiv preprint arXiv:2305.07304*, 2023.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
- [22] Dongze Lian, Jing Li, Jia Zheng, Weixin Luo, and Shenghua Gao. Density map regression guided detection network for rgb-d crowd counting and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1821–1830, 2019.
- [23] 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [24] Chang Liu, Yujie Zhong, Andrew Zisserman, and Weidi Xie. Countr: Transformer-based generalised visual counting. *arXiv preprint arXiv:2208.13721*, 2022.
- [25] Erika Lu, Weidi Xie, and Andrew Zisserman. Class-agnostic counting. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 669–684. Springer, 2019.
- [26] Zhiheng Ma, Xiaopeng Hong, and Qinnan Shangguan. Can sam count anything? an empirical study on sam counting. *arXiv preprint arXiv:2304.10817*, 2023.
- [27] T Nathan Mundhenk, Goran Konjevod, Wesam A Sakla, and Kofi Boakye. A large contextual dataset for classification, detection and counting of cars with deep learning. In *Computer*

- Vision-ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III 14*, pages 785–800. Springer, 2016.
- [28] Thanh Nguyen, Chau Pham, Khoi Nguyen, and Minh Hoai. Few-shot object counting and detection. In *European Conference on Computer Vision*, pages 348–365. Springer, 2022.
- [29] 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, 2021.
- [30] Viresh Ranjan and Minh Hoai Nguyen. Exemplar free class agnostic counting. In *Proceedings of the Asian Conference on Computer Vision*, pages 3121–3137, 2022.
- [31] Viresh Ranjan, Udbhav Sharma, Thu Nguyen, and Minh Hoai. Learning to count everything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3394–3403, 2021.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, 2015.
- [33] Min Shi, Hao Lu, Chen Feng, Chengxin Liu, and Zhiguo Cao. Represent, compare, and learn: A similarity-aware framework for class-agnostic counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9529–9538, 2022.
- [34] Zenglin Shi, Ying Sun, and Mengmi Zhang. Training-free object counting with prompts. *arXiv preprint arXiv:2307.00038*, 2023.
- [35] Qi Wang, Junyu Gao, Wei Lin, and Xuelong Li. Nwpu-crowd: A large-scale benchmark for crowd counting and localization. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):2141–2149, 2020.
- [36] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*, 2020.
- [37] Yujie Wei, Shiwei Zhang, Zhiwu Qing, Hangjie Yuan, Zhiheng Liu, Yu Liu, Yingya Zhang, Jingren Zhou, and Hongming Shan. Dreamvideo: Composing your dream videos with customized subject and motion. *arXiv preprint arXiv:2312.04433*, 2023.
- [38] Jingyi Xu, Hieu Le, Vu Nguyen, Viresh Ranjan, and Dimitris Samaras. Zero-shot object counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15548–15557, 2023.
- [39] Hongyu Yang, Di Huang, Yunhong Wang, and Anil K Jain. Learning face age progression: A pyramid architecture of gans. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 31–39, 2018.
- [40] Zhao Yunlong, Deng Xiaoheng, Liu Yijing, Pei Xinjun, Xia Jiazhi, and Chen Wei. Fully exploiting every real sample: Super-pixel sample gradient model stealing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [41] 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.
- [42] 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.
- [43] Hao Zhang, Feng Li, Xuayan Zou, Shilong Liu, Chunyuan Li, Jianwei Yang, and Lei Zhang. A simple framework for open-vocabulary segmentation and detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1020–1031, 2023.
- [44] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 589–597, 2016.
- [45] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022.
- [46] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, pages 350–368. Springer, 2022.
- [47] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.