

OmniCount: Multi-label Object Counting with Semantic-Geometric Priors

Anindya Mondal^{1*}, Sauradip Nag^{2*}, Xiatian Zhu¹, Anjan Dutta¹

¹University of Surrey

²Simon Fraser University

{a.mondal, xiatian.zhu, anjan.dutta}@surrey.ac.uk, snag@sfu.ca

Abstract

Object counting is pivotal for understanding the composition of scenes. Previously, this task was dominated by class-specific methods, which have gradually evolved into more adaptable class-agnostic strategies. However, these strategies come with their own set of limitations, such as the need for manual exemplar input and multiple passes for multiple categories, resulting in significant inefficiencies. This paper introduces a more practical approach enabling simultaneous counting of multiple object categories using an open-vocabulary framework. Our solution, OmniCount, stands out by using semantic and geometric insights (priors) from pre-trained models to count multiple categories of objects as specified by users, all without additional training. OmniCount distinguishes itself by generating precise object masks and leveraging varied interactive prompts via the Segment Anything Model for efficient counting. To evaluate OmniCount, we created the OmniCount-191 benchmark, a first-of-its-kind dataset with multi-label object counts, including points, bounding boxes, and VQA annotations. Our comprehensive evaluation in OmniCount-191, alongside other leading benchmarks, demonstrates OmniCount’s exceptional performance, significantly outpacing existing solutions.

Web — <https://mondalanindya.github.io/OmniCount>

Introduction

Understanding object distribution across multiple categories is crucial for comprehensive scene analysis, driving increased interest in object counting research. It aims to estimate specific object counts in natural scenes. Traditionally, object counting has focused on class-specific methods for categories such as human crowds (Li, Zhang, and Chen 2018; Song et al. 2021; Han et al. 2023; Li et al. 2023; Liang et al. 2023; Liu et al. 2023a), cells (Khan 2016), fruits (Rahmoonfar and Sheppard 2017), and vehicles (Bui, Yi, and Cho 2020). However, these methods require extensive training data and are limited to predefined categories.

Recent efforts have shifted towards class-agnostic counting, using exemplars (cropped images and class names) to count arbitrary categories (Chattopadhyay et al. 2017; Ranjan et al. 2021; Ranjan and Nguyen 2022; Jiang, Liu, and

Chen 2023). Some operate in low-shot settings (You et al. 2022; Xu et al. 2023a), but they still require substantial training data and separate processing for each category, increasing computational demands for multi-category scenes. Further, detection and instance segmentation methods (Chattopadhyay et al. 2017; Cholakkal et al. 2022) can count multiple categories by name, but struggle with small or non-atomic objects like grapes or bananas, which are hard to detect individually.

We present **OmniCount**, a training-free method for simultaneous open-vocabulary object counting. Unlike previous approaches (Xu et al. 2023a; Dai, Liu, and Cheung 2024), OmniCount leverages vision-language models to count diverse object categories without training (*i.e.*, training-free), using semantic and geometric cues from pre-trained foundation models to partition images, identify occluded objects, and ensure precise counting through object recovery. Segmentation models often struggle with dense scenes, missing distant/occluded objects (Xie et al. 2022; Phillion and Fidler 2020). Our approach separates semantic and geometric priors, using metric depth to rectify dense scenes where semantic estimation fails. Through k -nearest-neighbors search, we refine missed instances and estimate class-specific reference points for similar objects (see Fig. 3). This enables detection across varying object properties, utilizing SAM (Kirillov et al. 2023) for point-prompted, fine-grained instance segmentation and counting.

Object counting is an under-studied area lacking a dataset with diverse annotations for multiple generic categories per image. Existing datasets like PASCAL VOC (Everingham et al. 2009) and MS COCO (Lin et al. 2014) are sparsely populated and inadequate for real-world counting challenges. The REC-8K dataset (Dai, Liu, and Cheung 2024) focuses on fine-grained counting within the same category but doesn’t support counting across different coarse-grained categories.

To address this gap, we create the **OmniCount-191** benchmark. It includes 302,300 object instances across 191 categories in 30,230 images, featuring multiple categories per image and detailed annotations such as counts, points, and bounding boxes for each object (Fig. 4).

We make the following contributions: (1) We re-promote multi-label object counting that bypasses the conventional reliance on object detection and semantic segmentation

*Authors have equal contributions.

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

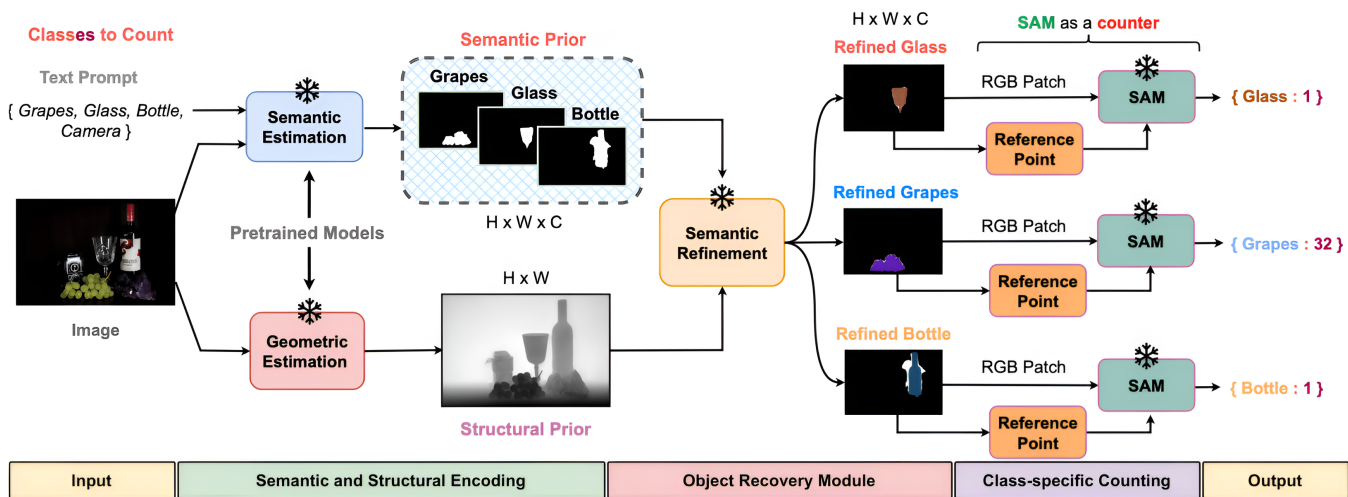


Figure 1: **OmniCount pipeline:** OmniCount processes the input image and target object classes using Semantic Estimation (SAN) and Geometric Estimation (Marigold) modules to generate class-specific masks and depth maps. These initial semantic and geometric priors are then refined through an Object Recovery module, producing precise binary masks. The refined masks help extract RGB patches and reference points, reducing over-counting. SAM then uses these RGB patches and reference points to generate instance-level masks, resulting in accurate object counts. (❄️ denotes pre-trained, frozen models)

models, addressing common accuracy issues such as over- and under-counting; (2) We introduce a novel, efficient, and user-friendly framework **OmniCount** for multi-label object counting by leveraging semantic and geometric cues without necessitating additional training; (3) We create a new multi-label object counting dataset, **OmniCount-191**, with rich annotations for fostering the development of this newly introduced setting; (4) We conduct extensive experiments to demonstrate OmniCount’s superior performance over existing methods on our dataset and establish benchmarks.

Related Work

Learning-based object counting: Traditional methods focus on specific categories, such as crowds (Li, Zhang, and Chen 2018; Song et al. 2021; Han et al. 2023; Huang et al. 2023; Liang et al. 2023; Liu et al. 2023a; Peng and Chan 2024; Guo et al. 2024), cells (Khan 2016), fruits (Rahmehoonfar and Sheppard 2017), and vehicles (Bui, Yi, and Cho 2020), using regression-based techniques to create density maps from point annotations (Lempitsky and Zisserman 2010; Zhang et al. 2016; Xu et al. 2021). These approaches predict counts by summing pixel values in density maps but are limited to their trained categories. Class-agnostic counting, in contrast, uses exemplars for flexible counting across categories (Lu, Xie, and Zisserman 2019; Zhang et al. 2019; Ranjan et al. 2021; Shi et al. 2022; Gupta et al. 2021; Zhang et al. 2021; Ranjan and Nguyen 2022; Shi, Mettes, and Snoek 2024) and includes data-efficient zero-shot (Xu et al. 2023a; Xu, Le, and Samaras 2023; Jiang, Liu, and Chen 2023; Dai, Liu, and Cheung 2024) and few-shot (You et al. 2022; Yang et al. 2021) variants. However, these methods require extensive data and training. We propose an open-vocabulary counter that leverages prompts like points, boxes, or text to count objects without training, en-

abling broader applicability without the data burden.

Multi-label object counting: Real-world scenarios often involve multiple coexisting object classes (You et al. 2022). Prior works (Cholakal et al. 2019, 2022; Chattopadhyay et al. 2017) addressed multi-label counting in sparse settings, focusing on global counts within human-discernible ranges but struggled with non-atomic or densely clustered objects like grapes. Few-shot methods (Ranjan et al. 2021) typically handle one category per image. Recently, (Dai, Liu, and Cheung 2024) introduced GrREC, a model for fine-grained multi-class counting, trained on predefined categories, alongside the REC-8K dataset with referring expressions. In contrast, our open-vocabulary model leverages semantic and geometric cues from pre-trained models without additional training, emphasizing the need for datasets like OmniCount-191 that capture real-world, dense, multi-class interactions.

Prompt-based foundation models: LLMs like GPT (Brown et al. 2020) and foundation models like CLIP (Radford et al. 2021) have revolutionized NLP and vision tasks, enabling effective zero- and few-shot transfer through textual prompts. In segmentation, the Segment Anything Model (SAM) (Kirillov et al. 2023) generates precise masks from diverse prompts (points, boxes, text) and demonstrates strong zero-shot capabilities. An ideal object counter should be *visually promptable*, *interactive*, and capable of *open-set* counting. While SAM meets these criteria, it struggles with occlusions (Ji et al. 2023) and multi-class counting (Shi, Sun, and Zhang 2024) due to its class-agnostic nature. We enhance SAM by integrating depth and semantic priors, improving its performance in complex counting tasks involving occlusions and multiple object classes.

OmniCount

In this work, we aim to achieve open-vocabulary, multi-label training-free object counting within a given image and with a set of labels to be counted in that image. Our proposed model is illustrated in Fig. 1.

Problem formulation

The problem of multi-label object counting can be defined as obtaining an object counter $\mathcal{F}_{\text{count}}$ using a training set $\mathcal{D}_{\text{train}} = \{(I_1, \mathcal{P}_1, \mathcal{C}_1), \dots, (I_N, \mathcal{P}_N, \mathcal{C}_N)\}$, where each $I_i \in \mathbb{R}^{H \times W \times 3}$ represents an RGB image, $\mathcal{P}_i = \{p_1, \dots, p_{m_i}\}$ is a set of class labels and $\mathcal{C}_i = \{c_1, \dots, c_{m_i}\}$ are the corresponding object counts (i.e. object with label p_k occurs c_k times in I_i), with m_i being the number of unique objects in the i -th image and N the total number of training data points in $\mathcal{D}_{\text{train}}$. For an image I_k and a subset of labels $\{p_1, \dots, p_{k_l}\} \subseteq \mathcal{P}_k$, the function $\mathcal{F}_{\text{count}}$ should result in:

$$\{c_1, \dots, c_{k_l}\} = \mathcal{F}_{\text{count}}(I_k, \{p_1, \dots, p_{k_l}\}) \quad (1)$$

where c_{k_l} is the number of occurrences of the object with label p_{k_l} in the image I_k . Our goal is to develop an open-vocabulary multi-label object counting model $\mathcal{F}_{\text{count}}$, such that it generalizes well to $\mathcal{D}_{\text{test}}$, a held-out test set of data points with classes not in $\mathcal{D}_{\text{train}}$, i.e., $\mathcal{D}_{\text{train}} \cap \mathcal{D}_{\text{test}} = \phi$. To achieve this, we introduce OmniCount, a multi-label object counting model that utilizes semantic and geometric priors, avoiding training that requires large datasets and expensive computational resources. Since our model is training-free, we do not use $\mathcal{D}_{\text{train}}$ and only evaluate our model on $\mathcal{D}_{\text{test}}$.

Semantic and structural encoding

Semantic estimation module: To count multiple objects in a single forward pass, we segment the image into relevant semantic regions. While any standard open-vocabulary segmentation model can be used, we employ the Side Adapter Network (SAN) (Xu et al. 2023b) as a semantic segmentation model \mathcal{E}_{sem} that takes an image I and a set of class labels $\mathcal{P} = \{p_1, \dots, p_m\}$ as input and results in $\mathbf{S}_{\mathcal{P}} = \{S_1, \dots, S_m\}$, a set of binary semantic masks corresponding to the classes in \mathcal{P} as follows:

$$\mathbf{S}_{\mathcal{P}}, F_{\mathcal{P}} = \mathcal{E}_{\text{sem}}(I, \mathcal{P}) \quad (2)$$

where $F_{\mathcal{P}} \in \mathbb{R}^{\frac{H}{K} \times \frac{W}{K} \times C}$ is an intermediate low-resolution feature activations, with K and C being integers depending on the design of \mathcal{E}_{sem} . We use $\mathbf{S}_{\mathcal{P}}$ as a semantic prior, bridging the gap between multi-label counting and semantic awareness. While segmenting 2D RGB images, SAM (Kirillov et al. 2023) primarily relies on texture information, such as colour. Combined with occlusion, this reliance can result in over-segmented masks (see Fig. 1 for the ‘‘bottle’’ class) and consequently, over-counting. To address this, we incorporate geometric information to achieve fine-grained segmentation, mitigating over-segmentation and over-counting. We refine the segmentation mask using geometric priors to be discussed in the next paragraph.

Geometric estimation module: Classical segmentation models (Long, Shelhamer, and Darrell 2015; Kirillov et al.

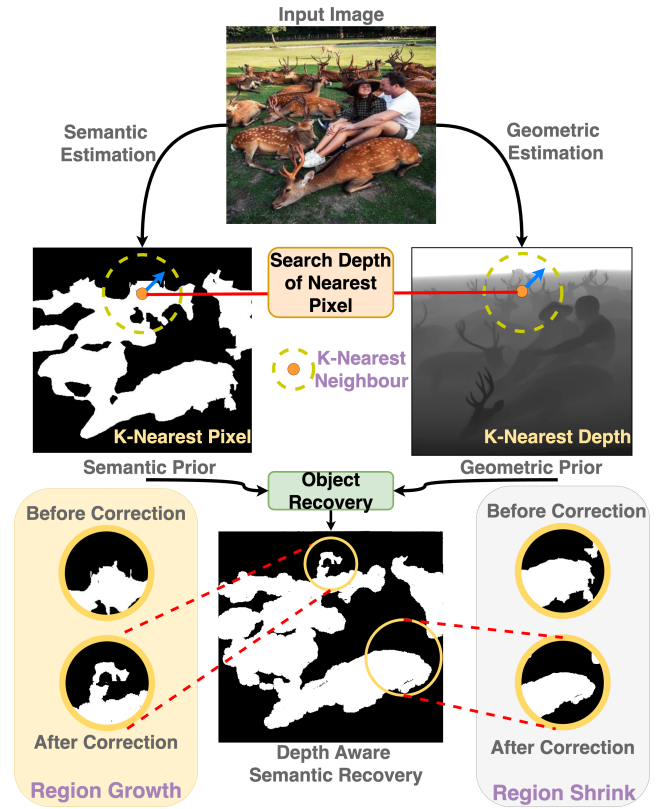


Figure 2: **Geometry aware Object Recovery:** We refine semantic masks with geometric priors using k-nearest neighbor searches to filter edge pixels by category uniqueness and depth alignment, enhancing mask precision through depth-integrated segmentation.

2023) primarily rely on texture information for object delineation, which often fails under significant occlusion.

So, counting dense interacting or overlapping objects requires information beyond RGB. Therefore, similar to how density maps have been utilized in classical object counting by providing a clearer representation of object distribution, we leverage depth maps to enhance segmentation accuracy. Depth maps, like density maps, ignore texture and focus on structural information, aiding in the segmentation of objects regardless of their distance from the camera, as shown in Fig. 1. This structural prior helps recover hidden objects and refine object semantics.

Using the off-the-shelf depth map rendering model ‘Marigold’ (Ke et al. 2024) denoted by $\mathcal{E}_{\text{depth}}$, we generate depth map D for the image I as $D = \mathcal{E}_{\text{depth}}(I)$, which serve as geometric prior for our model. Notably, D is a matrix of the same spatial dimension as the image I , with pixels having normalized depth values in $[0, 1]$. This implies that for a pixel $x \in I$, $D(x) \in [0, 1]$ indicates the normalized depth value of x . We utilize this pixel-wise depth information to refine the coarse segmentation mask $\mathbf{S}_{\mathcal{P}}$, discussed in the following subsection.

Geometry aware object recovery

Providing precise reference points is crucial for guiding SAM towards accurate counting and minimizing over-counting. We inject geometric priors into semantic priors before passing them into SAM (see Fig. 3) to obtain such reference points.

To refine the semantic prior $S_{\mathcal{P}}$ with geometric insights from the depth prior D , we conduct a k-nearest-neighbor (kNN) (refer to Fig. 2) search centering around each of the pixels of the individual semantic mask to ensure two conditions are met:

(1) A pixel must exclusively belong to its designated object category, preventing any overlapping with masks of other categories. For a pixel x in mask S_j :

$$C_1(x) : x \notin S_k, \forall k \neq j$$

(2) The absolute difference between a pixel’s depth and the mean depth of its object category must fall below a specified tolerance τ , ensuring depth consistency (for objects with curved edges). For a pixel x in S_n with mean depth $D_{\mu}(S_n)$:

$$C_2(x) : |D(x) - D_{\mu}(S_n)| < \tau$$

where $D(x)$ denotes the depth of pixel x and $D_{\mu}(S_n) = \frac{1}{s} \sum_y D(y), \forall y \in S_n$ with s as the total number of pixels in S_n . Pixels that fulfil both conditions $C_1(x)$ and $C_2(x)$ are integrated into their appropriate object class, leading to a refined semantic prior $S'_{\mathcal{P}}$, computed as $S'_{\mathcal{P}} = \mathcal{R}_{\text{geom}}(S_{\mathcal{P}}, D)$, where $\mathcal{R}_{\text{geom}}$ is the geometry-aware semantic refinement function, enhancing the precision of semantic masks by considering depth information. This depth-aware refined mask (see Fig. 2) minimizes the risk of over-segmentation (see Table 3) or recovers undiscovered objects in occluded scenes.

Reference point guided counting

We use SAM (Kirillov et al. 2023) as an object counter, which employs a point grid generator to place uniform points across the image and generate masks. This can lead to overcounting due to points falling on both foreground and background. To prevent this, we propose a reference point selection procedure that focuses solely on the foreground.

Reference point selection:

We select reference points (refer to Fig. 3) for the SAM decoder using the feature activation $F_{\mathcal{P}}$ from semantic priors (see Eq. (2)) to enhance text-image similarity accuracy. A set of reference points $\mathbf{P} = \{\rho_1, \dots, \rho_s\}$ are identified as local maxima within $F_{\mathcal{P}}$, but direct upsampling can misalign (Zhang et al. 2020) them due to quantization errors (refer to Table 3). To address this, we apply Gaussian refinement to the low-resolution reference points \mathbf{P} (Zhang et al. 2020), resulting in corrected reference points $\mathbf{P}' = \{\rho'_1, \dots, \rho'_s\}$. To specifically target foreground objects and avoid background segmentation, we compute the Hadamard product between the refined semantic mask S'_m for class p_m and the corrected reference points \mathbf{P}' as $\mathbf{Q}_m = S'_m \circ \mathbf{P}'$, where \circ represents the Hadamard product, and $\mathbf{Q}_m \subseteq \mathbf{P}'$ denotes the set of resulting reference points for objects belonging to the label p_m , serving as a guide for identifying regions likely to contain the target objects, as illustrated in Fig. 3. This

ensures the reference points \mathbf{Q}_m to guide the segmentation to the target objects. These per-class reference points act as density maps (see Fig. 7(b)), allowing for accurate object counting across varying densities. This automated selection can also be replaced by manual point or box annotations, making our model interactive for multiple user inputs.

SAM mask generator: By incorporating the reference object activation $F_{\mathcal{P}}$ and modifying the mask generation process, SAM’s mask decoder can better focus on the reference object features. This additional contextual information from $F_{\mathcal{P}}$ helps the mask generator accurately distinguish and segment target objects. Since SAM’s encoder requires an RGB image, we extract an RGB patch I_m of the target object by multiplying the input image I with the refined semantic mask S'_m (see Fig. 2). The resultant mask \mathbf{M}_m from SAM is obtained as $\mathbf{M}_m = \text{SAM}(I_m, \mathbf{Q}_m)$, where $\mathbf{M}_m = \{M_{m_1}, \dots, M_{m_n}\}$ is the set of individual object masks segmented by SAM. We count these masks to determine the total number of objects for class label p_m . This approach focuses on target objects without segmenting unrelated entities, enhancing efficiency and accuracy beyond the standard “segment everything” strategy. Finally, masks that are empty or cover an insignificantly small area are discarded, creating a refined subset $\mathbf{N}_m \subseteq \mathbf{M}_m$ containing only significant masks for the final object count. The cardinality of \mathbf{N}_m , $\text{card}(\mathbf{N}_m)$, denotes the final count of objects of class-label m in image I .

OmniCount-191 Dataset

To effectively evaluate OmniCount across open-vocabulary, supervised, and few-shot counting tasks, a dataset catering to a broad spectrum of visual categories and instances featuring various visual categories with multiple instances and classes per image is essential. The current datasets, primarily designed for object counting (Ranjan et al. 2021) focusing on singular object categories like humans and vehicles, fall short for multi-label object counting tasks. Despite the presence of multi-class datasets like MS COCO (Lin et al. 2014), PASCAL VOC (Everingham et al. 2009), and REC-8K (Dai, Liu, and Cheung 2024), their utility is limited for counting due to the sparse nature of object appearance and fine-grained referencing. Addressing this gap, we created a new dataset with 30,230 images spanning 191 diverse categories, including kitchen utensils, office supplies, vehicles, and animals. This dataset features a wide range of object instance counts per image, ranging from 1 to 160, with an average of 10, bridging the existing gap and setting a benchmark for assessing counting models in varied scenarios.

Dataset statistics: The OmniCount-191 benchmark presents images with small, densely packed objects from multiple classes, reflecting real-world object counting scenarios. This dataset encompasses 30,230 images, with dimensions averaging 700×580 pixels. Each image contains an average of 10 objects, totalling 302,300, with individual images ranging from 1 to 160. We use the same annotation rules defined in existing counting datasets (Ranjan et al. 2021). To ensure diversity, the dataset is split into training and testing sets, with no overlap in object categories – 118 categories for training and 73 for testing,

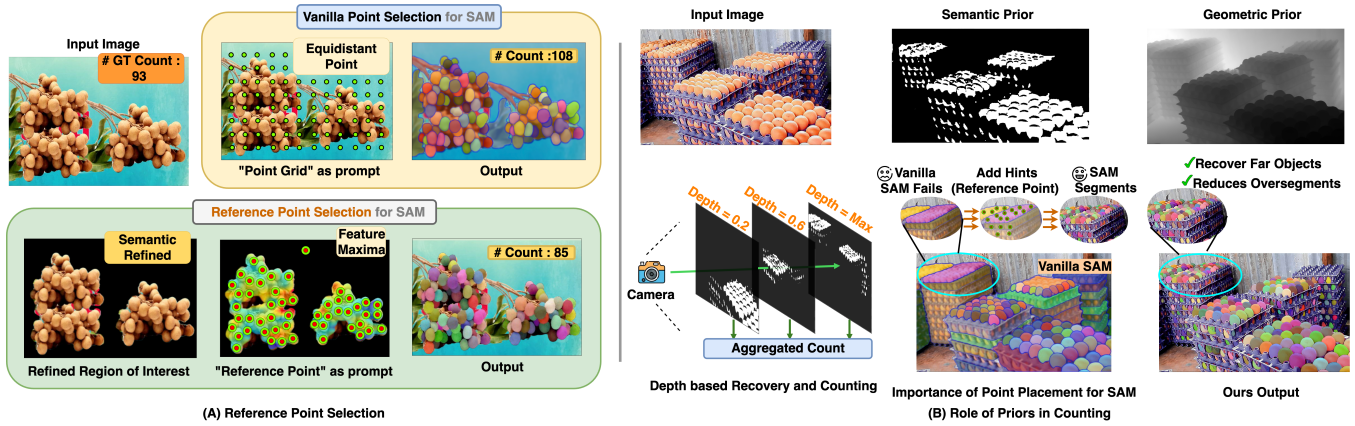


Figure 3: **Reference Point Selection:** SAM’s segmentation accuracy is enhanced by refining reference point selection. Panel (A) shows how integrating semantic priors, identifying local maxima, and applying Gaussian refinement improve reference point accuracy, focusing them on foreground objects for better segmentation and counting. Panel (B) demonstrates the benefits of incorporating semantic and geometric priors, where depth-based recovery and precise reference points help SAM recover distant or occluded objects, reducing over-segmentation issues found in the default “everything mode”.




Figure 4: **OmniCount-191 Annotations:** A collection of images with 191 classes across nine domains, annotating each image with captions, VQA, boxes, and points.

Methods	Training	PASCAL VOC		OmniCount-191	
		mRMSE ↓	mRMSE-nz ↓	mRMSE ↓	mRMSE-nz ↓
ILC (Cholakkal et al.)	✓	<u>0.29</u>	<u>1.14</u>	4.56	9.39
CEOES (Chattopadhyay et al.)	✓	0.42	1.65	-	-
Grounding-DINO (Liu et al.)	✗	0.0066	0.05	1.29	3.27
CLIPSeg (Lüdtke and Ecker)	✗	0.0091	0.08	1.54	4.28
TFOC (Shi, Sun, and Zhang)	✗	0.0084	0.03	0.95	2.89
GrREC (Dai, Liu, and Cheung)	✓	-	-	<u>0.50</u>	<u>1.87</u>
OmniCount	✗	0.0023	0.009	0.70	2.00

Table 1: **Performance comparison in multi-label object counting using text prompts.** Results on the PASCAL VOC and OmniCount-191 datasets. Methods requiring training are marked (✓). The best results are in bold, while the best scores among the learning-based methods are underlined. Zero-shot models are marked blue.

images for training and 3,252 for testing. Class splits are available for few and zero-shot settings for specific applications, detailed in the supplementary.

Experiments

Datasets: We evaluate OmniCount on our proposed OmniCount-191 benchmark, tailored for multi-class scenarios. Following prior multi-label counting methods (Chattopadhyay et al. 2017), we also benchmark on the PASCAL VOC dataset (Everingham et al. 2009), which contains 9963 images across 20 classes, with 4952 for testing. For single-class counting, we use the test sets from FSC-147 (Ranjana et al. 2021) (1190 images, 29 categories) and CARPK (Hsieh, Lin, and Hsu 2017) (1014 test images). FSC-147 and CARPK include point and box annotations compatible with our model, while PASCAL VOC provides box annotations only.

corresponding to a 60%-40% split. This results in 26,978

Models	Training	Prompt	FSC-147				CARPK			
			MAE ↓	RMSE ↓	NAE ↓	SRE ↓	MAE ↓	RMSE ↓	NAE ↓	SRE ↓
CFOCNet+ (Yang et al. 2021)	✓	box	22.10	112.71	-	-	-	-	-	-
GMN (Lu, Xie, and Zisserman 2019)	✓	box	26.52	124.57	-	-7.48	9.90	-	-	-
BMNet+ (Shi et al. 2022)	✓	box	<u>14.62</u>	91.83	<u>0.25</u>	<u>2.74</u>	<u>5.76</u>	<u>7.83</u>	-	-
Vanilla SAM (Kirillov et al. 2023)	✗	N.A.	42.48	137.50	1.14	8.13	16.97	20.57	0.70	5.30
PSeCo (Huang 2024)	✓	N.A.	16.58	129.77	-	-	-	-	-	-
TFOC (Shi, Sun, and Zhang 2024)	✗	box	19.95	132.16	0.29	3.80	10.97	14.24	0.48	3.70
OmniCount	✗	box	18.63	112.98	0.14	2.99	9.92	12.15	0.23	2.11
TFOC (Shi, Sun, and Zhang 2024)	✗	point	20.10	132.83	0.30	3.87	11.01	14.34	0.51	3.89
OmniCount	✗	point	19.24	115.27	0.25	3.21	10.66	13.15	0.31	2.45
ZSOC (Xu et al. 2023a)	✓	text	22.09	115.17	0.34	3.74	-	-	-	-
TFOC (Shi, Sun, and Zhang 2024)	✗	text	24.79	137.15	0.37	4.52	-	-	-	-
GrREC (Dai, Liu, and Cheung 2024)	✓	text	<u>10.12</u>	<u>107.19</u>	-	-	-	-	-	-
OmniCount	✗	text	21.46	133.28	0.32	0.39	-	-	-	-

Table 2: **Results in single-label object counting** setting using text, point, and box prompts. The **bold** denotes the best among training-free methods, while the underlined font is the best among learning-based methods. Zero-shot models are marked blue.

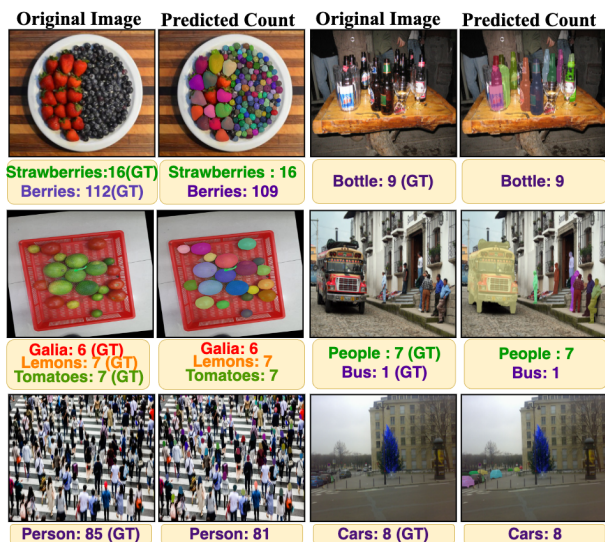


Figure 5: **Qualitative Results using OmniCount:** OmniCount-191 (left), PASCAL VOC (right).

Multi-label object counting

Competitors: We compare OmniCount with state-of-the-art methods by adapting the single-label counting model TFOC (Shi, Sun, and Zhang 2024) for multi-label counting, running it per class to obtain counts. We replicate and evaluate ILC (Cholakkal et al. 2019) and GrREC (Dai, Liu, and Cheung 2024), though replicating CEOES (Chattopadhyay et al. 2017) was constrained by its Lua implementation. Additionally, we employ open-vocabulary baselines: Grounding-DINO (Liu et al. 2023b) for object detection and CLIPSeg (Lüddecke and Ecker 2022) for semantic segmentation. Grounding-DINO counts objects by enumerating detected bounding boxes per category, while CLIPSeg uses a ViT encoder and spectral clustering to estimate category

counts via connected components.

Results: Table 1 compares OmniCount with existing multi-label counting methods, showing its strong performance as a training-free model. While GrREC achieves slightly better results, OmniCount remains competitive, leveraging geometric priors to handle occlusions and count multiple categories in a single pass, unlike traditional models requiring separate passes.

OmniCount also outperforms CLIPSeg and Grounding-DINO baselines, demonstrating SAM’s effectiveness for counting tasks. Qualitative results (Fig. 5, Fig. 6) highlight OmniCount’s robustness in handling objects of various sizes, from large items like seals to small entities like pulses and berries.

Single-label Counting

Competitors: We evaluate training-based methods, including CFOCNet+ (Yang et al. 2021), GMN (Lu, Xie, and Zisserman 2019), BMNet (Shi et al. 2022), ZSOC (Xu et al. 2023a), PSeCo (Huang 2024), and GrREC (Dai, Liu, and Cheung 2024), alongside training-free approaches like TFOC (Shi, Sun, and Zhang 2024) and a SAM-based baseline (Vanilla SAM (Kirillov et al. 2023)) with uniform point layout.

Results: As shown in Table 2, OmniCount outperforms both training-based and training-free models across text, box, and point prompts in all four key metrics, demonstrating its robustness and efficiency. These results highlight the importance of incorporating semantic and geometric priors, as Vanilla SAM performs poorly in comparison.

Further Analysis

Impact of Depth Refinement: OmniCount uses semantic (SP) and geometric priors (GP) to improve SAM’s segmentation for object counting. On the OmniCount-191 dataset (Table 3), removing GP increases m-RMSE and m-RMSE-nz errors by 56% and 8%, as SAM over-segments and over-counts without structural information. GP helps recover boundaries and reduce errors caused by occlusions.

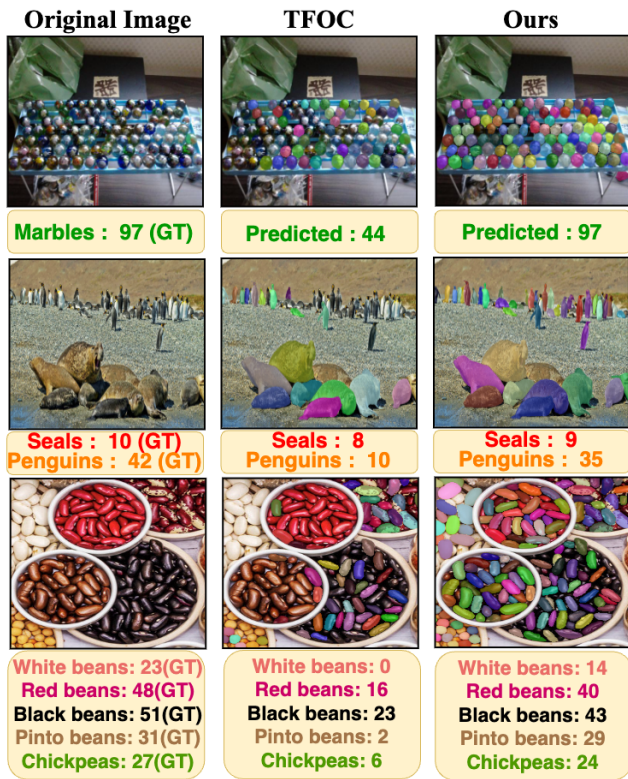


Figure 6: **Qualitative comparisons** with TFOC on the OmniCount-191 dataset.

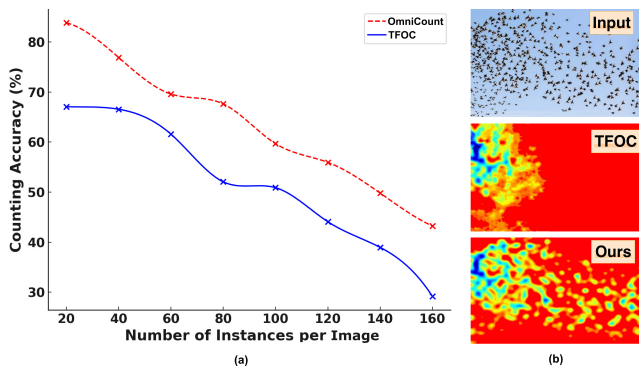


Figure 7: Performance comparison in dense scenes. (a) Counting accuracy vs number of instances per image (b) Counting heatmap in varying depth image

Importance of Reference Points: Reference point (RP) selection significantly reduces over-counting compared to SAM’s default “everything mode.” As shown in Table 3, RP selection lowers m-RMSE and m-RMSE-nz errors by 67% and 21%, ensuring more accurate segmentation and counting.

Counting in Dense Scenarios: Figure 7 shows OmniCount’s resilience in dense scenes compared to TFOC. OmniCount’s depth priors mitigate errors from occlusion and clustering, allowing it to detect objects at varying depths, in-

SP	GP	RP	m-RMSE ↓	m-RMSE-nz ↓
✓	✗	✗	7.02	5.89
✓	✗	✓	1.62	2.17
✓	✓	✗	2.12	2.54
✓	✓	✓	0.70	2.00

Table 3: Ablation of Semantic Prior (SP), Geometric Prior (GP) and Reference Point (RP) on OmniCount-191 dataset.

cluding distant ones, outperforming TFOC in complex scenarios.

Conclusion

We present OmniCount, an open-vocabulary, multi-label counting model that processes multiple categories in a single pass, leveraging semantic and geometric cues without training. Unlike category-specific models constrained by datasets, OmniCount integrates pre-trained foundation models for segmentation and depth estimation, tackling occlusions and enabling precise counting. To address the lack of a suitable dataset, we introduce OmniCount-191, comprising 30,230 images across 191 categories. Extensive testing on benchmarks and OmniCount-191 demonstrates OmniCount’s superior performance, efficiency, and scalability, highlighting its potential for real-world applications and establishing multi-label counting as a practical tool.

References

- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; and et al. 2020. Language models are few-shot learners. *NeurIPS*.
- Bui, K.-H. N.; Yi, H.; and Cho, J. 2020. A Vehicle Counts by Class Framework using Distinguished Regions Tracking at Multiple Intersections. In *CVPRW*.
- Chattopadhyay, P.; Vedantam, R.; Selvaraju, R. R.; Batra, D.; and Parikh, D. 2017. Counting Everyday Objects in Everyday Scenes. In *CVPR*.
- Cholakkal, H.; Sun, G.; Khan, F. S.; and Shao, L. 2019. Object counting and instance segmentation with image-level supervision. In *CVPR*.
- Cholakkal, H.; Sun, G.; Khan, S.; Khan, F. S.; Shao, L.; and Gool, L. V. 2022. Towards Partial Supervision for Generic Object Counting in Natural Scenes. *IEEE TPAMI*.
- Dai, S.; Liu, J.; and Cheung, N.-M. 2024. Referring Expression Counting. In *CVPR*.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2009. The PASCAL Visual Object Classes (VOC) Challenge. *IJCV*.
- Guo, M.; Yuan, L.; Yan, Z.; Chen, B.; Wang, Y.; and Ye, Q. 2024. Regressor-Segmenter Mutual Prompt Learning for Crowd Counting. In *CVPR*.
- Gupta, S. K.; Zhang, M.; Wu, C.-C.; Wolfe, J.; and Kreiman, G. 2021. Visual search asymmetry: Deep nets and humans share similar inherent biases. In *NeurIPS*.

- Han, T.; Bai, L.; Liu, L.; and Wanli, O. 2023. STEERER: Resolving Scale Variations for Counting and Localization via Selective Inheritance Learning. In *ICCV*.
- Hsieh, M.-R.; Lin, Y.-L.; and Hsu, W. H. 2017. Drone-based object counting by spatially regularized regional proposal network. In *ICCV*.
- Huang, Z. e. a. 2024. Point Segment and Count: A Generalized Framework for Object Counting. In *CVPR*.
- Huang, Z.-K.; Chen, W.-T.; Chiang, Y.-C.; Kuo, S.-Y.; and Yang, M.-H. 2023. Counting Crowds in Bad Weather. In *ICCV*.
- Ji, G.-P.; Fan, D.-P.; Xu, P.; Cheng, M.-M.; Zhou, B.; and Van Gool, L. 2023. SAM Struggles in Concealed Scenes – Empirical Study on “Segment Anything”. *SCIS*.
- Jiang, R.; Liu, L.; and Chen, C. 2023. CLIP-Count: Towards Text-Guided Zero-Shot Object Counting. In *ACM MM*.
- Ke, B.; Obukhov, A.; Huang, S.; Metzger, N.; Daudt, R. C.; and Schindler, K. 2024. Marigold: Repurposing Diffusion-Based Image Generators for Monocular Depth Estimation. In *CVPR*.
- Khan, A. 2016. Deep Convolutional Neural Networks for Human Embryonic Cell Counting. In *ECCV*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *ICCV*.
- Lempitsky, V.; and Zisserman, A. 2010. Learning to count objects in images. In *NeurIPS*.
- Li, C.; Hu, X.; Abousamra, S.; and Chen, C. 2023. Calibrating Uncertainty for Semi-Supervised Crowd Counting. In *ICCV*.
- Li, Y.; Zhang, X.; and Chen, D. 2018. CSRNet: Dilated convolutional neural networks for understanding the highly congested scenes. In *CVPR*.
- Liang, D.; Xie, J.; Zou, Z.; Ye, X.; Xu, W.; and Bai, X. 2023. CrowdCLIP: Unsupervised Crowd Counting via Vision-Language Model. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common objects in context. In *ECCV*.
- Liu, C.; Lu, H.; Cao, Z.; and Liu, T. 2023a. Point-Query Quadtree for Crowd Counting, Localization, and More. In *ICCV*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023b. Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection. In *ECCV*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*.
- Lu, E.; Xie, W.; and Zisserman, A. 2019. Class-agnostic counting. In *ACCV*.
- Lüddecke, T.; and Ecker, A. 2022. Image segmentation using text and image prompts. In *CVPR*.
- Peng, Z.; and Chan, S.-H. G. 2024. Single Domain Generalization for Crowd Counting. In *CVPR*.
- Phillion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Rahmehoonfar, M.; and Sheppard, C. 2017. Deep Count: Fruit Counting Based on Deep Simulated Learning. In *Sensors*.
- Ranjan, V.; and Nguyen, M. H. 2022. Exemplar Free Class Agnostic Counting. In *ACCV*.
- Ranjan, V.; Sharma, U.; Nguyen, T.; and Hoai, M. 2021. Learning To Count Everything. In *CVPR*.
- Shi, M.; Lu, H.; Feng, C.; Liu, C.; and Cao, Z. 2022. Represent, compare, and learn: A similarity-aware framework for class-agnostic counting. In *CVPR*.
- Shi, Z.; Mettes, P.; and Snoek, C. G. 2024. Focus for Free in Density-Based Counting. *IJCV*.
- Shi, Z.; Sun, Y.; and Zhang, M. 2024. Training-free Object Counting with Prompts. In *WACV*.
- Song, Q.; Wang, C.; Jiang, Z.; Wang, Y.; Tai, Y.; Wang, C.; Li, J.; Huang, F.; and Wu, Y. 2021. Rethinking Counting and Localization in Crowds: A Purely Point-Based Framework. In *ICCV*.
- Xie, E.; Yu, Z.; Zhou, D.; Phillion, J.; Anandkumar, A.; Fidler, S.; Luo, P.; and Alvarez, J. M. 2022. M2BEV: Multi-Camera Joint 3D Detection and Segmentation with Unified Bird’s-Eye View Representation. *arXiv preprint arXiv:2204.05088*.
- Xu, J.; Le, H.; Nguyen, V.; Ranjan, V.; and Samaras, D. 2023a. Zero-Shot Object Counting. In *CVPR*.
- Xu, J.; Le, H.; and Samaras, D. 2023. Zero-Shot Object Counting with Language-Vision Models. *arXiv preprint arXiv:2309.13097*.
- Xu, M.; Zhang, Z.; Wei, F.; Hu, H.; and Bai, X. 2023b. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*.
- Xu, Y.; Zhong, Z.; Lian, D.; Li, J.; Li, Z.; Xu, X.; and Gao, S. 2021. Crowd counting with partial annotations in an image. In *ICCV*.
- Yang, S.-D.; Su, H.-T.; Hsu, W. H.; and Chen, W.-C. 2021. Class-agnostic few-shot object counting. In *WACV*.
- You, Z.; Yang, K.; Luo, W.; Lu, X.; Cui, L.; and Le, X. 2022. Few-shot Object Counting with Similarity-Aware Feature Enhancement. In *WACV*.
- Zhang, F.; Zhu, X.; Dai, H.; Ye, M.; and Zhu, C. 2020. Distribution-aware coordinate representation for human pose estimation. In *CVPR*.
- Zhang, L.; Shi, Z.; Cheng, M.-M.; Liu, Y.; Bian, J.-W.; Zhou, J. T.; Zheng, G.; and Zeng, Z. 2019. Nonlinear regression via deep negative correlation learning. *IEEE TPAMI*.
- Zhang, M.; Armendariz, M.; Xiao, W.; Rose, O.; Bendtz, K.; Livingstone, M.; Ponce, C.; and Kreiman, G. 2021. Look Twice: A Generalist Computational Model Predicts Return Fixations across Tasks and Species. *PLOS Comp. Bio.*

Zhang, Y.; Zhou, D.; Chen, S.; Gao, S.; and Ma, Y. 2016. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*.