

VLCounter: Text-Aware Visual Representation for Zero-Shot Object Counting

Seunggu Kang, WonJun Moon, Euiyeon Kim, Jae-Pil Heo*

Sungkyunkwan University
 {seunggu35, wjun0830, keywi9811, jaepilheo}@g.skku.edu

Abstract

Zero-Shot Object Counting (ZSOC) aims to count referred instances of arbitrary classes in a query image without human-annotated exemplars. To deal with ZSOC, preceding studies proposed a **two-stage** pipeline: discovering exemplars and counting. However, there remains a challenge of vulnerability to error propagation of the sequentially designed two-stage process. In this work, we propose an **one-stage** baseline, Visual-Language Baseline (VLBase), exploring the implicit association of the semantic-patch embeddings of CLIP. Subsequently, we extend the VLBase to Visual-language Counter (VLCounter) by incorporating three modules devised to tailor VLBase for object counting. First, we introduce Semantic-conditioned Prompt Tuning (SPT) within the image encoder to acquire target-highlighted representations. Second, Learnable Affine Transformation (LAT) is employed to translate the semantic-patch similarity map to be appropriate for the counting task. Lastly, we transfer the layer-wisely encoded features to the decoder through Segment-aware Skip Connection (SaSC) to keep the generalization capability for unseen classes. Through extensive experiments on FSC147, CARPK, and PUCPR+, we demonstrate the benefits of our end-to-end framework, VLCounter. Code is available at <https://github.com/seunggu0305/VLCounter>

1 Introduction

Object counting, which was initially studied for specific targets, e.g., crowds (Zhang et al. 2016), cells (Xie, Noble, and Zisserman 2018), animals (Arteta, Lempitsky, and Zisserman 2016), and cars (Mundhenk et al. 2016), has shown that the number of objects can be counted even within a dense image. Furthermore, recent works have shown significant advances in inferring the number of arbitrary objects with several human-annotated exemplar patches. However, such a strong prerequisite that every cumbersome guidance must be equipped is undoubtedly the main challenge to overcome to grant applicability to object counting methods. In this context, Zero-Shot Object Counting (ZSOC) was proposed to mitigate the need for human labor.

Current ZSOC approaches commonly adopt a two-stage pipeline as illustrated in Fig. 1. These works primarily focus




Figure 1: Comparison between two-stage pipeline and one-stage pipeline (ours). The two-stage pipeline requires training the exemplar discoverer (orange) before the counter (blue), along with the need for an extra training dataset to optimize the discoverer. In contrast, our one-stage pipeline is designed to be simpler and does not necessitate any additional data or training stage.

on identifying exemplar patches within the image and subsequently adopt the counting framework from the literature of few-shot object counting (Shi et al. 2022; Ranjan et al. 2021). To identify the exemplar patches, RepRPN (Ranjan and Nguyen 2022) considered the repetition score to detect object patches that frequently appear within the image. Requirement for counting the desired classes over frequent ones, ZSC (Xu et al. 2023) utilized the class names to enable the class specification. They localize exemplars by identifying the k -nearest neighbors of the class name embeddings among randomly cropped patches. Despite their progress, the potential localization error propagation in the two-stage training pipeline (Nag et al. 2022) is an untapped problem in ZSOC frameworks. Indeed, they utilized additional datasets to train decent exemplar discovery networks.

This paper pursues a simplified zero-shot object counting framework. We instantiate an end-to-end ZSOC counter namely Visual-Language Baseline (VLBase), which consists of a CLIP (Radford et al. 2021) encoder and counting decoder. By leveraging the embedding space of CLIP which enables the implicit association of the semantic and patch embeddings to localize the target object (Zhou, Loy, and Dai 2022; Li et al. 2023), VLBase eliminates the need for an exemplar discovery process.

Additionally, we introduce VLCounter which is built upon VLBase by incorporating three modules devised to tailor VLBase for object counting. First, we propose Semantic-

*Corresponding author

conditioned Prompt Tuning (SPT) which extends the visual prompt tuning (VPT) to efficiently finetune CLIP for the counting task. Instead of utilizing naïve learnable prompts, SPT employs conditioning via semantic embedding to generate patch embeddings that emphasize the region of interest. Subsequently, based on our observation that the similarity maps between patch embeddings obtained using SPT and semantic embeddings already provide a decent approximation of object locations, we employ simple Learnable Affine Transformation (LAT) to adjust only the finer details. Finally, to equip the decoder with the generalization capability and provide rich clues, we exploit intermediate features across different encoding layers of CLIP through Segment-aware Skip Connections (SaSC). With all components combined, our simple end-to-end one-stage framework records new state-of-the-art results on the FSC147 (Ranjan et al. 2021) dataset validating its superiority over the previous ZSOC methods. Moreover, we provide additional evidence of cross-dataset generalization by evaluating performance on the car counting dataset CARPK and PUCPR+ (Hsieh, Lin, and Hsu 2017).

Our contributions are three-fold:

- We instantiate an end-to-end baseline for ZSOC, VL-Base, by exploiting the vision-language association capability of CLIP.
- We propose a VLCounter consisting of SPT, LAT, and SaSC that allows the model to utilize the generalization capability of CLIP in a counting-specific manner.
- Our experiments on FSC147 and cross-dataset validation verify the effectiveness of VLCounter.

2 Related Works

2.1 Object Counting

Class-specific Object Counting focuses on quantifying specific class samples, e.g., crowds (Abousamra et al. 2021; Idrees et al. 2018; Lian et al. 2019; Sindagi, Yasarla, and Patel 2019), cars (Mundhenk et al. 2016; Hsieh, Lin, and Hsu 2017), animals (Arteta, Lempitsky, and Zisserman 2016), and cells (Xie, Noble, and Zisserman 2018). Most works fall into two main categories each employing detection (Chattopadhyay et al. 2017; Hsieh, Lin, and Hsu 2017; Laradji et al. 2018) or regression (Chan, Liang, and Vasconcelos 2008; Cholakkal et al. 2019; Liu, Salzmann, and Fua 2019; Wang et al. 2020) mechanism to measure the number of instances. The former predicts the bounding box for every instance using an object detector, whereas the latter predicts the density distribution of the image instead, thereby being recognized as a more robust stream against partially occluded objects (Lu, Xie, and Zisserman 2019).

Few-shot Object Counting To overcome the lack of generality of being constrained to a specific class, Generic Matching Network (GMN) (Lu, Xie, and Zisserman 2019) first formalized class-agnostic object counting to count the desired objects provided by the human-annotated exemplar patches. They introduced a two-stream architecture to encode each image and exemplar to handle the difference in their resolution. Following them, CFOcNet (Yang et al.

2021) and BMNet (Shi et al. 2022) also adopted and enhanced the two-stream approach by adding a layer-wise matching procedure and bilinear similarity metric. Other works adhere to single-stream architecture. To be specific, FamNet (Ranjan et al. 2021) and RCAC (Gong et al. 2022) use ROI pooling after feature extraction to obtain exemplar prototypes. However, the aforementioned studies suffer from the limitation that every inference requires human-annotated exemplars.

Zero-shot Object Counting has been proposed by RepRPN (Ranjan and Nguyen 2022) to discard the duty of annotating target exemplars for counting. To be specific, they trained the region proposal network (RPN) to capture the patches containing the most frequently appeared objects to replace human-annotated exemplars. Then, to further grant more applicability to exemplar-free object counter, ZSC (Xu et al. 2023) presented a method that takes guidance from semantic information. By matching semantic information to randomly generated patches, they sampled the most semantically relevant patches to obtain target exemplars. Our work shares the goal with ZSC in that we aim to train the counter that can count user-specified classes with only class names. Yet, as the mentioned methods adopt a two-stage pipeline that is prone to error propagation, we focus on mitigating such issues by proposing an end-to-end framework that localizes and counts at once.

2.2 Prompt Tuning

Prompt tuning is a popular strategy to adapt pre-trained large models for downstream tasks due to its efficiency compared to conventional fine-tuning methods (Wang et al. 2022; Gu et al. 2021; Brown et al. 2020; Jia et al. 2022). Whereas fine-tuning updates all parameters, prompt tuning freezes the pre-trained large models and introduces only a small set of learnable prompts to optimize (Li and Liang 2021; Jia et al. 2022). Following these works, we utilize prompt tuning to efficiently exploit the quality of the visual-language understanding capability of pre-trained CLIP. Yet, our work differs in using semantic information from the semantic embeddings to condition the prompts in the visual encoder to concentrate more on specification-relevant information.

3 Preliminaries

3.1 Problem Formulation: ZSOC

ZSOC aims to predict the density map $D \in \mathbb{R}^{H \times W \times 1}$ for image $I \in \mathbb{R}^{H \times W \times 3}$ that belongs to unseen classes C^u ($f : (I, C^u) \mapsto D$) without any visual exemplar clues. In the training stage, the model is trained with $\mathcal{D}_{\text{train}} = \{(I_i, C_i^s, D_i)\}_{i=1}^N$ where C_i^s denotes the seen class names during training. Then in the testing stage, the model is to yield a density map for $\mathcal{D}_{\text{test}} = \{(I_i, C_i^u, D_i)\}_{i=N+1}^M$, where $C^s \cap C^u = \emptyset$.

3.2 Overview of CLIP

This section introduces the underlying motivation behind our proposed method. CLIP is composed of two encoders: an image encoder $\phi_V(\cdot)$ and a text encoder $\phi_T(\cdot)$. The text

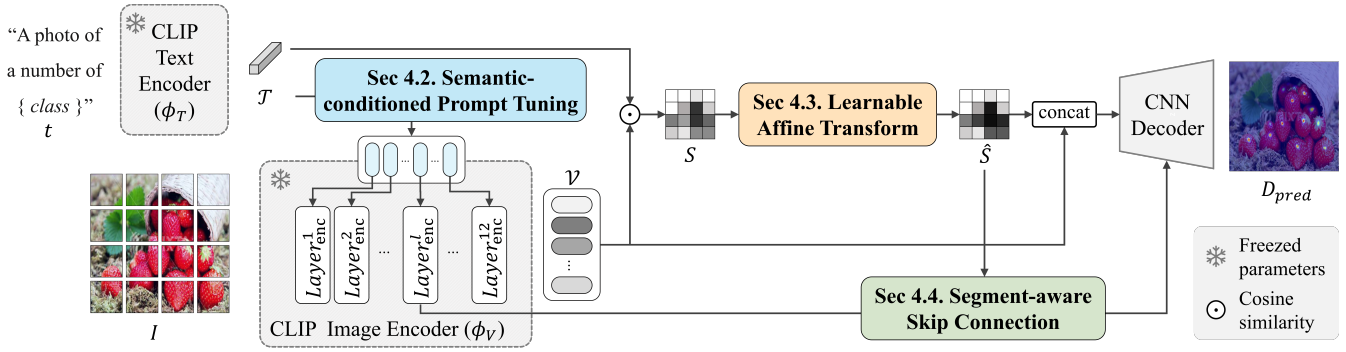


Figure 2: Overview of VLBase and VLCounter: each without and with colored components. The end-to-end baseline, VLBase, employs CLIP encoders to extract both image and text embeddings. Then, the decoder processes the image-text similarity map along with visual embeddings to count the number of specified objects. With three colored modules, VLCounter leverages the generalization capability of VLBase to be tailored for object counting.

encoder takes prompted class name t e.g., *A photo of [kiwi]* and produces a semantic embedding $\mathcal{T} \in \mathbb{R}^{1 \times d}$ where d represents an embedding dimension. The image encoder takes a learnable class token $[cls]$ along with embedded patch sequences V as inputs and encodes global and local semantics in the class token $[cls]$ and patch tokens \mathcal{V} respectively. Note that $V = [v_1, v_2, \dots, v_N] \in \mathbb{R}^{N \times (P^2 \cdot d)}$ where N is the number of embedded patches, and $(P \cdot P)$ is the resolution of each patch. Formally, this process can be expressed as follows:

$$\mathcal{T} = \phi_{\mathcal{T}}(t); \quad [[cls], \mathcal{V}] = \phi_{\mathcal{V}}([cls], V). \quad (1)$$

These encoders are trained collaboratively to map \mathcal{T} and $[cls]$ into a shared representation space.

Recently, there exist studies suggesting the implicit localization capability of CLIP, where each patch embedding preserves local image semantics (Zhou, Loy, and Dai 2022; Li et al. 2023). This property, coupled with the powerful image-text joint embedding space of CLIP, has provided a clear motivation for utilizing CLIP as a robust tool for zero-shot segmentation (localization). (Li et al. 2022; Rao et al. 2022; Lüddecke and Ecker 2022). Taking similar inspiration yet focused on object counting, we aim to leverage the implicit localization capability of CLIP to achieve precise and efficient object counting in an end-to-end manner.

4 Visual-Language Counter: End-to-End Framework for Zero-Shot Object Counting

This section presents Visual-Language Counter (VLCounter), an efficient end-to-end ZSOC framework. We first establish a baseline model referred to as Vision-Language Baseline (VLBase), which exploits the visual-language localization capacity of CLIP in Sec. 4.1. Then, we bring three improvements on top of VLBase to introduce VLCounter. Specifically, we emphasize the regions of interests (Sec. 4.2), learn task-specific visual-language similarity (Sec. 4.3), and exploit semantic-relevant information across the multi-level representations (Sec. 4.4). The overall architectures of the two models are illustrated in Fig. 2.

4.1 Visual-Language Baseline

VLBase is a standalone baseline, eliminating the need for few-shot counting techniques that previous ZSOC approaches heavily rely on. Given input query image I and class name C , VLBase obtains patch embedding \mathcal{V} and semantic embedding \mathcal{T} using CLIP encoders $\phi_{\mathcal{V}}(\cdot)$ and $\phi_{\mathcal{T}}(\cdot)$, respectively. By calculating the cosine similarity between \mathcal{T} and \mathcal{V} , the similarity map $S \in \mathbb{R}^{H \times W}$ is yielded:

$$S_{ij}(\mathcal{V}, \mathcal{T}) = \frac{v_{ij} \mathcal{T}^{\top}}{\|v_{ij}\| \|\mathcal{T}\|}, \quad (2)$$

where S_{ij} corresponds to the value at position (i, j) in matrix S and v_{ij} represents the embedding at position (i, j) of 2D-reshaped \mathcal{V} .

As mentioned in prior studies (Zhou, Loy, and Dai 2022; Li et al. 2023), we observed that the similarity map between CLIP-encoded semantic and patch embeddings provides an adequate indication of the degree of semantic similarity between the patch and semantic embedding. We find that this similarity map is a decent clue for a decoder to localize the target objects. Consequently, the CNN-based counting decoder predicts the density map D_{pred} by utilizing features of \mathcal{V} and S :

$$D_{\text{pred}} = \phi_{\text{decoder}}([\mathcal{V}, S]), \quad (3)$$

where $[\cdot, \cdot]$ denotes channel-wise concatenation. Finally, the object count prediction is derived by summing all values in D_{pred} .

Counting Loss For training, we adopt a conventional MSE loss:

$$\mathcal{L}_{\text{count}} = \|D_{\text{pred}} - D_{\text{gt}}\|_2^2, \quad (4)$$

where D_{gt} denotes the ground truth density map.

4.2 Semantic-conditioned Prompt Tuning (SPT)

To grant task-specificity to the CLIP image encoder without sacrificing its generalization capability, a straightforward approach is to employ visual prompt tuning (VPT) (Jia et al. 2022). However, the naïve VPT, which simply concatenates

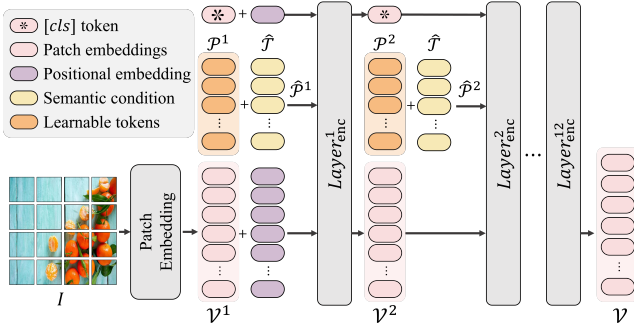


Figure 3: Illustration for Semantic-conditioned Prompt Tuning (SPT). In addition to learnable visual prompts (orange) in the image encoder, text features (yellow) are integrated to specify the desired semantics.

a few learnable tokens to the input sequence of each encoding layer does not take the semantic information into account. Hence, we introduce Semantic-conditioned Prompt Tuning (SPT), which utilizes semantic information along with the learnable tokens to assist the image encoder in extracting target-semantic-highlighted visual features.

Specifically, as illustrated in Fig. 3, SPT has new learnable tokens for each encoding layer. Learnable tokens for l^{th} layer are defined as $\mathcal{P}^l = [p_1^l, p_2^l, \dots, p_M^l]$ where the number of learnable tokens is denoted as M . These tokens are then, supplemented with the linearly projected semantic embedding $\hat{\mathcal{T}}$ to generate semantic-conditioned prompts $\hat{\mathcal{P}}$. The semantic-conditioned prompts for the l^{th} layer are defined as follows:

$$\hat{\mathcal{P}}^l = [p_1^l + \hat{\mathcal{T}}, p_2^l + \hat{\mathcal{T}}, \dots, p_M^l + \hat{\mathcal{T}}], \quad (5)$$

where $\hat{\mathcal{T}} = \phi_c(\mathcal{T})$ and ϕ_c denotes the parameters of the projection layer. Consequently, with the conditioned prompts $\hat{\mathcal{P}}$, the patch embedding process in l^{th} layer of the image encoder can be expressed as:

$$[[cls], _, \mathcal{V}^{l+1}] = Layer_{enc}^l([[cls], \hat{\mathcal{P}}^l, \mathcal{V}^l]), \quad (6)$$

where initial input $\mathcal{V}^1 = [v_1^1, v_2^1, \dots, v_N^1]$ is a sequence of embedded patches through the patch embedding layer prior to the encoder. Be aware that we follow VPT (Jia et al. 2022) to discard output tokens of $\hat{\mathcal{P}}$ (represented as $_$) and do not propagate to the subsequent layer.

4.3 Learnable Affine Transformation (LAT)

Through the adoption of the SPT, we obtain visual representations in which the corresponding regions of the target class are highlighted. Nevertheless, due to the nature of object counting, discovering the central points of the objects rather than encompassing the entire object area, a discrepancy might arise between the information contained in the similarity map S and the loss that needs to be backpropagated during training.

In light of this, we propose learnable affine transformation matrix (LAT) to facilitate the conversion of similarity

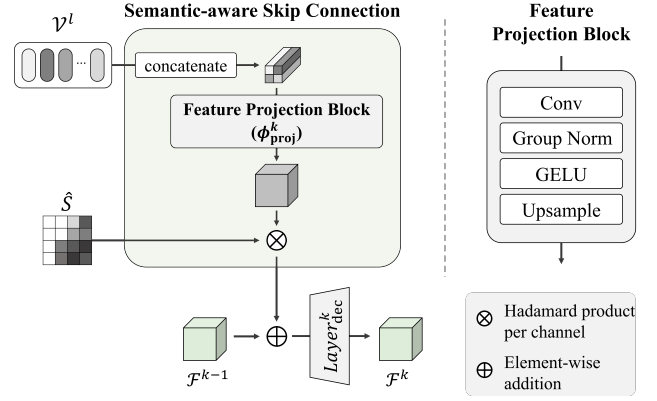


Figure 4: The flow of Semantic-aware Skip Connection (SaSC) and architecture of feature projection block. Intermediate visual features are projected and filtered with an object-aware counting map \hat{S} to produce object-relevant encoder features. Consequently, these are integrated into its counterpart in the decoder.

map S to counting map \hat{S} and establish a more task-specific visual-semantic linkage centered around individual objects as follows:

$$\hat{S} = W \otimes S + B, \quad (7)$$

where $W, B \in \mathbb{R}^{H \times W}$ are learnable matrices for affine transformation and \otimes indicates element-wise multiplication. In addition, we directly optimize the counting map \hat{S} with the rank-aware contrastive loss to learn the proper degree of activation for object counting. Details of rank-aware contrastive loss are elaborated in Sec. 4.5. With LAT, the input to the decoder $[\mathcal{V}, S]$ in Eq. 3 of VLBase is replaced by $[\mathcal{V}, \hat{S}]$.

4.4 Segment-aware Skip Connection (SaSC)

For ZSOC, where the model encounters unseen classes during inference, it is important to train a decoder that is tailored for object counting while maintaining a generalization ability. Sharing the motivation with VLBase in Sec. 4.1 that CLIP features inherently preserve local semantics, we adopt skip connections that incorporate intermediate features of the encoder to its counterpart in the decoder.

As shown in Fig. 4, the l^{th} encoder patch features are spatially concatenated and projected to yield decoder-assistive representations. Then, we multiply the affine transformed similarity map \hat{S} to emphasize the object-relevant patches. Finally, these patch features are added to the corresponding k^{th} layer features of the decoder. Formally, the k^{th} decoding layer with SaSC, receiving l^{th} encoder features, operates as follows:

$$\mathcal{F}^k = Layer_{dec}^k(\mathcal{F}^{k-1} + \phi_{proj}^k(\mathcal{V}^l) \otimes \hat{S}), \quad (8)$$

where $\phi_{proj}^k(\cdot)$, \mathcal{F}^k , and \otimes stand for the parameter of feature projection block, the output of the k -th decoding layer, and Hadamard products per channel, respectively.

Methods	Stage	Class Name	Train Dataset	Val set		Test set		Inference speed (s) ↓
				MAE ↓	RMSE ↓	MAE ↓	RMSE ↓	
<i>few-shot</i>								
GMN (Lu, Xie, and Zisserman 2019)	1	×	FSC147	29.66	89.81	26.52	124.57	-
FamNet (Ranjan et al. 2021)	1	×	FSC147	24.32	70.94	22.56	101.54	0.82
BMNet (Shi et al. 2022)	1	×	FSC147	19.06	67.95	16.71	103.31	0.86
BMNet+ (Shi et al. 2022)	1	×	FSC147	15.74	58.53	14.62	91.83	1.59
<i>zero-shot</i>								
RepRPN-Counter (Ranjan and Nguyen 2022)	2	×	FSC147 + MS COCO	31.69	100.31	28.32	128.76	-
ZSC (Xu et al. 2023)	2	✓	FSC147 + MS COCO	26.93	88.63	22.09	115.17	0.86+ α
VLBase (Ours)	1	✓	FSC147	31.82	98.89	32.20	130.51	0.81
VLCounter (Ours)	1	✓	FSC147	18.06	65.13	17.05	106.16	0.82

Table 1: Quantitative comparison to state-of-the-art approaches on the FSC147 dataset. α in the rightmost column indicates an additional cost necessary for the exemplar discovery process in the context of the two-stage pipeline.

4.5 Training Objectives

In addition to the counting loss described in Eq. 4, VLCounter additionally employs rank-aware contrastive loss (Hoffmann et al. 2022; Moon et al. 2023). Whereas the $\mathcal{L}_{\text{count}}$ trains the whole model to learn the counting objective, our focus in SPT and LAT is learning to yield the counting-tailored similarity map in the encoder. In this regard, we adopt rank-aware contrastive loss in the counting map \hat{S} to assign higher activations on the patches that are nearby the object centers. To design the hierarchical guidance for a rank-aware contrastive loss, we first normalize the ground truth density map D_{gt} to be mapped between 0 and 1. Then, we iterate the batch for K times with different thresholds to prepare positive and negative sets; patches are gathered as positive if the corresponding patch in D_{gt} has a higher value than the threshold, and if not, as negative. Formally, the rank contrastive loss with the positive set \hat{S}_r^{pos} and the negative set \hat{S}_r^{neg} is formulated as follows:

$$\mathcal{L}_{\text{rank}} = - \sum_{k=1}^K \log \frac{\sum_{\hat{S}_i \in \hat{S}_r^{\text{pos}}} \exp(\hat{S}_i / \tau)}{\sum_{\hat{S}_j \in (\hat{S}_r^{\text{pos}} \cup \hat{S}_r^{\text{neg}})} \exp(\hat{S}_j / \tau)}, \quad (9)$$

where τ is a temperature scaling parameter.

With the objectives in Eq. 4 and Eq. 9 combined, VLCounter’s final objective is as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{count}} + \lambda \cdot \mathcal{L}_{\text{rank}}, \quad (10)$$

where λ is a hyperparameter to balance between the losses.

5 Experiments

In this section, we provide a comprehensive explanation of experimental details. First, we delve into the implementation details, datasets, and evaluation metrics in Sec. 5.1, followed by a comparison of our model with existing state-of-the-art methods in Sec. 5.2. Then, we conduct an in-depth exploration of each component in Sec. 5.3 and qualitative analysis in Sec. 5.4.

Methods	CARPK		PUCPR+	
	MAE	RMSE	MAE	RMSE
<i>few-shot</i>				
FamNet	28.84	44.47	87.54	117.68
BMNet	14.61	24.60	103.18	112.42
BMNet+	10.44	13.77	62.42	81.74
<i>zero-shot</i>				
VLBase	20.47	24.33	90.82	104.01
VLCounter	6.46	8.68	48.94	69.08

Table 2: Cross-dataset validation performance on the CARPK and PUCPR+ dataset.

5.1 Experimental Details

Implementation Details. For all experiments, we employed CLIP ViT-B/16 as our encoders followed by a decoder consisting of 4 repeated units. Each of these units consists of one feature projection block in Fig. 4 and one additional convolutional layer. Regarding the image input, each image is resized to 384×384 , and augmentations such as Gaussian noise, Gaussian blur, affine transformation, horizontal flip, and color jittering were applied. We trained the model using AdamW (Loshchilov and Hutter 2017) optimizer with a learning rate of $1e^{-4}$ and weight decay of $1e^{-2}$ for 200 epochs with a batch size of 16 on a single NVIDIA RTX A6000. For temperature scaling and loss-balancing hyperparameter λ and τ , we used $1e^{-6}$ and 1.

Datasets. To explore the counting capability of models, we use FSC147 (Ranjan et al. 2021), the first large-scale dataset for class-agnostic counting. It includes 6135 images from 147 categories mainly composed of foods, animals, kitchen utensils, and vehicles. We also utilize CARPK and PUCPR+ (Hsieh, Lin, and Hsu 2017) datasets. These datasets exhibit different properties from the images in FSC147, so we use them for cross-dataset validation which is to test the model’s generality. To be specific, CARPK consists of 1,448 parking lot images with nearly 90,000 cars taken in a drone view at 40 meters height on average. On

No.	SPT	LAT	SaSC	Val set		Test set	
				MAE	RMSE	MAE	RMSE
M1	×	×	×	31.82	98.89	32.20	130.51
M2	✓	×	×	20.61	75.36	17.58	112.89
M3	×	✓	×	29.97	96.59	28.26	127.44
M4	×	×	✓	24.88	81.28	24.16	113.01
M5	✓	✓	✓	18.06	65.13	17.05	106.16

Table 3: Ablation study on each component of VLCounter.

the other hand, PUCPR+ contains nearly 16,456 cars in total which have 10th-floor-view images.

5.2 Comparison with State-of-the-art Methods

We compare VLBase and VLCounter against previous class-agnostic counting methods in Tab. 1. Despite its simple design, the performances of VLBase are comparable to the two-stage methods that even utilize additional training data. On the other hand, VLcounter clearly surpasses other ZSOC baselines. Particularly, when compared to ZSC, VLCounter achieves a relative improvement of 32.94% and 22.81% in terms of validation MAE and test MAE, respectively. Moreover, we remark on the comparable results to the state-of-the-art few-shot counting method: BMNet. This is an especially notable milestone for ZSOC since few-shot methods are generally seen as the upper bound of two-stage ZSOC methods; the counting framework in two-stage works is usually adopted from few-shot methods.

On the rightmost columns, we provide the inference speed per image. As our one-stage approaches (VLBase and VLCounter) only require the time to count the objects, it is shown that their inference speeds are much faster than a two-stage method (ZSC) which needs extra time to discover exemplars (denoted as α since the implementation is not fully publicized). In addition to the inference time, VLBase and VLCounter have much fewer parameters to learn, having their strength in shorter training time (Training time for VLCounter is approximately $2\times$ faster than BMNet+).

Following previous class-agnostic counting methods (Ranjan et al. 2021; Shi et al. 2022), we verify the generalization capability of VLBase and VLCounter by conducting a cross-dataset evaluation on CARPK and PUCPR+ datasets in Tab. 2, and VLBase and VLCounter demonstrate their benefits in generalization. Whereas the performance gaps between few-shot methods and VLBase is reduced, we observe the superiority of VLCounter to other methods by boosting MAE up to 38.12% and 27.54% in CARPK and PUCPR+ datasets compared to BMNet+. In particular, we emphasize the single-digit results of VLCounter in terms of both MAE and RMSE are derived without any fine-tuning (The average number of cars in each image of CARPK is 62). We attribute such success in cross-dataset validation to adapting the generality of CLIP to counting-specific and incorporating multi-level features to provide rich semantics into the prediction, each approximately taking 54% and 46% portions in the increase in CARPK MAE.

Condition		Val set		Test set	
		MAE	RMSE	MAE	RMSE
VLCounter		18.06	65.13	17.05	106.16
SPT	w/o \mathcal{T}'	19.07	65.72	17.19	107.54
SaSC	w/o \hat{S}	20.28	65.54	19.38	105.69

Table 4: Analysis of semantic-conditioning techniques in SPT and SaSC.

Text prompts	Val set		Test set	
	MAE	RMSE	MAE	RMSE
Singular	20.08	67.92	19.18	105.04
Plural	18.06	65.13	17.05	106.16

Table 5: Analysis of pluralized context to prompt the class names.

5.3 Ablation Studies on VLCounter

Component Analysis. To validate the effectiveness of individual components, we conducted an ablation study as presented in Tab. 3. Starting with VLBase (M1), we add SPT, LAT, and SaSC in M2, M3, and M4, respectively. Among the individual components, the effectiveness of SPT demonstrated in M2 is the most pronounced. This significant improvement demonstrates the importance of fine-tuning incorporated with the semantic condition. LAT in M3 is another important component. While it can be seen as not incurring a dramatic increase in performance, the counting map \hat{S} derived from LAT is also an essential element in SaSC. Lastly, M4 shows that SaSC not only boosts generalization capability but also task-specific predictions. This is because layer-wise intermediate representations in CLIP encoder are also semantically meaningful (Li et al. 2023) and SaSC aggregates them to aid counting prediction.

Effect of conditioning semantic information. We further conduct ablation studies on semantic conditioning. In Tab. 4, we compare conventional VPT with SPT and test the semantic conditioning in SaSC. Along with the benefits of VPT of granting task-specificity, utilizing semantic conditions in VPT allows the prompts to be more semantically specific. In addition, using semantic conditions in filtering the knowledge that is passed to the decoder with residual paths clearly benefits SaSC. We think that the semantic conditioning with the counting map \hat{S} suppresses the object-irrelevant information, thereby contributing to the improvements.

Effect of plural text prompts. We followed CLIP (Radford et al. 2021) to use different context prompts to encode the semantic embeddings. Yet, since the counting task mainly assumes the existence of multiple instances in every image, we modified text prompts to be in plural form. In Tab. 5, we compare the results between using singular and plural forms of text prompts, and text prompts in plural form have the advantage in the counting task.

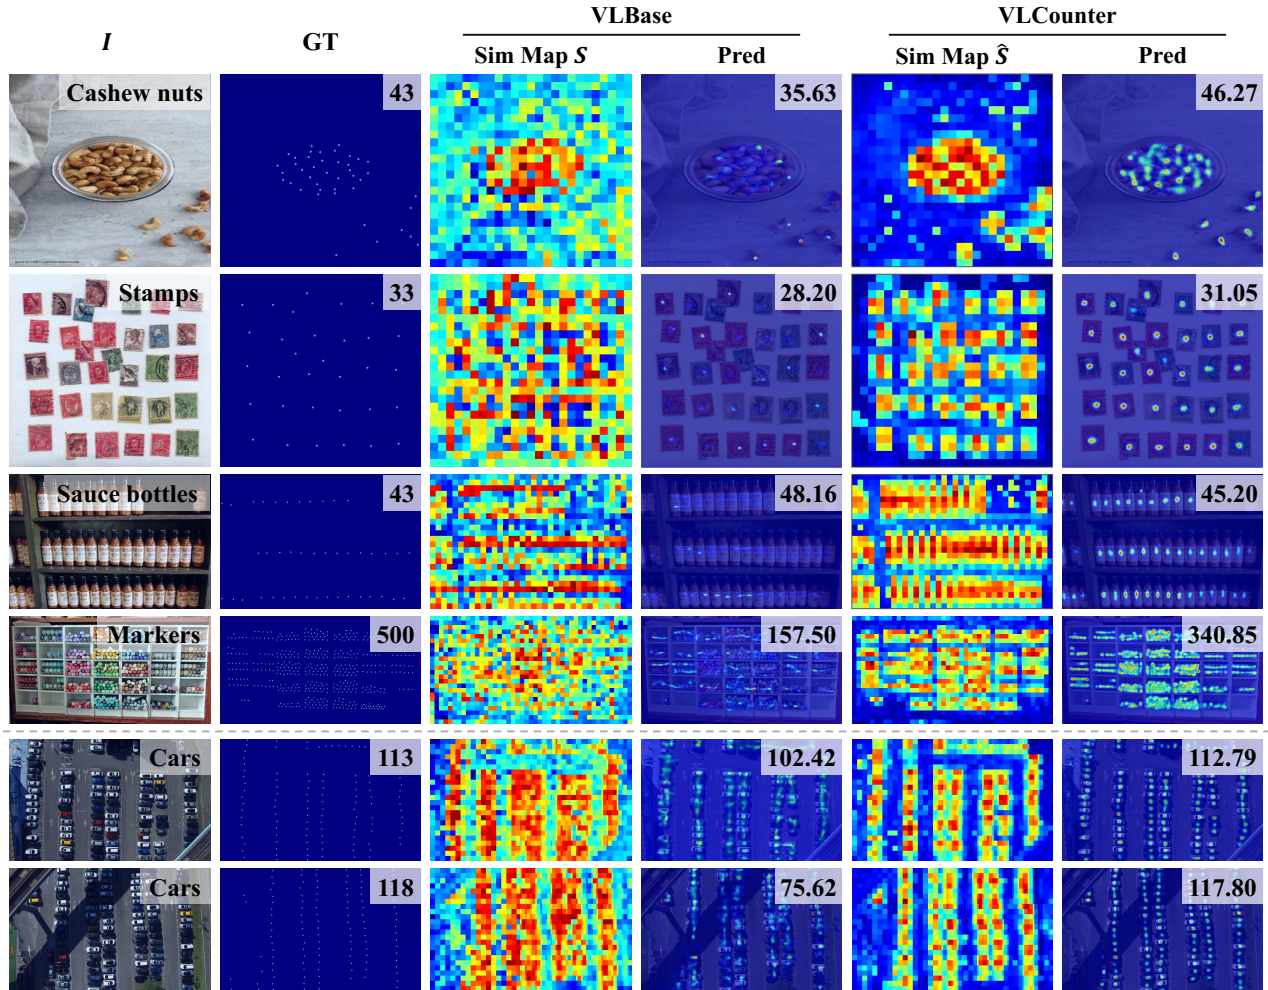


Figure 5: Qualitative comparison of VLBase and VLCOUNTER on the FSC-147 (Top 4 rows) and CARPK (Bottom 2 rows). Class names and counting values are shown at the right top of the query image (*I*) and the predicted density map, respectively.

5.4 Qualitative Results

Along with the quantitative results, we study how the components of VLCOUNTER affect class-specificity. In Fig. 5, we compare both the similarity map and the density map of VLBase and VLCOUNTER. By delivering the semantic condition and fine-tuning the similarity map, we find the similarity map to retain more compact salient regions; the activations in the background are suppressed (1st, 2nd rows) and object regions are clearly localized (2nd, 3rd rows). Then, by aggregating multi-level representations of rich semantics with these similarity maps in the decoder, we observe the clear discrepancy between the predicted density maps from VLBase and VLCOUNTER, especially for densely populated images (4th row).

Furthermore, we provide the cross-dataset results in the last two rows in Fig. 5. Similar to what we discussed with predictions for FSC147, we verify that VLCOUNTER is a counting-tailored and generalizable model across new categories, shapes, and densities of objects. These results ver-

ify the advantage of employing a pretrained vision-language model for capturing the semantics of newly seen objects, i.e., cars. Refer to the appendix for more visualizations.

6 Conclusion

In this work, we present a simple end-to-end framework VLBase and VLCOUNTER for zero-shot object counting that eliminates the need for the process of discovering exemplars. Simply put, VLBase is built upon the pre-trained vision-language CLIP model. Then, VLCOUNTER introduces three key components that bring task-specificity and object-specificity. Whereas the semantic-conditioned prompt tuning and learnable affine transformation fine-tune the encoding process to obtain counting-tailored representations, the segment-aware skip connection is designed to learn the generalizable decoder with the knowledge. Our thorough experiments on FSC147 and cross-dataset benchmarks validate the effectiveness and efficiency of VLCOUNTER.

Acknowledgments

This work was supported in part by MSIT/IITP (No. 2022-0-00680, 2019-0-00421, 2020-0-01821, 2021-0-02068), and MSIT&KNPA/KIPoT (Police Lab 2.0, No. 210121M06).

References

- Abousamra, S.; Hoai, M.; Samaras, D.; and Chen, C. 2021. Localization in the crowd with topological constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 872–881.
- Arteta, C.; Lempitsky, V.; and Zisserman, A. 2016. Counting in the wild. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, 483–498. Springer.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chan, A. B.; Liang, Z.-S. J.; and Vasconcelos, N. 2008. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *2008 IEEE conference on computer vision and pattern recognition*, 1–7. IEEE.
- Chattopadhyay, P.; Vedantam, R.; Selvaraju, R. R.; Batra, D.; and Parikh, D. 2017. Counting everyday objects in everyday scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1135–1144.
- Cholakkal, H.; Sun, G.; Khan, F. S.; and Shao, L. 2019. Object counting and instance segmentation with image-level supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12397–12405.
- Gong, S.; Zhang, S.; Yang, J.; Dai, D.; and Schiele, B. 2022. Class-agnostic object counting robust to intraclass diversity. In *European Conference on Computer Vision*, 388–403. Springer.
- Gu, Y.; Han, X.; Liu, Z.; and Huang, M. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Hoffmann, D. T.; Behrmann, N.; Gall, J.; Brox, T.; and Noroozi, M. 2022. Ranking info noise contrastive estimation: Boosting contrastive learning via ranked positives. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 897–905.
- Hsieh, M.-R.; Lin, Y.-L.; and Hsu, W. H. 2017. Drone-based object counting by spatially regularized regional proposal network. In *Proceedings of the IEEE international conference on computer vision*, 4145–4153.
- Idrees, H.; Tayyab, M.; Athrey, K.; Zhang, D.; Al-Maadeed, S.; Rajpoot, N.; and Shah, M. 2018. Composition loss for counting, density map estimation and localization in dense crowds. In *Proceedings of the European conference on computer vision (ECCV)*, 532–546.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *European Conference on Computer Vision*, 709–727. Springer.
- Laradji, I. H.; Rostamzadeh, N.; Pinheiro, P. O.; Vazquez, D.; and Schmidt, M. 2018. Where are the blobs: Counting by localization with point supervision. In *Proceedings of the european conference on computer vision (ECCV)*, 547–562.
- Li, B.; Weinberger, K. Q.; Belongie, S.; Koltun, V.; and Ranftl, R. 2022. Language-driven Semantic Segmentation. In *International Conference on Learning Representations*.
- Li, X. L.; and Liang, P. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Li, Y.; Wang, H.; Duan, Y.; and Li, X. 2023. Clip surgery for better explainability with enhancement in open-vocabulary tasks. *arXiv preprint arXiv:2304.05653*.
- Lian, D.; Li, J.; Zheng, J.; Luo, W.; and Gao, S. 2019. 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*, 1821–1830.
- Liu, W.; Salzmann, M.; and Fua, P. 2019. Context-aware crowd counting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5099–5108.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, E.; Xie, W.; and Zisserman, A. 2019. 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*, 669–684. Springer.
- Lüddecke, T.; and Ecker, A. 2022. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7086–7096.
- Moon, W.; Hyun, S.; Park, S.; Park, D.; and Heo, J.-P. 2023. Query-dependent video representation for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23023–23033.
- Mundhenk, T. N.; Konjevod, G.; Sakla, W. A.; and Boakye, K. 2016. 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*, 785–800. Springer.
- Nag, S.; Zhu, X.; Song, Y.-Z.; and Xiang, T. 2022. Zero-shot temporal action detection via vision-language prompting. In *European Conference on Computer Vision*, 681–697. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ranjan, V.; and Nguyen, M. H. 2022. Exemplar free class agnostic counting. In *Proceedings of the Asian Conference on Computer Vision*, 3121–3137.

- Ranjan, V.; Sharma, U.; Nguyen, T.; and Hoai, M. 2021. Learning to count everything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3394–3403.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18082–18091.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9529–9538.
- Sindagi, V. A.; Yasarla, R.; and Patel, V. M. 2019. Pushing the frontiers of unconstrained crowd counting: New dataset and benchmark method. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1221–1231.
- Wang, B.; Liu, H.; Samaras, D.; and Nguyen, M. H. 2020. Distribution matching for crowd counting. *Advances in neural information processing systems*, 33: 1595–1607.
- Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 139–149.
- Xie, W.; Noble, J. A.; and Zisserman, A. 2018. Microscopy cell counting and detection with fully convolutional regression networks. *Computer methods in biomechanics and biomedical engineering: Imaging & Visualization*, 6(3): 283–292.
- Xu, J.; Le, H.; Nguyen, V.; Ranjan, V.; and Samaras, D. 2023. Zero-shot Object Counting. *arXiv preprint arXiv:2303.02001*.
- Yang, S.-D.; Su, H.-T.; Hsu, W. H.; and Chen, W.-C. 2021. Class-agnostic few-shot object counting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 870–878.
- Zhang, Y.; Zhou, D.; Chen, S.; Gao, S.; and Ma, Y. 2016. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 589–597.
- Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract free dense labels from clip. In *European Conference on Computer Vision*, 696–712. Springer.