

Generalized-Scale Object Counting with Gradual Query Aggregation

Jer Pelhan, Alan Lukežič, Matej Kristan

Faculty of Computer and Information Science, University of Ljubljana, Slovenia
jer.pelhan@fri.uni-lj.si

Abstract

Few-shot detection-based counters estimate the number of instances in the image specified only by a few test-time exemplars. A common approach to localize objects across multiple sizes is to merge backbone features of different resolutions. Furthermore, to enable small object detection in densely populated regions, the input image is commonly upsampled and tiling is applied to cope with the increased computational and memory requirements. Because of these ad-hoc solutions, existing counters struggle with images containing diverse-sized objects and densely populated regions of small objects. We propose GeCo2, an end-to-end few-shot counting and detection method that explicitly addresses the object scale issues. A new dense query representation gradually aggregates exemplar-specific feature information across scales that leads to high-resolution dense queries that enable detection of large as well as small objects. GeCo2 surpasses state-of-the-art few-shot counters in counting as well as detection accuracy by 10% while running 3x faster at smaller GPU memory footprint.

Code — <https://github.com/jerpelhan/GECO2>

Introduction

Few-shot object counters estimate the number of objects, whose category was not observed in training, using only a few annotated exemplars. Driven by benchmarks (Ranjan et al. 2021; Nguyen et al. 2022), the research has initially focused on global counters, which report counts as summation over an image-wide estimated density map (Liu et al. 2022; Djukic et al. 2023). The continuous density is well suited for crowded regions, but cannot provide exact object locations, which makes the counts unexplainable and unsuitable for applications like biomedical analysis (Zavrtanik, Vodopivec, and Kristan 2020), where precise object positions and sizes estimates are required.

Detection-based counters, which produce per-object bounding boxes thus emerged. The current state-of-the-art is based on DETR-like (Carion et al. 2020) detection paradigms with fixed-set pre-trained object queries (Nguyen et al. 2022), region proposals (Pelhan et al. 2024a; Zhizhong et al. 2024) or dense object query maps (Pelhan et al. 2024b). However, these struggle with small objects in dense regions

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

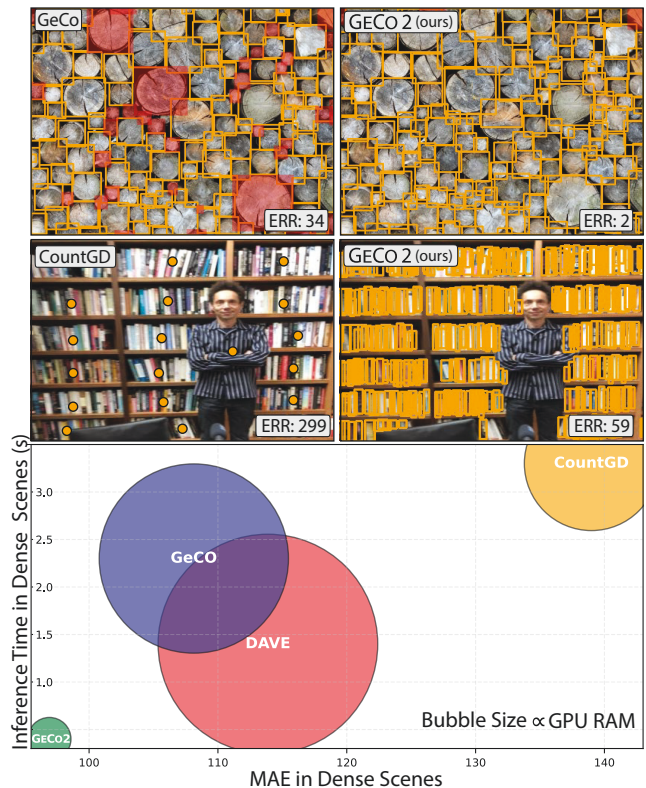


Figure 1: GeCo (Pelhan et al. 2024b) struggles with diverse object scales and CountGD (Amini-Naieni, Han, and Zisserman 2024) fails in dense scenes with small objects. GECO2 overcomes these issues by the new scale-generalized dense queries and achieves superior accuracy, fast inference and a substantially lower memory footprint.

(Figure 1). A major cause is the backbone feature resolution reduction, which causes several objects to occupy the same latent pixel. A common solution is to upsample the input image (Pelhan et al. 2024b; Liu et al. 2022; Amini-Naieni, Han, and Zisserman 2024), which however is limited by the GPU memory. To achieve sufficient resolution, tiling is employed (Liu et al. 2022; Amini-Naieni, Han, and Zisserman 2024), i.e., running inference on sub-images and

reassembling the results. Which leads to several issues: (i) exemplars cut by tiles cannot be easily extracted, (ii) transformer-based backbones yield tile-specific features incompatible with exemplars on tiles with different visual content and (iii) detection accuracy for objects at tile borders is low, often leading to double detections at tile reassembly. Moreover, since the scaling-based methods do not consider objects at multiple resolutions, their performance degrades when the object sizes in the image vary substantially (see Figure 1).

We propose a new Generalized-Scale Counter GECO2, that explicitly addresses the aforementioned challenges of diverse object sizes and dense regions. The core innovation is high-resolution dense query map construction, where prototypes are formed at each scale and query interactions occur independently at each scale, followed by cross-scale aggregation. While multi-scale features are widely used in vision, the proposed gradual high-resolution query construction through scale-specific encoding with class-defining exemplars on each separate scale is unique and outperforms the classical multi-scale formulations as demonstrated in experiments. In contrast to the current state-of-the-art, GECO2 thus avoids the need for the ad-hoc heuristic input upscaling and tiling for small object detection, enjoys a low memory footprint, and maintains excellent cross-scale detection capability.

GECO2 outperforms all few-shot density- and detection-based counters on the counting benchmark (Ranjan et al. 2021) by $\sim 25\%$ RMSE. It further delivers superior detection accuracy, surpassing the best method by 8% AP. On detection-field-adapted FSCD-LVIS (Nguyen et al. 2022), GECO2 outperforms state-of-the-art by 15% MAE in counting and by 23% AP in detection. On the recent multi-class benchmark (Hobley and Prisacariu 2025), it dominates by 25% MAE advantage, demonstrating excellent discrimination capabilities. Notably, GECO2 runs at least $3\times$ faster than the current detection-based counters, while using less than one-third of their GPU memory, underscoring its practical applicability.

Related Work

Traditional object counters originate from object detectors trained for predefined categories such as vehicles (Dai et al. 2019), cells (Falk et al. 2019), and people (Liu, Salzmann, and Fua 2019), and rely on extensively annotated datasets. Few-shot counting, introduced with FSC147 (Ranjan et al. 2021), shifts the paradigm by adapting to arbitrary and unseen categories at test time. Few-shot counters primarily follow two paradigms: (i) density map regression and (ii) detection-based counting. Early density-based methods, like FamNet (Ranjan et al. 2021) employed a Siamese network in combination with test-time fine-tuning, BMNet+ (Shi et al. 2022) formulated similarity metric learning, while SAFE-Count (You et al. 2023) introduced a new feature enhancement module. CounTR (Liu et al. 2022) employed an attention mechanism for exemplar-image feature interaction and introduced test-time normalization to adjust the total counts by the predicted counts within the exemplar bounding boxes.

Furthermore, when the given exemplars are small, indicating the presence of small objects, it divides the input image into sub-images, on which inference is run separately. Both test-time ad-hoc solutions lead to notable performance gains. LOCA (Djukic et al. 2023) improved generalization and performance by iteratively enhancing object prototypes with shape and appearance with attention.

Detection-based methods, in contrast, generate discrete bounding boxes, offering more informative outputs. DETR-based counters (Nguyen et al. 2022; Amini-Naieni, Han, and Zisserman 2024) rely on a fixed number of object queries (e.g., 600), limiting the counting capacity and forcing usage of ad-hoc tiling for dense scenes, which in practice restricts these counters to global count prediction, as detection merging is unfeasible in dense scenes. CountGD (Amini-Naieni, Han, and Zisserman 2024) introduces the SAM-based test-time normalization, a refined global count normalization based on the number of detections within the SAM (Kirillov et al. 2023) segmentation mask, which is inferred from exemplar bounding boxes. This normalization prevents reporting object location. A recent detection-based counter GeCo (Pelhan et al. 2024b) introduced a dense non-parametric prototype formulation, removing the fixed query constraint. However, to cope with small objects, these methods require input upsampling, which significantly increases inference time and memory usage due to the high cost of self-attention over dense representations.

Object scale variation is another fundamental challenge in few-shot counting, as areas of single objects vary from 20 pix^2 to $50,000 \text{ pix}^2$ (Ranjan et al. 2021). Most state-of-the-art approaches (Liu et al. 2022; You et al. 2023; Shi et al. 2022) do not address scale and regress on features from a single resolution. Other approaches provide only partial solutions. Famnet (Ranjan et al. 2021) scales the exemplar features, i.e., template, to obtain multiple correlation maps. To cope with memory limitations, LOCA (Djukic et al. 2023), DAVE (Pelhan et al. 2024a) and GeCo (Pelhan et al. 2024b) operate on low-resolution features, while GeCo (Pelhan et al. 2024b) upsamples the output using higher-resolution features. In contrast, our GECO2 introduces a new object query map encoder that comfortably scales over different resolutions and proposes a gradual latent query construction that leads to improved detection performance at a lower memory footprint.

GECO2

Given an image $I \in \mathbb{R}^{H_0 \times W_0 \times 3}$ and a set of k exemplar boxes $\mathbf{B}^E = \{\mathbf{b}_i\}_{i=1:k}$, the goal is to detect all target instances in I as a set of boxes $\mathbf{B}^P = \{\mathbf{b}_j\}_{j=1:N}$, and their count estimated as $N = |\mathbf{B}^P|$. To cope with dense object regions, we cast the detection problem as dense query map construction (Pelhan et al. 2024b), which is decoded into the set of final bounding boxes. There is a key difference between (Pelhan et al. 2024b), which predicts a low-resolution query map and upscales it with high-resolution backbone features at the decoding stage, and our approach. In particular, our approach directly builds a high-resolution generalized-scale dense query map, exploiting exemplars at

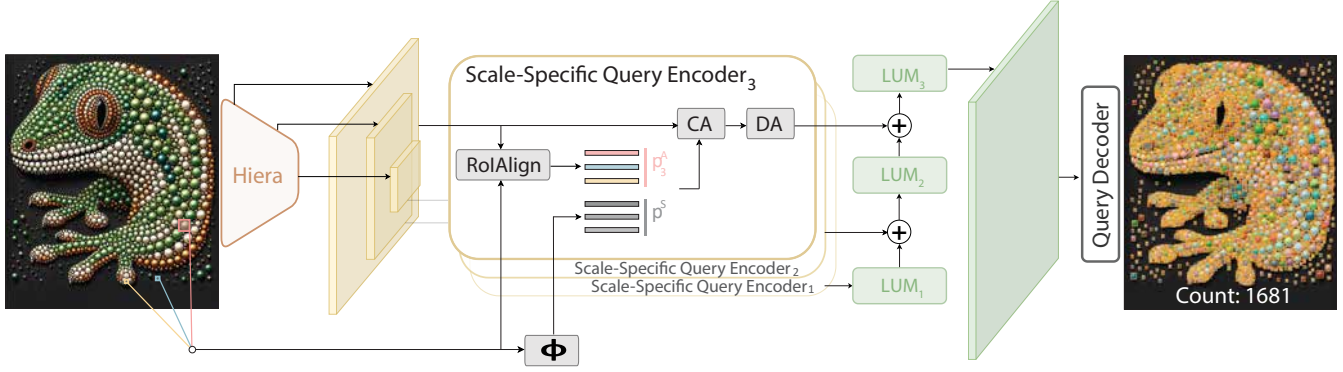


Figure 2: The architecture of the proposed generalized-scale detection-based object counter GECO2.

multiple scales and avoiding the need for image upscaling.

The proposed GECO2 architecture is outlined in Figure 2. Hiera backbone (Ryali et al. 2023) extracts multi-resolution features $\mathbf{f}^l \in \{\mathbf{C}_L\}_{L=1,2,3}$, where $\mathbf{C}_L \in \{\mathbb{R}^{h_L \times w_L \times d}\}$, $h_L = H_0/r_L$, $w_L = W_0/r_L$, $r \in \{16, 8, 4\}$ and d is number of feature channels. Each scale is processed separately by a scale-specific query encoder, employing scale-specific appearance and shape prototypes. A cross-scale query aggregation module then gradually transforms and merges scale-specific dense queries into a generalized single-scale dense query map, from which object detections are inferred. The detected instances are further refined by the SAM2 (Ravi et al. 2024) segmentation head. GECO2 can thus be seen as a new SAM2 head, adding a counting capability to the SAM2 swiss-knife backbone.

Scale-Specific Query Encoder

The task of the scale-specific query encoder is to extract exemplar-guided object presence information for the objects corresponding to the scale L of the considered backbone features \mathbf{C}_L . To focus on the relevant object types, scale-specific object appearance and object shape prototypes are extracted from the exemplars. The object shape prototypes $\mathbf{p}^S \in \mathbb{R}^{k \times d}$ are constructed from their exemplar width and height (W_{b_i}, H_{b_i}) in the original image resolution, i.e., $\mathbf{p}_i^S = \Phi([W_{b_i}, H_{b_i}])$, where $\Phi(\cdot)$ is a trainable MLP, as in (Djukic et al. 2023), with same parameters shared across all scales.

In addition to shape, the appearance prototypes $\mathbf{p}_L^A \in \mathbb{R}^{k \times d}$ are extracted separately at each scale by the RoiAlign (He et al. 2017) operation from their bounding boxes. The two sets of prototypes (\mathbf{p}_L^A and \mathbf{p}^S) are concatenated, forming scale-specific prototypes $\mathbf{p}_L \in \mathbb{R}^{2k \times d}$. The scale-specific dense object queries map $\mathbf{Q}_L \in \{\mathbb{R}^{h_L \times w_L \times d}\}$ is then constructed as follows. First, exemplar information is transferred to the scale-specific features by a sequence of cross-attention operations,

$$\mathbf{Q}'_L^{(i)} = \text{CA}(\mathbf{Q}'_L^{(i-1)}, \mathbf{p}_L, \mathbf{p}_L), \quad (1)$$

where $\text{CA}(\mathbf{q}, \mathbf{k}, \mathbf{v})$ is the standard cross-attention (Vaswani et al. 2017) with query/key/value tensors $(\mathbf{q}, \mathbf{k}, \mathbf{v})$, including

a skip connection, followed by a layer normalization, with $\mathbf{Q}'_L^{(0)} = \mathbf{C}_L$ and $i = 1 : N_{CA}$.

The queries are then refined by a sequence of local transformations, employing deformable attention $\text{DA}(\cdot)$ (Zhu et al. 2021), which attends to only a small number of nearby queries with a learnable displacement prediction. This enables processing low- and high-resolution scales without quadratic computational complexity restrictions of full self-attention, i.e.,

$$\mathbf{Q}''_L^{(j)} = \text{DA}(\mathbf{Q}''_L^{(j-1)}), \quad (2)$$

where $\mathbf{Q}''_L^{(0)} = \mathbf{Q}'_L^{(N_{CA})}$ and $j = 1 : N_{DA}$. The output of the scale-specific query encoder at L -th scale is thus $\mathbf{Q}_L = \mathbf{Q}''_L^{(N_{DA})}$.

Cross-Scale Query Aggregation

The scale-specific object queries \mathbf{Q}_L are gradually aggregated into generalized-scale high-resolution dense object queries \mathbf{Q} . The aggregation is performed by consecutive lightweight upsampling and fusion modules (denoted as LUM), each composed of $2 \times$ bilinear upsampling, 3×3 convolution, and a GeLU activation. Starting from the coarsest resolution, queries are progressively generalized across scales to a higher resolution, i.e.,

$$\mathbf{Q} = \text{LUM}_3(\text{LUM}_2(\text{LUM}_1(\mathbf{Q}_1) + \mathbf{Q}_2) + \mathbf{Q}_3). \quad (3)$$

Note that separate trainable parameters are used in each LUM_L , to enable scale-specific fusion. Since we consider three scales, the output dense object query map from the cross-scale query aggregation module is $\mathbf{Q} \in \mathbb{R}^{H \times W \times d}$, where $H = H_0/2$, $W = W_0/2$.

Dense Query Decoder

The dense query decoder transforms the generalized-scale high-resolution dense object queries \mathbf{Q} into object detections. In particular, two heads are employed to predict dense objectness score map and per-location bounding boxes. The objectness score map $\mathbf{y}^O \in \mathbb{R}^{HW \times 1}$ is obtained by applying a linear-layer transformation, followed by a Leaky ReLU activation,

$$\mathbf{y}^O = \text{LReLU}(\mathbf{W}^O \mathbf{Q}), \quad (4)$$

where \mathbf{W}^O is a learned matrix. The bounding boxes $\mathbf{y}^{\text{BB}} \in \mathbb{R}^{HW \times 4}$ are estimated in the *tlrb* format (Tian et al. 2019) using a three-layer MLP,

$$\mathbf{y}^{\text{BB}} = \sigma(\text{MLP}(\mathbf{Q})), \quad (5)$$

where $\sigma(\cdot)$ denotes the sigmoid function. Note that \mathbf{Q} in (4) and (5) is reshaped into $HW \times d$ for computation purposes. The final detections \mathbf{B}^P are obtained from $\mathbf{y}^O \in \mathbb{R}^{H \times W \times 1}$ and $\mathbf{y}^{\text{BB}} \in \mathbb{R}^{H \times W \times 4}$ as follows. Bounding boxes on spatial locations with objectness scores larger than their immediate 8 neighbors (a local maximum), surpassing a threshold set on the training set, are retrieved from \mathbf{y}^{BB} . During inference, these bounding boxes are fed as prompts to SAM2 (Ravi et al. 2024) decoder on the already computed backbone features \mathbf{f}^l to predict segmentation masks. Each bounding box is then refined by min-max fit to the corresponding mask. Note that in cases with over 800 objects detected, the average object size is approximately 25 pixels, at which point SAM2 masks become unreliable and are ignored. Finally, a non-maxima suppression is applied to the bounding boxes to remove duplicate detections, yielding the output bounding boxes \mathbf{B}^P and masks \mathbf{M}^P .

Training

The recent detection-oriented counting loss \mathcal{L}' with automatic hard-negative mining on \mathbf{y}^O and \mathbf{y}^{BB} is applied (Pelhan et al. 2024b), without the need for initial warm-up period, which we attribute to the new dense query map robustness. To enhance small object detection, an auxiliary loss computed from the highest feature resolution level is added. The dense object queries \mathbf{Q}_3 are transformed by a separate LUM module, i.e., $\mathbf{Q}^{\text{AUX}} = \text{LUM}_{\text{AUX}}(\mathbf{Q}_3)$, decoded by a separate Query Decoder, and a detection-oriented counting loss $\mathcal{L}'_{\text{AUX}}$ is computed. The auxiliary loss is computed only on images with average object size lower than $\theta_{\text{size}} = 25$ pixels. The final loss is thus $\mathcal{L} = \mathcal{L}' + \alpha \mathcal{L}'_{\text{AUX}}$.

Experiments

Implementation Details

GECO2 uses a pre-trained Hiera backbone (Ryali et al. 2023) from SAM2 (Ravi et al. 2024), which represents a unique selection due to its high-resolution multi-scale output. The backbone encodes the input image into three levels of features, all with channel depth $d = 256$. During training, the number of exemplars is set to $k = 3$. In the Scale-Specific Query Encoder, GECO2 performs $N_{\text{CA}} = 3$ iterations of prototype-to-image interactions, and $N_{\text{DA}} = 2$ query refinement iterations. GECO2 follows the standard test-time practice, where the input image is scaled to fit the width and height of the average exemplar under 80 pixels. The image is then zero-padded to $W_0 = H_0 = 1024$, which is the *constant* input resolution to GECO2.

The Hiera backbone parameters are frozen during training, while the other parameters of GECO2 are trained for 200 epochs using the AdamW optimizer (Loshchilov and

Hutter 2019) with a batch size of 8, the initial learning rate of 10^{-4} , and a weight decay of 5×10^{-5} . The auxiliary loss weight is set to $\alpha = 0.3$. Training runs on two A100 GPUs with the standard scale augmentations (Pelhan et al. 2024a; Djukic et al. 2023; Pelhan et al. 2024b) and images zero-padded to 1024×1024 pixels.

Experimental Protocols and Benchmarks

GECO2 is evaluated on two well-established few-shot counting and detection benchmarks FSCD147 (Nguyen et al. 2022) and FSCD-LVIS (Nguyen et al. 2022), and the most recent MCAC (Hobley and Prisacariu 2025), a multiclass synthesized benchmark. FSCD147 (Nguyen et al. 2022) extends FSC147 (Ranjan et al. 2021), a dataset comprising 6135 images across 147 object categories, with bounding box annotations. The dataset is split into 3659 training, 1286 validation, and 1190 test images, with disjoint categories ensuring that objects in the test set are unseen during training. FSCD-LVIS (Nguyen et al. 2022) features 377 object categories. We use the unseen split, which contains 3959 training and 2242 test images, ensuring no overlap in object categories between training and testing. The Multi-class, Class-Agnostic Counting dataset MCAC (Hobley and Prisacariu 2025), contains 1–4 object classes per image. For every preset category, separate counting runs are conducted, ensuring that models are rigorously tested on their ability to specialize at inference time for the target category.

Counting performance is evaluated by following the standard protocol (Ranjan et al. 2021), using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE). Detection accuracy is measured by Average Precision (AP) and AP at IoU=50 (AP50) (Nguyen et al. 2022). Both GPU usage and inference time are tracked using PyTorch on a single A100, measuring reserved memory, and execution time, excluding data loading. For additional insights, ablations and evaluations, please see the supplementary material.

State-of-the-Art Comparison

GECO2 is compared with the currently best state-of-the-art detection-based counters C-DETR (Nguyen et al. 2022), SAM-C (Ma, Hong, and Shanguan 2023), PSECO (Zhizhong et al. 2024), DAVE (Pelhan et al. 2024a), GeCo (Pelhan et al. 2024b), which provide object locations by bounding boxes, and CountGD (Amini-Naieni, Han, and Zisserman 2024) that provides point localization if test-time normalization or tiling is not activated. For completeness, we include state-of-the-art global counters, which only estimate the total count without localization, i.e., LOCA (Djukic et al. 2023), CACViT (Wang et al. 2024), CounTR (Liu et al. 2022), SAFECount (You et al. 2023), BMNet+ (Shi et al. 2022), VCN (Ranjan and Hoai 2022), CFOCFNet (Yang et al. 2021), MAML (Finn, Abbeel, and Levine 2017), FamNet (Ranjan et al. 2021) and CFOCFNet (Yang et al. 2021). Results are summarized in Table 1¹.

GECO2 outperforms all state-of-the-art detection-based counters on the FSCD147 (Nguyen et al. 2022) dataset, including the most recent GeCo (Pelhan et al. 2024b), with

¹Detailed qualitative analysis in the supplementary material

Method	Validation set				Test set			
	MAE(↓)	RMSE(↓)	AP(↑)	AP50(↑)	MAE(↓)	RMSE(↓)	AP(↑)	AP50(↑)
FamNet	23.75	69.07	-	-	22.08	99.54	-	-
CFOCNet	21.19	61.41	-	-	22.10	112.71	-	-
BMNet+	15.74	58.53	-	-	14.62	91.83	-	-
VCN	19.38	60.15	-	-	18.17	95.60	-	-
SAFEC	15.28	47.20	-	-	14.32	85.54	-	-
CounTR	13.13	49.83	-	-	11.95	91.23	-	-
LOCA	10.24	32.56	-	-	10.79	56.97	-	-
CACViT	10.63	37.95	-	-	9.13	48.96	-	-
CountGD	8.12	38.97	-	-	8.35	89.80	-	-
C-DETR	20.38	82.45	17.27	41.90	16.79	123.56	22.66	50.57
SAM-C	31.20	100.83	20.08	39.02	27.97	131.24	27.99	49.17
PSECO	15.31	68.36	32.12 ^③	60.02	13.05	112.86	42.98 ^③	73.33 ^③
DAVE	9.75 ^③	40.30 ^③	24.20	61.08 ^③	10.45 ^③	74.51 ^③	26.81	62.82
GeCo	9.52 ^②	43.00 ^②	33.51 ^②	62.51 ^②	7.91 ^②	54.28 ^②	43.42 ^②	75.06 ^②
GECo2	9.40^①	33.28^①	34.08^①	63.21^①	7.64^①	39.39^①	46.81^①	75.89^①

Table 1: Performance of few-shot global counting methods (top part) and detection-based methods (bottom part) on FSCD147.

Method	Count		Detection	
	MAE(↓)	RMSE(↓)	AP(↑)	AP50(↑)
AttRPN-PB	39.16	46.09	3.15	7.87
C-DETR	23.50	35.89	3.85	11.28
DAVE	15.47 ^③	25.95 ^②	4.12 ^③	14.16 ^③
GeCo	15.26 ^②	28.80 ^③	11.47 ^②	24.49 ^②
GECo2	13.01^①	23.43^①	14.08^①	28.84^①

Table 2: Few-shot counting and detection on the FSCD-LVIS "unseen" split.

improvements of 22% and 27% in RMSE on the test and validation splits, respectively, and 3% MAE on the test split (see Table 1). Notably, GECo2 is the only detection-based few-shot counter that surpasses all global counters, outperforming CACViT by 16% MAE and by a remarkable 20% RMSE on the test set. This advancement is driven by its scale-generalized query formulation, effectively addressing the challenges in dense regions with small objects, without resorting to ad-hoc scaling heuristics like the current state-of-the-art. GECo2 achieves superior detection accuracy, outperforming all state-of-the-art methods, e.g., the three-stage PSECO (Zhizhong et al. 2024), which leverages both SAM (Kirillov et al. 2023) and CLIP (Radford et al. 2021), and two-stage DAVE (Pelhan et al. 2024a), which employs an additional verification step. Furthermore, the current state-of-the-art GeCo (Pelhan et al. 2024b), which also exploits SAM, is outperformed by 8% AP on FSCD147 dataset (Nguyen et al. 2022). By excelling in both total count estimation as well as in localization accuracy, GECo2 demonstrates its superiority in low-shot counting.

GECo2 is further evaluated on FSCD-LVIS (Nguyen et al. 2022), unseen split, which follows the few-shot counting setup, where categories across splits are disjoint. Results in Table 2 show that GECo2 outperforms all state-of-the-art

detection counters by 15% MAE and 19% RMSE in counting, but also by 23% AP and 18% AP50 in detection. These results on the detection-field-adapted FSCD-LVIS support the conclusions drawn from the FSCD147 experiments.²

Multi-Class Performance. As observed in (Pelhan et al. 2024a; Ciampi et al. 2025), state-of-the-art methods often prioritize high recall, leading to overcounting by indiscriminately detecting all objects in the image, even categories not specified by the exemplars. To further evaluate GECo2, we test it on a recently proposed synthetic multi-class dataset MCAC (Hobley and Prisacariu 2025), where 52% of the images contain multiple object categories, requiring separate count predictions for each category. GECo2 is compared against state-of-the-art global counters from (Ciampi et al. 2025), including LOCA (Djukic et al. 2023), CounTR (Liu et al. 2022), BMNet+(Shi et al. 2022), and FamNet(Ranjan et al. 2021). Results in Table 3 show that GECo2 outperforms all few-shot counters, surpassing the benchmark top-performer, LOCA, by 27% in MAE and 22% in RMSE on the test set. These results indicate that GECo2 successfully differentiates between object categories, a critical ability that is overlooked (Pelhan et al. 2024a; Ciampi et al. 2025) in experimental evaluation on FSC147 (Ranjan et al. 2021).

Comparison with Scaling-Heuristic-Free Sota

As already mentioned, the current best-performing few-shot counters rely on ad-hoc resolution upscaling strategies. In particular GeCo (Pelhan et al. 2024b) upscales the input images with small objects to 1536×1536 , whereas CountGD (Amini-Naieni, Han, and Zisserman 2024) performs tiled inference, raising the effective resolution up to 3465×2400 . These techniques are crucial for achieving high performance, but they significantly increase computation and GPU memory requirements, thereby constraining their applicability. In contrast, GECo2 performs inference on fixed input resolution of 1024×1024 , with all im-

²Evaluation on CARPK in the supplementary material.

Method	Validation		Test	
	MAE(↓)	RMSE(↓)	MAE(↓)	RMSE(↓)
Famnet+	24.76	41.12	26.40	45.52
Bmnet+	15.83	27.07	17.29	29.83
GeCo	15.31	28.42	19.91	33.25
CounTR	15.07 ^③	26.26 ^③	16.12 ^③	29.28 ^③
LOCA	10.45 ^②	20.81 ^②	10.91 ^②	22.04 ^②
GECO2	9.38^①	18.80^①	7.93^①	17.05^①

Table 3: Few-shot counting on MCAC benchmark.

Method	MAE (↓)	RMSE (↓)	AP (↑)	AP50 (↑)
PSECO	13.05	112.86	42.98	73.33
LOCA	10.79	56.97	-	-
CountGD	9.44 \downarrow 1.0	98.84 \downarrow 9.0	-	-
DAVE	12.21 \downarrow 1.8	94.93 \downarrow 20.4	26.44 \downarrow 0.4	61.04 \downarrow 1.8
GeCo	8.88 \downarrow 1.0	75.67 \downarrow 21.4	43.15 \downarrow 0.3	73.40 \downarrow 1.7
GECO2	7.64	39.4	46.81	75.89

Table 4: Few-shot counters on the FSCD147 without upscaling and tiling. Performance drops are marked in red.

ages zero-padded if needed to fit this size. To emphasize the object-size generalization issues of existing methods, we turned off their ad-hoc scaling heuristics.

The results are shown in Table 4. GECO2 maintains its performance, since it does not apply the heuristics. On the other hand, the current state-of-the-art performance falls below the classic density-based counter LOCA (Djukic et al. 2023), underscoring their significant reliance on image upscaling. The performance gap between GECO2 and GeCo on the FSCD147 (Ranjan et al. 2021) test set further increases to 13% and 48% in MAE and RMSE, respectively. Additionally, GECO2 surpasses CountGD by 19% in MAE and a significant 151% in RMSE.

The observed RMSE drop of 39% for GeCo (Pelhan et al. 2024b), 27% for DAVE (Pelhan et al. 2024a), and 10% for CountGD highlights the poor generalization of state-of-the-art methods across different scales. In contrast, GeCo2 handles multiple scales in a principled way without reliance on ad-hoc pre-processing and without computational overhead.

Memory and Inference Speed

To assess the practical applicability of few-shot counters, we measured the average inference time per any image and the average inference time per image for images containing at least 300 objects on the validation and test set from FSCD147 (Ranjan et al. 2021), and reported the MAE error for this subset of dense scenes images. Additionally, we recorded the per-image GPU memory usage when ad-hoc scaling or tiling is applied in the counters.

The results are summarized in Table 5. Related methods require significantly more resources than GECO2, consuming at least 210% more GPU memory. On average, GECO2 infers the total count in 0.4 seconds, whereas the recent GeCo (Pelhan et al. 2024b) and CountGD (Amini-Naieni, Han, and Zisserman 2024) are 3 \times and 6 \times slower,

Method	mem(Gb)	time(s)	time(s) [†]	MAE [†]
DAVE	25.33	0.3 \downarrow 1.5 \times	1.4 \downarrow 3.5 \times	113.88
CountGD	15.47	1.3 \downarrow 6.5 \times	3.3 \downarrow 8.3 \times	138.99
GeCo	23.83	0.6 \downarrow 3 \times	2.3 \downarrow 5.8 \times	108.13
GECO2	4.98	0.2	0.4	96.92

Table 5: Inference time and GPU memory per image on FSCD147; [†] for scenes with over 300 objects.

respectively. Note that the gap between GECO2 and other counters increases even further on images with more than 300 objects, when the average width/height of objects is \sim 50 px. In particular, GeCo (Pelhan et al. 2024b) and CountGD (Amini-Naieni, Han, and Zisserman 2024) become 6 \times and 8 \times slower than GECO2.

To further analyze the source of the performance efficiency gain for GECO2, we measure the average inference times, excluding the backbones. We observe a 3.8 \times speedup for GECO2 (0.13 s) compared to GeCo (0.49 s). This conclusively validates the efficiency of the new multi-scale query map construction that eliminates the need for ad-hoc pre/post processing, while leading to a substantial \sim 10% error reduction in challenging dense scenarios.

To isolate the effect of the backbone, we trained GeCo (Pelhan et al. 2024b) with the Hiera backbone from SAM2, denoted GeCo+Hiera (see Table 6). This variant requires 17GB RAM and averages 0.3s per FSCD147 image—1.5 \times slower than GECO2. On dense scenes, inference time increases to 1.9s, nearly 5 \times slower, confirming that the efficiency gains of GECO2 stem from the new query map construction method, rather than from the backbone.

Ablation Study

Ablation results on FSC147 (Ranjan et al. 2021) are summarized in Table 6. First, to assess the contribution of the Scale-Specific Query Encoder, we ablate the multi-scale query formulation by limiting the prototype-image interaction only to the lowest scale Q_1 . The query generation process can be defined by rewriting (3) as follows

$$\mathbf{Q} = \text{LUM}_3 \left(\text{LUM}_2 \left(\text{LUM}_1(\mathbf{Q}_1) + \mathbf{C}_2 \right) + \mathbf{C}_3 \right). \quad (6)$$

This version (denoted GECO2_{FP}) can be viewed as a common pipeline with interaction performed on the lowest scale, followed by a feature pyramid to get more details from high-resolution features. This variant incurs a substantial 11% increase in MAE and 51% in RMSE. The qualitative comparison in Figure 3 further confirms this degradation, demonstrating the necessity of query generation at multiple scales. The proposed multi-scale query formulation is key to robust scale generalization, enabling accurate detections of small objects in dense regions (rows 1 and 2) without reliance on ad-hoc upscaling strategies, as well as objects with varying sizes within a single image (row 3). Furthermore, to demonstrate the impact of individual feature layers, GECO2 is modified to use only features from the first, second, or third backbone layers. These versions are denoted by GECO2_{Q₁}, GECO2_{Q₂} and GECO2_{Q₃}, respectively. The results clearly

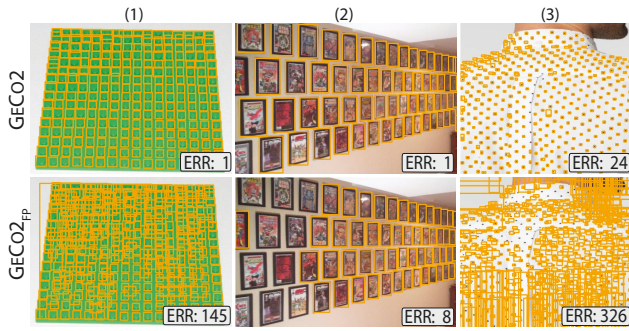


Figure 3: GECO2 and GECO2_{FP} with removed multi-scale query generation, using C_2 and C_3 for upsampling.

show that both versions result in a significant performance drop in both counting and detection, which implies that using only a single level of high-resolution features does not deliver good performance. Deep low-resolution features capture more semantic context and perform well in complex scenes, whereas, high-resolution features are critical for detecting small objects and handling densely populated regions. The proposed query map construction leverages both: it individually generates scale-specific queries, which are progressively refined across scales, enabling strong semantic reasoning without sacrificing spatial precision.

We also ablate N_{CA} , which when set to $N_{CA} = 2$ results in a minor drop (3% MAE, 10% RMSE). Setting number of deformable attentions $N_{DA} = 1$, for query refinement, results in on-pair performance considering MAE and 30% RMSE performance drop. This justifies the proposed prototype formulation and dense query transformation steps.

To assess the importance of auxiliary supervision for small objects, we set the auxiliary loss weight $\alpha = 0$. An increase of 8% in MAE and 35% in RMSE, along with reduced detection metrics is observed. This underscores the importance of the additional training signal on the highest resolution of features Q_3 for guiding the network to learn accurate small object detection. Alternatively, larger values, e.g., $\alpha = 0.5$, also deteriorate the performance.

To give insight into the shape prototype encoder that is shared over the scales, we re-trained a variant that uses a separate encoder for each scale (GECO2_{3 Φ (\cdot)}). This led to 16% MAE and 60% RMSE performance drop, confirming that integrating the same shape prototypes at multiple scales is crucial for providing the same object size cues when generating object queries at all feature levels. Collectively, ablation experiments validate that the architecture design – comprising multi-scale dense query formulation of GECO2, deformable attention for query refinement, auxiliary supervision for small objects, and object-size-cue-providing shape prototypes – is designed robustly, and efficiently.

We also evaluate the impact of SAM2 (Ravi et al. 2024) mask refinement. Removing it (GECO2_{REF}) does not increase global count errors (MAE/RMSE), but results in a 30% AP and a modest 4% AP50 drops. This means, that although SAM2 indeed improves individual box quality, the objects are already well localized before refinement.

Method	Counting		Detection	
	MAE(\downarrow)	RMSE(\downarrow)	AP(\uparrow)	AP50(\uparrow)
GECO2	9.40	33.28	34.08	63.21
GECO2 _{FP}	10.48	50.49	34.02	62.43
GECO2 _{Q₁}	12.71	60.29	34.05	62.10
GECO2 _{Q₂}	22.64	58.15	24.98	49.86
GECO2 _{Q₃}	19.97	55.33	24.75	47.97
GECO2 _{N_{DA}=1}	9.45	43.19	34.03	63.05
GECO2 _{N_{CA}=2}	11.06	45.68	33.58	62.15
GECO2 _{$\alpha=0$}	10.15	44.89	33.61	61.19
GECO2 _{$\alpha=0.5$}	9.39	36.74	33.47	63.02
GECO2 _{REF}	9.13	31.76	24.15	60.22
GECO2 _{3Φ(\cdot)}	10.93	53.08	34.63	63.82
GeCo+Hiera	10.99	47.99	33.75	62.84

Table 6: Ablation study on the FSCD147 val split.

Lastly, we ablate the effect of the SAM2 backbone by integrating it with GeCo, to verify whether the performance gains can be attributed to the better backbone, rather than the proposed multi-scale dense query construction. The resulting variant, GeCo+Hiera, demonstrates improved detection performance over the original GeCo, yet GECO2 consistently outperforms it in both counting and detection metrics. This confirms that the gains of GECO2 are primarily driven by its scale-specific query design rather than the backbone substitution.

Conclusion

We proposed GECO2, a novel end-to-end low-shot detection-based object counter that explicitly models scale through exemplar-conditioned, scale-specific query encoders and a generalized-scale aggregation strategy, all while maintaining a low memory footprint. Unlike other counters that rely on computationally expensive upsampling or tiling, GECO2 delivers accurate counting even in dense scenes. Extensive evaluations on multiple benchmarks show that GECO2 outperforms state-of-the-art detection-based and global counters, reducing MAE by $\sim 10\%$ and RMSE by $\sim 20\%$. Furthermore, GECO2 surpasses detection-based counters by $\sim 10\%$ AP while running $3\times$ faster with a lower memory footprint. By eliminating the ad-hoc heuristics, GECO2 provides a principled architecture and sets a new state-of-the-art in few-shot counting.

GECO2 relies on backbones that support high-resolution input and yield high-quality multi-scale features, which may be seen as its limitation. Yet, it naturally aligns it with the SAM2 Hiera (Ravi et al. 2024) encoder, which positions GECO2 as an additional head on SAM2 swiss-army knife toolbox, seamlessly extending SAM2 with unprecedented counting capabilities.

Acknowledgements

This work was supported by the Slovenian Research Agency program P2-0214 and projects L2-3169, J2-60054, as well as the supercomputing network SLING (ARNES, EuroHPC Vega - IZUM).

References

- Amini-Naieni, N.; Han, T.; and Zisserman, A. 2024. CountGD: Multi-Modal Open-World Counting. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Ciampi, L.; Messina, N.; Pierucci, M.; Amato, G.; Avvenuti, M.; and Falchi, F. 2025. Mind the Prompt: A Novel Benchmark for Prompt-based Class-Agnostic Counting.
- Dai, Z.; Song, H.; Wang, X.; Fang, Y.; Yun, X.; Zhang, Z.; and Li, H. 2019. Video-based vehicle counting framework. *IEEE Access*, 7: 64460–64470.
- Djukic, N.; Lukezic, A.; Zavrtnik, V.; and Kristan, M. 2023. A Low-Shot Object Counting Network With Iterative Prototype Adaptation. In *ICCV*.
- Falk, T.; Mai, D.; Bensch, R.; Çiçek, Ö.; Abdulkadir, A.; Marrakchi, Y.; Böhm, A.; Deubner, J.; Jäckel, Z.; Seiwald, K.; et al. 2019. U-Net: deep learning for cell counting, detection, and morphometry. *Nature methods*, 16(1): 67–70.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Hobley, M.; and Prisacariu, V. 2025. Abc easy as 123: A blind counter for exemplar-free multi-class class-agnostic counting. In *European Conference on Computer Vision*, 304–319. Springer.
- 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Liu, C.; Zhong, Y.; Zisserman, A.; and Xie, W. 2022. CounTR: Transformer-based Generalised Visual Counting. In *BMVC*. BMVA Press.
- Liu, W.; Salzmann, M.; and Fua, P. 2019. Context-Aware Crowd Counting. In *CVPR*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled weight decay regularization. *ICLR*.
- Ma, Z.; Hong, X.; and Shangquan, Q. 2023. Can SAM Count Anything? An Empirical Study on SAM Counting. arXiv:2304.10817.
- Nguyen, T.; Pham, C.; Nguyen, K.; and Hoai, M. 2022. Few-shot Object Counting and Detection. In *ECCV*, 348–365. Springer.
- Pelhan, J.; Lukežič, A.; Zavrtnik, V.; and Kristan, M. 2024a. DAVE – A Detect-and-Verify Paradigm for Low-Shot Counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Pelhan, J.; Lukežič, A.; Zavrtnik, V.; and Kristan, M. 2024b. A Novel Unified Architecture for Low-Shot Counting by Detection and Segmentation. In *Advances in Neural Information Processing Systems*, volume 37. Curran Associates, Inc.
- Radford, A.; Sutskever, I.; Kim, J. W.; Krueger, G.; and Agarwal, S. 2021. CLIP: connecting text and images. *OpenAI*. <https://openai.com/blog/clip/>.
- Ranjan, V.; and Hoai, M. 2022. Vicinal Counting Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 4221–4230.
- Ranjan, V.; Sharma, U.; Nguyen, T.; and Hoai, M. 2021. Learning to count everything. In *CVPR*, 3394–3403.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Ryali, C.; Hu, Y.-T.; Bolya, D.; Wei, C.; Fan, H.; Huang, P.-Y.; Aggarwal, V.; Chowdhury, A.; Poursaeed, O.; Hoffman, J.; et al. 2023. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *International Conference on Machine Learning*, 29441–29454. PMLR.
- 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*, 9529–9538.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9627–9636.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Z.; Xiao, L.; Cao, Z.; and Lu, H. 2024. Vision transformer off-the-shelf: A surprising baseline for few-shot class-agnostic counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5832–5840.
- Yang, S.-D.; Su, H.-T.; Hsu, W. H.; and Chen, W.-C. 2021. Class-agnostic few-shot object counting. In *WACV*, 870–878.
- You, Z.; Yang, K.; Luo, W.; Lu, X.; Cui, L.; and Le, X. 2023. Few-shot object counting with similarity-aware feature enhancement. In *WACV*, 6315–6324.
- Zavrtnik, V.; Vodopivec, M.; and Kristan, M. 2020. A segmentation-based approach for polyp counting in the wild. *Engineering Applications of Artificial Intelligence*, 88: 103399.
- Zhizhong, H.; Mingliang, D.; Yi, Z.; Junping, Z.; and Hongming, S. 2024. Point, Segment and Count: A Generalized Framework for Object Counting. In *CVPR*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.