Few-shot Object Counting with Similarity-Aware Feature Enhancement

Zhiyuan You¹, Kai Yang², Wenhan Luo³, Xin Lu², Lei Cui⁴, Xinyi Le^{1*} ¹Shanghai Jiao Tong University, ²SenseTime, ³Sun Yat-sen University, ⁴Tsinghua University zhiyuanyou@foxmail.com, {yangkai, luxin}@sensetime.com, whluo.china@gmail.com

cuill9@mails.tsinghua.edu.cn, lexinyi@sjtu.edu.cn, *Corresponding Author

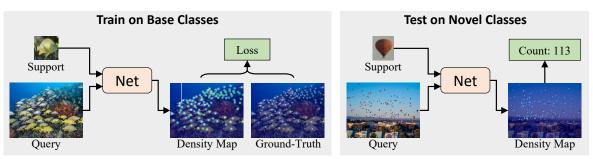


Figure 1. **Illustration of few-shot object counting**, where we would like to find how many exemplar objects described by *a few* support images occur in the query image. Besides the objects included in the training phase, we also expect the model to handle novel classes at the test stage *without retraining*.

Abstract

This work studies the problem of few-shot object counting, which counts the number of exemplar objects (i.e., described by one or several support images) occurring in the query image. The major challenge lies in that the target objects can be densely packed in the query image, making it hard to recognize every single one. To tackle the obstacle, we propose a novel learning block, equipped with a similarity comparison module and a feature enhancement module. Concretely, given a support image and a query image, we first derive a score map by comparing their projected features at every spatial position. The score maps regarding all support images are collected together and normalized across both the exemplar dimension and the spatial dimensions, producing a reliable similarity map. We then enhance the query feature with the support features by employing the developed point-wise similarities as the weighting coefficients. Such a design encourages the model to inspect the query image by focusing more on the regions akin to the support images, leading to much clearer boundaries between different objects. Extensive experiments on various benchmarks and training setups suggest that we surpass the state-of-the-art methods by a sufficiently large margin. For instance, on a recent large-scale FSC-147 dataset, we surpass the state-of-theart method by improving the mean absolute error from 22.08 to 14.32 (35%). Code has been released in https: //github.com/zhiyuanyou/SAFECount.

1. Introduction

Object counting [3,4], which aims at investigating how many times a certain object occurs in the query image, has received growing attention due to its practical usage [8,13, 17,45]. Most existing studies assume that the object to count at the test stage is covered by the training data [1,10,11,17, 26,44,45]. As a result, each learned model can only handle a specific object class, greatly limiting its application.

To alleviate the generalization problem, few-shot object counting (FSC) is recently introduced [21]. Instead of predefining a common object that is shared by all training images, FSC allows users to customize the object of their own interests with a few support images, as shown in Fig. 1. In this way, we can use a single model to unify the counting of various objects, and even adapt the model to novel classes (*i.e.*, unseen in the training phase) without any retraining.

A popular solution to FSC is to first represent both the exemplar object (*i.e.* the support image) and the query image with expressive features, and then pinpoint the candidates via analyzing the feature correlation [18,21,40]. Active attempts roughly fall into two folds. One is featurebased [18], as shown in Fig. 2a, where the pooled support feature is concatenated onto the query feature, followed by a regress head to recognize whether the two features are close enough. However, the spatial information of the support image is omitted by pooling, leaving the feature comparison unreliable. The other is similarity-based [21,40],

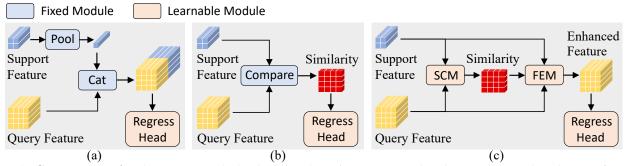


Figure 2. **Concept comparison** between our method and existing alternatives. (a) Feature-based approach [18], where the query feature is concatenated with the pooled support feature for regression. (b) Similarity-based approach [21,40], where a similarity map is developed from raw features for regression. (c) Our proposed *similarity-aware feature enhancement* block, consisting of a similarity comparison module (SCM) and a feature enhancement module (FEM). Concretely, the reliable feature similarity developed by SCM is exploited as the guidance of FEM to enhance the query feature with the support feature. The details of SCM and FEM can be found in Sec. 3.2 and Fig. 3.

as shown in Fig. 2b, where a similarity map is developed from raw features as the regression object. Nevertheless, the similarity is far less informative than feature, making it hard to identify clear boundaries between objects (see Fig. 5). Accordingly, the counting performance heavily deteriorates when the target objects are densely packed in the query image, like the shoal of fish in Fig. 1.

In this work, we propose a Similarity-Aware Feature Enhancement block for object Counting (SAFECount). As discussed above, feature is more informative while similarity *better captures the support-query relationship.* Our novel block adequately integrates both of the advantages by exploiting similarity as a guidance to enhance the features for regression. Intuitively, the enhanced feature not only carries the rich semantics extracted from the image, but also gets aware of which regions within the query image are similar to the exemplar object. Specifically, we come up with a similarity comparison module (SCM) and a feature enhancement module (FEM), as illustrated in Fig. 2c. On one hand, different from the naive feature comparison in Fig. 2b, our SCM learns a feature projection, then performs a comparison on the projected features to derive a score map. This design helps select from features the information that is most appropriate for object counting. After the comparison, we derive a reliable similarity map by collecting the score maps with respect to all support images (i.e., few-shot) and normalizing them along both the exemplar dimension and the spatial dimensions. On the other hand, the FEM takes the point-wise similarities as the weighting coefficients, and fuses the support features into the query feature. Such a fusion is able to make the enhanced query feature focus more on the regions akin to the exemplar object defined by support images, facilitating more precise counting.

Experimental results on a very recent large-scale FSC dataset, FSC-147 [21], and a car counting dataset, CARPK [10], demonstrate our *substantial improvement* over state-of-the-art methods. Through visualizing the intermediate similarity map and the final predicted density

map, we find that our SAFECount substantially benefits from the clear boundaries learned between objects, even when they are densely packed in the query image.

2. Related Work

Class-specific object counting counts objects of a specific class, such as people [17, 26, 44, 45], animals [1], cars [10], among which crowd counting has been widely explored. For this purpose, traditional methods [14, 28, 33] count the number of people occurring in an image through person detection. However, object detection is not particularly designed for the counting task and hence shows unsatisfying performance when the crowd is thick. To address this issue, recent work [32] employs a deep model to predict the density map from the crowd image, where the sum over the density map gives the counting result [15]. Based on this thought, many attempts have been made to handle more complicated cases [2, 16, 20, 23–25, 36, 38, 41, 42, 45]. Some recent studies [26, 31] propose effective loss functions that help predict the position of each person precisely. However, all of these methods can only count objects regarding a particular class (e.g., person), making them hard to generalize. There are also some approaches targeting counting objects of multiple classes [13, 19, 27, 37]. In particular, Stahl et al. [27] propose to divide the query image into regions and regress the counting results with the inclusion-exclusion principle. Laradji et al. [13] formulate counting as a segmentation problem for better localization. Michel et al. [19] detect target objects and regress multi-class density maps simultaneously. Xu et al. [37] mitigate the mutual interference across various classes by proposing categoryattention module. Nevertheless, they still can not handle the object classes beyond the training data.

Few-shot object counting (FSC) has recently been proposed [18,21,40] and presents a much stronger generalization ability. Instead of pre-knowing the type of object to count, FSC allows users to describe the exemplar object

of their own interests with one or several support images. This setting makes the model highly flexible in that it does not require the test object to be covered by the training samples. In other words, a well-learned model could easily make inferences on novel classes (i.e., unseen in the training phase) as long as the support images are provided. To help the model dynamically get adapted to an arbitrary class, a great choice is to compare the object and the query image in feature space [18, 21, 40]. GMN [18] pools the support feature, and concatenates the pooling result onto the query feature, then learns a regression head for pointwise feature comparison. However, the comparison built on concatenation is not as reliable as the similarity [40]. Instead, CFOCNet [40] first performs feature comparison with dot production, and then regresses the density map from the similarity map derived before. FamNet [21] further improves the reliability of the similarity map through multi-scale augmentation and test-time adaptation. But similarities are far less informative than features, hence regressing from the similarity map fails to identify clear boundaries between the densely packed objects. In this work, we propose a similarityaware feature enhancement block, which integrates the advantages of both features and similarities.

Few-shot learning has received popular attention in the past few years thanks to its high data efficiency [6,7,34,35,39,43]. The rationale behind this is to adapt a well-trained model to novel test data (*i.e.*, having a domain gap to the training data) with a few support samples. In the field of image classification [7, 34], MAML [7] proposes to fit parameters to novel classes at the test stage using a few steps of gradient descent. FRN [34] formulates few-shot classification as a reconstruction problem. As for object detection [6, 35], Fan *et al.* [6] exploit the similarity between the input image and the support images to detect novel objects. Wu et al. [35] create multi-scale positive samples as the object pyramid for prediction refinement. When the case comes to semantic segmentation [39, 43], CANet [43] iteratively refines the segmentation results by comparing the query feature and the support feature. Yang *et al.* [39] aim to alleviate the problem of feature undermining and enhance the embedding of novel classes. In this work, we explore the usage of few-shot learning on the object counting task.

3. Method

3.1. Preliminaries

Few-shot object counting (FSC) [21] aims to count the number of exemplar objects occurring in a query image with only a few support images describing the exemplar object. In FSC, object classes are divided into base classes C_b and novel classes C_n , where C_b and C_n have no intersection. For each query image from C_b , both a few support images and the ground-truth density map are provided. While, for query

images from C_n , only a few support images are available. FSC aims to count exemplar objects from C_n using only a few support images by leveraging the generalization knowledge from C_b . If we denote the number of support images for one query image as K, the task is called K-shot FSC.

3.2. Similarity-Aware Feature Enhancement

Overview. Fig. 3 illustrates the core block in our framework, termed as the similarity-aware feature enhancement block. We respectively denote the support feature and the query feature as $f_{S} \in \mathbb{R}^{K \times C \times H_{S} \times W_{S}}$ and $f_{Q} \in \mathbb{R}^{C \times H_{Q} \times W_{Q}}$, where K is the number of support images. The similarity comparison module (SCM) first projects f_S and f_Q to a comparison space, then compares these projected features at every spatial position, deriving a score map, R_0 . Then, R_0 is normalized along both the exemplar dimension and the spatial dimensions, resulting in a reliable similarity map, R. The following feature enhancement module (FEM) first obtains the similarity-weighted feature, f_R , by weighting f_S with R, and then manages to fuse f_R into f_Q , producing the enhanced feature, f_Q' . By doing so, the features regarding the regions similar to the support images are "highlighted", which could help the model get distinguishable borders between densely packed objects. Finally, the density map is regressed from f'_{O} .

Similarity Comparison Module (SCM). As discussed above, similarity can better characterize how a particular image region is alike the exemplar object. However, we find that the conventional feature comparison approach (*i.e.*, using the vanilla dot production) used in prior arts [21,40] is not adapted to fit the FSC task. By contrast, our proposed SCM develops a reliable similarity map from the input features with the following three steps.

Step-1: Learnable Feature Projection. Before performing feature comparison, f_S and f_Q are first projected to a comparison space via a 1×1 convolutional layer. This projection asks the model to automatically select suitable information from the features. We also add a shared layer normalization after the projection to make these two features subject to the same distribution as much as possible.

Step-2: Feature Comparison. The point-wise feature comparison is realized with convolution. In particular, we convolve the projected f_Q with the projected f_S as kernels, which gives us the score map, $R_0 \in \mathbb{R}^{K \times 1 \times H_Q \times W_Q}$, as

$$\boldsymbol{R_0} = \operatorname{conv}(g(\boldsymbol{f_Q}), \operatorname{kernel} = g(\boldsymbol{f_S})), \quad (1)$$

where $g(\cdot)$ denotes the feature projection described in *Step-1*, *i.e.*, a 1×1 convolutional layer followed layer normalization.

Step-3: Score Normalization. The values of the score map, R_0 , are normalized to a proper range to avoid some unusual (*e.g.*, too large) entries from dominating the learning. Here, we propose Exemplar Normalization (ENorm) and

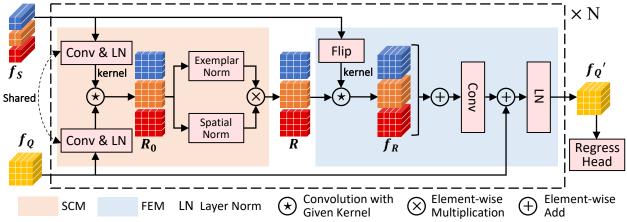


Figure 3. Illustration of the similarity-aware feature enhancement block under the 3-shot case. Given features, f_S , f_Q , that are extracted from the support images and the query image respectively, the similarity comparison module (SCM) first develops a score map, R_0 , by comparing the projected features, and then produces a similarity map, R, via normalizing R_0 along both the exemplar dimension and the spatial dimensions. Here, feature projection is implemented with a 1×1 convolution. The following feature enhancement module (FEM) weights f_S with R to derive a similarity-weighted feature, f_R , and manages to fuse f_R into f_Q as a feature enhancement. Such a block can be stacked for multiple times in the training framework.

Spatial Normalization (SNorm). On the one hand, ENorm normalizes R_0 along the exemplar dimension as

$$\boldsymbol{R_{EN}} = \texttt{softmax}_{dim=0} (\frac{\boldsymbol{R_0}}{\sqrt{H_S W_S C}}), \qquad (2)$$

where $softmax_{dim}(\cdot)$ is the softmax layer operated along a specific dimension. On the other hand, R_0 is also normalized along the spatial dimensions (*i.e.*, the height and width) with SNorm, as

$$\boldsymbol{R_{SN}} = \frac{\exp(\boldsymbol{R_0}/\sqrt{H_S W_S C})}{\max_{dim=(2,3)}(\exp(\boldsymbol{R_0}/\sqrt{H_S W_S C}))}, \quad (3)$$

where $\max_{dim}(\cdot)$ finds the maximum value from the given dimensions. After SNorm, the score value of the most support-relevant position would be 1, and others would be among [0, 1]. Finally, the similarity map, \boldsymbol{R} , is obtained from $\boldsymbol{R_{EN}}$ and $\boldsymbol{R_{SN}}$ with

$$\boldsymbol{R} = \boldsymbol{R}_{\boldsymbol{E}\boldsymbol{N}} \otimes \boldsymbol{R}_{\boldsymbol{S}\boldsymbol{N}} \in \mathbb{R}^{K \times 1 \times H_Q \times W_Q}, \qquad (4)$$

where \otimes is the element-wise multiplication. The studies of the effect of ENorm and SNorm can be found in Sec. 4.4.

Feature Enhancement Module (FEM). Recall that, compared to similarity, feature is more informative in representing the image yet less accurate in capturing the supportquery relationship. To take sufficient advantages of both, we propose to use the similarity developed by SCM as the guidance for feature enhancement. Specifically, our FEM integrates the support feature, f_S , into the query feature, f_Q , with similarity values in R as the weighting coefficients. In this way, the model can inspect the query image by paying more attention to the regions that are akin to the support images. This module consists of the following two steps. Step-1: Weighted Feature Aggregation. In this step, we aggregate the support feature, f_s , by taking the pointwise similarity, R, into account. Namely, the feature point corresponding to a higher similarity score should have larger voice to the final enhancement. Such a weighted aggregation is implemented with convolution, which outputs the similarity-weighted feature,

$$f'_{R} = \operatorname{conv}(R, \operatorname{kernel} = \operatorname{flip}(f_{S})) \in \mathbb{R}^{K \times C \times H_{Q} \times W_{Q}},$$
(5)

$$\boldsymbol{f_R} = \operatorname{sum}_{dim=0}(\boldsymbol{f'_R}) \in \mathbb{R}^{C \times H_Q \times W_Q}, \tag{6}$$

where $\operatorname{sum}_{dim}(\cdot)$ accumulates the input tensor along specific dimensions, $\operatorname{flip}(\cdot)$ denotes the flipping operation, which flips the input tensor both horizontally and vertically. Flipping helps f'_R preserve the spatial structure of f_S . The intuitive illustration and the performance improvement of flipping can be found in *Supplementary Material*.

Step-2: Learnable Feature Fusion. The similarityweighted feature, f_R , is fused into the query feature, f_Q , via an efficient network. It contains a convolutional block and a layer normalization, as shown in Fig. 3. Finally, we obtain the enhanced feature, f'_Q , with

$$f'_{Q} = \texttt{layer_norm}(f_{Q} + h(f_{R})) \in \mathbb{R}^{C \times H_{Q} \times W_{Q}},$$
 (7)

where $h(\cdot)$ is implemented with two convolutional layers.

Comparison with Attention. A classical attention module [30] involved with *query*, *key*, and *value* (denoted as q, k, v) is represented as similarity(q, k)v. The key idea is employing the similarity values between q and kas weighting coefficients to aggregate v as an information aggregation. Our SAFECount is similar with the similarityguided aggregation of existing attention modules. However, existing attention modules omit the spatial information as they need to flatten a feature map $(C \times H \times W)$ to a collection of feature vectors $(C \times HW)$. Instead, in all processes of our SCM and FEM, the feature maps are designed to *maintain their spatial structure* $(C \times H \times W)$, which plays a vital role in *learning clear boundaries between objects*. The ablation study in Sec. 4.4 confirms our significant advantage over the classical attention module.

3.3. Training Framework

Sec. 3.2 describes the core block of our approach, SAFECount. In practice, such a block should work together with a feature extractor, which feeds input features into the block, and a regression head, which receives the enhanced feature for object counting. Moreover, it is worth mentioning that our SAFECount allows stacking itself for continuous performance improvement. In this part, we will introduce these assistant modules, whose detailed structures are included in *Supplementary Material*.

Feature Extractor. When introducing our SAFECount block, we start with the support feature, f_S , and the query feature, f_Q , which are assumed to be well prepared. Specifically, we use a *fixed* ResNet-18 [9] pre-trained on ImageNet [5] as the feature extractor. In particular, given a query image, we resize the outputs of the first three stages of ResNet-18 to the same size, $H_Q \times W_Q$, and concatenate them along the channel dimension as the query feature. Besides, given a support image, which is usually cropped from a large image so as to contain the exemplar object only, the support feature is obtained by applying ROI pooling [22] on the feature extracted from its parent before cropping. Here, the ROI pooling size is the size of f_S , *i.e.*, $H_S \times W_S$.

Regression Head. After getting the enhanced feature, f'_Q , we convert it to a density map, $D \in \mathbb{R}^{H \times W}$, with a regression head. Following existing methods [18, 21, 40], the regression head is implemented with a sequence of convolutional layers, followed by Leaky ReLU activation and bi-linear upsampling.

Multi-block Architecture. Recall that our proposed SAFECount block enhances the input query feature, f_Q , with the support features, f_S . The enhanced feature, f'_Q , is with exactly the same shape as f_Q . As a result, it can be iteratively enhanced simply by stacking more blocks. The ablation study on the number of blocks can be found in Sec. 4.4, where we verify that *adding one block is already enough to boost the performance substantially*.

Objective Function. Most counting datasets are annotated with the coordinates of the target objects within the query image [3,4,45]. However, directly regressing the coordinates is hard [15,45]. Following prior work [21], we generate the ground-truth density map, $D_{GT} \in \mathbb{R}^{H \times W}$, from the labeled coordinates, using Gaussian smoothing with adaptive

Table 1. Statistics of the four fold splits from FSC-147 [21].

Fold	Class Indices	#Classes	#Images
0	0-35	36	2033
1	36-72	37	1761
2	73-109	37	1239
3	110-146	37	1113

window size. Our model is trained with the MSE loss as

$$\mathcal{L} = \frac{1}{H \times W} ||\boldsymbol{D} - \boldsymbol{D_{GT}}||_2^2.$$
(8)

4. Experiments

4.1. Metrics and Datasets

Metrics. We choose Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) to measure the performance of counting methods following [8,21]:

$$MAE = \frac{1}{N_Q} \sum_{i=1}^{N_Q} |C^i - C_{GT}^i|,$$

$$RMSE = \sqrt{\frac{1}{N_Q} \sum_{i=1}^{N_Q} (C^i - C_{GT}^i)^2},$$
(9)

where N_Q is the number of query images, C^i and C^i_{GT} are the predicted and ground-truth count of the i^{th} query image, respectively.

FSC-147. FSC-147 [21] is a multi-class, 3-shot FSC dataset with 147 classes and 6135 images. Each image has 3 support images to describe the target objects. Note that the training classes share no intersection with the validation classes and test classes. The training set contains 89 classes, while validation set and test set both contain another disjoint 29 classes. The number of objects per image varies extremely from 7 to 3701 with an average of 56.

Cross-validation of FSC-147. In original FSC-147 [21], the dataset split and the shot number are both fixed, while other few-shot tasks including classification [34], detection [6], and segmentation [39] all contain multiple dataset splits and shot numbers. Therefore, we propose to evaluate FSC methods with multiple dataset splits and shot numbers by incorporating FSC-147 and cross-validation. Specifically, we split all images in FSC-147 to 4 folds, whose class indices, class number, and image number are shown in Tab. 1. The class indices ranging from 0 to 146 are obtained by sorting the class names of all 147 classes. Note that these 4 folds share no common classes. When fold-i (i = 0, 1, 2, 3) serves as the test set, the remaining 3 folds form the training set. Also, we evaluate FSC methods in both 1-shot and 3shot cases. For 3-shot case, the original three support images in FSC-147 are used. For 1-shot case, we randomly sample one from the original three support images.

CARPK. A car counting dataset, CARPK [10], is used to test our model's ability of cross-dataset generality. CARPK

Table 2. **Quantitative results** on FSC-147 dataset [21], where we surpass other competitors by a sufficiently large margin.

Method	Val Set		Test Set	
	MAE	RMSE	MAE	RMSE
GMN [18]	29.66	89.81	26.52	124.57
MAML [7]	25.54	79.44	24.90	112.68
FamNet [21]	24.32	70.94	22.56	101.54
FamNet+ [21]	23.75	69.07	22.08	99.54
CFOCNet [40]	21.19	61.41	22.10	112.71
SAFECount (ours)	15.28	47.20	14.32	85.54

contains 1448 images and nearly 90, 000 cars from a drone perspective. These images are collected in various scenes of 4 different parking lots. The training set contains 3 scenes, while another scene is used for test.

4.2. Class-agnostic Few-shot Object Counting

Our method is evaluated on the FSC dataset FSC-147 [21] under the original setting and the cross-validation setting.

Setup. The sizes of the query image, the query feature map, and the support feature map, $H \times W$, $H_Q \times W_Q$, and $H_S \times W_S$, are selected as 512×512 , 128×128 , and 3×3 , respectively. The dimension of the projected features are set as 256. The multi-block number is set as 4. The model is trained with Adam optimizer [12] for 200 epochs with batch size 8. The learning rate is set as 2e-5 initially, and it is dropped by 0.25 every 80 epochs.

FSC-147. Quantitative results on FSC-147 are given in Tab. 2. Our method is compared with GMN [18], MAML [7], FamNet [21], and CFOCNet [40]. Our approach outperforms all counterparts with a quite large margin. For example, we surpass FamNet+ by 8.47 MAE and 21.87 RMSE on validation set, 7.76 MAE and 14.00 RMSE on test set. Note that FamNet+ needs test-time adaptation for novel classes, while our SAFECount needs no test-time adaptation. These significant advantages demonstrate the effectiveness of our method. In Fig. 4, we show some qualitative results of SAFECount. Compared with Famnet+ [21], our SAFECount has much stronger ability to separate each independent object within densely packed objects, thus helps obtain an accurate count. Especially, for densely packed green peas (Fig. 4b), we not only exactly predict the object count, but also localize target objects with such a high precision that every single object could be clearly distinguished.

Cross-validation of FSC-147. The dataset split and shot number are both fixed in FSC-147 benchmark, which could not provide a comprehensive evaluation. Therefore, we incorporate FSC-147 with cross-validation to evaluate FSC methods with 4 dataset splits and 2 shot numbers. Our approach is compared with FSC baselines including GMN [18] and FamNet [21]. These baselines are trained and evaluated by ourselves with the official code. The cross-validation results are shown in Tab. 3, where foldi (i = 0, 1, 2, 3) indicates the test set. Under all dataset splits and shot numbers, our method significantly outperforms baseline methods with both MAE and RMSE. Averagely, we outperform FamNet by 8.82 MAE and 20.18 RMSE in 1-shot case, 9.67 MAE and 21.74 RMSE in 3-shot case. Moreover, from 1-shot case to 3-shot case, our approach gains more performance improvement than two baseline methods, reflecting the superior ability of our SAFECount to utilize multiple support images.

4.3. Cross-dataset Generalization

Following FamNet [21], we test our model's generalization ability on the car counting dataset, CARPK [10]. The models are first pre-trained on FSC-147 [21] (the "car" category is excluded), then fine-tuned on CARPK. The results are shown in Tab. 4. For the models pre-trained on FSC-147, we significantly surpass FamNet by 42.23% in MAE and 45.85% in RMSE. When it comes to the finetuning scenario, our method still consistently outperforms all baselines. For instance, we surpass GMN [18] by 28.74% in MAE and 28.89% in RMSE. Therefore, our SAFECount has much better ability in cross-dataset generalization.

4.4. Ablation Study

To verify the effectiveness of the proposed modules and the selection of hyper-parameters, we implement extensive ablation studies on FSC-147 [21].

SCM and FEM. To demonstrate the effectiveness of the proposed SCM and FEM modules, we conduct diagnostic experiments. The substitute of FEM is to concatenate the similarity map and query feature together, then recover the feature dimension through a 1×1 convolutional layer. To invalidate SCM, we replace the score normalization to a naive maximum normalization (dividing the maximum value). The results are shown in Tab. 5a. Both SCM and FEM are necessary for our SAFECount. Specifically, when we replace FEM, the performance drops remarkably by 9.05 MAE in test set. This reflects that FEM is of vital significance in our SAFECount, since the core of our insight, *i.e.* feature enhancement, is completed in FEM. Besides, when we remove SCM, the performance also drops by 1.81 MAE in test set. This indicates that SCM derives a similarity map with a proper value range, promoting the performance. Score Normalization. We conduct ablation experiments regarding ENorm and SNorm in Tab. 5b. A naive maximum normalization (dividing the maximum value) serves as the baseline when both normalization methods are removed. Even if without score normalization, we still stably outperform all baselines in Tab. 2. Adding either ENorm or SNorm improves the performance greatly (≥ 4 MAE in test set), indicating the significance of score normalization. ENorm together with SNorm brings the best performance, reflecting that the two normalization methods could cooperate together for further performance improvement.

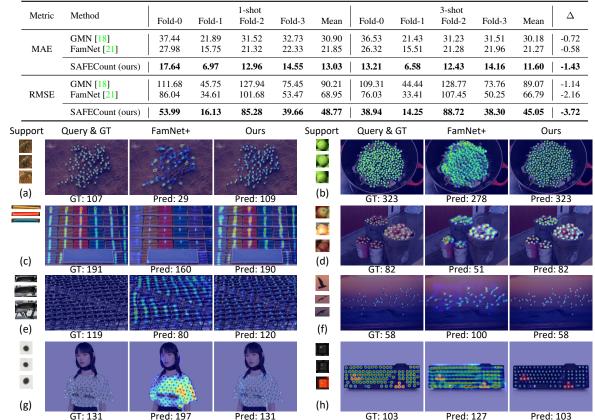


Table 3. Counting performance with cross-validation setting on FSC-147 dataset [21]. Fold-*i* (i = 0, 1, 2, 3) indicates the test set. Δ stands for the averaged improvement of the 3-shot case over the 1-shot case.

Figure 4. **Qualitative results** on the FSC-147 dataset [21] under the 3-shot case. From left to right: support images, query image overlaid by the ground-truth density map, predicted density map by FamNet+ [21], and our prediction. The numbers bellow are the counting results.

Table 4. **Cross-dataset generalization** on the car counting dataset CARPK [10]. The models are first pre-trained on FSC-147 [21] (the "car" category is excluded), then fine-tuned on CARPK.

	Method	MAE	RMSE
Pre-trained on FSC-147 [21]	GMN [18] FamNet [21]	32.92 28.84	39.88 44.47
[]	SAFECount (ours)	16.66	24.08
Fine-tuned on CARPK [10]	GMN [18] FamNet [21]	7.48 18.19	9.90 33.66
	SAFECount (ours)	5.33	7.04

Block Number. It is described in Sec. 3.3 that our SAFE-Count could be formulated to a multi-block architecture. Here we explore the influence of the block number. As shown in Tab. 5c, only 1-block SAFECount has achieved state-of-the-art performance by a large margin, which illustrates the effectiveness of our designed SAFECount architecture. Furthermore, the performance gets improved gradually when the block number increases from 1 to 4, and decreased slightly when the block number is added to 5. As proven in [9], too many blocks could hinder the training process, decreasing the performance. Finally, we set the block number as 4 for FSC-147.

Regressing from Similarity Map v.s. Enhanced Feature.

The density map can be regressed from either the enhanced feature or the similarity map. We compare these two choices in Tab. 5d. *Raw Similarity* is similar to FamNet [21] (without test-time adaptation), predicting the density map directly from the raw similarity. The rest 3 methods follow our design, where *i-block Similarity* and *i-block Feature* mean that the density map is regressed from the similarity map and enhanced feature of the *i*th block, respectively. Obviously, *I-block Feature* and *4-block Feature* significantly outperform *Raw Similarity* and *4-block Similarity*, respectively. The reason may be that the enhanced feature contains rich semantics and can filter out some erroneous high similarity values, *i.e.* the high similarity values that do not correspond to target objects, as proven in [29].

Comparison with Attention. In Tab. 5e, when the similarity derivation in SCM and the feature aggregation in FEM are replaced by an vanilla attention [30], the performance drops dramatically. As stated in Sec. 3.2, our method could better utilize the *spatial structure* of features than vanilla attention, which helps find more accurate boundaries between objects and brings substantial improvement.

Kernel Flipping in FEM. The kernel flipping in FEM could help the similarity-weighted feature, f_R , inherit the spatial

Table 5. **Ablation studies** on (a) the effect of the similarity comparison module (SCM) and feature enhancement module (FEM), (b) the score normalization in SCM, (c) the stacked number of our SAFECount block, and (d) the place to regress density map, (e) comparison with attention, (f) kernel flipping, (g) training or freezing backbone.

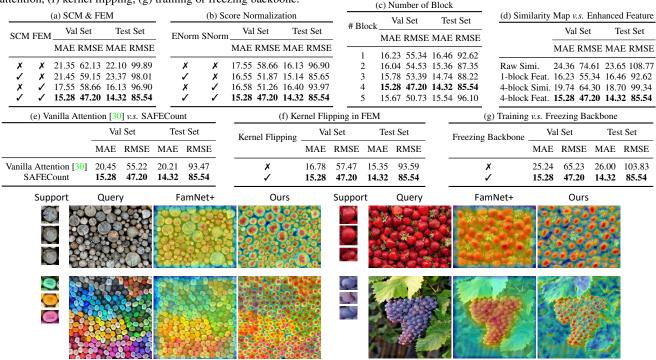


Figure 5. Visualization of the similarity maps developed by FamNet+ [21] and our SAFECount. Benefiting from the proposed SAFECount block, our approach recognizes *much clearer boundaries* between densely packed objects.

structure from the support feature, f_S (see *Supplementary Material* for details). The effectiveness of adding the flipping is proven by Tab. 5f. Adding the flipping could improve the performance stably (≥ 1 MAE), reflecting that preserving the spatial structure of f_R benefits the counting performance.

Training *v.s.***Freezing Backbone.** The comparison results are provided in Tab. 5g. The frozen backbone significantly surpasses the trainable backbone. Considering that the testing classes are different from training classes in FSC-147 [21], training backbone will lead the backbone to extract more relevant features to training classes, which decreases the performance in the validation and test sets.

4.5. Visualization

We visualize and compare the intermediate similarity map in FamNet [21] and SAFECount in Fig. 5, which intuitively explains why SAFECount surpasses FamNet substantially. Here the similarity map in SAFECount means the one in the last block. In FamNet, the similarity map is derived by direct comparison between the raw features of the query image and support images. However, the similarity map is far less informative than features, making it hard to identify clear boundaries within densely packed objects. In contrast, we weigh the support feature based on the similarity values, then integrate the similarity-weighted feature into the query feature. This design encodes the support-query relationship into features, while keeping the rich semantics extracted from the image. Also, our similarity comparison module is learnable. Benefiting from these, our SAFECount gets clear boundaries between densely packed objects in the similarity map, which is beneficial to regress an accurate count.

5. Conclusion

In this work, to tackle the challenging few-shot object counting task, we propose the similarity-aware feature enhancement block, composed of a similarity comparison module (SCM) and a feature enhancement module (FEM). Our SCM compares the support feature and the query feature to derive a score map. Then the score map is normalized across both the exemplar and spatial dimensions, producing a reliable similarity map. The FEM views these similarity values as weighting coefficients to integrate the support features into the query feature. By doing so, the model will pay more attention to the regions similar to support images, bringing distinguishable borders within densely packed objects. Extensive experiments on various benchmarks and training settings demonstrate that we achieve state-of-the-art performance by a considerably large margin.

Acknowledgement. This work is sponsored by the National Key Research and Development Program of China (2021YFB1716000) and National Natural Science Foundation of China (62176152).

References

- Carlos Arteta, Victor Lempitsky, and Andrew Zisserman. Counting in the wild. In *Eur. Conf. Comput. Vis.*, pages 483–498, 2016. 1, 2
- [2] Deepak Babu Sam, Shiv Surya, and R Venkatesh Babu. Switching convolutional neural network for crowd counting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5744–5752, 2017. 2
- [3] Antoni B Chan, Zhang-Sheng John Liang, and Nuno Vasconcelos. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1–7, 2008. 1, 5
- [4] Ke Chen, Chen Change Loy, Shaogang Gong, and Tony Xiang. Feature mining for localised crowd counting. In *Brit. Mach. Vis. Conf.*, page 3, 2012. 1, 5
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009. 5
- [6] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Fewshot object detection with attention-rpn and multi-relation detector. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4013–4022, 2020. 3, 5
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Modelagnostic meta-learning for fast adaptation of deep networks. In *Int. Conf. Mach. Learn.*, pages 1126–1135, 2017. 3, 6
- [8] Eran Goldman, Roei Herzig, Aviv Eisenschtat, Jacob Goldberger, and Tal Hassner. Precise detection in densely packed scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5227–5236, 2019. 1, 5
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. 5, 7
- [10] Meng-Ru Hsieh, Yen-Liang Lin, and Winston H Hsu. Dronebased object counting by spatially regularized regional proposal network. In *Int. Conf. Comput. Vis.*, pages 4145– 4153, 2017. 1, 2, 5, 6, 7
- [11] Ersin Kilic and Serkan Ozturk. An accurate car counting in aerial images based on convolutional neural networks. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–10, 2021. 1
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [13] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vazquez, and Mark Schmidt. Where are the blobs: Counting by localization with point supervision. In *Eur. Conf. Comput. Vis.*, pages 547–562, 2018. 1, 2
- [14] Bastian Leibe, Edgar Seemann, and Bernt Schiele. Pedestrian detection in crowded scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 878–885, 2005. 2
- [15] Victor Lempitsky and Andrew Zisserman. Learning to count objects in images. In Adv. Neural Inform. Process. Syst., pages 1324–1332, 2010. 2, 5
- [16] Jiang Liu, Chenqiang Gao, Deyu Meng, and Alexander G Hauptmann. DecideNet: Counting varying density crowds

through attention guided detection and density estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5197–5206, 2018. 2

- [17] Lingbo Liu, Jiaqi Chen, Hefeng Wu, Guanbin Li, Chenglong Li, and Liang Lin. Cross-modal collaborative representation learning and a large-scale rgbt benchmark for crowd counting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4823–4833, 2021. 1, 2
- [18] Erika Lu, Weidi Xie, and Andrew Zisserman. Class-agnostic counting. In Asian Conf. Comput. Vis., pages 669–684, 2018. 1, 2, 3, 5, 6, 7
- [19] Andreas Michel, Wolfgang Gross, Fabian Schenkel, and Wolfgang Middelmann. Class-aware object counting. In *IEEE Winter Conf. Appl. Comput. Vis.*, 2022. 2
- [20] Viresh Ranjan, Hieu Le, and Minh Hoai. Iterative crowd counting. In *Eur. Conf. Comput. Vis.*, pages 270–285, 2018. 2
- [21] Viresh Ranjan, Udbhav Sharma, Thu Nguyen, and Minh Hoai. Learning to count everything. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3394–3403, 2021. 1, 2, 3, 5, 6, 7, 8
- [22] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, pages 91–99, 2015. 5
- [23] Deepak Babu Sam and R Venkatesh Babu. Top-down feedback for crowd counting convolutional neural network. In Assoc. Adv. Artif. Intell., 2018. 2
- [24] Zenglin Shi, Le Zhang, Yun Liu, Xiaofeng Cao, Yangdong Ye, Ming-Ming Cheng, and Guoyan Zheng. Crowd counting with deep negative correlation learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5382–5390, 2018. 2
- [25] Vishwanath A Sindagi and Vishal M Patel. Generating highquality crowd density maps using contextual pyramid cnns. In *Int. Conf. Comput. Vis.*, pages 1861–1870, 2017. 2
- [26] Qingyu Song, Changan Wang, Zhengkai Jiang, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yang Wu. Rethinking counting and localization in crowds: A purely point-based framework. In *Int. Conf. Comput. Vis.*, pages 3365–3374, 2021. 1, 2
- [27] Tobias Stahl, Silvia L Pintea, and Jan C Van Gemert. Divide and count: Generic object counting by image divisions. *IEEE Trans. Image Process.*, pages 1035–1044, 2019. 2
- [28] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2325–2333, 2016.
- [29] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8934–8943, 2018. 7
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inform. Process. Syst.*, pages 5998–6008, 2017. 4, 7, 8
- [31] Jia Wan, Ziquan Liu, and Antoni B Chan. A generalized loss function for crowd counting and localization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1974–1983, 2021. 2

- [32] Jia Wan, Wenhan Luo, Baoyuan Wu, Antoni B Chan, and Wei Liu. Residual regression with semantic prior for crowd counting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4036–4045, 2019. 2
- [33] Meng Wang and Xiaogang Wang. Automatic adaptation of a generic pedestrian detector to a specific traffic scene. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3401–3408, 2011.
 2
- [34] Davis Wertheimer, Luming Tang, and Bharath Hariharan. Few-shot classification with feature map reconstruction networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8012–8021, 2021. 3, 5
- [35] Jiaxi Wu, Songtao Liu, Di Huang, and Yunhong Wang. Multiscale positive sample refinement for few-shot object detection. In *Eur. Conf. Comput. Vis.*, pages 456–472, 2020. 3
- [36] Feng Xiong, Xingjian Shi, and Dit-Yan Yeung. Spatiotemporal modeling for crowd counting in videos. In *Int. Conf. Comput. Vis.*, pages 5151–5159, 2017. 2
- [37] Wei Xu, Dingkang Liang, Yixiao Zheng, Jiahao Xie, and Zhanyu Ma. Dilated-scale-aware category-attention convnet for multi-class object counting. *IEEE Signal Processing Letters*, 2021. 2
- [38] Zhaoyi Yan, Yuchen Yuan, Wangmeng Zuo, Xiao Tan, Yezhen Wang, Shilei Wen, and Errui Ding. Perspective-guided convolution networks for crowd counting. In *Int. Conf. Comput. Vis.*, pages 952–961, 2019. 2
- [39] Lihe Yang, Wei Zhuo, Lei Qi, Yinghuan Shi, and Yang Gao. Mining latent classes for few-shot segmentation. In *Int. Conf. Comput. Vis.*, pages 8721–8730, 2021. 3, 5
- [40] Shuo-Diao Yang, Hung-Ting Su, Winston H Hsu, and Wen-Chin Chen. Class-agnostic few-shot object counting. In *IEEE Winter Conf. Appl. Comput. Vis.*, pages 870–878, 2021. 1, 2, 3, 5, 6
- [41] Yifan Yang, Guorong Li, Zhe Wu, Li Su, Qingming Huang, and Nicu Sebe. Reverse perspective network for perspectiveaware object counting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4374–4383, 2020. 2
- [42] Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. Cross-scene crowd counting via deep convolutional neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 833–841, 2015. 2
- [43] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. CANet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5217–5226, 2019. 3
- [44] Qi Zhang, Wei Lin, and Antoni B Chan. Cross-view crossscene multi-view crowd counting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 557–567, 2021. 1, 2
- [45] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 589–597, 2016. 1, 2, 5