

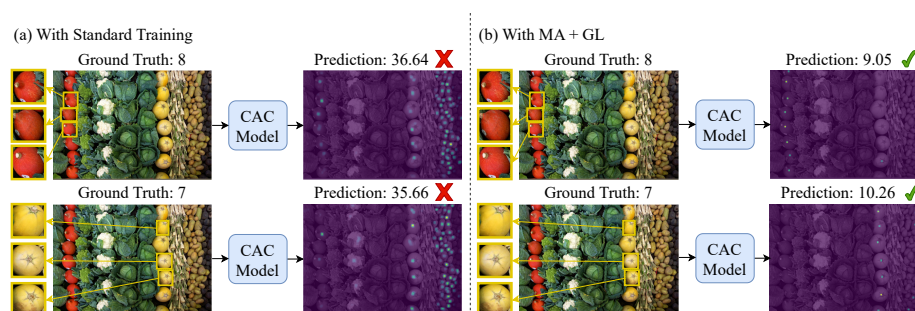
# A Recipe for CAC: Mosaic-based Generalized Loss for Improved Class-Agnostic Counting

Tsung-Han Chou<sup>1</sup>, Brian Wang<sup>2</sup>, Wei-Chen Chiu<sup>1</sup>, and Jun-Cheng Chen<sup>3</sup>

<sup>1</sup> National Yang Ming Chiao Tung University, Taiwan

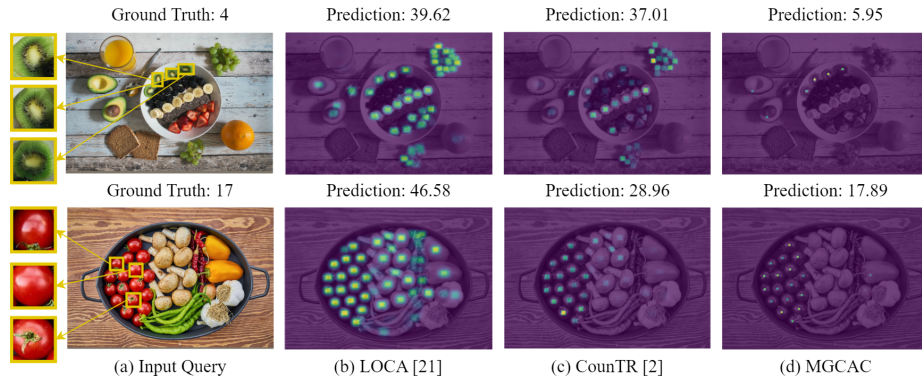
<sup>2</sup> University of Michigan, Ann Arbor

<sup>3</sup> Research Center of Information Technology Innovation, Academia Sinica, Taiwan



**Fig. 1: Comparison of CAC Models with Different Training Strategies in Real-World Images.** Given the same query image with different references (annotated by yellow bounding boxes), (a) the state-of-the-art CAC models (e.g., BMNet [18], LOCA [21]) adopting the standard training setting upon the FSC-147 dataset fail to distinguish the appearance of objects and gives wrong counts (i.e. the model does not fire at the objects of the same class as references on the query image, but the objects of majority. Noting that the dilation is applied on the density maps for best viewing). Conversely, (b) our suggested recipe, which utilizes novel technique of integrating Mosaic Augmentation (MA) and Generalized Loss (GL), enables the models to effectively discriminate different objects, thereby resulting in significantly better count prediction.

**Abstract.** Class agnostic counting (CAC) is a vision task that can be used to count the total occurrence number of any given reference objects in the query image. The task is usually formulated as a density map estimation problem through similarity computation among a few image samples of the reference object and the query image. In this paper, we point out a severe issue of the existing CAC framework: Given a multi-class setting, models don't consider reference images and instead blindly match all dominant objects in the query image. Moreover, the current evaluation metrics and dataset cannot be used to faithfully assess the model's generalization performance and robustness. To this end, we discover that the combination of mosaic augmentation with generalized



**Fig. 2: More qualitative results of various CAC models upon real images.** The existing state-of-the-art CAC models such as (b) LOCA [21] and (c) CountTR [2] have problems on either counting all objects in the query image (and completely disregarding the target objects showcased in reference images) or unintentionally detecting irrelevant objects. In contrast, our proposed MGCAC model, featuring the multi-class training scenario and the more feasible objective to CAC task, faithfully follows the guidance from the references to achieve superior results of counting.

loss is essential for addressing the aforementioned issue of CAC models to count objects of majority (i.e. dominant objects) regardless of the references. Furthermore, we introduce a new evaluation protocol and metrics for resolving the problem behind the existing CAC evaluation scheme and better benchmarking CAC models in a more fair manner. Besides, extensive evaluation results demonstrate that our proposed recipe can consistently improve the performance of different CAC models. The code is available at <https://github.com/littlepenguin89106/MGCAC>.

## 1 Introduction

Object counting is a popular research topic in the vision community with a wide spread of applications, including visual surveillance, intelligent agriculture, etc. It aims to count the occurrence number of target objects in an image. Object counting methods can be classified into two major categories: class-specific object counting and class-agnostic counting. Class-specific object counting usually focuses on counting a specific category such as cars, animals, people, etc, whereas among these applications, crowd counting, which aims to count the number of people in an image, is well studied by [3, 4, 10, 19]. However, class-specific counting requires training an individual model for each category with tremendous efforts on collecting thousands of training images with annotations and fails to work for unseen classes.

In contrast, class-agnostic counting (CAC) studied by [13, 17, 23] arises recently and aims to count any novel objects within the query image, especially for those objects of unseen classes during the training stage. Ideally, CAC should consider the appearance of  $K$  reference images and count occurrences of similar objects in the query image. Nevertheless, in the existing evaluation protocol, query images typically only contain objects from a single class (even with clean backgrounds), which significantly differ from real-world multi-class cases. By our examination of real-world examples, we observed that most prior works degenerate back to blindly capture objects without considering the references (Figure 1(a) and Figure 2). We term the problem as Reference Oversight (RO). To verify RO, we propose a novel multi-class mosaic evaluation dataset for CAC.

Secondly, we disclose that current benchmarking protocols are not fair indications of model generalization. Given FSC-147’s greatly skewed per-image-object-count distribution (noting that FSC-147 [17] is a popular and representative dataset of CAC), we show that evaluation metrics such as Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) tend to strongly encourage models to only optimize counting on images with high object counts while overlooking the remaining images. To ensure fair and comprehensive comparisons, we also incorporate metrics such as Normalized Absolute Error (NAE) and Squared Relative Error (SRE).

Through in-depth analysis, we pinpoint the root causes of RO problem: the lack of variety of classes in each query image during the training stage incapacitates the model’s need to differentiate among object classes. Additionally, using pixel-wise losses (e.g., MSE) fails to accurately localize target objects and still contributes to the model’s inability to differentiate objects effectively. To address these challenges, we propose the Mosaic-based Generalized-Loss (MG) recipe for CAC, which consistently allows different CAC models to effectively capture variations in reference appearances and precisely localize target objects as illustrated in Figure 1(b) and demonstrated by our extensive experimental results. Finally, we also provide a strong CAC baseline by exploiting more advanced cross-attention modules from [5] as the backbone with our proposed MG loss recipe for training, denoted as MGCAC. The proposed MGCAC achieves the state-of-the-art results for CAC upon the original FSC-147 [17] and our proposed mosaic evaluation dataset.

Our contributions are summarized as follows: Firstly, our research is pioneering in exposing the performance gap of CAC methods between the existing evaluation protocol and real-world, multi-class scenarios. Our proposed mosaic evaluation dataset emphasizes RO problem on all existing CAC models, even the state-of-the-art methods (e.g. LOCA [21] or CounTR [2]) have such critical flaws (Figure 2(b) and 2(c)). Secondly, to improve the robustness and localization of CAC, we introduce MG, a recipe that proves generalizable to various CAC models, and build a strong baseline, MGCAC. Our work underscores the importance of effective detection and discrimination of multi-class objects and hopes to reveal the “devil” hidden under the details of the current CAC framework.

## 2 Related Works

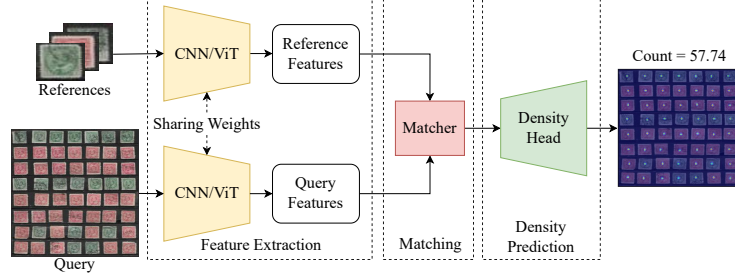
### 2.1 Class-specific Object Counting

Class-specific counting focuses on the task of counting the occurrence number of the objects of specific classes, such as crowd counting. In addition, the task is usually approached in two paradigms: detection-based and regression-based methods. The detection-based methods, including the methods proposed by [7,9], perform explicit object detection over the input image using visual object detectors and then get the count. However, similar to the object detection methods, their performances are also sensitive to situations when the objects are overlapped, occluded, or crowded. To address these problems, the regression methods proposed by [3,14,20] instead predict the density map of input images where each pixel value can be interpreted as the fraction or the confidence level of the target object present in the query image. The sum of these values is then used as the estimated object count. In addition, the ground truth density maps for training are generated by convolving point annotations of the training images with properly selected Gaussian kernels. Although class-specific counting methods have achieved satisfactory results, they usually require training an individual network for each class of objects and can only be applied for counting the classes that have been seen in the training data.

### 2.2 Class-Agnostic Counting

Concerning class-specific object counting, class-agnostic counting aims to perform object counting for every class. In its setting, the models usually take a query image and several reference images from the same class as inputs and predict the count of the class that appears in the query image through the density map regression. Lu et al. [13] propose GMNet which is the first CAC framework, where the reference and query feature maps are extracted independently through the feature extractor based on the ResNet [6]. These features are then directly exploited as a pixel-wise regression task to perform object counting. To better capture the interactions between the query and the references during the matching stage, Ranjan et al. [17] propose FamNet which introduces template matching for the CAC task. FamNet enhances the matching framework by convolving the reference feature maps across query feature maps. Each pixel value in the resulting query feature map represents the similarity between the query image and reference images at that specified location and therefore supports better localization. Building upon the existing structures, CFOCNet proposed by Yang et al. [23] matches query and references feature maps from different stages of backbone for multi-scale strategy. Moreover, CFOCNet integrates self-attention to strengthen the distinct features of the query image. Shi et al. [18] propose BMNet which adopts bilinear similarity for matching. The bilinear similarity is a special case of generalized inner product and provides more flexibility for matching than previous methods using convolution. SPDCN [12] introduces Scale-Prior Deformable Convolution and Scale-Sensitive Generalized Loss to enhance the





**Fig. 3: CAC Framework.** The general pipeline of CAC models consists of a feature extractor, matcher, and density head. Given a query image and  $K$  reference images (i.e.,  $K = 3$  in this diagram), the model learns to predict a density map for counting.

model scale-awareness. CounTR proposed in [2] introduces transformer-based architecture into class-agnostic counting, which utilizes attention to capture the similarity between the query and the references. Djukic et al. [21] propose a low-shot object counting network with iterative prototype adaptive (LOCA) which introduces the object prototype extraction module (OPE). The OPE module focuses on constructing strong object prototypes to improve object localization accuracy and iteratively adapts shape information and object appearances into object prototypes by recursively performing cross-attention operations.

With the progress of multi-modality, recent works [1] and [8] adopt CLIP [16] model to perform CAC with input prompts. Given a class description and a query, by matching the extracted text features and the query features, the objects that match the description in the query will be identified.

### 3 Deficiencies of Current CAC Training and Evaluation

In this section, we first introduce the general CAC framework. After that, we outline the blindly counting problem and the biases of the prevalent benchmark. We provide our solutions to resolve the problem in Section 5 and show their effectiveness in Section 6.

#### 3.1 CAC Framework

As shown in Figure 3, standard CAC models comprise three components: (1) a feature extractor for visual feature extraction from reference and query images, (2) a cross-matching module for similarity map computation, and (3) a density map estimator for counting. Given a query image  $X \in \mathbb{R}^{H_X \times W_X \times 3}$ ,  $K$  reference images  $Z = \{z_i\}_{i=1}^K$ ,  $z_i \in \mathbb{R}^{H_{z_i} \times W_{z_i} \times 3}$ , and a ground truth label  $Y \in \mathbb{R}^{H_X \times W_X}$  (in the form of density map), the CAC model minimizes the regression problem to predict the number of the references occurring in the query image:

$$\min_{\theta, \phi, \psi} \mathbb{E}_{(X, Z, Y)} [\mathcal{L}(R_\theta(S_\phi(F_\psi(X), F_\psi(Z))), Y)], \quad (1)$$

where  $F_\psi(\cdot)$  is a pre-trained feature extractor,  $S_\phi(\cdot, \cdot)$  computes the similarity map,  $R_\theta(\cdot)$  is the regressor or density map estimator, and  $\mathcal{L}(\cdot, \cdot)$  is the loss function (e.g., pixel-wise L2 loss).

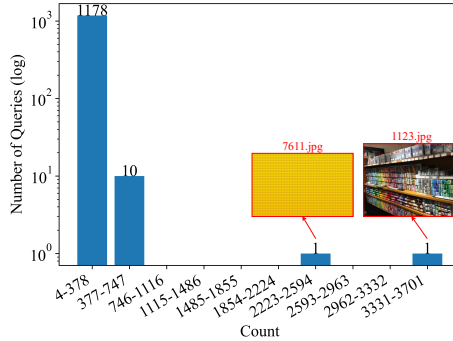
### 3.2 Reference Oversight

A robust CAC model should accurately count and localize objects corresponding to reference images in diverse query scenarios. However, most models are evaluated on widely used CAC datasets (e.g., FSC-147 [18] and CARPK [7]), which primarily feature a single-class. Such datasets lack real-life variability in object distributions and backgrounds. This setting raises concerns about the model’s ability to precisely match target objects in multi-class scenarios. As shown in Figure 2, we observe that most CAC models, including the current state-of-the-art (SOTA) LOCA, tend to annotate all dominant objects in query images, neglecting the similarity between objects in query and reference images. In contrast, CounTR leverages mosaic augmentation and therefore better identifies target objects. However, CounTR fails to calibrate the correct count for reference images due to inadequate localization information. With regards to supervised CAC models, we hypothesize that: (1) The dominant issues are caused by training with queries dominated by objects of a single class, which CounTR uses mosaic augmentation that alleviates the issue (2) Supervision with pixel-wise L2 loss is not sensitive to localization on the object-level and leads to unsatisfied performance for CounTR.

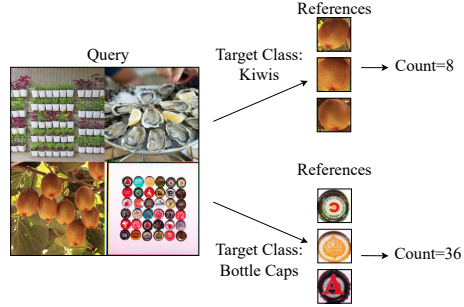
### 3.3 Benchmarking

A well-designed evaluation protocol should encompass diverse scenarios and utilize fair metrics. However, existing CAC datasets, as discussed in Section 3.2, often lack diversity of categories within a single query. Conversely, detection datasets like COCO [11] include multiple classes but predominantly have a small count of objects per image, failing to comprehensively assess the model’s ability to identify corresponding objects in queries.

Furthermore, the current CAC evaluation protocol evaluates model counting performance by aggregating counting errors across all images. Given groundtruth count  $c_l$  and predicted count  $\tilde{c}_l$  for the  $l$ th image in the dataset with  $L$  images,  $\text{MAE} = \frac{1}{L} \sum_{i=1}^L |\tilde{c}_l - c_l|$  and  $\text{RMSE} = \sqrt{\frac{1}{L} \sum_{i=1}^L (\tilde{c}_l - c_l)^2}$  are usually used. The challenge arises from the long-tailed distribution of per-image object counts, introducing bias towards large-count query images. Figure 4 illustrates that only two images featuring large counts in FSC-147 test dataset. By excluding the first two highest object-count query images from RMSE calculations on the FSC-147 test dataset (Table 1), we show significant impacts on the evaluation results (details in Appendix). Recently SOTA models (i.e., LOCA and CounTR) get worse RMSE than BMNet+ when removing large-count queries. Those phenomena reinforce the need for an evaluation strategy that accounts for the distributional characteristics of object counts, especially in datasets with varying count sizes.



**Fig. 4: Distribution of object count in FSC-147 test dataset.** We split the test dataset into 10 bins. The x-axis is the count range of the corresponding bin, y-axis is the number of queries in log-scale.



**Fig. 5: Mosaic evaluation.** Given different references of the target class, the objective is to count the corresponding objects in the query.

**Table 1: Comparison of the performance of CAC models on the full FSC-147 test dataset and the dataset without outliers (i.e., 1123.jpg and 7611.jpg).**

(a) Full test set.						(b) Test set after excluding outliers					
	MAE	RMSE	NAE	SRE			MAE	RMSE	NAE	SRE	
BMNet+ [18]	14.62	91.83	0.27	6.20			12.04	30.13	0.27	6.04	
CounTR [2]	11.95	91.23	0.23	7.44			9.93	33.61	0.23	6.47	
LOCA [21]	10.79	56.97	0.19	2.19			9.46	33.03	0.18	1.94	

## 4 Proposed Evaluation Protocol

### 4.1 Mosaic Evaluation Dataset

Currently, there is a lack of real-world multi-class counting datasets for CAC evaluation. We leverage the FSC-147 dataset as a foundation to create a multi-class setting. As shown in Figure 5, for a synthesized mosaic collage, we randomly select 4 query images of distinct classes from the validation and test sets of FSC-147. These images are then randomly cropped to the size of  $384 \times 384$  and put together as size  $768 \times 768$ . Each class in the collage will individually serve as the reference class. Given the associated reference images, the model has to accurately predict the number of target objects for the collage (i.e. predicted count should equate to the ground truth count of the query image containing the target class). In total, we generate 1,000 queries (i.e., 4,000 data pairs), facilitating a comprehensive evaluation of the object recognition capabilities within the context of mosaic collages. We name this dataset as FSC-Mosaic, which stands for Few-Shot Counting Dataset with Mosaic Evaluation.

## 4.2 Additional Metric

As discussed in Section 3.3, using MAE and RMSE is heavily influenced by query images with large object counts. For balancing the contribution of each image to the performance score, we additionally benchmark with  $\text{NAE} = \frac{1}{L} \sum_{l=1}^L \frac{|\tilde{c}_l - c_l|}{c_l}$  and  $\text{SRE} = \sqrt{\frac{1}{L} \sum_{l=1}^L \frac{(\tilde{c}_l - c_l)^2}{c_l}}$ , which provide a better reflection of the model’s overall performance on the entire dataset. As shown in Table 1, the use of NAE and SRE provides a more stable assessment of model performance on datasets with skewed object count distributions.

## 5 Proposed MG Recipe for CAC Training

To improve the class discrimination of CAC models and their matching performance on targets, we tackle it in two ways: (1) exposing multi-class setting through mosaic augmentation, and (2) leveraging Generalized Loss (GL) that is more localization-aware and robust. We further elucidate how the combination of these strategies improves the RO problem.

### 5.1 Mosaic Augmentation

The main issue for the CAC domain is its single-class per-image data setting, which hinders the model’s generalizability under multi-class setting. To directly transform single-class setting to multi-class setting, we apply mosaic augmentation (MA) on FSC-147 images. We construct it by (1) randomly selecting 4 images, each randomly cropping to dimensions  $H_Z \times W_Z$ , (2) arranging it into  $2H_Z \times 2W_Z$  mosaiced image. During training, we randomly select 1 out of the  $M$  reference classes in the mosaic for prediction. As compared to the original single-class image, the model is now exposed to  $M$  object classes at once, therefore acquiring the capability to distinguish among different class objects while recognizing reference objects that slightly vary in appearance. Our experiments show the non-trivial impact of this data augmentation technique that greatly strengthens the model’s robustness on both our mosaic evaluation protocol and real-life multi-class scenarios.

### 5.2 Generalized Loss

Currently, most CAC approaches based on density map estimation employ the pixel-wise  $L_2$  loss (or MSE loss) to train the models where the ground truth density map is generated by converting the point annotations of each training image through convolving with a Gaussian kernel. However,  $L_2$  loss only offers pixel-wise supervision. It penalizes the same no matter whether the predictions are near or far from the ground-truth dot annotation. This makes  $L_2$  loss inappropriate for accurate localization. Recently, Wan et al. [22] propose a Generalized Loss (GL) that uses unbalanced optimal transport (OT) to measure the transport cost between the predicted density map and the ground truth point

annotations directly. GL penalizes more when the predictions are farther from the ground truth. This suffices the need for precise localization for single and multi-class settings. In addition, it also has shown that the  $L_2$  loss is a special case of GL and usually results in suboptimal solutions. Thus, we mainly exploit GL to train the proposed models shown as follows:

$$L_C^T = \min_P \langle C, P \rangle - \varepsilon H(P) + \tau \|P \mathbf{1}_m - a\|_2^2 + \tau \|P^\top \mathbf{1}_n - b\|_1, \quad (2)$$

where  $C$  is the transport cost of moving predicted density to ground truth dot annotation,  $P$  is the corresponding transport plan,  $H(\cdot)$  is the entropic regularization term,  $n$  is the number of pixels,  $m$  is the number of annotation points,  $a$  is the predicted density map, and  $b$  is the ground truth dot map. Furthermore, to better encode perspective information, it proposes the Perspective-Guided Transport Cost:

$$C_{ij} = \exp\left(\frac{1}{\eta(x_i, y_i)} \|x_i - y_j\|_2\right), \quad (3)$$

where  $\eta(x_i, y_i)$  is a adaptive perspective factor. We simply choose a fixed  $\eta$  for our experiment.

### 5.3 Combination of MA and GL

While both MA and GL are not new in counting tasks, previous studies have not fully exploited the potential of MA and GL in multi-class scenarios. For instance, CounTR utilizes MA for resolving long-tailed problems (i.e. few images with large object counts), which differs from our intention of creating multi-class setting. Similarly, SPDCN uses GL for better localization in single-class setting, which can not distinguish different classes. By combining MA and GL, models have stronger awareness of reference locations, allowing them to discern the appearance of reference objects in multi-class settings. Furthermore, they also generalize well to count unseen real-world multi-class images as illustrated in Figure 1(b). In the following experiment, we show the necessity of our approach in bridging the domain gap between single and multi-class settings.

## 6 Experiment

In this section, we present the details of the experimental setup. Then, we show extensive evaluation results and ablation studies of the proposed training recipe on the widely used CAC benchmarks. We further evaluate CAC models on our proposed evaluation protocol to examine CAC models' ability to localize reference targets under multi-class scenarios.

### 6.1 Implementation Details

**Datasets.** We utilize three datasets to evaluate our model performance: (1) FSC-147 dataset [17], a standard class-agnostic object counting dataset composed of

147 categories and 6,135 images across different scales, in which each image is annotated with three reference objects. Given that the dataset contains decent annotations, we use the FSC-147 dataset to train and evaluate our models. (2) We evaluate FSC-147-trained models on our proposed FSC-Mosaic dataset to validate their generalizability in multi-class scenarios. (3) Moreover, we evaluate FSC-147-trained models on CARPK [7], a dataset that contains 1,448 photos of cars in parking lots from a bird-view perspective, to ensure the generalizations of our models on unseen datasets (see Appendix).

**Evaluation Metrics.** To compare performance against existing CAC models, we preserve the previous evaluation protocol [17] based on MAE and RMSE. We additionally benchmark with NAE and SRE to ensure fair consideration of each image in the evaluation process.

**Model Comparison.** We utilize MixFormer [5] architecture design with a multi-scale density head to better localize features of different levels as our baseline, denoted as MGCAC as mentioned in Section 1. (Please refer to Appendix for more detailed descriptions.) It is worth noting that our training recipe can be applied to other CAC model architectures as well. To demonstrate the effectiveness and general applicability of our proposed training recipe, we compare BMNet+, CounTR, and LOCA with their original training approach against our proposed recipe on the FSC-Mosaic dataset.

**Data Preprocessing.** For our baseline training preprocessing, query images are (1) randomly flipped horizontally, and (2) randomly cropped to the same size of  $384 \times 384$ . When applying mosaic augmentation, 4 images with different classes are randomly cropped to  $192 \times 192$ , so the synthesized mosaic collage keeps the same image size in a batch for efficient training. We randomly select 1 out of the 4 reference classes as the target, using the corresponding references and synthetic mosaic collage as input, with the corresponding annotations as ground truth. During the evaluation stage, the processing of reference images is similar to training, but query images are fully used.

**Optimizer and Training Hyperparameters.** Our model is trained with 300 epochs, batch size of 8, optimized with AdamW optimizer, and weight decay of  $5e-4$ . For MGCAC, the backbone learning rate is  $1e-5$ , and the non-backbone learning rate is  $1e-4$ . We use PyTorch [15] as the implementation framework. We use GL to supervise the final output that is  $\frac{1}{4}$  of the original query image size. The hyperparameters of GL are set as  $\varepsilon = 0.01$ ,  $\tau = 0.5$ , and  $\eta = 0.6$ .

**Inference time.** The running time of our MGCAC takes 38ms for one evaluation on average using a single NVidia RTX 3090 and an Intel i7-8700 CPU.

## 6.2 Evaluation Results with Other SOTAs

**Evaluation on Existing FSC-147 Setting.** As shown in Table 2, our model surpasses most of models on NAE and SRE, which are metrics that indicate counting errors across images, unweighted by object counts. As compared to the SOTA model LOCA, our performance yields an improvement of 45.4% on VAL NAE, 1.9% on VAL SRE, and 15.7% on TEST NAE. These results indicate that our model performs better on average across each image. When considering the

**Table 2:** Comparisons with SOTA CAC models on the FSC-147 dataset. Notation ‘\*’ indicates that we further adopt the test-time normalization [2] during inference.

Method	VAL (validation)				TEST			
	MAE↓	RMSE↓	NAE↓	SRE↓	MAE↓	RMSE↓	NAE↓	SRE↓
GMN [13]	29.66	89.81	—	—	25.52	124.57	—	—
FamNet [17]	23.75	69.07	0.51	4.24	22.08	99.54	0.44	6.45
CFOCNet [23]	21.19	61.41	—	—	22.10	112.71	—	—
BMNet+ [18]	15.74	58.53	0.25	2.73	14.62	91.83	0.27	6.20
CounTR [2]	13.13	49.83	0.23	2.59	11.95	91.23	0.23	7.44
SAFECount [24]	15.28	48.20	0.26	3.99	14.32	85.54	0.25	6.37
SPDCN [12]	14.59	49.97	0.22	2.79	13.51	96.80	0.22	6.70
LOCA [21]	10.24	<b>32.56</b>	0.22	2.09	10.79	56.97	0.19	<b>2.19</b>
<b>MGCAC (Ours)</b>	11.00	51.42	<b>0.12</b>	2.05	10.46	96.60	<b>0.16</b>	6.17
<b>MGCAC* (Ours)</b>	<b>9.93</b>	41.08	<b>0.12</b>	<b>1.91</b>	<b>9.10</b>	<b>54.21</b>	<b>0.16</b>	5.68

existing metrics MAE and RMSE, our MAE is comparable to LOCA’s while having a gap in RMSE. As shown in Section 3.3, RMSE is not an indicative metric of model performance. We also demonstrate MGCAC\* that intentionally overfits and optimizes the performance on a few large-count images by adopting test-time normalization from [2], which acquires SOTA on the FSC-147 dataset. This validates the unbalanced contribution of RMSE in large-count images and our model superiority on non-outlier images. In the next subsection, we discuss the evaluation of multi-class settings, where the model’s awareness of reference context is challenged.

**Mosaic Evaluation.** Seen from Table 3, models under the CAC typical training process, such as BMNet+ and LOCA, exhibit unsatisfactory performance under mosaic evaluation. The introduction of mosaic augmentation, as demonstrated by CounTR, yields a better performance. After employing our training recipe, BMNet+, CounTR, and LOCA exhibit substantial error decreases in all metrics. Specifically, BMNet+ yields an improvement of 47.5% in RMSE and over 75% across the remaining three metrics. These improvements demonstrate the importance and efficacy of applying the combination of mosaic augmentation and generalized loss on CAC models in discriminating reference classes and accurately matching targets. Notably, MGCAC achieves SOTA on the FSC-Mosaic dataset with 37% improvement on all metrics against all models, showing our superiority in counting target objects.

**Real-world examples.** To verify whether the CAC models are general for real-world scenarios, we test them with some real-world examples as well. As shown in Figure 2(d), MGCAC excels in accurately localizing target objects and providing relatively precise prediction counts. Notably, as shown in Figure 7, LOCA and CounTR, trained with our MG recipe, efficiently locate target objects and mitigate the RO problem, validating the effectiveness and generability of our training recipe.

**Table 3:** Evaluation results on the FSC-Mosaic dataset of the CAC models trained with the proposed training recipe (**Red**) against those using their original training recipe where we use their publicly available pre-trained models (**Green**).

MA	GL	Model	MAE	RMSE	NAE	SRE
×	×	<b>FamNet</b> [17]	37.09	66.63	1.65	10.57
×	×	<b>BMNet+</b> [18]	81.69	101.75	5.28	24.73
×	×	<b>LOCA</b> [21]	42.24	67.19	2.45	13.33
✓	×	BMNet+ [18]	24.38	52.81	1.17	7.81
✓	×	<b>CounTR</b> [2]	26.07	65.04	0.51	4.55
×	✓	<b>SPDCN</b> [12]	40.60	85.38	0.99	6.36
✓	✓	<b>BMNet+</b> [18]	20.51	53.37	0.58	5.65
✓	✓	<b>CounTR</b> [2]	23.77	63.44	0.43	4.40
✓	✓	<b>LOCA</b> [21]	27.69	64.36	0.64	4.66
✓	✓	<b>MGCAC (Ours)</b>	<b>8.58</b>	<b>33.28</b>	<b>0.17</b>	<b>1.93</b>

### 6.3 Ablation Studies

We perform ablation studies on MGCAC to examine the functionality of our proposed framework in the FSC-147 dataset and FSC-Mosaic dataset. Based on the results, we have the following observations:


**Mosaic Augmentation.** In Table 4a, comparing S1 and S2, the introduction of mosaic augmentation (S2) worsens the performance on FSC-147 dataset. In contrast, S2 gets significant improvement on the FSC-Mosaic dataset (Table 4b). From Figure 6(c), we observe that mosaic augmentation allows S2 to focus on the correct region. Nevertheless, S2 still assigns density values on irrelevant objects and gives inaccurate counts.

**Generalized Loss.** With only applying GL (comparing S1 and S3), although S3 gets huge improvement on FSC-147 dataset (Table 4a), its performance becomes worse on FSC-Mosaic dataset (Table 4b). As shown in Figure 6(d), although S3 captures the target objects, it also includes irrelevant objects. These observations indicate that while the use of GL enhances the localization ability, it inadvertently leads to model overfitting in single-class settings, hindering its effectiveness in differentiating classes.


**Combination of MA and GL.** Comparing S2 and S4 quantitatively, the introduction of GL yields a relative improvement of 10% in MAE, 26% in NAE, and 7% in SRE for our MGCAC. By combining MA and GL, the model has stronger supervision under multi-class setting, allowing them to discern the appearance of reference objects. The large improvement in object localization is shown in Figure 6(e). While the count is not accurate, the crucial point is that S4 effectively captures only the target objects, excluding irrelevant ones.

**Multi-scale Feature with U-Net-like Fusion.** Comparing S4 and S5 quantitatively, the incorporation of multi-scale feature aggregation (i.e. the multi-scale features from the feature extractor and the matcher are fused via a U-Net-like density head) allows MGCAC to be competitive on the FSC-147 dataset while still demonstrating its ability to distinguish reference classes on the FSC-





**Fig. 6: Ablation of MGCAC (cf. Section 6.3).** We visualize the predicted density map results of our MGCAC on the real image. The top-left corner is the input query image with references (annotated with bounding boxes), while its caption provides information about the number and the class of target objects.



**Fig. 7: Comparison of CAC models with original training strategy (official pre-trained models) and our proposed MG training recipe.**

Mosaic dataset. The predicted count closely aligns with the ground truth in Figure 6(f).

## 7 Conclusion

In this work, we present a Mosaic-based Generalized Loss recipe that resolves the CAC model’s failure to distinguish among reference classes. We analyze the skewed distribution phenomenon that hides under the optimization of existing metrics of MSE/RMSE. With our proposed evaluation protocol, we can effectively quantify the model’s ability to learn discriminative reference features and accurately locate target objects in a multi-class setting. We also utilize different metrics (NAE, SRE) to represent the model’s overall performance across images. Our mosaic augmentation and supervision of generalized loss have been

**Table 4:** The ablation results of the proposed MGCAC on FSC-147 and FSC-Mosaic dataset, where ‘MA’ denotes using mosaic augmentation, ‘GL’ denotes replacing the MSE loss with the generalized loss, ‘MS’ means that the multi-scale features from the feature extractor and the matcher are fused via a U-Net-like density head.

(a) Ablation results on FSC-147 dataset.

Setting	MA	GL	MS	VAL				TEST			
				MAE	RMSE	NAE	SRE	MAE	RMSE	NAE	SRE
S1	×	×	×	14.83	63.70	0.17	2.60	14.30	101.89	0.22	7.39
S2	✓	×	×	16.38	70.93	0.18	2.83	16.16	116.32	0.23	6.89
S3	×	✓	×	13.23	62.15	0.13	2.42	12.29	100.75	0.18	6.51
S4	✓	✓	×	14.68	67.40	0.13	2.62	13.61	119.97	0.17	<b>6.10</b>
S5	✓	✓	✓	<b>11.00</b>	<b>51.42</b>	<b>0.12</b>	<b>2.05</b>	<b>10.46</b>	<b>96.60</b>	<b>0.16</b>	6.17

(b) Ablation results on FSC-Mosaic dataset.

Setting	MA	GL	MS	MAE	RMSE	NAE	SRE
S1	×	×	×	37.42	70.58	2.22	15.94
S2	✓	×	×	12.17	41.23	0.35	2.79
S3	×	✓	×	45.03	76.91	2.88	17.28
S4	✓	✓	×	10.45	40.40	0.17	2.24
S5	✓	✓	✓	<b>8.58</b>	<b>33.28</b>	<b>0.17</b>	<b>1.93</b>

empirically proven to significantly improve CAC models’ robustness. With these contributions, we hope to set up a better standard for CAC domain that reflects model generalizability and robustness in real-life settings.

## 8 Acknowledgements

This research work is supported in part by National Science and Technology Council (NSTC) under grants 112-2222-E-001-001-MY2, 112-2634-F-002-006, 112-2221-E-A49-087-MY3, 111-2628-E-A49-018-MY4, and by Academia Sinica under grant AS-CDA-110-M09. We also thank to National Center for High-performance Computing (NCHC) of National Applied Research Laboratories (NARLabs) in Taiwan for providing computational and storage resources.

## References

1. Amini-Naieni, N., Amini-Naieni, K., Han, T., Zisserman, A.: Open-world text-specified object counting. In: British Machine Vision Conference (BMVC) (2023)
2. Chang, L., Yujie, Z., Andrew, Z., Weidi, X.: Countr: Transformer-based generalised visual counting. In: British Machine Vision Conference (BMVC) (2022)
3. Cheng, Z.Q., Dai, Q., Li, H., Song, J., Wu, X., Hauptmann, A.G.: Rethinking spatial invariance of convolutional networks for object counting. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19638–19648 (2022)

4. Cheng, Z.Q., Li, J.X., Dai, Q., Wu, X., Hauptmann, A.G.: Learning spatial awareness to improve crowd counting. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6152–6161 (2019)
5. Cui, Y., Jiang, C., Wang, L., Wu, G.: Mixformer: End-to-end tracking with iterative mixed attention. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13608–13618 (2022)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
7. Hsieh, M.R., Lin, Y.L., Hsu, W.H.: Drone-based object counting by spatially regularized regional proposal network. In: IEEE International Conference on Computer Vision (ICCV). pp. 4145–4153 (2017)
8. Jiang, R., Liu, L., Chen, C.: Clip-count: Towards text-guided zero-shot object counting. In: Proceedings of ACM International Conference on Multimedia (ACM MM). p. 4535–4545. MM '23 (2023)
9. Leibe, B., Seemann, E., Schiele, B.: Pedestrian detection in crowded scenes. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). vol. 1, pp. 878–885 (2005)
10. Li, Y., Zhang, X., Chen, D.: Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In: IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1091–1100 (2018)
11. Lin, T.Y., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision (ECCV) (2014)
12. Lin, W., Yang, K., Ma, X., Gao, J., Liu, L., Liu, S., Hou, J., Yi, S., Chan, A.: Scale-prior deformable convolution for exemplar-guided class-agnostic counting. In: British Machine Vision Conference (BMVC) (2022)
13. Lu, E., Xie, W., Zisserman, A.: Class-agnostic counting. In: Asian Conference on Computer Vision (ACCV). pp. 669–684 (2018)
14. Ma, Z., Hong, X., Wei, X., Qiu, Y., Gong, Y.: Towards a universal model for cross-dataset crowd counting. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3205–3214 (2021)
15. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)* **32** (2019)
16. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (ICML) (2021)
17. Ranjan, V., Sharma, U., Nguyen, T., Hoai, M.: Learning to count everything. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3394–3403 (2021)
18. Shi, M., Lu, H., Feng, C., Liu, C., Cao, Z.: Represent, compare, and learn: A similarity-aware framework for class-agnostic counting. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9529–9538 (2022)
19. Song, Q., Wang, C., Jiang, Z., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Wu, Y.: Rethinking counting and localization in crowds: A purely point-based framework. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3365–3374 (2021)

20. Thanasutives, P., Fukui, K.i., Numao, M., Kijsirikul, B.: Encoder-decoder based convolutional neural networks with multi-scale-aware modules for crowd counting. In: IEEE International Conference on Pattern Recognition (ICPR). pp. 2382–2389 (2021)
21. Đukić, N., Lukežič, A., Zavrtanik, V., Kristan, M.: A low-shot object counting network with iterative prototype adaptation. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 18872–18881 (October 2023)
22. Wan, J., Liu, Z., Chan, A.B.: A generalized loss function for crowd counting and localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
23. Yang, S.D., Su, H.T., Hsu, W.H., Chen, W.C.: Class-agnostic few-shot object counting. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 870–878 (2021)
24. You, Z., Yang, K., Luo, W., Lu, X., Cui, L., Le, X.: Few-shot object counting with similarity-aware feature enhancement. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 6315–6324 (2023)