

Interactive Class-Agnostic Object Counting

Yifeng Huang¹, Viresh Ranjan^{2*}, Minh Hoai^{1,3}

¹Stony Brook University, Stony Brook, NY, USA

²Amazon, USA ³VinAI Research, Hanoi, Vietnam

Abstract

We propose a novel framework for interactive class-agnostic object counting, where a human user can interactively provide feedback to improve the accuracy of a counter. Our framework consists of two main components: a user-friendly visualizer to gather feedback and an efficient mechanism to incorporate it. In each iteration, we produce a density map to show the current prediction result, and we segment it into non-overlapping regions with an easily verifiable number of objects. The user can provide feedback by selecting a region with obvious counting errors and specifying the range for the estimated number of objects within it. To improve the counting result, we develop a novel adaptation loss to force the visual counter to output the predicted count within the user-specified range. For effective and efficient adaptation, we propose a refinement module that can be used with any density-based visual counter, and only the parameters in the refinement module will be updated during adaptation. Our experiments on two challenging class-agnostic object counting benchmarks, FSCD-LVIS and FSC-147, show that our method can reduce the mean absolute error of multiple state-of-the-art visual counters by roughly 30% to 40% with minimal user input. Our project can be found at <https://yifehuag97.github.io/ICACountProjectPage/>.

1. Introduction

The need for counting objects in images arises in many applications, and significant progress has been made for both class-specific [17, 30, 13, 46, 47, 9, 3, 24, 44, 25, 19, 38, 23, 16, 34, 36, 1] and class-agnostic [49, 35, 33, 41, 51, 26, 31, 32] counting. However, unlike in many other computer vision tasks where the predicted results can be verified for reliability, visual counting results are difficult to validate, as illustrated in Fig. 1. Mistakes can be made, and often there are no mechanisms to correct them. To enhance the practicality of visual counting methods, the results need




Figure 1. Given an input image and several exemplar objects, a class-agnostic counter will output a density map and the total object count. It is often challenging to validate these outputs, making it difficult to adopt automatic visual counting in practice. To improve the practicality of a visual counter, we propose an interactive framework that allows a human user to quickly detect mistakes and improve performance based on the identified errors.

to be more intuitive and verifiable, and feedback mechanisms should be incorporated to allow errors to be corrected. This necessitates a human-in-the-loop framework that can interactively display the predicted results, collect user feedback, and adapt the visual counter to reduce counting errors.

It is, however, challenging to develop an interactive framework for visual counting. The first challenge is to provide the user with an intuitive visualizer for the counting result. Current state-of-the-art visual counting methods typically generate a density map and then sum the density values to obtain the final count. However, as shown in Fig. 1, verifying the final predicted count can be difficult, as can verifying the intermediate density map, due to the mismatch between the continuous nature of the density map and the discrete nature of the objects in the image. The second challenge is to design an appropriate user interaction method that requires minimal user effort while being suited for providing feedback on object counting. The third challenge is developing an effective adaptation scheme for the selected interaction type that can incorporate user feedback and improve the performance of visual counters. In this paper, we address all three aforementioned challenges to develop an interactive framework for visual counting.

For the first challenge, we propose a novel segmentation method that segments a density map into non-overlapping regions, where the sum of density values in each region is a near-integer value that can be easily verified. This provides the user with a more natural and understandable interpretation of the predicted density map. Notably, developing such

*Work done at Stony Brook University, prior to joining Amazon

an algorithm that must also be suitably fast for an interactive system is challenging, which constitutes a technical contribution of our paper.

For the second challenge, we propose a novel type of interaction that enables the user to provide feedback with just two mouse clicks: the first click selects the region, and the second click selects the appropriate range for the number of objects in the chosen region. The proposed user interaction method is unique as it is specifically tailored for object counting and requires minimal user effort. Firstly, the auto-generated segmentation map allows the user to select an image region using just one mouse click, which is faster compared to drawing a polygon or scribbles. Secondly, by leveraging the humans' subitizing ability, which allows them to estimate the number of objects in a set quickly without counting them individually, we can obtain an approximate count with just another mouse click, which is quicker than one by one counting using dot annotations.

For the third challenge, we develop an interactive adaptation loss based on range constraints. To update the visual counter efficiently and effectively and to reduce the disruption of the learned knowledge in the visual counter, we propose the refinement module that directly refines the spatial similarity feature in the regression head. Furthermore, we propose a technique to estimate the user's feedback confidence and use this confidence to adjust the learning rate and gradient steps during the adaptation process.

In this paper, we primarily focus on class-agnostic counting, and we demonstrate the effectiveness of our framework with experiments on FSC-147 [35] and FSCD-LVIS [31]. However, our framework can be extended to category-specific counting, as will be seen in our experiments on several crowd-counting and car-counting benchmarks, including ShanghaiTech [52], UCF-QNRF [11], and CARPK [8]. We also conduct a user study to investigate the practicality of our method in a real-world setting.

In short, the main contribution of our paper is a framework that improves the accuracy and practicality of visual counting. Our technical contributions include: (1) a novel segmentation method that quickly segments density maps into non-overlapping regions with near-integer density values, which enhances the interpretability of predicted density maps for users; (2) an innovative user feedback scheme that requires minimal user effort for object counting by utilizing subitizing ability and auto-generated segmentation maps; and (3) an effective adaptation approach that incorporates the user's feedback into the visual counter through a refinement module and a confidence estimation method.

2. Related Works

Visual counting. Various visual counting methods have been proposed, e.g., [17, 30, 13, 46, 47, 9, 3, 24, 44, 25,

19, 38, 23, 16], but most of them are class-specific counters, requiring large amounts of training data with hundreds of thousands of annotated objects. To address this limitation and enable counting of objects across multiple categories, several class-agnostic counters have been proposed [49, 35, 33, 41, 51, 26, 31]. These methods work by regressing the object density map based on the spatial correlation between the input image and the provided exemplars. However, in many cases, a limited number of exemplars are insufficient to generalize over object instances with varying shapes, sizes, and appearances.

Interactive counting. There exists only one prior method for interactive counting [2]. This method uses low-level features and ridge regression to predict the density map. To visualize the density map, it uses MSER [29] or spectral clustering [40] to generate some candidate regions, then seeks a subset of candidate regions that can keep the integrality of each region in the subset. At each iteration, the user must draw a region of interest and mark the objects in this region with dot annotations. Additionally for the first iteration, the user has to specify the diameter of a typical object. This method [2] has two drawbacks. First, it requires significant effort from the user to draw a region of interest, specify the typical object size, and mark all objects in the region. Second, MSER and spectral clustering may not generate suitable candidate regions for dense scenes, as will be shown in Sec. 3.2. To alleviate the user's burden, the counting results should be easy to verify, and feedback should be simple to provide. In this paper, we propose a density map visualization method that can generate regions by finding and expanding density peaks. Unlike MSER and spectral clustering, our approach works well on dense density maps.

Interactive methods for other computer vision tasks. Various interactive methods have been developed for other computer vision tasks, such as object detection [50], tracking [39], and segmentation [12, 42, 5, 22, 20, 10, 28, 27, 18, 21, 48, 43]. While the success of these methods is inspiring, none of them are directly applicable to visual counting due to unique technical challenges. Unlike object detection, tracking, and segmentation, the immediate and final outputs of visual counting are difficult to visualize and verify. Designing an interactive framework for visual counting requires addressing the technical challenges discussed in the introduction section, none of them has been considered in previous interactive methods.

3. Proposed Approach

We propose an interactive framework for visual counting, as illustrated in Fig. 2. Each interactive iteration consists of two phases. The first phase visualizes the predicted density map to collect user feedback. The second phase uses the provided feedback to improve the visual counter.

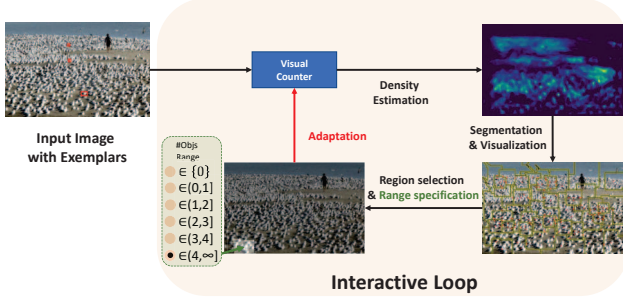


Figure 2. We propose a practical approach for visual counting based on interactive user’s feedback. In each iteration: (1) the visual counter estimates the density map for the input image; (2) the density map is segmented and visualized; (3) the user selects a region and provides a range for the number of objects in the region; (4) an objective function is defined based for the provided region and count range, and the parameters of a refinement module are updated by optimizing this objective function.

3.1. Overview of the two phases

In the first phase, we will visualize the density map by segmenting it into regions $\{R_1, \dots, R_n\}$ with the following desiderata:

- C1. Non-overlapping: $R_i \cap R_j = \emptyset$ for all $i \neq j$,
- C2. Total coverage: $\bigcup_{i=1}^n R_i =$ the predicted density map,
- C3. Moderate size: each region is not too big or too small,
- C4. Near-integer and small integral: the sum of density values within each region should be close to an integer and smaller than a verifiable counting limit.

The above desiderata are for visualization and easy verification of the results. The last desideratum is motivated by humans’ subitizing ability, which is the ability to identify the number of objects in an image simply by quickly looking at them, not by counting them one by one.

In the second phase of each iteration, the user is prompted to pick a region and specify the range for the number of objects in that region. Let R denote the selected region and $c = (c_l, c_u]$ the range specified by the user for the number of objects in R , a loss $\mathcal{L}(R, c)$ will be generated and used to adapt the counting model. For efficient and effective adaptation, instead of adapting the whole counting network, we propose a refinement module that directly refines the feature map in the regression head and we only adapt the parameters of this module using gradient descent.

3.2. Density map segmentation algorithm

One technical contribution of this paper is the development of a fast segmentation algorithm called Iterative Peak Selection and Expansion (IPSE) that satisfies the desiderata described in Sec. 3.1. The input to this algorithm is

a smoothed density map, and the output is a set of non-overlapping regions. IPSE is an iterative algorithm where the output set will be grown one region at a time, starting from an empty set. To yield a new region for the output, it starts from the pixel p with the highest density value (called the peak) among the remaining pixels that have not been included in any previously chosen region. IPSE seeks a region R containing p that minimizes the below objective:

$$h(R) = \frac{|R_s - \lceil R_s - \frac{1}{2} \rceil|}{\max(1, \lceil R_s - \frac{1}{2} \rceil)} + \frac{\max(0, T_l - R_a)}{T_l} + \lceil \max(0, R_s - C) \rceil, \quad (1)$$

where R_s denote the sum of density values in R and R_a the area of R . T_l is the preferred minimum area for the region R , and C is the preferred maximum number of objects in the region. The first term of the objective function encourages the sum of the densities to be close to an integer. It also encourages the region not to have too small count. The second term penalizes small regions. The region cannot be too big either. The expansion algorithm will stop when the area reaches a predefined upper bound. The last term penalizes regions with a total density greater than C .

Because there are exponentially many regions containing a given peak p , finding the optimal region R that minimizes $h(R)$ is an intractable problem. Fortunately, we only need to obtain a sufficiently good solution. We therefore restrict the search space to a smaller list of expanding regions $S_0 \subset S_1 \subset \dots \subset S_m$ and perform an exhaustive search among this list. This list can be constructed by starting from the seed p , i.e., $S_0 = \{p\}$ and constructing S_{i+1} from S_i by adding a pixel q selected from neighboring pixels of S_i that have not been part of any existing output regions. If multiple such neighboring pixels are available, we prioritize pixels with positive density values and select the one closest to p . The process terminates when any of the following conditions is met: (1) all neighboring pixels of S_i have been included in an output region; (2) the area or the sum of density values in S_i has reached a predefined limit; or (3) the proportion of zero-density pixels in S_i has exceeded a predefined threshold.

The above peak selection and expansion process is used repeatedly to segment the density map, with each iteration commencing from the seed pixel p that has the highest density value among the remaining pixels that have not been included in any prior output region. If all the remaining pixels have zero-density values, a random seed location is selected. The process continues until complete coverage of the density map is achieved, at which point all small regions are merged into their neighboring regions.

Comparison with other density segmentation methods.

Our algorithm for segmenting the density map bares some similarity to gerrymandering or political redistricting.

Gerrymandering involves dividing a large region into several smaller regions while adhering to certain constraints related to the population and contiguity of the regions. Most methods for this problem are based on weighted graph partitioning or heuristic algorithms [6, 37, 14, 4, 7]. However, an object density map contains several hundred thousand pixels, making these methods too slow for interactive systems. For example, the time and iteration limit of [6] is 600 seconds and 1000 iterations, which cannot meet the real-time requirements of interactive systems. In contrast, our method takes less than one second, as reported in Sec. 4.3.

Another approach for visualizing a density map is to use MSER [29] or spectral clustering[40] to generate some candidate regions, as used in [2]. MSER and spectral clustering, however, often fail to generate suitable candidate regions for dense scenes, as shown in Fig. 3.

3.3. Interactive feedback and adaptation

Upon presenting the segmented density map, the user will be prompted to pick a region R and choose a numeric range c for the number of objects in R , from a list of range options, $c \in \{(-\infty, 0], (0, r], (r, 2r], \dots, (C - r, C], (C, \infty)\}$, where r is the range interval and C is the counting limit. This method of user interaction is innovative and specifically tailored for object counting, necessitating only two mouse clicks per iteration. The reason for using a range instead of an exact number is because it can be ambiguous to determine an exact number for a region. Despite our segmentation efforts to ensure that each region contains an integral number of objects, some regions may still contain partial objects, making it more challenging for the user to provide an accurate number quickly.

An important technical contribution of our paper is the creation of an adaptation technique capable of leveraging the user’s weak supervision feedback. This feedback is weak in terms of both object localization and count. Specifically, it does not provide exact locations of object instances, only indicating the number of objects present in an image region. Additionally, the number of objects is provided only as a range, rather than an exact count. Below, we will detail our adaptation technique, starting from the novel refinement module for effective and efficient adaptation.

3.3.1 Refinement module

We aim for an adaptation method that works for all class-agnostic counting networks [49, 35, 33, 41, 51]. Most of them contain three components: a feature extractor f , a spatial-similarity module g (e.g., convolution [35] or learnable bilinear transform [41]), and a regression head h consisting of upsampling layers and convolution layers. Let \mathbf{I} be the input image, \mathbf{E} the exemplars, $\mathbf{S} = g(f(\mathbf{I}), f(\mathbf{E}))$ the correlation map for the spatial similarity between the input

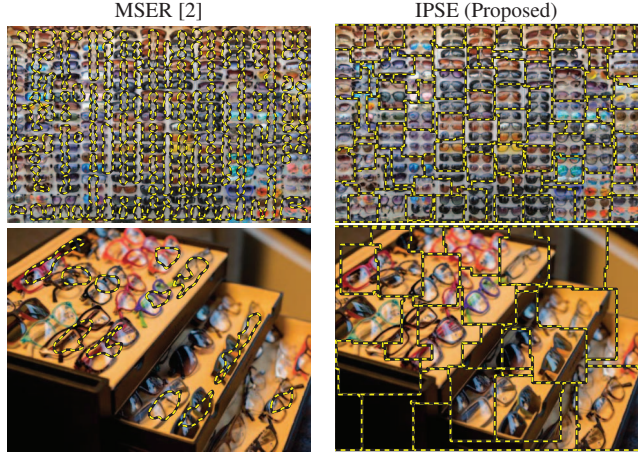


Figure 3. Density map segmentation comparison with MSER [2]. These two examples are from the FSC-147 dataset.

image and exemplars, and $\mathbf{D} = h(\mathbf{S})$ the predicted density map. The predicted count is obtained by summing over the predicted density map \mathbf{D} .

The correlation map serves as input to the regression head, which applies several convolution and upsampling layers to generate the output object density map. We observe that if the correlation map or any intermediate feature map between the input and output maps accurately represents the spatial similarity between the input image and exemplar objects, the final output density map and the predicted count will be correct. Therefore, the adaptation process need not revert to layers earlier than the correlation map. To minimize disruption to learned knowledge and accelerate adaptation for user interaction, we propose a lightweight refinement module that is integrated only within the regression head.

The refinement module, depicted in Fig. 4, can be applied to any intermediate feature map \mathbf{F} between the input correlation map \mathbf{S} and the output density map \mathbf{D} : $\mathbf{F}' = \mathcal{R}(\mathbf{F})$, where \mathbf{F}' is the refined feature map. Our refinement module consists of two components: channel-wise refinement and spatial-wise refinement.

The channel-wise refinement is illustrated in Fig. 5, and it refines the feature of each channel by multiplying with a scale parameter and adding with a bias parameter: $\mathcal{R}_{ch}(\mathbf{F}) = \theta_{ch}^{scale} \odot \mathbf{F} + \theta_{ch}^{bias}$, where $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$ is the input feature map, $\theta_{ch}^{scale} \in \mathbb{R}^C$ is the vector of scale parameters, $\theta_{ch}^{bias} \in \mathbb{R}^C$ is the vector of bias parameters.

The spatial-wise refinement is illustrated in Fig. 6, and it also refines the feature at each spatial position: $\mathcal{R}_{sp}(\mathbf{F}) = \theta_{sp}^{scale} \odot \mathbf{F} + \theta_{sp}^{bias}$, where $\theta_{sp}^{scale} \in \mathbb{R}^{H \times W}$, $\theta_{sp}^{bias} \in \mathbb{R}^{H \times W}$.

The overall refinement module is the successive application of channel-wise refinement and spatial-wise refinement: $\mathbf{F}' = \mathcal{R}(\mathbf{F}) = \mathcal{R}_{sp}(\mathcal{R}_{ch}(\mathbf{F}))$. The set of adaptable




Figure 4. The refinement module can be integrated into the **regression head** of any density estimation visual counter.




Figure 5. Channel-wise refinement refines the feature of each channel by multiplying with scale and adding with a bias.




Figure 6. Spatial-wise refinement refines the feature at each spatial position by multiplying with scale and adding with a bias.

parameters of the two refinement modules are: $\theta^{scale} = [\theta_{ch}^{scale}; \theta_{sp}^{scale}]$ and $\theta^{bias} = [\theta_{ch}^{bias}; \theta_{sp}^{bias}]$. At the beginning of each adaptation iteration, the scale parameters are reset to one and the bias parameters to zero, so that the refined feature map \mathbf{F}' is the same as the input feature map \mathbf{F} .

3.3.2 Adaptation schemes

Given a selected region R and a specified range $c = (c_l, c_u]$ for the number of objects in region R , a weekly-supervised adaptation loss is defined as:

$$\mathcal{L}_I(R, c) = \text{ReLU}(c_l - R_s) + \text{ReLU}(R_s - c_u), \quad (2)$$

If the sum of predicted density values in the selected region is outside the count range provided by the user, the above loss will be positive.

To account for the scenario where the user can provide feedback for multiple regions either in a single iteration or through multiple iterations, we extend the adaptation loss to use multiple regions to update the counter. Let $\Omega = (R, c)$ denote the user's selected regions and their corresponding specified count ranges. We use the following combined loss for adaptation:

$$\mathcal{L}(\Omega) = \mathcal{L}_L(\Omega) + \mathcal{L}_G(\Omega) + \eta(\|\theta^{scale} - 1\| + \|\theta^{bias}\|), \quad (3)$$

where $\mathcal{L}_L(\Omega)$ is the sum of regional losses, with each loss corresponding to an individual region separately:

$$\mathcal{L}_L(\Omega) = \sum_{(R,c) \in \Omega} \mathcal{L}_I(R, c), \quad (4)$$

and $\mathcal{L}_G(\Omega)$ is the single loss for all the regions combined:

$$\mathcal{L}_G(\Omega) = \mathcal{L}_I\left(\sum_{(R,c) \in \Omega} R_s, \sum_{(R,c) \in \Omega} c\right). \quad (5)$$

Here, we combine all the selected regions and view them as one big region, and then use Eq. (2) on the big region to adapt the visual counter. Hereafter, we will refer to $\mathcal{L}_L(\Omega)$ as Local Loss and $\mathcal{L}_G(\Omega)$ as Global Loss.

The last term of Eq. (3) is a regularization term to discourage large changes to the scale and bias parameters of the refinement module and η is the weight for the regularization term. In our experiments, $\eta = 0.002$.

Our adaptation loss is based on the sum of predicted density values within one or multiple regions, which provides a weak supervision signal as it lacks penalty terms related to individual values. This type of supervision signal can be problematic if a large learning rate is used. Meanwhile, using a small learning rate would require numerous gradient steps, leading to a prolonged convergence process. To overcome this issue, we propose determining the impact value of the user's feedback and using it to adjust the learning rate and gradient steps, resulting in a smoother adaptation process. When the uncertainty level is higher, such as having too few regions or inconsistent error types (i.e., under-counting in some regions and over-counting in others), we use a lower learning rate with more gradient steps. Conversely, we apply a larger learning rate with fewer gradient steps to shorten the adaptation process when the uncertainty level is lower. We define the user feedback value as follows:

$$F_C(\Omega) = 0.5F_i(\Omega) + 0.5F_s(\Omega), \quad (6)$$

The first term of the above measures the informativeness of the feedback while the second term measures the inconsistency of the feedback for multiple regions. The informativeness term can be defined as:

$$F_i(\Omega) = \min\left(1, \exp\left(\frac{|\Omega| - t}{T}\right)\right), \quad (7)$$

where $|\Omega|$ is the size of the set Ω , T is the temperature, and t is the informativeness threshold. Specifically, in our experiments $t = 3, T = 2$.

Let Ω_o and Ω_u be the sets of over-counting and under-counting regions. Let $p = \frac{|\Omega_o|}{|\Omega_o| + |\Omega_u|}$, the feedback consistency value is defined based on negative entropy:

$$F_s(\Omega) = 1 + p \log p + (1 - p) \log(1 - p). \quad (8)$$

Based on the estimated value of the feedback, the learning rate and the number of gradient steps will be scaled accordingly as follows: $\gamma' = \gamma F_C(\Omega)$, $N' = \frac{N}{F_C(\Omega)}$, where γ and N are the default values for the learning rate and the number of gradient steps.

4. Experiments

4.1. Class-agnostic counting

4.1.1 Experiment settings

Datasets. We evaluate our approach on two challenging class-agnostic counting benchmarks: FSC-147 [35] and FSCD-LVIS [31].

Class-agnostic visual counters. Our interactive framework is applicable to many different class-agnostic counters, and we experiment with FamNet [35], SAFECOUNT [51], and BMNet+ [41] in this paper.

Evaluation metrics. We use Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) as performance metrics, which are widely used for evaluating visual counting performance [49, 35, 33, 41, 51, 26, 31].

User feedback simulation. To simulate user feedback, we randomly select a displayed region and provide the counting range for that region. We repeat each experiment five times with different random seeds (5, 10, 15, 20, 25) and report the **average** and **standard error**.

Count limit and pre-defined ranges. According to [45], people can quickly and correctly estimate the number of objects without one-to-one counting if the number of objects is less or equal to four. Therefore, we set the count limit C to 4 and the pre-defined ranges to $\{(-\infty, 0], (0, 1], (1, 2], (2, 3], (3, 4], (4, \infty)\}$.

Implementation details. We insert the refinement module after the **first convolution layer** in the regression head. For faster computation, we first downsample the density map by a factor of four, then perform the IPSE, and finally upsample the density map segmentation result to its original size. We adapt a FamNet with an Adam optimizer [15]. On FSC-147, the default learning rate is 0.02 while the default number of gradient steps is 10. On FSCD-LVIS, FamNet does not converge as well, so the default learning rate and number of gradient steps are set to 0.01 and 20, respectively.

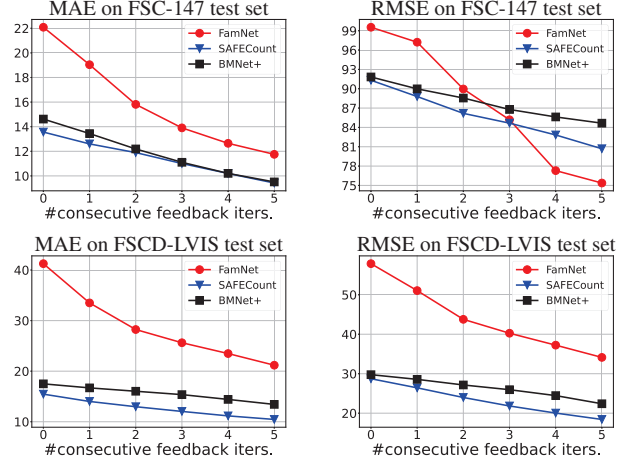


Figure 7. The proposed framework can be used to improve the performance of various visual counters. This shows the MAE and RMSE values of FamNet, SAFECOUNT, and BMNet+ on FSC-147 and FSCD-LVIS test data, as the number of feedback iterations is increased from zero (without any adaption) to five.

4.1.2 Experiment results

The proposed interactive framework can be used to improve the performance of various types of visual counters, including FamNet [35], SAFECOUNT [51], and BMNet+ [41]. As can be seen from Fig. 7, the benefits are consistently observed in multiple experiments and metrics (three visual counters, two datasets, and two performance metrics). Significant error reduction is already achieved even after one feedback iteration, as also shown in Fig. 8. After five iterations, the amount of error reduction is huge, with an average value of 30%.

The proposed framework requires minimal user input, but it should not be viewed as a competitor to few-shot counting methods. Rather, it is an interactive framework that offers complementary benefits to few-shot methods. Notably, most class-agnostic visual counters, including FamNet, SAFECOUNT, and BMNet+, are few-shot methods that can count with just a few exemplars. As shown earlier, the proposed framework enhances the performance of these counters. However, there is another approach to improve these counters, which is to provide more exemplars. Rather than using these visual counters in our interactive framework, we could offer them additional exemplars. As our framework requires two mouse clicks for each iteration and drawing a bounding box also requires the same effort, we compare the performance of our framework with five iterations to the performance of visual counters with five additional exemplars, and the results are shown in Table 1. As shown, our framework produces greater improvement with the same level of user effort. This may be due to several reasons: (1) supplying extra exemplars does not immediately highlight prediction errors, resulting in a weaker supervi-



Figure 8. Qualitative results of our approach for test images in FSC-147 with FamNet as the visual counter. The brighter region is the selected region, and the red dot is our approximate location of each region generated by peak selection and non-maximum suppression on each region. The selected region is highlighted in this example, and the input range is $(-\infty, 0]$, since the object of interest is cake. With one single interaction, our method can improve the counting result locally and globally. **More qualitative results and a demo video are in our supplementary.**

| | FSC-147 Test set | | FSCD-LVIS Test set | |
|-----------------|------------------|------------|--------------------|------------|
| | MAE | RMSE | MAE | RMSE |
| FamNet [35] | 22.08 | 99.54 | 41.26 | 57.87 |
| + 5 exemplars | 21.52 ↓2% | 98.10 ↓1% | 40.36 ↓2% | 57.85 ↓0% |
| + our framework | 11.75 ↓47% | 75.37 ↓24% | 21.18 ↓49% | 34.13 ↓41% |
| SAFECount [51] | 13.56 | 91.31 | 15.45 | 28.73 |
| + 5 exemplars | 13.01 ↓4% | 94.22 ↑3% | 14.83 ↓4% | 28.01 ↓2% |
| + our framework | 9.42 ↓31% | 80.69 ↓12% | 10.45 ↓32% | 18.42 ↓36% |
| BMNet+ [41] | 14.62 | 91.83 | 17.49 | 29.76 |
| + 5 exemplars | 14.40 ↓2% | 91.56 ↓0% | 17.27 ↓1% | 29.60 ↓1% |
| + our framework | 9.51 ↓35% | 84.66 ↓8% | 13.43 ↓23% | 22.39 ↓25% |

Table 1. Comparing the performance of the proposed interactive framework with **five feedback iterations** to a few-shot baseline approach that uses the base counters with **5 additional exemplars**.

sion signal; (2) an exemplar provides information for only one object, less informative than a region containing multiple objects; (3) most class-agnostic counting methods are trained with a predefined number of exemplars (e.g., three), so the model may not be able to fully utilize the additional exemplars to improve the performance.

One technical contribution of our method is the innovative segmentation algorithm. Table 2 compares this algorithm with four segmentation methods: MSER [2], K-means, Watershed, and DBSCAN. For K-means and DBSCAN, we use the spatial coordinate and the density value as the feature to perform clustering for segmentation. For K-means, K is set to $\min(S_u, \max(S_l, \frac{Sum(D)}{C}))$, where S_u and S_l are the pre-defined upper bound and lower bound, $Sum(D)$ is the summed density, and C is the count limit. For these density map segmentation baselines, we also use our feature refinement adaptation scheme to update the visual counter. As shown in Table 2, our proposed algorithm surpasses the other methods by a wide margin. Fig. 8 shows some qualitative results.

4.1.3 Ablation study

We perform some experiments to evaluate the contribution of different components, including the refinement module,

| | FSC-147 Test set | | FSCD-LVIS Test set | |
|-----------------|-------------------|-------------------|--------------------|-------------------|
| | MAE | RMSE | MAE | RMSE |
| Initial error | 22.08 | 99.54 | 41.26 | 57.87 |
| MSER [2] | 16.46±0.13 | 86.06±3.28 | 30.61±0.37 | 44.57±0.77 |
| Watershed | 18.95±0.09 | 74.08±4.66 | 27.35±0.25 | 43.21±0.79 |
| K-means | 15.32±0.23 | 86.49±2.43 | 32.70±0.04 | 47.80±0.09 |
| DBSCAN | 19.69±0.24 | 78.26±8.08 | 41.26±0.00 | 57.87±0.00 |
| IPSE (proposed) | 11.75±0.12 | 75.37±5.21 | 21.18±0.28 | 34.13±0.88 |

Table 2. Results of different segmentation methods under the same adaptation scheme with **five feedback iterations**, when FamNet is the base visual counter. Each experiment is repeated five times, and the **mean** and **standard error** are reported.



Figure 9. Channel-wise contributes more to the improvement than spatial-wise refinement since channel-wise refinement corrects global errors, while spatial-wise focuses on local errors.

the adaptation loss, and the setting of user feedback simulation. All ablation studies are conducted on the FSC-147 validation set.

Refinement module. The results of the ablation study for the refinement module are shown in Table 3. Both the channel-wise and spatial-wise refinement steps are important. The channel-wise refinement contributes more to the improvement than the spatial-wise refinement step. This is perhaps because the spatial-wise refinement refines locally while the channel-wise refinement refines globally, as shown in Fig. 9. Also, the order of these two refinement steps has little effect on the final result.

Adaptation loss. Table 4 shows the ablation study on the adaptation loss. Both the Local Loss $\mathcal{L}_L(\Omega)$ and the Global Loss $\mathcal{L}_G(\Omega)$ contribute to the reduction in MAE and RMSE. The confidence scaling is less important.

| Refinement Component | FamNet | | SAFECount | |
|----------------------|--------------|--------------|-------------|--------------|
| | MAE | RMSE | MAE | RMSE |
| Spatial | 21.63 | 67.84 | 13.90 | 10.67 |
| Channel | 13.72 | 52.99 | 9.46 | 13.01 |
| Spatial + Channel | 12.84 | 45.85 | 9.98 | 36.15 |
| Channel + Spatial | 12.79 | 47.21 | 9.29 | 33.83 |

Table 3. Analyzing the contribution of the spatial and refinement components of the refinement module. The result is on FSC-147 validation set with FamNet and SAFECount as the visual counter.

| | | | | |
|-------------------------------------|------------|------------|------------|-------------------|
| Local Loss $\mathcal{L}_L(\Omega)$ | ✓ | ✗ | ✓ | ✓ |
| Global Loss $\mathcal{L}_G(\Omega)$ | ✗ | ✓ | ✓ | ✓ |
| Confidence scaling | ✗ | ✗ | ✗ | ✓ |
| MAE | 15.41±0.29 | 14.21±0.17 | 13.25±0.21 | 12.79±0.16 |
| RMSE | 55.72±1.57 | 52.94±1.56 | 48.44±1.69 | 47.21±2.05 |

Table 4. Ablation on the adaptation loss, the result is on FSC-147 validation set with FamNet as the visual counter.

| Refinement Component | FamNet | | SAFECount | |
|----------------------|--------------|--------------|-------------|--------------|
| | MAE | RMSE | MAE | RMSE |
| LocalCorrection | 20.74 | 64.58 | 13.01 | 49.15 |
| AllParamAdapt | 17.07 | 61.61 | 9.98 | 36.15 |
| Proposed | 12.79 | 47.21 | 9.29 | 33.83 |

Table 5. Analyzing the effectiveness of our adaptation scheme.

Adaptation scheme. Table 5 shows the ablation study on the adaptation scheme. LocalCorrection is the method that only corrects the prediction of the selected region, and will not adapt the visual counter. AllParamAdapt is the method that updates all the parameters in the regression head.

User feedback simulation. To further assess the efficiency and efficacy of various region selection strategies, we consider two methods that may offer advantages over random selection used in previous experiments. These two strategies are: 1) prioritizing background regions containing no objects, and 2) selecting regions with largest errors. The comparison of these region selection strategies is presented in Table 6. The error-based approach emerges as the most successful, whereas prioritizing background regions has a detrimental impact on performance. Concerning time efficiency, random selection is the quickest, followed by background prioritization, with error-based selection being the slowest due to the need for error estimation for each region.

In the primary experiment, we set interactions at five, given consecutively. To validate this five-interaction approach, we conducted extra trials on the FSC-147 dataset with varying interaction counts. Results are shown in Fig. 10, revealing continued performance improvement beyond five interactions, though at a slower rate.

We also compared consecutive and non-consecutive in-

| Region Selection Strategy | MAE | RMSE |
|---------------------------|--------------|--------------|
| Background Prior | 14.02 | 55.31 |
| Error Based | 10.61 | 46.85 |
| Random Selection | 12.79 | 47.21 |

Table 6. Result on different region selection strategies with FamNet on FSC-147 validation set.

| | FamNet | | SAFECount | |
|-----------------|--------|-------|-----------|-------|
| | MAE | RMSE | MAE | RMSE |
| Consecutive | 12.79 | 47.21 | 9.29 | 33.83 |
| Non Consecutive | 16.87 | 54.42 | 12.36 | 46.97 |

Table 7. Comparison of consecutive and non-consecutive interactions, the result is on the FSC-147 validation set.




Figure 10. MAE with varying numbers of interactions on the FSC-147 validation set.

teractions to confirm the importance of sequential engagement. Table 7 outlines these results, indicating better performance with sequential interactions. However, this approach increases time consumption due to added segmentation and adaptation time.

4.2. Class-specific counting

Our interactive counting framework has primarily focused on class-agnostic counters, with its foundation being the correlation between the exemplar objects and the image. However, it is natural to wonder if this framework can be extended to class-specific visual counters. The main concern is that the user feedback collected in this manner may be too weak for class-specific counters that are trained on hundreds of thousands or even millions of annotated objects. In this section, we report our experiments on the crowd-counting and car-counting task, where we have found that our method can be used to reduce the errors of a crowd counter if the counting ranges for the user feedback are adjusted.

Specifically, we apply our framework to DM-Count [46] for crowd counting on two crowd-counting benchmarks: ShanghaiTech [52] and UCF-QNRF [11]. We set the counting limit to 50, and the range interval to 10, so the counting ranges are $\{[-\infty, 0], (0, 10], \dots, (40, 50], (50, \infty)\}$. We insert the refinement module after the first convolution layer. We adapt the DM-Count with 20 gradient steps with a learning rate of 5×10^{-4} . We compare with the other baseline

| | ShanghaiTech A | | UCF-QNRF | |
|-----------------|-------------------|-------------------|-------------------|--------------------|
| | MAE | RMSE | MAE | RMSE |
| Initial error | 59.60 | 95.56 | 85.65 | 148.35 |
| MSER [2] | 57.90±0.32 | 94.84±0.50 | 81.38±0.58 | 142.73±0.93 |
| Watershed | 58.84±0.71 | 90.17±2.06 | 82.78±0.46 | 144.70±0.39 |
| K-means | 48.20±0.75 | 79.23±2.46 | 73.29±1.77 | 131.37±3.98 |
| DBSCAN | 56.99±0.21 | 92.73±0.28 | 80.62±0.93 | 142.29±1.66 |
| LocalCorrection | 43.86±0.26 | 74.76±0.84 | 78.27±0.23 | 135.07±0.53 |
| AllParamAdapt | 31.03±0.33 | 59.47±0.96 | 62.30±2.81 | 123.03±7.52 |
| Proposed | 33.85±0.78 | 57.50±1.94 | 58.13±1.04 | 102.32±2.63 |

Table 8. Results of interactive adaptation methods for a crowd-counting network (DM-Count) using five feedback iterations.

| | Initial Error | | Five Interactions | |
|-----------|---------------|-------|-------------------|-------|
| | MAE | RMSE | MAE | RMSE |
| FamNet | 18.34 | 35.77 | 13.91 | 20.14 |
| SAFECount | 4.91 | 6.32 | 4.16 | 5.91 |

Table 9. Result on CARPK with FamNet and SAFECount using five feedback iterations.

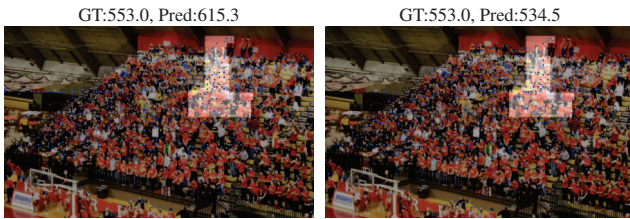


Figure 11. Qualitative results of our approach for test images in ShanghaiTech A with DM-Count as the visual counter. The left figure is before interaction, and the right is after one interaction. The selected region is highlighted, and the input range is (20, 30).

methods, and the feedback simulation is identical to those reported in Sec. 4.1. The quantitative results are shown in Table 8, and the qualitative results are shown in Fig. 11. Our approach reduces the MAE and RMSE by approximately 40% and 30% on ShanghaiTech A, and approximately 30% and 25% on UCF-QNRF. For car counting, we apply our framework to FamNet and SAFECount on CARPK [8]. The results are shown in Table 9. The MAE decreased more than 15% on CARPK.

4.3. User study

To assess the feasibility of the proposed interactive counting framework, we conducted a user study with eight participants. We selected 20 images with high counting errors from the FSC-147 dataset and used FamNet as the visual counter. Each participant was allowed a maximum of five iterations for each image, but they could terminate the process if they felt the prediction was accurate enough. The average number of iterations for one image is 3.08, and the variance is 0.41. Additionally, we carried out an experi-




Figure 12. Per-iteration statistics. The visualization time (including segmentation time) and the adaptation time are both less than one second, sufficiently fast for interactive systems.

| #Feedback Iterations | MAE | RMSE |
|------------------------------|-------------------|--------------------|
| Initial error | 93.41 | 125.57 |
| Real User (avg. 3.08 iters) | 45.11 ± 2.63 ↓52% | 90.64 ± 2.68 ↓28% |
| Simulated feedback (3 iters) | 59.97 ± 3.34 ↓34% | 110.67 ± 5.96 ↓12% |
| Simulated feedback (5 iters) | 43.86 ± 3.02 ↓52% | 92.93 ± 2.90 ↓26% |

Table 10. Comparison between real user’s feedback and the simulated feedback used in our quantitative experiments.

ment involving simulated user feedback on the same set of images, to compare the results with those obtained from real user feedback. As shown in Table 10, the users were able to improve the performance of the counter using our framework, demonstrating its potential for practical usage. Moreover, the benefits achieved from real user feedback were comparable to those obtained from simulated feedback, indicating that many of our analyses using simulated feedback can be extrapolated to real-world scenarios.

The user study was conducted on an RTX3080 machine, and several time statistics of a single iteration are presented in Fig. 12. It took less than a second to segment any density image and display it. The average time for a user to select a region and specify a range was four seconds, and the adaptation time for a single iteration was less than one second. All operations are sufficiently fast for interactive systems.

5. Conclusions

We have proposed an interactive framework primarily for class-agnostic counting, but can also be extended to class-specific counting. It uses a novel method for density map segmentation to generate an intuitive display for the user, enabling them to visualize the results and provide feedback. Additionally, we have developed an adaptation loss and a refinement module to efficiently and effectively update the visual counter with the user’s feedback. Experiments on two class-agnostic counting datasets and two crowd-counting benchmarks with four different visual counters demonstrate the effectiveness and general applicability of our framework.

Acknowledgement: This project was supported by US National Science Foundation Award NSDF DUE-2055406.

References

- [1] Shahira Abousamra, Minh Hoai, Dimitris Samaras, and Chao Chen. Localization in the crowd with topological constraints. In *Proceedings of AAAI Conference on Artificial Intelligence*, 2021.
- [2] Carlos Arteta, Victor Lempitsky, J Alison Noble, and Andrew Zisserman. Interactive object counting. In *Proceedings of the European Conference on Computer Vision*, 2014.
- [3] Shuai Bai, Zhiqun He, Yu Qiao, Hanzhe Hu, Wei Wu, and Junjie Yan. Adaptive dilated network with self-correction supervision for counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [4] Sean L Forman and Yading Yue. Congressional districting using a tsp-based genetic algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2003.
- [5] Prasad Gabbur, Manjot Bilkhu, and Javier Movellan. Probabilistic attention for interactive segmentation. *arXiv preprint arXiv:2106.15338*, 2021.
- [6] Alex Gliesch, Marcus Ritt, Arthur HS Cruz, and Mayron CO Moreira. A heuristic algorithm for districting problems with similarity constraints. In *IEEE Congress on Evolutionary Computation (CEC)*, 2020.
- [7] Sebastian Goderbauer and Jeff Winandy. Political districting problem: Literature review and discussion with regard to federal elections in germany. *consultado em outubro de*, 2018.
- [8] Meng-Ru Hsieh, Yen-Liang Lin, and Winston H Hsu. Drone-based object counting by spatially regularized regional proposal network. In *Proceedings of the IEEE international conference on computer vision*, 2017.
- [9] Yutao Hu, Xiaolong Jiang, Xuhui Liu, Baochang Zhang, Jungong Han, Xianbin Cao, and David Doermann. Nas-count: Counting-by-density with neural architecture search. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [10] Yang Hu, Andrea Soltoggio, Russell Lock, and Steve Carter. A fully convolutional two-stream fusion network for interactive image segmentation. *Neural Networks*, 109:31–42, 2019.
- [11] Haroon Idrees, Muhmmad Tayyab, Kishan Athrey, Dong Zhang, Somaya Al-Maadeed, Nasir Rajpoot, and Mubarak Shah. Composition loss for counting, density map estimation and localization in dense crowds. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [12] Won-Dong Jang and Chang-Su Kim. Interactive image segmentation via backpropagating refinement scheme. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [13] Xiaoheng Jiang, Li Zhang, Mingliang Xu, Tianzhu Zhang, Pei Lv, Bing Zhou, Xin Yang, and Yanwei Pang. Attention scaling for crowd counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [14] Myung Jin Kim. Multiobjective spanning tree based optimization model to political redistricting. *Spatial Information Research*, 26(3):317–325, 2018.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vazquez, and Mark Schmidt. Where are the blobs: Counting by localization with point supervision. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [17] Yuhong Li, Xiaofan Zhang, and Deming Chen. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [18] Zhuwen Li, Qifeng Chen, and Vladlen Koltun. Interactive image segmentation with latent diversity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [19] Dongze Lian, Jing Li, Jia Zheng, Weixin Luo, and Shenghua Gao. Density map regression guided detection network for rgb-d crowd counting and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [20] JunHao Liew, Yunchao Wei, Wei Xiong, Sim-Heng Ong, and Jiashi Feng. Regional interactive image segmentation networks. In *Proceedings of the IEEE international conference on computer vision*, 2017.
- [21] Jun Hao Liew, Scott Cohen, Brian Price, Long Mai, and Jiashi Feng. Deep interactive thin object selection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021.
- [22] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive image segmentation with first click attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [23] Chenchen Liu, Xinyu Weng, and Yadong Mu. Recurrent attentive zooming for joint crowd counting and precise localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [24] Xiyang Liu, Jie Yang, Wenrui Ding, Tieqiang Wang, Zhijin Wang, and Junjun Xiong. Adaptive mixture regression network with local counting map for crowd counting. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [25] Yuting Liu, Miaoqing Shi, Qijun Zhao, and Xiaofang Wang. Point in, box out: Beyond counting persons in crowds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [26] Erika Lu, Weidi Xie, and Andrew Zisserman. Class-agnostic counting. In *Proceedings of the Asian conference on computer vision*, 2018.
- [27] Sabarinath Mahadevan, Paul Voigtlaender, and Bastian Leibe. Iteratively trained interactive segmentation. *arXiv preprint arXiv:1805.04398*, 2018.
- [28] Soumajit Majumder and Angela Yao. Content-aware multi-level guidance for interactive instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

- [29] Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pa-jdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*, 22(10):761–767, 2004.
- [30] Yunqi Miao, Zijia Lin, Guiguang Ding, and Jungong Han. Shallow feature based dense attention network for crowd counting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [31] Thanh Nguyen, Chau Pham, Khoi Nguyen, and Minh Hoai. Few-shot object counting and detection. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [32] Viresh Ranjan and Minh Hoai. Exemplar free class agnostic counting. In *Proceedings of the Asian Conference on Computer Vision*, 2022.
- [33] Viresh Ranjan and Minh Hoai. Vicinal counting networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop*, 2022.
- [34] Viresh Ranjan, Hieu Le, and Minh Hoai. Iterative crowd counting. In *Proceedings of the European Conference on Computer Vision*, 2018.
- [35] Viresh Ranjan, Udbhav Sharma, Thu Nguyen, and Minh Hoai. Learning to count everything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [36] Viresh Ranjan, Boyu Wang, Mubarak Shah, and Minh Hoai. Uncertainty estimation and sample selection for crowd counting. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [37] Eric Alfredo Rincón-García, Miguel Ángel Gutiérrez-Andrade, Sergio Gerardo de-los Cobos-Silva, Roman Anselmo Mora-Gutiérrez, Antonin Ponsich, and Pedro Lara-Velazquez. A comparative study of population-based algorithms for a political districting problem. *Kybernetes*, 2017.
- [38] Deepak Babu Sam, Skand Vishwanath Peri, Mukuntha Narayanan Sundararaman, Amogh Kamath, and Venkatesh Babu Radhakrishnan. Locate, size and count: Accurately resolving people in dense crowds via detection. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [39] Minmin Shen, Chen Li, Wei Huang, Paul Szyszka, Kimiaki Shirahama, Marcin Grzegorzec, Dorit Merhof, and Oliver Deussen. Interactive tracking of insect posture. *Pattern Recognition*, 48(11):3560–3571, 2015.
- [40] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- [41] Min Shi, Hao Lu, Chen Feng, Chengxin Liu, and Zhiguo Cao. Represent, compare, and learn: A similarity-aware framework for class-agnostic counting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [42] Konstantin Sofiiuk, Ilia Petrov, Olga Barinova, and Anton Konushin. f-brs: Rethinking backpropagating refinement for interactive segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [43] Gwangmo Song, Heesoo Myeong, and Kyoung Mu Lee. Seednet: Automatic seed generation with deep reinforcement learning for robust interactive segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [44] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [45] Lana M Trick and Zenon W Pylyshyn. Why are small and large numbers enumerated differently? a limited-capacity preattentive stage in vision. *Psychological review*, 101(1):80, 1994.
- [46] Boyu Wang, Huidong Liu, Dimitris Samaras, and Minh Hoai. Distribution matching for crowd counting. In *Advances in Neural Information Processing Systems*, 2020.
- [47] Chenfeng Xu, Kai Qiu, Jianlong Fu, Song Bai, Yongchao Xu, and Xiang Bai. Learn to scale: Generating multipolar normalized density maps for crowd counting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [48] Ning Xu, Brian Price, Scott Cohen, Jimei Yang, and Thomas S Huang. Deep interactive object selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [49] Shuo-Diao Yang, Hung-Ting Su, Winston H Hsu, and Wen-Chin Chen. Class-agnostic few-shot object counting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021.
- [50] Angela Yao, Juergen Gall, Christian Leistner, and Luc Van Gool. Interactive object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2012.
- [51] Zhiyuan You, Kai Yang, Wenhan Luo, Xin Lu, Lei Cui, and Xinyi Le. Few-shot object counting with similarity-aware feature enhancement. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023.
- [52] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.