ABC Easy as 123: A Blind Counter for Exemplar-Free Multi-Class Class-agnostic Counting

Michael Hobley and Victor Prisacariu

Active Vision Laboratory, University of Oxford [mahobley, victor]@robots.ox.ac.uk

Abstract. Class-agnostic counting methods enumerate objects of an arbitrary class, providing tremendous utility in many fields. Prior works have limited usefulness as they require either a set of examples of the type to be counted or that the query image contains only a single type of object. A significant factor in these shortcomings is the lack of a dataset to properly address counting in settings with more than one kind of object present. To address these issues, we propose the first Multi-class, Class-Agnostic Counting dataset (MCAC) and A Blind Counter (ABC123), a method that can count multiple types of objects simultaneously without using examples of type during training or inference. ABC123 introduces a new paradigm where instead of requiring exemplars to guide the enumeration, examples are found after the counting stage to help a user understand the generated outputs. We show that ABC123 outperforms contemporary methods on MCAC without needing human in-theloop annotations. We also show that this performance transfers to FSC-147, the standard class-agnostic counting dataset. MCAC is available at MCAC.active.vision and ABC123 is available at ABC123.active.vision

1 Introduction

Given an image and told to 'count', a person would generally understand the intended task and complete it with accuracy even if there are multiple previously unseen classes of object present. This natural human ability to count arbitrarily has not been modelled by today's methods. Most automated counting methods are class-specific [1,6], counting objects of classes that were present during training. These methods are not generalisable and require retraining for each new type of object. *Class-agnostic* methods [19,21] can count objects of an arbitrary type removing the need for retraining. However, they usually require an exemplar image as a prior on the class to count and only count a single class at a time. Recently, *exemplar-free* or *zero-shot* class-agnostic counting methods [7,18] have been developed that do away with the need for exemplars to define type removing the need for human intervention during deployment. These methods either perform poorly [18] or are bounded to images that only contain a single class of object [4,7].



Fig. 1: ABC123 counts objects of multiple unseen types. Not only does our method not need exemplars to define the type to count, it finds examples of each type it has counted.



Fig. 2: MCAC Contains images with up to 4 classes and up to 300 instances per class. All objects have associated instance labels, class labels, bounding boxes, centre points, and occlusion percentages.

We propose ABC123, a transformer-based, multi-class class-agnostic counter which does not need exemplars during training or inference. ABC123 achieves this by first regressing density maps for each type present then enumerating the instances using integration. As it is sometimes difficult to interpret what has been counted given only a density map and count, we design an example discovery stage which locates instances of the counted object.

A significant factor in the limitations of the current methods is the lack of a dataset for class-agnostic counting that includes images with more than one class present, as currently they all focus on the single-class scenario. In order to train and evaluate our method, as well as other methods in multi-class settings we introduce MCAC, a new synthetic multi-class class-agnostic counting dataset. We show that methods previously assumed to work in multi-class settings perform poorly on MCAC and that ABC123 significantly outperforms them while also generalising to other datasets.

Our main contributions are:

- We propose ABC123, the first exemplar-free multi-class class-agnostic counter and show it tackles multi-class counting effectively.
- We introduced MCAC, the first multi-class class-agnostic counting dataset and use it to demonstrate prior methods do not perform as expected in multi-class settings.
- We introduce the idea of *example finding* to exemplar-free counting and demonstrate its utility in aiding a user in understanding what has been counted.

In the remainder of this paper, we outline the relevant prior work in Section 2 before introducing MCAC in Section 3. In Section 4, we introduce ABC and detail our experimental setup. Section 5 presents our state-of-the-art results in single- and multi-class class-agnostic settings. Finally, we summarise the impact of the work as a whole.

2 Related Work

Class-specific counting methods aim to enumerate the instances of a single or small set of known classes [1,6,9,26]. These methods struggle to adapt to novel classes, needing specific data and training for each type of object. To address these issues, Lu et al. [14] proposed class-agnostic counting, a framework where inference-time classes are not present during training. Still, most class-agnostic methods [19, 21, 22, 28], require exemplar images of the test-time class. These methods generally work by creating a sufficiently general feature space and applying some form of matching to regress a density map of the counted objects.

Recent works, RepRPN [18], CounTR [13], ZSC [27], CLIP-count [10], RCC [7] and LOCA [4] do away with exemplar images at inference-time, removing the need for intervention during deployment. RepRPN is a two-step method which proposes regions likely to contain an object of interest and then uses them for an exemplar-based density map regression method. It proposes more than one bounding box and enumerates them separately. RepRPN performs poorly in comparison to other contemporary methods. ZSC [27] uses a multi-stage process in which a text input is used to generate a generic image of the type to be counted that is then used to find exemplar patches. These exemplar patches then act as the input to an exemplar-based method [21]. CounTR [13] uses a large vision-transformer encoder-decoder to regress a density map of instance locations. It is trained in a mixed few/zero-shot way, applying understanding gained from exemplar-based examples to exemplar-free cases. LOCA [4] also uses a vision transformer backbone and can perform both few- and zero-shot counting. LOCA separately extract the shape and appearance of exemplar and non-exemplar objects to create a more informed object prototype. This prototype is then matched to areas of the image to generate a density map prediction.

It has been assumed that current exemplar-based methods can function in multi-class settings. However, this has not been proven rigorously as the main dataset for class-agnostic counting (FSC-133/147 [7,19]) contains only one labelled class per image. In fact, we show in Sec. 5 that these methods perform poorly in contexts with multiple types present. FSC-133/147 being single-class has also explicitly motivated work such as RCC, which regresses a single scalar count from an image. It is trained without exemplar images and uses scalar supervision instead of density maps. Even with the constraint of only counting one kind of object and with no further direction on the type to count, RCC achieves competitive results with exemplar-based methods on FSC-147, showing the limitations of this dataset.

While large models with image inputs [16] like SAM [11] would seem to be able to effectively count objects of arbitrary types, in fact these methods have poor numerical understanding [15], and perform unsatisfactorily on counting tasks especially when images have small objects or a high density of objects. Recently, Paiss *et al.* [17] attempted to utilise multi-modal deep learning to count. Specifically, they introduce a method that teaches a CLIP model a coherent understanding of words related to counting. However, they are only able to achieve this up to 10 objects.

3 MCAC Dataset

There are currently no datasets suitable for class class-agnostic counting problems with multiple types of object present at once. This significantly impacts the research into addressing these tasks. In order to facilitate the development of multi-class class-agnostic counting methods as well as the evaluation of prior work, we introduce MCAC, the first multi-class class-agnostic counting dataset.

While the deployment query scenario, 'counting given an unlabelled image of objects', is natural, the training and quantitative evaluation of methods to address it is not. To facilitate training and evaluation of methods in multi-class settings, we need images with multiple objects of multiple types. To evaluate a methods generalisability to unseen object types, the classes present in the images need to be mutually exclusive between training, validation and testing. It is infeasible to gather natural images with (a) a wide variety of classes, (b) a wide variety of the number of times an object appears in an image, and (c) no repetition of the types of object between the train, test, and validation splits. Using synthetic images allows the above constraints to be satisfied while also providing a high level of precision and accuracy in the labels for each image. As shown in Sec. 5.2, the understanding gained from training on synthetic data is general enough to apply to standard photographic datasets.

MCAC contains images with between 1 and 4 classes of object (mean of 1.75) and between 1 and 300 instances per class (mean of 47.66). MCAC has three data splits: training with 4756 images drawn from 287 classes; validation 2413 images drawn from 37 classes, and testing with 2114 images drawn from 19 classes. MCAC-M1 is the single-class subset of the MCAC images which have only one class present per image. MCAC-M1 totals 4259 images, with a mean of 114.89 instances per image. These distributions were designed to replicate that of real-world counting tasks.

All instances in an image have associated class labels, model labels, center coordinates, bounding box coordinates, segmentation maps, unoccluded segmentation maps, and occlusion percentages. The occlusion percentage is calculated as $1 - \frac{A_0}{A_1}$, where A_0 is the number of pixels in the final image and A_1 is the number of pixels that would be seen if the object was unoccluded and completely within the bounds of the image.

Objects are 'dropped' into the scene, ensuring random locations and orientations. As objects in real settings often vary in size, we vary the size of objects by $\pm 50\%$ from a random nominal size. We also vary the number, location, and intensity of lights present. Models and textures are drawn from ShapeNetSem [20].

Both exemplar-based and exemplar-free methods bump into problems of ambiguity. If there are objects of varied levels of generality, which boundary should be used? For example, on a chess board with a single white pawn as the exemplar, should the count be of all the pieces, all the white pieces, all the white pawns, all the pawns, and so on? Figure 8 shows examples from FSC-147 of cases with an ambiguity of what is to be counted.

Given the infeasibility of defining every possible way of grouping the objects present in an image, we define a single way of grouping the objects: an identical



Fig. 3: The ABC123 pipeline. Our method learns to count objects of multiple novel classes without needing exemplar images. During training and quantitative evaluation, the matcher aligns the unguided predictions to the ground truth labels. To aid a user in understanding the results, the example prediction stage locates instances associated with each generated count.

mesh and texture, independent of size or orientation. We do, however, acknowledge the existence of other *valid-but-unknown* counts, the unlabelled ways of grouping the objects.

4 Method

Our method, ABC123, takes an image with multiple instances of objects of multiple types and regresses the count of each type. This is achieved *blind*, i.e. on objects of arbitrary classes with no requirement to have seen the object class during training or to have an exemplar image to define the type during inference. We achieve this by first regressing density maps for each type then enumerating the instances using integration. To facilitate training and evaluating ABC123 in an exemplar-free way, we propose a matching stage. To further increase the interpretability of the outputs of ABC123, we design an example discovery stage which finds specific instances of the counted object. The pipeline of ABC123 is presented in Fig. 3.

4.1 Density Map Regression

For each image, there are m classes present, each with an associated ground truth count y and density map d. We regress \hat{m} counts and density map predictions, \hat{y} and \hat{d} respectively. \hat{m} acts as upper-bound of the number of counts ABC123 can regress.

$$\hat{y} = \sum_{h,w} \hat{d}_{(h,w)} \tag{1}$$

where $\hat{d}_{(h,w)}$ denotes the density value for pixel (h, w). We achieve this by using \hat{m} convolutional up-sampling heads on top of a vision transformer backbone [5]. We use a vision transformer backbone due to its globally receptive field and self-attention mechanism, which Hobley and Prisacariu [7] showed is helpful to generating a complex understanding in counting settings. Each head regresses a single pixel-wise density map prediction and count prediction from a patch-wise low-resolution high-dimensional feature space. Similar to other contemporary methods [12, 13, 19, 28], we use the pixel-wise error $||d - \hat{d}||_1$ as our loss, where d and \hat{d} are the ground truth and predicted density maps.

4.2 Matching

In single-class or exemplar-based settings, there is a single prediction-label pair. However, in multi-class exemplar-free settings, there are multiple predictions as well as multiple labels, without a clearly defined pairing. This resembles other open-set problems like class-discovery [24, 29] and clustering [8, 25], where the number and cardinality of new classes is not necessarily known. In keeping with these fields and to facilitate training and quantitative evaluation, we find correspondences between the set of m known counts and the set of \hat{m} predicted counts. The correspondence matrix is defined as $\mathcal{X} = \{0, 1\}^{\hat{m} \times m}$ where $\mathcal{X}_{i,j} = 1$ iff prediction i is assigned to label j. A problem instance is described by an $\hat{m} \times m$ cost matrix \mathcal{C} , where $\mathcal{C}_{i,j}$ is the cost of matching prediction i and ground truth label j. The goal is to find the complete assignment of predictions to labels of minimal cost. Formally, the optimal assignment has cost

$$\min_{\mathcal{X}} \sum_{i=0}^{m} \sum_{j=0}^{\hat{m}} \mathcal{C}_{i,j} \cdot \mathcal{X}_{i,j}$$
(2)

Specifically, our cost function is defined as the the pixel-wise distance of the normalised ground truth density map d_i and the predicted density map \hat{d}_j .

$$C_{i,j} = \left\| \frac{d_i}{||d_i||_2} - \frac{\hat{d}_j}{||\hat{d}_j||_2} \right\|_2$$
(3)

The normalisation ensures the matching is done on the locality of the counted objects rather than the magnitude of the prediction itself. We use the Hungarian algorithm, specifically the Jonker-Volgenant algorithm outlined in Crouse [3], to solve for \mathcal{X} robustly.

The supervision loss for each image is the sum of the L_1 difference of the ground truth density maps and their matched predictions as:

$$\mathcal{L} = \sum_{i,j}^{m,\hat{m}} ||d_i - \hat{d}_j||_1 \cdot \mathcal{X}_{i,j}$$
(4)

It should be noted that every label has an associated prediction, but the inverse is not the case as generally $\hat{m} > m$. This means we do not impose a



Fig. 4: Example Finding. Instead of using exemplars to define the count, we count 'blind' and then find meaningful bounding boxes to aids a user in understanding what has been counted. The examples are found using the query image, learnt features, our regressed density maps and a SAM network.

loss on the unmatched density maps. This allows the network to generate more nuanced count definitions as it does not punish valid-but-unknown counts which are likely present in any counting setting. As is usual [24,25,29], we use the same matching procedure to evaluate our performance at inference-time. However, during deployment, when there are no ground-truth density maps, the matching is both unnecessary and impossible. We instead combine any similar predictions, remove predictions of zero, and present the user with the predictions.

4.3 Example Discovery

While exemplar-free counting saves a user time, as no manual intervention is required, it does require the user to interpret the results. A set of scalar counts or density maps can be unclear as it is not always obvious which count corresponds to which type of object in the input image, especially in high density situations; see Figure 4. To aid the user in understanding to which class a generated count corresponds, we propose flipping the usual exemplar-based paradigm. Instead of using exemplar images to define the type to count, we find examples of the type that was counted.

To find examples corresponding to a given count, we first find example points which are high in the corresponding density map while low in the others. To increase the diversity of the found examples we select the points with the largest latent feature distance. We use these points as seed inputs for a pre-trained segmentation method, SAM [11]. The end user is presented with cropped areas of the query image centred on these segmentations, as in Fig. 4.

4.4 Experiments

We use a ViT-Small [5] backbone due to its lightweight nature and for comparison to methods that use the ResNet-50 backbone, such as FamNet and BM-Net [19,21]. ViT-S has a similar number of parameters (21M vs 23M), throughput (1237im/sec vs 1007im/sec), and supervised ImageNet performance (79.3% vs 79.8%) as ResNet-50 [23]. Our choice of ViT-S for ABC123 limits the input



Fig. 5: Comparison to other methods on MCAC. ABC123 produces more accurate results than the exemplar-based methods without using exemplar images. The ground truth (GT) and predicted counts are shown in the top right corner of their respective density maps.

resolution to (224×224) as opposed to the $(\geq 384 \times 384)$ resolution used by the method we compare to with ResNet-50 or larger ViT backbones. The effect of using ResNet or ViT-S backbones in various counting methods was discussed in detail in Hobley and Prisacariu [7] where they found that neither was inherently superior and the performance varied dependant on the method.

Since vision transformers typically demand substantial training data, we initialise our transformer backbone with weights from Caron *et al.* [2]. This selfsupervised pre-training endows the network with an understanding of meaningful image features prior to exposure to our dataset and without supervision. This reduces the risk of overfitting when the model is then trained on MCAC.

Our counting heads, comprised of 3 Conv-ReLU-Upsample blocks, increase the patch-wise resolution of the trained counting features from $k \times (28 \times 28)$ to a pixel-wise density map prediction of $\hat{m} \times (224 \times 224)$, where k is the dimensionality of the transformer features and \hat{m} is the number of predicted counts. For ABC123, k = 384. We set $\hat{m} = 5$ to ensure there is the capacity to generate a count per defined class in MCAC and at least one valid-but-unknown count.

ABC123 trains in less than eight hours using two 1080Tis. It takes less than two hours to train just the head with a frozen backbone (ABC123*). During training we use an Adam optimiser, a batch-size of 2 and a learning rate of $3 * 10^{-5}$ which halves every 35 epochs for a total of 100 epochs. The example discovery stage uses a frozen pretrained ViT-B SAM model [11].

Here we lay out our usage of MCAC, which we used to generate our results. We recommend future works use a similar approach and we will release code for a PyTorch dataset to enable easy adoption. We exclude objects that are more than 70% occluded by either other objects or the edge of the frame.

MCAC enables the use of true pixel-wise density maps. While this increases the accuracy of our method, we found this significantly decreased the performance of other methods, especially those with test-time adaptations. For fairness, we use the standard [12, 13, 19, 28] pseudo-density-map for all methods,



Fig. 6: Results of our method on images with 4 classes. ABC123 is able to generate accurate counts and meaningful density maps from images with four novel classes. MCAC has between one and four classes of object per image.

including our own. These are generated by placing a Gaussian kernel centred on the center pixel of each object.

When training exemplar-based methods, we take bounding boxes randomly from instances with less than 30% occlusion. We evaluate these methods using the bounding boxes of the three least occluded instances.

5 Results

We evaluate our method against two trivial baseline methods, predicting the training-set mean or median count for all inference images. As there are no previous multi-class exemplar-free class-agnostic counting methods, we compare ABC123 to exemplar-based methods using separate exemplars from each of the classes present. We compare our method to FamNet [19], BMNet [21] CounTR [13] and LOCA [4] on MCAC and FSC-133/147 as these are the current state-of-the-art methods with publicly available implementations. Additionally, we also compare to RCC [7] and CounTR in its zero-shot configuration on MCAC-M1, the subset of MCAC with only a single type of object present per image.

As in Xu *et al.* [27], we use Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Normalized Absolute Error (NAE), and Squared Relative Error (SRE) to evaluate the performance of each method.

$$MAE = \frac{1}{nm} \sum_{j=1}^{n} \sum_{i=1}^{m_j} |y_i - \hat{y}_i|, \qquad RMSE = \sqrt{\frac{1}{nm} \sum_{j=1}^{n} \sum_{i=1}^{m_j} (y_i - \hat{y}_i)^2} \qquad (5)$$

$$NAE = \frac{1}{nm} \sum_{j=1}^{n} \sum_{i=1}^{m_j} \frac{|y_i - \hat{y}_i|}{y_i}, \qquad SRE = \sqrt{\frac{1}{nm} \sum_{j=1}^{n} \sum_{i=1}^{m_j} \frac{(y_i - \hat{y}_i)^2}{y_i}} \qquad (6)$$

where n is the number of test images, m_j is the number of classes in image j, y_i and \hat{y}_i are the ground truth and predicted number of objects of class i in image j.

5.1 MCAC

We achieve significantly better result compared to FamNet, BMNet, CounTR and LOCA on MCAC both quantitatively and qualitatively without needing exemplars; see Tab. 1 for results and Fig. 5 for comparative examples. As seen in Fig. 5, FamNet often fails to discriminate between objects of different classes when they are visually similar or in high density applications. Both quantitatively and qualitatively, BMNet and CounTR outperform FamNet. However, in many cases, they appear to count the 'most obvious' objects in the image regardless of the provided exemplar images. This behaviour is present but is less prevalent with LOCA. Our method performs well on images with up to 4 classes even when they have high intra-class appearance variation, such as having different colours on different sides, and low inter-class variation; see Fig. 6. A downside to current exemplar-based class-agnostic counting methods is that while they have some multi-class capabilities, they all take a single exemplar at a time and produce only one count. This is slow and inefficient as compared to our approach which generates all counts simultaneously.

As would be expected, the performance of all methods improves when evaluating on MCAC-M1, the images from MCAC with only a single class present; see Tab. 2. This is due to a lack of ambiguity of the type to be counted. This was more significant when the methods were trained on MCAC-M1 instead of MCAC. In this training configuration, the methods generally learnt a broader definition of similarity as there was no chance they would accidentally combine classes or count instances from another class. RCC performs well on MCAC-M1, showing the strength of the simple count-wise loss in cases where there is little ambiguity as to what is to be counted. In contrast to other methods, ABC123 trained on MCAC-M1 has similar performance to when it is trained on the full MCAC dataset, demonstrating that it avoids issues concerning intra-class variance and combining classes. Training ABC123 with only a single head ($\hat{m} = 1$) and no matching stage has very similar performance to using its default ($\hat{m} = 5$) configuration with a matching stage. This increases our belief that the matching head does not provide an unfair advantage to our method's quantitative results.

5.2 Applicability to FSC-147/133

ABC123 trained on only MCAC, a synthetic dataset, produces accurate results and outperforms other contemporary methods when applied to FSC-133/147,

		Val Set				Test Set				
Method	Shots	MAE	RMSE	NAE	SRE	MAE	RMSE	NAE	SRE	
Mean Median	N/A N/A	$39.87 \\ 36.25$	$53.56 \\ 58.15$	$\begin{array}{c} 3.07\\ 1.51 \end{array}$	$11.40 \\ 6.70$	$42.67 \\ 39.81$	$59.68 \\ 65.36$	$2.79 \\ 1.38$	$10.93 \\ 6.73$	
$\begin{array}{l} Exemplar-based\\ FamNet+ [19]\\ BMNet+ [21]\\ CounTR [13]\\ LOCA [4] \end{array}$	3 3 3 3	24.76 15.83 15.07 10.45	$\begin{array}{c} 41.12 \\ 27.07 \\ 26.26 \\ 20.81 \end{array}$	$1.12 \\ 0.71 \\ 0.63 \\ 0.43$	$6.86 \\ 4.97 \\ 4.79 \\ 4.18$	$26.40 \\ 17.29 \\ 16.12 \\ 10.91$	45.52 29.83 29.28 22.04	$1.04 \\ 0.75 \\ 0.67 \\ 0.37$	$6.87 \\ 6.08 \\ 5.71 \\ 4.05$	
Exemplar-free ABC123 ☆ ABC123	$\begin{array}{c} 0 \\ 0 \end{array}$	14.64 8.96	23.67 15.93	0.46 0.29	2.97 2.02	15.76 9.52	25.72 1 7.64	0.45 0.28	3.11 2.23	

Table 1: Comparison to SOTA methods on MCAC. We significantly outperform methods which use exemplar images and test-time adaptation without requiring them. ABC123* denotes our method trained with a frozen pre-trained backbone.

the standard, more complex, photographic dataset, as seen in Figure 7 and Table 3. As the standard benchmark evaluation metrics rely on absolute count error they are all very sensitive to even a small number of very dense images. This phenomenon was discussed in Hobley and Prisacariu [7]. We found that the errors on the few images with counts between 300 (the largest count in MCAC) and 3000 (the largest count in FSC) corrupted the metrics, making comparison to other literature more difficult. For this reason these images are excluded from the quantitative evaluation in Table 3. These exclusions amount to 3.0% of the validation set and 1.1% of the test set. It should be noted that the relative rankings of the methods with and without this exclusion remain the same.

While ABC123 performs well on FSC, it often finds valid-but-unknown counts. As seen in Fig. 8, the generated counts are correct for the type of object counted, but the type counted may not be aligned with the labels in the original dataset. Classes are often divided into sub-classes, and unlabelled classes are discovered.

This is due to a difference in definition of what is similar. MCAC associates a count with objects of the same mesh and texture, however, FSC is labelled by hand uses high-level semantic understanding so often groups objects with significantly different geometries, colours, or textures. Interestingly, an unguided segmentation method [11], which identifies instances' relations often finds the same class divisions as ABC123, shown in Figure 8. To generate quantitative results, we borrow the approach of other open set methods [8,11,15], combining sub-classes. We perform this mapping by combining the density maps of subclass counts, either by the summation of both the separate counts or by trying to combine the density maps using the maximum density at a given point and then counting the instances. The maximum density map configuration produces results which are competitive with other contemporary methods while the density map summation is clearly SOTA, as presented in Table 3. The two approaches only differ in cases where sub-class density maps overlap.

Table 2: Comparison to SOTA methods on MCAC-M1. MCAC-M1 is the subset of MCAC with only one class present per image. Methods are either trained on the full multi-class dataset (\checkmark) or MCAC-M1 (\bigstar). ' \hat{m} ' denotes the number of predictions the method generates per query and 'Shots' denotes the number of exemplar images per query at inference time. CounTR† is an exemplar-free adaption of CounTR as in Hobley and Prisacariu [7]. ABC123 outperforms other methods when trained in single or multi-class settings.

				Val Set			Test Set				
Method	Multi-Class Training	Shots	\hat{m}	MAE	RMSE	NAE	SRE	MAE	RMSE	NAE	SRE
Mean	N/A	N/A	N/A	53.36	67.14	3.53	13.46	58.54	75.58	3.37	13.27
Median	N/A	N/A	N/A	45.98	76.64	1.08	6.68	51.35	86.61	1.03	7.00
Exemplar-base	d Training										
FamNet+ [19]	1	3	1	24.97	48.63	0.36	3.79	28.31	54.88	0.35	3.97
FamNet+ [19]	X	3	1	12.54	24.69	0.37	4.71	13.97	26.19	0.25	2.12
BMNet+[21]	1	3	1	11.70	23.08	0.26	2.39	11.57	22.25	0.24	1.96
BMNet+[21]	X	3	1	6.82	12.84	0.25	2.95	8.05	14.57	0.19	1.43
CounTR [13]	1	3	1	11.44	21.37	0.33	2.36	10.91	21.70	0.29	2.01
CounTR [13]	1	0	1	13.57	25.53	0.30	2.48	13.09	25.72	0.29	2.41
CounTR [13]	X	3	1	9.00	16.91	0.41	3.56	9.96	18.92	0.38	2.93
CounTR [13]	X	0	1	9.16	17.13	0.42	3.56	10.10	19.10	0.40	3.02
LOCA [4]	1	3	1	5.62	12.24	0.15	1.73	6.25	13.09	0.12	1.14
LOCA [4]	×	3	1	5.01	11.47	0.22	3.35	6.52	13.37	0.15	1.36
$Exemplar\-free$	Training										
$CounTR^{\dagger}$ [13]	X	0	1	11.46	21.24	0.35	2.78	12.54	23.84	0.31	2.38
RCC [7]	X	0	1	7.78	15.40	0.24	2.71	8.81	16.92	0.19	1.73
LOCA [4]	X	0	1	5.46	11.74	0.22	2.90	6.94	14.58	0.19	1.70
ABC123	X	0	5	10.78	18.83	0.28	1.97	13.23	24.57	0.29	2.39
ABC123	X	0	1	11.38	19.73	0.40	3.51	14.31	25.40	0.37	2.79
ABC123	1	0	5	10.98	18.85	0.30	1.93	13.13	23.93	0.29	2.18
ABC123	×	0	5	5.82	11.74	0.15	1.22	7.54	15.30	0.21	1.87
ABC123	X	0	1	5.85	12.91	0.24	3.37	7.53	15.69	0.22	2.19
ABC123	1	0	5	6.08	12.62	0.16	1.22	6.82	14.70	0.16	1.51

5.3 Validating the Number of Predictions

It should be noted that as the matching stage uses the ground-truth density maps, it could be used to significantly benefit a method's quantitative results without improving its deployment capabilities. Specifically, a method could generate a high number of diverse counts and use the matching stage to select the best one. We found this to be the case with our method, see Tab. 4 for the complete results. We believe, however, that this does not align with a more useful method in a deployment situation. This numerical gain derives purely from the matching stage, which is not present during deployment. In fact, during deployment, this would correspond to a much more difficult to interpret output as a user would have to figure out which of the many outputs was most relevant. We limit ourselves to regressing 5 predictions to minimise this behaviour while still allowing the network to generate valid-but-unknown counts. We also found that when high numbers of predictions were generated, fewer than half were used,



Fig. 7: ABC123 trained on MCAC applied to FSC-147 produces accurate counts. The ground truth and predicted counts are in the top left and top right corners.

i.e. the outputs of some heads were rarely or never picked. This is likely due to these heads not being matched frequently during training so the loss is rarely propagated back through them. There is also significant redundancy between the heads. The predictions of certain heads over the whole dataset were clearly similar and could be grouped. Of the 39 utilised heads when $\hat{m} = 100$, there were three groups of, respectively, 13, 6, and 4 heads that were very similar, lowering the effective number of utilised heads to 19.

6 Conclusion

In this work, we present ABC123, a multi-class exemplar-free class-agnostic counter, and show that it is superior to prior exemplar-based methods in a multi-class setting. ABC123 requires no human input at inference-time, works in complex settings with more than one kind of object present, and outputs easy to understand information in the form of examples of the counted objects. Due to this, it has potential for deployment in various fields. We also propose MCAC, a multi-class class-agnostic counting dataset, and use it to train our method as well as to demonstrate that exemplar-based counting methods may not be as robust as previously assumed in multi-class settings.

Table 3: Comparison to SOTA methods when trained on MCAC and applied to the cases in FSC147 with fewer than 300 objects. Combining sub-class density maps with a sum rather than a max is more effective as it is more accurate in cases where instances of the sub classes are spatially close or overlapping as both instances are counted completely. CounTR † is an exemplar-free modification of CounTR.

		Sub-Class		Val S	let		Test Set				
Method	Shots	Combine	MAE	RMSE	NAE	SRE	MAE	RMSE	NAE	SRE	
Exemplar-based											
FamNet+	3	N/A	25.83	46.31	0.50	4.29	28.05	45.59	0.48	4.33	
BMNet+	3	N/A	29.47	53.15	0.51	4.72	30.74	52.00	0.47	4.68	
CounTR	3	N/A	21.22	40.28	0.47	3.85	21.09	40.79	0.38	3.66	
LOCA	3	N/A	25.70	48.64	0.45	4.30	29.93	49.89	0.48	4.62	
Exemplar-	free										
CounTR †	0	N/A	23.50	45.83	0.42	4.06	23.57	42.00	0.39	3.85	
LOCA	0	N/A	29.37	54.01	0.51	4.84	33.96	56.66	0.53	5.17	
ABC123	0	Max	19.56	46.71	0.20	3.54	22.43	47.35	0.22	3.70	
ABC123	0	Sum	11.13	34.47	0.12	2.44	11.75	33.41	0.11	2.38	

Table 4: Effect of using more prediction heads. Increasing the number of prematching predictions improves the quantitative results of our method. However, as the number of heads increases, the percentage of heads that are frequently (> 0.4%)matched decreases.

		Val S	Set			Test			
\hat{m}	MAE	RMSE	NAE	SRE	MAE	RMSE	NAE	SRE	Head Utilisation
4	9.43	17.42	0.31	2.57	10.19	19.44	0.33	2.81	100%
5	8.96	15.93	0.28	2.02	9.52	17.64	0.28	2.23	100%
10	8.39	14.93	0.28	1.97	9.08	16.96	0.27	2.15	100%
20	7.78	13.75	0.26	1.82	8.29	15.53	0.24	1.95	85%
50	7.26	12.80	0.25	1.69	7.99	15.14	0.24	1.81	68%
100	7.11	12.81	0.23	1.59	7.43	14.48	0.21	1.72	39%



Fig. 8: ABC123 trained on MCAC applied to Ambiguous Images in FSC-147. In cases with class ambiguity, ABC123 often finds valid-but-unknown counts, i.e. they don't align with the human annotations. Similar to an unguided segmentation method (right), ABC123 discovers unlabelled classes, divides labelled classes into sub classes, or counts component parts of objects. The unguided segmentations are coloured by latent feature to demonstrate how they would be grouped.

References

- Cao, X., Wang, Z., Zhao, Y., Su, F.: Scale aggregation network for accurate and efficient crowd counting. In: Proceedings of the European conference on computer vision (ECCV). pp. 734–750 (2018)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9650–9660 (2021)
- Crouse, D.F.: On implementing 2d rectangular assignment algorithms. IEEE Transactions on Aerospace and Electronic Systems 52(4), 1679–1696 (2016). https://doi.org/10.1109/TAES.2016.140952
- 4. Djukic, N., Lukezic, A., Zavrtanik, V., Kristan, M.: A low-shot object counting network with iterative prototype adaptation. arXiv preprint arXiv:2211.08217 (2023)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
- Go, H., Byun, J., Park, B., Choi, M.A., Yoo, S., Kim, C.: Fine-grained multiclass object counting. In: 2021 IEEE International Conference on Image Processing (ICIP). pp. 509–513. IEEE (2021)
- 7. Hobley, M., Prisacariu, V.: Learning to count anything: Reference-less classagnostic counting with weak supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023)
- Hobley, M.A., Prisacariu, V.A.: Dms: Differentiable mean shift for dataset agnostic task specific clustering using side information. arXiv preprint arXiv:2305.18492 (2023)
- Hoekendijk, J., Kellenberger, B., Aarts, G., Brasseur, S., Poiesz, S.S., Tuia, D.: Counting using deep learning regression gives value to ecological surveys. Scientific reports 11(1), 1–12 (2021)
- Jiang, R., Liu, L., Chen, C.: Clip-count: Towards text-guided zero-shot object counting. arXiv preprint arXiv:2305.07304 (2023)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
- Lin, H., Hong, X., Wang, Y.: Object counting: You only need to look at one. arXiv preprint arXiv:2112.05993 (2021)
- Liu, C., Zhong, Y., Zisserman, A., Xie, W.: Countr: Transformer-based generalised visual counting. arXiv preprint arXiv:2208.13721 (2022)
- Lu, E., Xie, W., Zisserman, A.: Class-agnostic counting. In: Asian Conference on Computer Vision (2018)
- Ma, Z., Hong, X., Shangguan, Q.: Can sam count anything? an empirical study on sam counting. arXiv preprint arXiv:2304.10817 (2023)
- 16. OpenAI: Gpt-4 technical report (2023)
- Paiss, R., Ephrat, A., Tov, O., Zada, S., Mosseri, I., Irani, M., Dekel, T.: Teaching clip to count to ten. arXiv preprint arXiv:2302.12066 (2023)
- Ranjan, V., Hoai, M.: Exemplar free class agnostic counting. arXiv preprint arXiv:2205.14212 (2022)
- Ranjan, V., Sharma, U., Nguyen, T., Hoai, M.: Learning to count everything. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3394–3403 (2021)

- 16 Hobley and Prisacariu
- Savva, M., Chang, A.X., Hanrahan, P.: Semantically-Enriched 3D Models for Common-sense Knowledge. CVPR 2015 Workshop on Functionality, Physics, Intentionality and Causality (2015)
- Shi, M., Lu, H., Feng, C., Liu, C., Cao, Z.: 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. pp. 9529– 9538 (2022)
- 22. Sokhandan, N., Kamousi, P., Posada, A., Alese, E., Rostamzadeh, N.: A few-shot sequential approach for object counting. arXiv preprint arXiv:2007.01899 (2020)
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. pp. 10347–10357. PMLR (2021)
- Troisemaine, C., Lemaire, V., Gosselin, S., Reiffers-Masson, A., Flocon-Cholet, J., Vaton, S.: Novel class discovery: an introduction and key concepts. arXiv preprint arXiv:2302.12028 (2023)
- 25. Xie, J., Girshick, R., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: Balcan, M.F., Weinberger, K.Q. (eds.) Proceedings of The 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 48, pp. 478–487. PMLR, New York, New York, USA (20–22 Jun 2016), https://proceedings.mlr.press/v48/xieb16.html
- Xie, W., Noble, J.A., Zisserman, A.: Microscopy cell counting and detection with fully convolutional regression networks. Computer methods in biomechanics and biomedical engineering: Imaging & Visualization 6(3), 283–292 (2018)
- Xu, J., Le, H., Nguyen, V., Ranjan, V., Samaras, D.: Zero-shot object counting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15548–15557 (2023)
- Yang, S.D., Su, H.T., Hsu, W.H., Chen, W.C.: Class-agnostic few-shot object counting. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 870–878 (2021)
- Yang, Y., Xu, D., Nie, F., Yan, S., Zhuang, Y.: Image clustering using local discriminant models and global integration. IEEE Transactions on Image Processing 19(10), 2761–2773 (2010). https://doi.org/10.1109/TIP.2010.2049235