

# Single Domain Generalization for Crowd Counting

Zhuoxuan Peng, S.-H. Gary Chan

The Hong Kong University of Science and Technology

{zpengac, gchan}@cse.ust.hk

## Abstract

Due to its promising results, density map regression has been widely employed for image-based crowd counting. The approach, however, often suffers from severe performance degradation when tested on data from unseen scenarios, the so-called “domain shift” problem. To address the problem, we investigate in this work single domain generalization (SDG) for crowd counting. The existing SDG approaches are mainly for image classification and segmentation, and can hardly be extended to our case due to its regression nature and label ambiguity (i.e., ambiguous pixel-level ground truths). We propose MPCount, a novel effective SDG approach even for narrow source distribution. MPCount stores diverse density values for density map regression and reconstructs domain-invariant features by means of only one memory bank, a content error mask and attention consistency loss. By partitioning the image into grids, it employs patch-wise classification as an auxiliary task to mitigate label ambiguity. Through extensive experiments on different datasets, MPCount is shown to significantly improve counting accuracy compared to the state of the art under diverse scenarios unobserved in the training data characterized by narrow source distribution. Code is available at <https://github.com/Shimmer93/MPCount>.

## 1. Introduction

Image-based crowd counting is to estimate the number of people (or objects) in an image. The recent mainstream counting approach is density map regression, where a density map with continuous pixel value is predicted from the input image, and the total crowd count is obtained by summing all these pixel values. During training, the ground-truth density map is generated by smoothing point annotations given by the locations of human heads. Density map regression has been demonstrated to achieve encouraging results in crowd counting.

Existing training models predominantly assume that the testing data share the same source distribution, i.e., fully supervised learning. Such assumption, however, often breaks

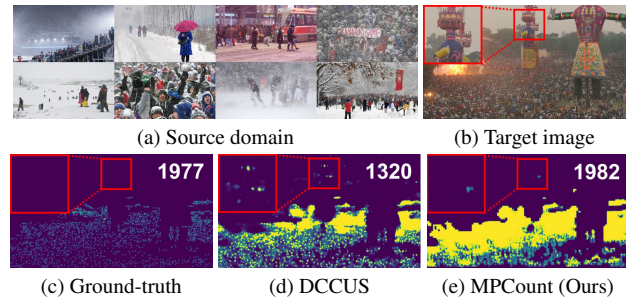


Figure 1. Single domain generalization for crowd counting. (a) Sample images from a single source domain  $\mathcal{S}$  with narrow distribution. (b) A challenging target image from an unseen scenario. (c) The ground-truth density map. (d) The predicted density map from the previous work DCCUS [6] trained on  $\mathcal{S}$ . (e) The predicted density map from our MPCount trained on  $\mathcal{S}$ . MPCount achieves much lower counting error than DCCUS and makes better predictions in the cropped area against the domain shift.

in practice due to changes in camera position, weather condition, etc. Such deviation from the training data (i.e. source domain) to unseen testing data (i.e. target domain) is known as *domain shift*, which severely degrades the performance of deep learning models in real-world deployment.

To mitigate the domain shift problem, domain adaptation (DA) approaches have been proposed to transfer knowledge from the source domain to the target one, typically by fine-tuning pre-trained models with target domain data [2, 10, 39, 44]. Despite promising results, the fine-tuning process of DA is usually laborious. Furthermore, target domain data may not be easily available in reality.

In contrast to DA, domain generalization (DG) aims to generalize models to any unobserved target domain without the need for target data. As data and model fine-tuning are not necessary at the target domain, DG is more convenient to be deployed in unseen or unpredictable domain. We consider in this work the challenging case of single domain generalization (SDG) for image-based crowd counting, where only one source domain, likely of narrow distribution (Fig. 1a), is used for training.

A common general SDG technique is to construct

domain-invariant features from different augmented versions of an image. Such approach has been shown promising for various tasks [5, 13, 34]. Independent of data augmentation, many other SDG approaches have been proposed for classification and segmentation to further boost performance. Unfortunately, most of them cannot be extended to crowd counting, a regression task without category information. Moreover, while labels in classification and segmentation tasks are precise and domain-invariant, the labels for density-based crowd counting are usually ambiguous. Such “label ambiguity” stems from the fact that the density values are generated from point annotations, and hence human heads and backgrounds may be assigned to similar value under different scenarios.

The recent work of [6] achieves DG for crowd counting based on dividing the source into multiple sub-domains. While encouraging, it requires a rather broad source domain distribution. We propose, for the first time, an SDG approach for crowd counting called MPCount without sub-domain division, thereof effective and extensible to narrow source distribution. We highlight its performance as compared with a state-of-the-art scheme in Fig. 1. MPCount achieves its superior performance with the following two novel components:

- *An attention memory bank (AMB) to tackle density regression:* MPCount employs only one attention memory bank (AMB) that takes a pair of features encoded from a crowd image and its data-augmented version as input. To cover continuous density values for regression with a finite memory size, AMB reconstructs each feature vector as an attention over memory vectors. These memory vectors learn domain-invariant representations from the pair of features of different styles but similar content (i.e. crowd densities). We present the content error mask (CEM) to eliminate domain-related content from input features, as characterized by a large discrepancy between instance-normalized feature element pairs. Furthermore, the newly employed attention consistency loss (ACL) enforces the similarity of attention scores produced from input features, ensuring the consistency of memory vectors.
- *Patch-wise classification (PC) to tackle label ambiguity:* To address the label ambiguity, MPCount employs a novel auxiliary task called patch-wise classification (PC). In this task, each crowd image is evenly divided into fixed-size square patches, say  $16 \times 16$ , classified into two classes, i.e., containing human heads or not. The density values in the areas classified as without head are filtered out during density regression, so that the pixel-level ambiguity is overcome by coarser but more accurate patch-wise binary labels.

To validate MPCount, we conduct extensive experiments on various crowd counting datasets including ShanghaiTech A (SHA) & B (SHB) and JHU-Crowd++.

We introduce a new challenging setting of narrow source distribution, where only images of the same category, such as “snow” (SN) and “fog/haze” (FH) in JHU-Crowd++, are used in training. We demonstrate that MPCount achieves excellent performance not only on traditional inter-dataset benchmarks, but also under our newly introduced narrow source setting. Let  $\mathcal{S} \rightarrow \mathcal{T}$  denote the case where the source domain is  $\mathcal{S}$  and the target domain is  $\mathcal{T}$ . As compared to the state of the art, MPCount is shown to reduce significantly counting error by 21.8% on SN  $\rightarrow$  FH, 18.6% on FH  $\rightarrow$  SN, 18.2% on SHB  $\rightarrow$  SHA and 9.5% on SHA  $\rightarrow$  SHB.

## 2. Related Works

### 2.1. Fully Supervised Crowd Counting

The mainstream approach to crowd counting is via density map regression, where each pixel of an image is assigned a count value, and the summation of these values yields an estimate of the total number of people. Most methods focus on improving counting accuracy under the fully supervised setting by employing novel network designs [1, 16, 17, 31, 45] or loss functions [20, 26, 36, 37]. While encouraging results have been shown, their performance usually drop significantly when evaluated on out-of-distribution data, due to their limited generalization ability. Besides, several works [4, 22, 38] have proposed auxiliary tasks in crowd counting for certain purposes. However, these tasks still require pixel-level predictions and contribute little to tackling label ambiguity.

### 2.2. Domain Adaptation (DA) for Crowd Counting

Domain adaptation (DA) has been widely studied in crowd counting to address domain shift by adapting source domain information to a particular target domain. While a few methods study supervised DA where labeled data from target domain are available [40, 44], most approaches tackle unsupervised DA, utilizing only unlabeled target domain data [2, 7, 18, 39, 41, 47]. Despite the notable advancements, DA methods normally require data from the target domain, which may not be easily available in practice.

### 2.3. Single Domain Generalization (SDG)

Domain Generalization (DG) approaches are designed to train a network with generalization ability solely utilizing source domain data, and single domain generalization (SDG) is the special case when only one source domain is available. There are various techniques for SDG, including 1) adversarial data generation [15, 25, 46, 50], 2) feature normalization/whitening transformation [5, 12, 23, 24, 43] and 3) domain-general network design, such as a convolution layer [35], vision transformer [32] and memory bank [3]. Among various SDG methods, a common general technique is to simulate domain shift with data augmenta-

tion and extract domain-invariant information from different augmented versions of features [5, 13, 43]. While SDG has attracted widespread attention in classification and segmentation, it is still a nascent topic in crowd counting due to its regression nature and label ambiguity.

DCCUS [6] studies SDG for crowd counting. In this method, the source domain is dynamically clustered into sub-domains simulated as meta-training and -test sets in a meta-learning strategy. Two types of memory modules and several losses are proposed to distinguish and record domain-invariant and -related information. DCCUS designs the memory mechanism specifically for regression, but the problem of ambiguous labels is not considered. Moreover, when the source domain follows a narrow distribution, the sub-domain division process may experience reduced effectiveness and affect the generalization ability of the counting model. MPCount, without sub-domain division, employs novel patch-wise classification to address the label ambiguity problem. Additionally, a single memory bank for regression is introduced without the need for sub-domain partitioning, enabling MPCount to generalize well even in source domains of narrow distributions.

### 3. MPCount

This section provides a detailed description of our MPCount scheme. We begin by reviewing the foundational concepts of crowd counting in Sec. 3.1. Next, in Sec. 3.2, we present a comprehensive overview of the entire scheme. We then elaborate the attention memory bank (Sec. 3.3) designed for density regression, followed by the content error mask (Sec. 3.4) and attention consistency loss (Sec. 3.5). After that, we discuss the patch-wise classification (Sec. 3.6) proposed to address the challenge of ambiguous labels. Finally, we detail the overall training loss in Sec. 3.7.

#### 3.1. Crowd Counting Preliminaries

First, we review the preliminary knowledge of deep learning-based crowd counting via density map regression. Our objective is to train a neural network  $\mathcal{N}$  with parameter  $\theta$  that takes in an input image  $\mathcal{I} \in \mathbb{R}^{H \times W \times 3}$  and outputs a density map  $\hat{\mathcal{D}} = \mathcal{N}(\mathcal{I}; \theta)$  of size  $H \times W$ . The estimated count  $\hat{c}$  is the sum of all pixel-wise density values in  $\hat{\mathcal{D}}$ , i.e.,

$$\hat{c} = \sum_{i=1}^H \sum_{j=1}^W \hat{\mathcal{D}}_{ij}. \quad (1)$$

Generation of ground-truth density maps usually follows the method in [28]. Given the human head point annotations  $\mathcal{H}$  where each  $h \in \mathcal{H}$  is the coordinate of a head location, the ground-truth density map  $\mathcal{D} \in \mathbb{R}^{H \times W}$  is calculated as the sum of the 2D Gaussian filter applied to each head coordinate  $h$ , i.e.,

$$\mathcal{D} = \sum_{h \in \mathcal{H}} \delta(x - h) \times G_{\sigma}(x), \quad (2)$$

where  $\delta(\cdot)$  is the discrete delta function, and  $G_{\sigma}(\cdot)$  is the Gaussian kernel with a fixed variance  $\sigma$ .

A common objective function to supervise the network  $\mathcal{N}$  can be written as

$$\mathbb{L}(\theta) = \|\mathcal{D}, \hat{\mathcal{D}}\|_2^2. \quad (3)$$

#### 3.2. Scheme Overview

The overall structure of our method is illustrated in Fig. 2. During training, photometric transformations are applied to the original image  $\mathcal{I}^{ori}$  to produce an augmented version  $\mathcal{I}^{aug}$ . The feature extractor encodes them into features  $\mathcal{F}^{ori}$  and  $\mathcal{F}^{aug}$ , and a content error mask (CEM)  $\mathcal{M}$  is applied to them. The masked features are reconstructed into  $\tilde{\mathcal{F}}^{ori}$  and  $\tilde{\mathcal{F}}^{aug}$  via an attention memory bank (AMB) and passed to the density head for crowd density prediction. Meanwhile, the highest-level features encoded by the feature extractor,  $\mathcal{Z}^{ori}$  and  $\mathcal{Z}^{aug}$ , are fed into the patch-wise classification (PC) head, and the predicted PC maps (PCMs)  $\hat{\mathcal{C}}^{ori}$  and  $\hat{\mathcal{C}}^{aug}$  are binarized and resized to serve as a mask to filter out areas without human heads in the estimated density maps  $\hat{\mathcal{D}}^{ori}$  and  $\hat{\mathcal{D}}^{aug}$ . The final predictions  $\mathcal{D}^{ori}$  and  $\mathcal{D}^{aug}$  are two density maps with certain areas masked out by the PCMs. The entire training process is supervised under a combination of density losses, PC losses and the attention consistency loss (ACL). The inference process solely employs  $\mathcal{I}^{ori}$  as input, and outputs a single density map  $\mathcal{D}^{ori}$ .

#### 3.3. Attention Memory Bank (AMB)

The attention memory bank (AMB) is designed to automatically learn domain-invariant features for density regression. [14] proposes a memory bank where memory vectors are updated individually, each corresponding to a certain category. However, this design cannot be applied to regression tasks, since a finite number of memory vectors cannot cover continuous density values. Inspired by [6], we reconstruct each feature vector as an attention over memory vectors, so that any density value may be represented as a linear combination of representative values in the memory bank.

The AMB  $\mathcal{V} \in \mathbb{R}^{M \times C}$  consists of  $M$  memory vectors of dimension  $C$ . Given a flattened input feature map  $\mathcal{F} \in \mathbb{R}^{HW \times C}$ , we first compute the attention scores  $\mathcal{A}$  between the feature  $\mathcal{F}$  (query) and the memory  $\mathcal{V}$  (key). Then the reconstructed feature map  $\tilde{\mathcal{F}}$  is calculated as a linear combination of  $\mathcal{V}$  (value) with  $\mathcal{A}$  as weights. The reconstruction process can be summarized as follows:

$$\mathcal{A}(\mathcal{F}, \mathcal{V}) = \text{Softmax} \left( \frac{\mathcal{F}\mathcal{V}^T}{\sqrt{C}} \right), \quad (4)$$

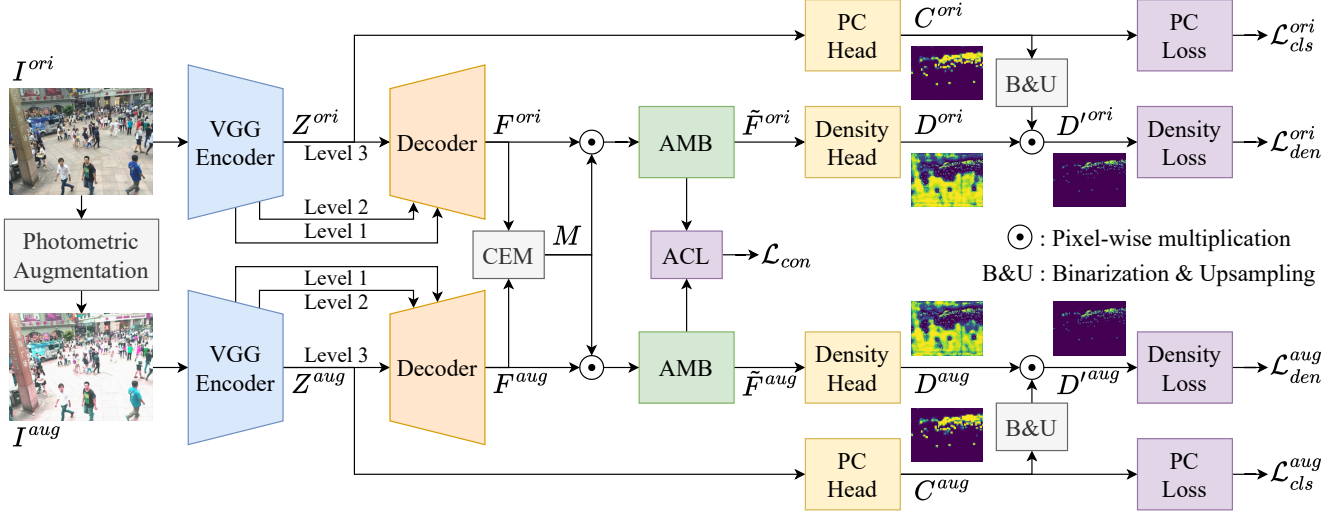


Figure 2. The overall training pipeline of our proposed MPCount. All identical modules in this diagram share the same weights. In the encoder-decoder structure, a higher level indicates a deeper feature.

and

$$\tilde{\mathcal{F}} = \mathcal{A}\mathcal{V}. \quad (5)$$

While [6] relies on multiple memory banks corresponding to different sub-domains to distinguish domain-invariant and -related information, we only employ a single AMB and propose the novel CEM (Sec. 3.4) and ACL (Sec. 3.5) to ensure that it stores domain-invariant representations. To this end, our feature reconstruction mechanism can work without sub-domain division, thereof is effective even with a narrow source distribution.

### 3.4. Content Error Mask (CEM)

CEM is proposed to guarantee the resemblance of content contained in the pair of input features by excluding possible domain-related content information. A common technique to disentangle style and content in a feature is instance normalization (IN) [33], where the feature statistics (channel-wise mean and variance) are considered to contain style information, while the instance normalized features retain content information. Based on this idea, we assume that the discrepancy of instance normalized features reflect effects of domain shift on feature content and may mislead AMB to learn domain-related information. Therefore, we filter out elements in the input features with difference between their instance normalized versions above a certain threshold, resulting in a pair of features with similar content information for domain-invariant feature reconstruction.

More specifically, given the features extracted from the original and augmented images,  $\mathcal{F}^{ori}$  and  $\mathcal{F}^{aug} \in \mathbb{R}^{H \times W \times C}$ , we define the content error mask  $\mathcal{M} \in \mathbb{R}^{H \times W \times C}$  as follows:

$$\mathcal{M}_{ijk} = \begin{cases} 1 & \text{if } |\text{IN}(\mathcal{F}^{ori})_{ijk} - \text{IN}(\mathcal{F}^{aug})_{ijk}| \leq \alpha \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where  $\alpha$  is the threshold indicating whether the error value reflects possible inconsistency of content information caused by domain shift.

Given  $\mathcal{F}$  to be one of  $\mathcal{F}^{ori}$  and  $\mathcal{F}^{aug}$ , we filter out the domain-related content information by

$$\mathcal{F}' = \text{Dropout2D}(\mathcal{F} \odot \mathcal{M}). \quad (7)$$

The random 2D dropout is applied to  $\mathcal{F}'$  to prevent the memory from relying solely on certain channels.

### 3.5. Attention Consistency Loss (ACL)

The style information of  $\mathcal{F}^{ori}$  and  $\mathcal{F}^{aug}$  preserved in Sec. 3.4 may produce a large discrepancy on their attention distributions over the memory vectors in AMB even if they contain identical content information. To ensure that each memory vector consistently store specific domain-invariant representations, we propose the attention consistency loss (ACL) to enforce the similarity of attention scores produced from  $\mathcal{F}^{ori}$  and  $\mathcal{F}^{aug}$ .

We consider  $\mathcal{A}$  in Eq. 4 as a distribution, and the ACL  $\mathbb{L}_{con}$  is calculated as the distance between the distribution  $\mathcal{A}^{ori}$  produced by  $\mathcal{F}^{ori}$  and  $\mathcal{A}^{aug}$  produced by  $\mathcal{F}^{aug}$ . Here the simple Euclidean distance is selected as the distance measure due to its computational stability during training:

$$\mathbb{L}_{con} = \|\mathcal{A}^{aug}, \mathcal{A}^{ori}\|_2^2. \quad (8)$$



Figure 3. Illustration of label ambiguity and how PCM tackles it. Point A is a typical example which belongs to a human head but is assigned the density value 0 as the background in the density map due to varying head sizes. In the PCM, the patch covering A is correctly classified as containing heads, thanks to the patch-level accuracy of the labels.

### 3.6. Patch-wise Classification (PC)

As discussed in Sec. 3.1, ground-truth density maps are usually calculated only based on point annotations. In different scenarios, pixels of a human head and the background might be assigned the same density value due to various head sizes (Fig. 3a), violating the common assumption that labels are accurate and domain-invariant, thereof resulting in the ambiguous labels in crowd counting.

To tackle the label ambiguity, we propose a novel auxiliary task named patch-wise classification (PC) and its corresponding supervision signal, patch-wise classification map (PCM), as illustrated in Fig. 3b. A PCM divides an image evenly into  $P \times P$  ( $P = 16$  empirically) patches and classifies each patch into two categories: containing human heads or not. Such patch-level predictions mitigate the uncertainty in pixel-level density maps by providing coarser but more accurate patch-level information.

In practice, a ground truth PCM  $\mathcal{C}$  with patch size  $P$  can be calculated as:

$$\mathcal{C}_{ij} = 1 - \delta \left( \sum_{k=iP}^{(i+1)P} \sum_{l=jP}^{(j+1)P} \mathcal{D}_{kl} \right), \quad (9)$$

where  $\mathcal{D}$  is the ground truth density map and  $\delta(\cdot)$  is the discrete delta function.

Our model predicts a PCM  $\hat{\mathcal{C}}$  supervised by the binary cross entropy (BCE) loss:

$$\mathbb{L}_{cls} = \text{BCE}(\mathcal{C}, \hat{\mathcal{C}}). \quad (10)$$

Subsequently,  $\hat{\mathcal{C}}$  is binarized and resized to obtain  $\mathcal{C}'$  which matches the dimensions of the predicted density map  $\hat{\mathcal{D}}$ .  $\mathcal{C}'$  is then used to mask the regions classified without crowds in  $\hat{\mathcal{D}}$ , resulting in the final density map  $\mathcal{D}'$ :

$$\mathcal{D}' = \hat{\mathcal{D}} \odot \mathcal{C}'. \quad (11)$$

Finally, the density regression head is supervised with the ground truth density map  $\mathcal{D}$ , and the common Euclidean distance in Eq. (3) is used as the objective function:

$$\mathbb{L}_{den} = \|\mathcal{D}, \mathcal{D}'\|_2^2. \quad (12)$$

### 3.7. Model Training

The overall training loss is a combination of density losses  $\mathbb{L}_{den}^{ori}, \mathbb{L}_{den}^{aug}$ , PC losses  $\mathbb{L}_{cls}^{ori}, \mathbb{L}_{cls}^{aug}$  and the attention consistency loss  $\mathbb{L}_{con}$ , which can be written as

$$\mathbb{L} = \mathbb{L}_{den}^{ori} + \mathbb{L}_{den}^{aug} + \lambda_{cls} (\mathbb{L}_{cls}^{ori} + \mathbb{L}_{cls}^{aug}) + \lambda_{con} \mathbb{L}_{con}, \quad (13)$$

where  $\lambda_{cls}$  and  $\lambda_{con}$  are weighting parameters to balance different loss terms.

## 4. Illustrative Experimental Results

In this section, we first introduce the datasets used in experiments in Sec. 4.1. Next, the implementation details are specified in Sec. 4.2, then evaluation metrics are described in Sec. 4.3. We present the results compared with state of the art in Sec. 4.4. Finally, we conduct ablation studies and other additional analysis in Sec. 4.5.

### 4.1. Datasets

We evaluate our method on four mainstream crowd counting datasets: ShanghaiTech Part A & B, UCF-QNRF and JHU-Crowd++.

- *ShanghaiTech* [45] is composed of two parts, SHA (A) and SHB (B). SHA contains 300 training images and 182 testing images, while SHB contains 400 training images and 316 testing images. Images in SHA are collected from the Internet and feature highly crowded scenes. In contrast, SHB is captured from several streets in Shanghai, and its images exhibit crowd densities generally lower than those in SHA.
- *UCF-QNRF* (Q) [11] consists of 1201 training images and 334 testing images. It is a challenging dataset with a large range of crowd densities, scenes, viewpoints and lighting conditions.
- *JHU-Crowd++* [29] is a large-scale dataset containing 4372 images, with 2722, 500 and 1600 images in the training, validation and testing set, respectively. In addition, image-level labels are also provided in this dataset, including 16 types of scenes and 4 types of weather conditions. For scene annotations, there are 879 images labeled “stadium” (SD) and 573 labeled “street” (ST), which are utilized as two source domains. For weather annotations, we use 201 images labeled “snow” (SN) and 168 images labeled “fog/haze” (FH) as source domains. In each domain, 80% of the images are selected as the training set while the rest 20% are for testing. Data with the same label belong to a narrower distribution than mainstream datasets, thus are more challenging for SDG. Illustrative samples in these scene-specific and weather-specific domains are available in Fig. 4.

Let  $\mathcal{S} \rightarrow \mathcal{T}$  denote the case where  $\mathcal{S}$  is the source domain and  $\mathcal{T}$  is the target domain. The datasets are utilized



Figure 4. Sample images with different instance-level labels in JHU-Crowd++.

in experiments under two distinct types of settings: 1) An entire dataset is regarded as a single domain:  $A \rightarrow B / Q$ ,  $B \rightarrow A / Q$  and  $Q \rightarrow A / B$ ; 2) A dataset is partitioned into subsets according to image-level labels, and one subset constitutes a single domain:  $SD \leftrightarrow SR$  and  $SN \leftrightarrow FH$ .

## 4.2. Implementation

We adopt VGG16-BN [27] as the feature extractor in our model. The density regression head is a single convolution layer with  $1 \times 1$  filters, while the classification head consists of a  $3 \times 3$  and a  $1 \times 1$  convolution layers. Batch normalization is employed after all convolution layers except the last ones. For data augmentation, we apply three types of photometric transformations, color jittering, Gaussian blurring and sharpening. We also randomly crop image patches with a size of  $320 \times 320$  and adopt random horizontal flipping to both the original and augmented images simultaneously. We select AdamW [19] as the optimizer and OneCycleLR [30] as the learning rate scheduler with maximum learning rate set to  $1e-3$  and a maximum epoch of 300. The memory size  $M$  is set to 1024 and the dimension  $C$  is 256. We use a content error threshold  $\alpha$  of 0.5, and the loss weights  $\lambda_{cls}$  and  $\lambda_{con}$  are both set to 10.

## 4.3. Evaluation Metrics

We evaluate our method with mean absolute error (MAE) and mean squared error (MSE), defined as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^N |c_i - \hat{c}_i|, MSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (c_i - \hat{c}_i)^2}, \quad (14)$$

where  $N$  is the number of testing images,  $c_i$  is the ground truth count of the  $i$ -th image and  $\hat{c}_i$  is the predicted count. Lower values of both metrics indicate better performance.

## 4.4. Comparison with State of the Art

In this subsection, we compare our MPCount with state-of-the-art methods on different benchmarks, as illustrated

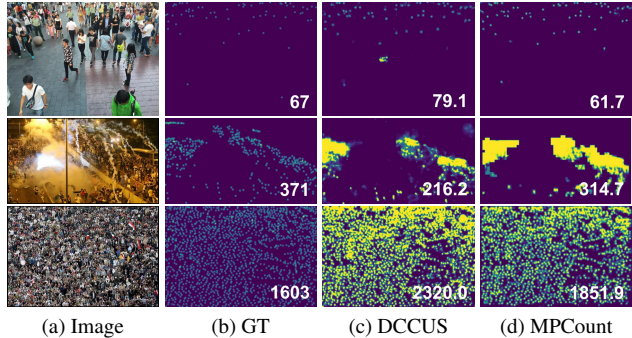


Figure 5. Visualization results of DCCUS [6] and MPCount under different DG settings. First row:  $A \rightarrow B$ ; second row:  $SN \rightarrow FH$ ; third row:  $B \rightarrow A$ .

in Fig. 5. Tab. 1 shows the results on  $A \rightarrow B / Q$ ,  $B \rightarrow A / Q$  and  $Q \rightarrow A / B$ . The selected baselines can be divided into three categories:

- *Fully supervised crowd counting*: BL [20], DM-Count [37], SASNet [31], ChfL [26] and MAN [17].
- *Domain adaptation for crowd counting*: RBT [18],  $C^2$ MoT [41], FGF [47], DAOT [48] and FSIM [49].
- *Single domain generalization*: IBN [23], SW [24], ISW [5], DG-MAN [21] and DCCUS [6].

All results are copied from previous papers except IBN, SW and ISW, which are feature normalization/whitening-based SDG methods originally designed for classification or segmentation and adapted to crowd counting by us. Domain adaptation methods are listed only for reference, as data from target domains are usually visible during adaptation.

Our MPCount outperforms all the DG methods under most of these settings, including a significant error reduction by 18.2% on  $B \rightarrow A$ . It is notable that MPCount exhibits outstanding performance in some cases even compared to DA methods, providing compelling evidence for the effectiveness of our design.

In Tab. 2, we further conduct experiments on scene-specific domains  $SD \leftrightarrow SR$  and weather-specific domains  $SN \leftrightarrow FH$ . BL [20], MAN [17], DAOT [48] and DCCUS [6] are trained using their released code, while IBN [23], SW [24] and ISW [5] are adapted by ourselves. We observe that the normalization/whitening-based methods cannot achieve satisfactory results, possibly because useful information is eliminated from feature statistics. DCCUS also obtains worse results than fully supervised methods, reflecting that the clustering procedure may fail to produce meaningful sub-domains when the source domain distribution is narrow. In contrast, our MPCount still performs the best among all tested methods, demonstrating its superiority in challenging conditions with narrow domain distributions.

Source $\rightarrow$ Target		A $\rightarrow$ B		A $\rightarrow$ Q		B $\rightarrow$ A		B $\rightarrow$ Q		Q $\rightarrow$ A		Q $\rightarrow$ B		
Method	DA	DG	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
BL [20]	$\times$	$\times$	15.9	25.8	166.7	287.6	138.1	228.1	226.4	411.0	-	-	-	-
DMCount [37]	$\times$	$\times$	23.1	34.9	134.4	252.1	143.9	239.6	203.0	386.1	73.4	135.1	14.3	27.5
SASNet [31]	$\times$	$\times$	21.3	33.2	211.2	418.6	132.4	225.6	273.5	481.3	73.9	116.4	13.0	22.1
ChfL [26]	$\times$	$\times$	18.7	29.1	122.3	218.0	121.3	200.8	197.1	357.9	68.7	118.5	14.0	27.0
MAN [17]	$\times$	$\times$	22.1	32.8	138.8	266.3	133.6	255.6	209.4	378.8	67.1	122.1	12.5	22.2
RBT [18]	$\checkmark$	$\times$	13.4	29.3	175.0	294.8	112.2	218.2	211.3	381.9	-	-	-	-
C <sup>2</sup> MoT [41]	$\checkmark$	$\times$	12.4	21.1	125.7	218.3	120.7	192.0	198.9	368.0	-	-	-	-
FGFD [47]	$\checkmark$	$\times$	12.7	23.3	124.1	242.0	123.5	210.7	209.7	384.7	70.2	118.4	12.5	20.3
DAOT [48]	$\checkmark$	$\times$	10.9	18.8	113.9	215.6	-	-	-	-	67.0	128.4	11.3	19.6
FSIM [49]	$\checkmark$	$\times$	11.1	19.3	105.3	191.1	120.3	202.6	194.9	324.5	66.8	111.5	11.0	19.7
IBN [23]	$\times$	$\checkmark$	19.1	30.8	280.2	561.0	125.9	202.3	183.5	317.4	105.9	174.6	16.0	25.5
SW [24]	$\times$	$\checkmark$	20.2	30.4	285.5	431.0	126.7	193.8	200.7	333.2	102.4	168.8	19.0	32.9
ISW [5]	$\times$	$\checkmark$	23.9	37.7	215.6	399.7	156.2	291.5	263.1	442.2	83.4	136.0	22.1	34.7
DG-MAN [21]	$\times$	$\checkmark$	17.3	28.7	129.1	238.2	130.7	225.1	182.4	325.8	-	-	-	-
DCCUS [6]	$\times$	$\checkmark$	12.6	24.6	119.4	216.6	121.8	203.1	179.1	316.2	67.4	112.8	<b>12.1</b>	<b>20.9</b>
<b>MPCount (Ours)</b>	$\times$	$\checkmark$	<b>11.4</b>	<b>19.7</b>	<b>115.7</b>	<b>199.8</b>	<b>99.6</b>	<b>182.9</b>	<b>165.6</b>	<b>290.4</b>	<b>65.5</b>	<b>110.1</b>	12.3	24.1

Table 1. Comparison with the state-of-the-art methods on SHA (A), SHB (B) and UCF-QNRF (Q). Domain adaptation methods are only for reference since data in target domain might be used.

Source $\rightarrow$ Target		SD $\rightarrow$ SR		SR $\rightarrow$ SD		SN $\rightarrow$ FH		FH $\rightarrow$ SN		
Method	DA	DG	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
BL [20]	$\times$	$\times$	42.1	79.0	262.7	1,063.9	48.1	129.5	343.8	770.5
MAN [17]	$\times$	$\times$	45.1	79.0	246.1	950.8	38.1	68.0	445.0	979.3
DAOT [48]	$\checkmark$	$\times$	45.3	88.0	278.7	1,624.3	42.3	73.0	151.6	273.9
IBN [23]	$\times$	$\checkmark$	92.2	178.0	318.1	1,420.4	109.7	267.7	491.8	1,110.4
SW [24]	$\times$	$\checkmark$	110.3	202.4	312.6	1,072.4	131.5	306.6	381.3	825.0
ISW [5]	$\times$	$\checkmark$	108.1	212.4	385.9	1,464.8	151.6	365.7	276.6	439.8
DCCUS [6]	$\times$	$\checkmark$	90.4	194.1	258.1	1,005.9	54.5	125.8	399.7	945.0
<b>MPCount (Ours)</b>	$\times$	$\checkmark$	<b>37.4</b>	<b>70.1</b>	<b>218.6</b>	<b>935.9</b>	<b>31.3</b>	<b>55.0</b>	<b>216.3</b>	<b>421.4</b>

Table 2. Comparison with the state of the art on data in JHU-Crowd++ with labels ‘‘Stadium’’(SD), ‘‘Street’’(SR), ‘‘Snow’’(SN) and ‘‘Fog/Haze’’(FH). Domain adaptation methods are only for reference since data in target domain might be used.

#### 4.5. Ablation Studies

In this section, we conduct ablation studies and other analysis under the setting A  $\rightarrow$  B / Q.

*Effects of model components.* We start from a simple encoder-decoder structured baseline model and add or remove each component individually to validate its effectiveness. The experimental results are displayed in Tab. 3.

- *Attention Memory Bank:* After adding AMB, the performance consistently improves on both target domains, as shown in Tab. 3. This validates that AMB can effectively help the model generalize by reconstructing domain-invariant features for density regression.
- *Content Error Mask & Attention Consistency Loss:* Next, we test the effects of CEM and ACL on AMB. The results

in Tab. 3 show that both components generally improves the performance, especially when PC is present.

- *Patch-wise Classification:* Finally, models with PC outperforms their counterparts without PC, as visible in Tab. 3. This suggests that such an auxiliary task is useful for DG by compensating for ambiguous pixel-level labels with accurate patch-level labels.

*Effects of  $\alpha$ .* The threshold of CEM  $\alpha$  mentioned in Eq. (6) controls the Portion of Diminished Elements (PDE) in features. According to the results in Tab. 4,  $\alpha = 0.5$  produces the best performance under all settings, with a PDE of 10.5%. It is also indicated that using a too large value of  $\alpha$  (0.9) may be inadequate to filter out domain-related information, while a too small value (0.1) would let useful information eliminated.

Components				A → B		A → Q	
AMB	CEM	ACL	PC	MAE	MSE	MAE	MSE
				21.6	39.1	200.9	325.5
✓				18.9	32.0	176.3	292.8
✓	✓			15.3	25.5	156.8	232.7
✓		✓		16.9	28.7	172.5	281.6
✓	✓	✓		14.1	24.3	170.6	286.8
			✓	14.4	30.5	144.2	262.6
✓			✓	13.4	25.2	133.4	224.0
✓	✓		✓	13.3	23.4	126.9	222.3
✓	✓	✓	✓	<b>11.4</b>	<b>19.7</b>	<b>115.7</b>	<b>199.8</b>

Table 3. Ablation studies on the model components.

Parameter	Property	A → B		A → Q	
$\alpha$	PDE(%)	MAE	MSE	MAE	MSE
0.1	28.6	16.7	25.9	157.3	338.5
0.3	14.6	12.4	20.8	132.1	223.6
0.5	10.5	<b>11.4</b>	<b>19.7</b>	<b>115.7</b>	<b>199.8</b>
0.7	5.9	11.9	23.8	123.8	214.1
0.9	3.6	13.8	27.2	124.7	209.6

Table 4. Ablation studies on the error threshold  $\alpha$  in Eq. (6).

Parameters		A → B		A → Q	
$M$	$C$	MAE	MSE	MAE	MSE
256	256	15.9	24.0	138.4	226.1
512	256	13.7	28.1	124.4	217.2
1024	256	<b>11.4</b>	<b>19.7</b>	<b>115.7</b>	<b>199.8</b>
2048	256	12.0	23.3	128.1	218.1
1024	128	16.3	28.6	127.4	223.3
1024	256	<b>11.4</b>	<b>19.7</b>	<b>115.7</b>	<b>199.8</b>
1024	512	14.3	27.7	130.1	250.1
1024	1024	14.2	24.8	128.4	220.2

Table 5. Ablation studies on the number of memory vectors  $M$  and vector size  $C$ .

*Effects of  $M$  and  $C$ .* The number of memory vectors  $M$  and dimension of each vector  $C$  are two important parameters of AMB. To demonstrate their effects independently, we adjust one parameter while keeping the other fixed. The results in Tab. 5 shows that among all variations of parameters,  $M = 1024, C = 256$  is the optimal choice.

*Visualizations of CEMs.* We present visualizations of CEMs in Fig. 6, each summed along the channel dimension. The results indicate that crowd information is generally more




Figure 6. Visualized CEMs. Yellower means less filtered.




Figure 7. Visualization of (a) sample images from SHB, (b) the ground-truth PCMs, (c) PCMs predicted by MPCount trained on SHA and (d) predicted PCMs after binarization.

sensitive to domain shift than environments.

*Visualizations of PCMs.* We visualize the predicted and binarized PCMs from MPCount for qualitative analysis. As shown in Fig. 7, MPCount predicts accurate PCMs that provide reliable information of crowd locations. In the raw predicted PCMs, there are unconfident predictions with low classification scores, which are filtered out in the binarized PCMs. The results verify that PC can effectively mitigate the label ambiguity problem and contribute to a more robust crowd counting model against domain shift.

## 5. Conclusion

This paper proposes MPCount to address single domain generalization for crowd counting with possibly narrow source distribution. MPCount tackles two unique challenges, namely, density regression and label ambiguity. We propose an attention memory bank to reconstruct domain-invariant features for regression, with the content error mask eliminating domain-related content information and the attention consistency loss ensuring the consistency of memory vectors. Patch-wise classification is a novel auxiliary task to mitigate the ambiguous pixel-level labels with accurate information, enhancing the robustness of density predictions. Extensive experiments on well-known datasets show that MPCount achieves significantly the best results on various benchmarks by outperforming the state of the art under different unseen scenarios and narrow source domain.

**Acknowledgement.** We would like to thank the Turing AI Computing Cloud (TACC) [42] and HKUST iSING Lab for providing us computation resources on their platform.



## References

- [1] Haoyue Bai, Hao He, Zhuoxuan Peng, Tianyuan Dai, and S.-H. Gary Chan. CounTr: An End-to-End Transformer Approach for Crowd Counting and Density Estimation. In *Proceedings of European Conference on Computer Vision (ECCV) International Workshop on Distributed Smart Cameras*, pages 207–222, Tel Aviv, Israel, 2022. Springer Lecture Notes in Computer Science. **2**
- [2] Yiqing Cai, Lianggangxu Chen, Haoyue Guan, Shaohui Lin, Changhong Lu, Changbo Wang, and Gaoqi He. Explicit invariant feature induced cross-domain crowd counting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37:259–267, 2023. **1, 2**
- [3] Yang Chen, Yu Wang, Yingwei Pan, Ting Yao, Xinmei Tian, and Tao Mei. A style and semantic memory mechanism for domain generalization\*. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9144–9153, 2021. **2**
- [4] Jian Cheng, Haipeng Xiong, Zhiguo Cao, and Hao Lu. Decoupled two-stage crowd counting and beyond. *IEEE Transactions on Image Processing*, 30:2862–2875, 2021. **2**
- [5] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne Taery Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11575–11585, 2021. **2, 3, 6, 7, 1**
- [6] Zhipeng Du, Jiankang Deng, and Miaoqing Shi. Domain-general crowd counting in unseen scenarios. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 561–570, 2023. **1, 2, 3, 4, 6, 7**
- [7] Junyu Gao, Tao Han, Yuan Yuan, and Qi Wang. Domain-adaptive crowd counting via high-quality image translation and density reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 34:4803–4815, 2019. **2**
- [8] Junyu Gao, Wei Lin, Bin Zhao, Dong Wang, Chenyu Gao, and Jun Wen. C<sup>3</sup> framework: An open-source pytorch code for crowd counting. *arXiv preprint arXiv:1907.02724*, 2019. **1**
- [9] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015. **1**
- [10] Tianlang He, Zhiqiu Xia, Jierun Chen, Haoliang Li, and S.-H. Gary Chan. Target-agnostic Source-free Domain Adaptation for Regression Tasks. In *Proceedings of 40th IEEE International Conference on Data Engineering (ICDE’24)*, Utrecht, Netherlands, 13-16 May 2024 (to appear). IEEE. **1**
- [11] Haroon Idrees, Muhammad Tayyab, Kishan Athrey, Dong Zhang, Somaya Ali Al-Maadeed, Nasir M. Rajpoot, and Mubarak Shah. Composition loss for counting, density map estimation and localization in dense crowds. In *European Conference on Computer Vision*, 2018. **5**
- [12] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Style normalization and restitution for domain generalization and adaptation. *IEEE Transactions on Multimedia*, 24:3636–3651, 2022. **2**
- [13] Hyeonseong Kim, Yoonsu Kang, Changgyoon Oh, and Kuk-Jin Yoon. Single domain generalization for lidar semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17587–17598, 2023. **2, 3**
- [14] Jin Kim, Jiyoung Lee, Jungin Park, Dongbo Min, and Kwanghoon Sohn. Pin the memory: Learning to generalize semantic segmentation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4340–4350, 2022. **3**
- [15] Lei Li, Ke Gao, Juan Cao, Ziyao Huang, Yepeng Weng, Xiaoyue Mi, Zhengze Yu, Xiaoya Li, and Boyang Xia. Progressive domain expansion network for single domain generalization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 224–233, 2021. **2**
- [16] 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*, pages 1091–1100, 2018. **2**
- [17] Hui Lin, Zhiheng Ma, Rongrong Ji, Yaowei Wang, and Xiaopeng Hong. Boosting crowd counting via multifaceted attention. In *CVPR*, 2022. **2, 6, 7**
- [18] Yuting Liu, Zheng Wang, Miaoqing Shi, Shin’ichi Satoh, Qijun Zhao, and Hongyu Yang. Towards unsupervised crowd counting via regression-detection bi-knowledge transfer. In *Proceedings of the 28th ACM International Conference on Multimedia*, page 129–137, New York, NY, USA, 2020. Association for Computing Machinery. **2, 6, 7**
- [19] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *ArXiv*, abs/1711.05101, 2017. **6**
- [20] Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. Bayesian loss for crowd count estimation with point supervision. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6141–6150, 2019. **2, 6, 7**
- [21] Lucas Mansilla, Rodrigo Echeveste, D.H. Milone, and Enzo Ferrante. Domain generalization via gradient surgery. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6610–6618, 2021. **6, 7**
- [22] Yanda Meng, Hongrun Zhang, Yitian Zhao, Xiaoyun Yang, Xuesheng Qian, Xiaowei Huang, and Yalin Zheng. Spatial uncertainty-aware semi-supervised crowd counting. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15529–15539, 2021. **2**
- [23] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *European Conference on Computer Vision*, 2018. **2, 6, 7, 1**
- [24] Xingang Pan, Xiaohang Zhan, Jianping Shi, Xiaoou Tang, and Ping Luo. Switchable whitening for deep representation learning. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. **2, 6, 7, 1**
- [25] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12556–12565, 2020. **2**

- [26] Weibo Shu, Jia Wan, Kay Chen Tan, Sam Kwong, and Antoni B Chan. Crowd counting in the frequency domain. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19618–19627, 2022. [2](#), [6](#), [7](#)
- [27] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [6](#)
- [28] Vishwanath A. Sindagi and Vishal M. Patel. Cnn-based cascaded multi-task learning of high-level prior and density estimation for crowd counting. *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6, 2017. [3](#)
- [29] Vishwanath A. Sindagi, Rajeev Yasarla, and Vishal M. Patel. Jhu-crowd++: Large-scale crowd counting dataset and a benchmark method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:2594–2609, 2020. [5](#)
- [30] Leslie N. Smith and Nicholay Topin. Super-convergence: very fast training of neural networks using large learning rates. In *Defense + Commercial Sensing*, 2017. [6](#)
- [31] Qingyu Song, Changan Wang, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Jian Wu, and Jiayi Ma. To choose or to fuse? scale selection for crowd counting. *The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021. [2](#), [6](#), [7](#)
- [32] Maryam Sultana, Muzammal Naseer, Muhammad Haris Khan, Salman Khan, and Fahad Shahbaz Khan. Self-distilled vision transformer for domain generalization. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 3068–3085, 2022. [2](#)
- [33] Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *ArXiv*, abs/1607.08022, 2016. [4](#)
- [34] Vidit Vidit, Martin Engilberge, and Mathieu Salzmann. Clip the gap: A single domain generalization approach for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3219–3229, 2023. [2](#)
- [35] Chaoqun Wan, Xu Shen, Yonggang Zhang, Zhiheng Yin, Xinmei Tian, Feng Gao, Jianqiang Huang, and Xiansheng Hua. Meta convolutional neural networks for single domain generalization. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4672–4681, 2022. [2](#)
- [36] Jia Wan, Ziquan Liu, and Antoni B. Chan. A generalized loss function for crowd counting and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1974–1983, 2021. [2](#)
- [37] Boyu Wang, Huidong Liu, Dimitris Samaras, and Minh Hoai. Distribution matching for crowd counting. In *Advances in Neural Information Processing Systems*, 2020. [2](#), [6](#), [7](#)
- [38] Mingjie Wang, Hao Cai, Xianfeng Han, Jun Zhou, and Minglun Gong. Stnet: Scale tree network with multi-level auxiliary for crowd counting. *IEEE Transactions on Multimedia*, 25:2074–2084, 2020. [2](#)
- [39] Qi Wang, Junyu Gao, Wei Lin, and Yuan Yuan. Learning from synthetic data for crowd counting in the wild. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8198–8207, 2019. [1](#), [2](#)
- [40] Qi Wang, Tao Han, Junyu Gao, and Yuan Yuan. Neuron linear transformation: Modeling the domain shift for crowd counting. *IEEE Transactions on Neural Networks and Learning Systems*, 33(8):3238–3250, 2022. [2](#)
- [41] Qiangqiang Wu, Jia Wan, and Antoni B. Chan. Dynamic momentum adaptation for zero-shot cross-domain crowd counting. In *Proceedings of the 29th ACM International Conference on Multimedia*, page 658–666, New York, NY, USA, 2021. Association for Computing Machinery. [2](#), [6](#), [7](#)
- [42] Kaiqiang Xu, Xinchun Wan, Hao Wang, Zhenghang Ren, Xudong Liao, Decang Sun, Chaoliang Zeng, and Kai Chen. Tacc: A full-stack cloud computing infrastructure for machine learning tasks. *arXiv preprint arXiv:2110.01556*, 2021. [8](#)
- [43] Qi Xu, Lili Yao, Zhengkai Jiang, Guannan Jiang, Wenqing Chu, Wenhui Han, Wei Zhang, Chengjie Wang, and Ying Tai. Dirl: Domain-invariant representation learning for generalizable semantic segmentation. In *AAAI Conference on Artificial Intelligence*, 2022. [2](#), [3](#)
- [44] Cong Zhang, Hongsheng Li, Xiaogang Wang, and Xiaokang Yang. Cross-scene crowd counting via deep convolutional neural networks. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 833–841, 2015. [1](#), [2](#)
- [45] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 589–597, 2016. [2](#), [5](#)
- [46] Yuxiang Zhang, Wei Li, Weidong Sun, Ran Tao, and Qian Du. Single-source domain expansion network for cross-scene hyperspectral image classification. *IEEE Transactions on Image Processing*, 32:1498–1512, 2022. [2](#)
- [47] Huilin Zhu, Jingling Yuan, Zhengwei Yang, Xian Zhong, and Zheng Wang. Fine-grained fragment diffusion for cross domain crowd counting. In *Proceedings of the 30th ACM International Conference on Multimedia*, page 5659–5668, New York, NY, USA, 2022. Association for Computing Machinery. [2](#), [6](#), [7](#)
- [48] Huilin Zhu, Jingling Yuan, Xian Zhong, Zhengwei Yang, Zheng Wang, and Shengfeng He. Daot: Domain-agnostically aligned optimal transport for domain-adaptive crowd counting. In *Proceedings of the 31st ACM International Conference on Multimedia*, page 4319–4329, New York, NY, USA, 2023. Association for Computing Machinery. [6](#), [7](#)
- [49] Huilin Zhu, Jingling Yuan, Xian Zhong, Liang Liao, and Zheng Wang. Find gold in sand: Fine-grained similarity mining for domain-adaptive crowd counting. *IEEE Transactions on Multimedia*, 26:3842–3855, 2024. [6](#), [7](#)
- [50] Su Zixian, Yao Kai, Yang Xi, Wang Qiufeng, Sun Jie, and Huang Kaizhu. Rethinking data augmentation for single-source domain generalization in medical image segmentation. In *AAAI*, 2023. [2](#)