

Counting Crowds in Bad Weather

Zhi-Kai Huang^{1*} Wei-Ting Chen^{1,2*} Yuan-Chun Chiang¹ Sy-Yen Kuo¹ Ming-Hsuan Yang^{3,4,5}

¹National Taiwan University ²Stanford University ³UC Merced ⁴Google Research ⁵Yonsei University

Abstract

Crowd counting has recently attracted significant attention in the field of computer vision due to its wide applications to image understanding. Numerous methods have been proposed and achieved state-of-the-art performance for real-world tasks. However, existing approaches do not perform well under adverse weather such as haze, rain, and snow since the visual appearances of crowds in such scenes are drastically different from those images in clear weather of typical datasets. In this paper, we propose a method for robust crowd counting in adverse weather scenarios. Instead of using a two-stage approach that involves image restoration and crowd counting modules, our model learns effective features and adaptive queries to account for large appearance variations. With these weather queries, the proposed model can learn the weather information according to the degradation of the input image and optimize with the crowd counting module simultaneously. Experimental results show that the proposed algorithm is effective in counting crowds under different weather types on benchmark datasets. The source code is available in our [project page](#).

1. Introduction

Crowd counting aims to estimate the number of persons in a scene. With the advances of deep learning and construction of large-scale datasets [69, 19, 45, 54], this topic has become an important topic with numerous applications. State-of-the-art approaches are mainly based on convolutional neural networks (CNNs) [22, 24, 31, 57, 50, 13, 1, 27, 29] and transformers [48, 47, 37, 26]. While these approaches count crowds effectively on normal images, they typically do not perform well under adverse weather conditions such as haze, rain, and snow. However, adverse weather is a common and inevitable scenario that causes large appearance variations of crowd scenes, thereby significantly affecting the performance of methods developed for clear weather. As shown in Figure 1, the state-of-the-art crowd counting method [29]

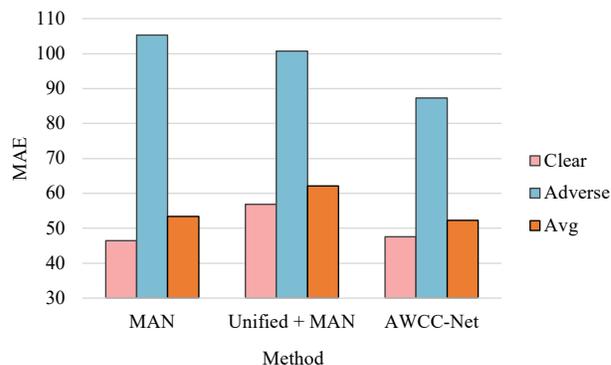


Figure 1. Performance of State-of-the-art crowd counting methods under adverse and clear weather on the JHU-Crowd++ [45] dataset using mean absolute error (MAE). The MAN method [29] achieves low MAE in clear scenes but high MAE in adverse weather. On the other hand, the two-stage method, based on Unified [8] and MAN [29]), performs slightly better in adverse weather but slightly worse in clear scenes. Overall, the proposed AWCC-Net performs favorably in both scenarios.

performs well only under good weather conditions. Thus, it is of great importance to develop robust crowd counting methods for adverse weather conditions.

To solve the crowd counting problem under adverse weather, a plausible solution is to consider a two-stage model. Specifically, this strategy pre-processes images using the state-of-the-art image restoration modules [8, 49] and then applies the state-of-the-art crowd counting method [29]. However, this two-stage method may not perform well due to several factors. First, adopting existing image restoration methods does not always facilitate the crowd counting task significantly since these methods are designed to restore image contents rather than visual classification or regression. As shown in Figure 1, a two-stage approach does not address this problem effectively. Second, this strategy requires to collect and label images under adverse weather conditions for the restoration process. In addition, a two-stage strategy may increase computational complexity significantly.

In this work, we propose a method based on a trans-

* indicates equal contribution.

former to robustly count crowds under adverse and clear scenes without exploiting typical image restoration modules. We develop the **adverse weather crowd counting network** (AWCC-Net) which leverages learned weather queries for robust crowd counting. In our model, the weather queries are adopted to match the keys and values extracted from the VGG encoder [43] based on the cross-attention design. However, the learned weather queries cannot well represent the weather information without appropriate constraints. To improve this situation, we present a contrastive weather-adaptive module to improve the learned weather queries. This adaptive module computes a weight vector which is combined with the learned weather bank to construct input-dependent weather queries. We use the proposed contrastive loss to enforce the learned weather queries to be weather-related. With this weather-adaptive module, more robust weather queries can be learned which help the network understand and adapt to the various weather degradation types in the input image.

Extensive experimental results show that the proposed AWCC-Net model performs robustly and favorably against the state-of-the-art schemes for crowd counting under adverse weather conditions without weather annotations. We make the following contributions in this work:

- We present the AWCC-Net for crowd counting under adverse weather. Our method is based on a transformer architecture that can learn weather-aware crowd counting.
- We propose a module to constrain the learned queries to be weather-relevant, thereby improving model robustness under adverse weather.

2. Related Work

2.1. Crowd Counting

Existing crowd counting methods are mainly based on: (i) detection [30, 40, 32, 57], (ii) density map [13, 19, 1, 14, 27, 60, 22, 44], and (iii) regression [22, 44, 18, 5, 4, 46].

Counting by detection. These approaches use detectors, e.g., Faster RCNN [39], for crowd counting. In [32], a method based on curriculum learning detects and learns to predict and count bounding boxes of persons via a locally-constrained regression loss. Sam *et al.* [40] leverages the nearest neighbor distance to generate pseudo bounding boxes and adopt the winner-take-all loss to optimize the box selection during the training stage. This technique can benefit the optimization of images with higher resolutions.

Counting by density map. In this mainstream approach, the crowd count is derived via summing over the estimated density map of a scene. In [19], the Gaussian Kernel is applied to construct the density map and propose the composition loss to optimize the crowd counting. The DACC method [13] leverages inter-domain features segregation to

generate coarse counter and then applies Gaussian prior to compute the counting results. Abousamra *et al.* [1] develops the topological constraint, which is achieved by a persistence loss to improve the spatial arrangement of dots for density map-based methods. Recently, the focal inverse distance transform map [27] shows better performance in representing person locations compared to the aforementioned Gaussian-based density map.

Counting by regression. The number of the crowd can be regressed from the contextual information among extracted features of cropped image patches. In [4], a closed-form approximation based on the Bayesian Poisson regression is proposed to compute the crowd counting. The Fourier analysis and SIFT features can be used to estimate the number of persons in a crowded scene [18]. On the other hand, the surrogate regression [46] achieves crowd counting and localization based on a set of point proposals.

Although the aforementioned methods can solve the crowd counting problem on clear images effectively, they do not perform well when the input images are taken in adverse weather. Thus, it is of great importance to develop a solution to cope with this problem.

2.2. Image Restoration

In recent years, numerous restoration algorithms have been proposed to handle images acquired in adverse weather, which can be categorized as: (i) single-weather removal; (ii) multi-degradation removal; and (iii) all-in-one weather removal.

Single-purpose image restoration. These methods are developed to restore image contents degraded by one specific degradation such as rain [21, 12, 63, 62, 59, 56, 20, 11, 52, 53, 64, 61], snow [70, 36, 58, 65, 25, 68, 7], and haze [15, 2, 67, 17, 10, 41, 38, 6]. While these methods are effective for specific conditions, they do not generate clear images when inputs are degraded by other or multiple weather types.

Multi-purpose image restoration. These methods aim to recover various weather types using a unified model. A general architecture is developed by Pan *et al.* [35] to estimate structures and details simultaneously in parallel branches. In [66], the MPRNet exploits a multi-stage strategy and an attention module to refine the incoming features at each stage for effective image restoration. Although it is based on a unified framework, different model weights need to be learned for each weather condition.

All-purpose image restoration. In recent years, much effort has been made to all-purpose image restoration since it only requires a set of pre-trained weights to recover images degraded by different factors. An end-to-end training scheme based on the neural architecture search is proposed by Li *et al.* [23] to investigate crucial features from multiple

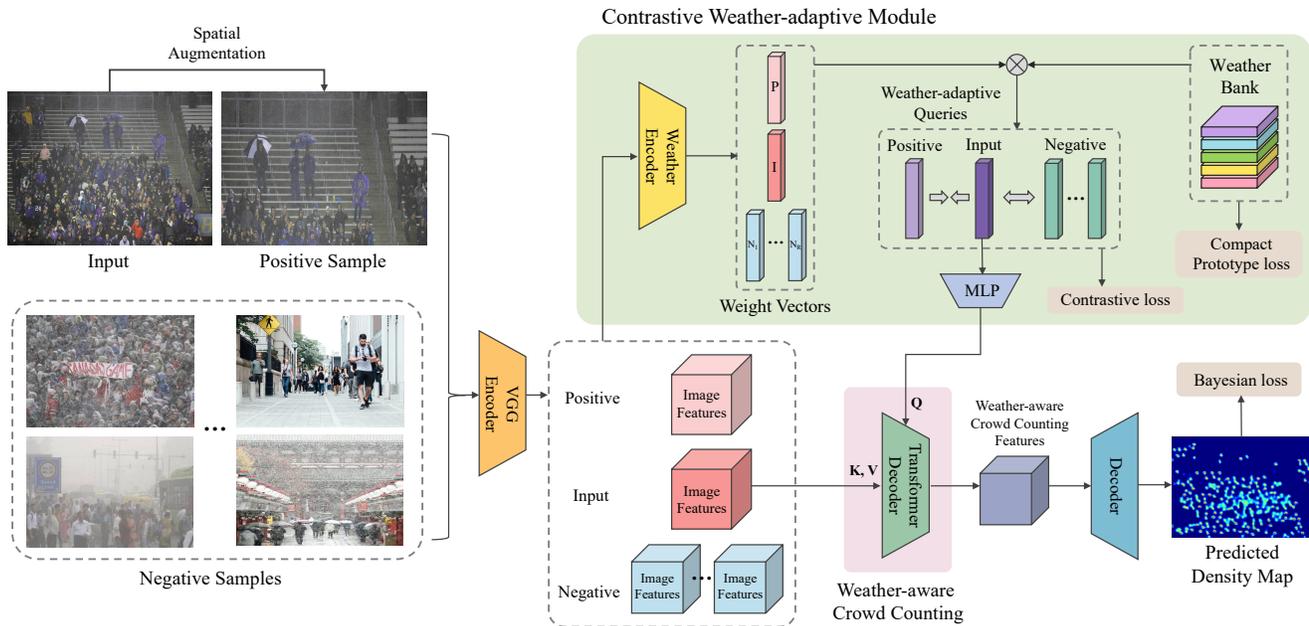


Figure 2. **AWCC-Net for adverse weather crowd counting.** A VGG encoder is adopted to extract image features, which are fed into the contrastive weather-adaptive module to generate patch-wise weather queries. Then, the transformer decoder leverages crowd counting features and weather-adaptive queries to generate weather-aware crowd counting features. The decoder computes the density map based on the weather-aware crowd counting features.

encoders for different weather types and then reconstruct the clear results. In [8], a two-stage learning strategy and the multi-contrastive regularization based on knowledge distillation [16] are developed to recover multi-weather types. Most recently, the TransWeather model [49] leverages the intra-patch visual information to extract fine detail features for the all-purpose weather removal.

Although these methods can be adopted in pre-process stage for crowd counting under adverse weather, they are designed to recover image content (visual appearance) rather than estimate head counts (visual classification or regression). Thus, adopting them with crowd counting may have limited performance gain under adverse weather.

3. Proposed Method

In this work, we tackle the crowd counting problem under adverse weather. As shown in Figure 2, image features are first extracted using a pre-trained backbone model, e.g., VGG-19 [43]. The contrastive weather-adaptive module uses the image features as input to generate the weather-adaptive queries, which are leveraged in the weather-aware crowd counting (WACC) model to generate the weather-aware crowd counting features. These features are then fed into the decoder to compute the density map for crowd estimation. We discuss the details of these modules in the following sections.

3.1. Weather-Aware Crowd Counting Model

In this work, we learn weather conditions as queries in our crowd counting model based on a transformer. Image features are first extracted via a pre-trained backbone model. The decoder is similar to that of a vision transformer [3], as shown in the pink region of Figure 2. The key (K) and value (V) are computed from the image features and the query (Q) is trainable weather queries learned with the network simultaneously. Through this operation, the output features of transformer decoder contain crowd counting features with weather information closest to the input scene. These features are then fed into a decoder which contains three 3×3 convolution blocks with ReLU to predict the density map D .

However, as shown in Table 6, this strategy may obtain limited performance gain since the weather queries are not adaptive to the weather types of the input image as no constraints are enforced. That is, the weather information is not learned by the weather queries, and thus they cannot well represent the weather type. Without effective weather queries, a transformer cannot generate effective features for crowd counting.

3.2. Contrastive Weather-adaptive Module

To enforce the learned queries are adaptive to weather conditions, we propose to construct input-dependent weather queries by computing a weight vector. The input-dependent

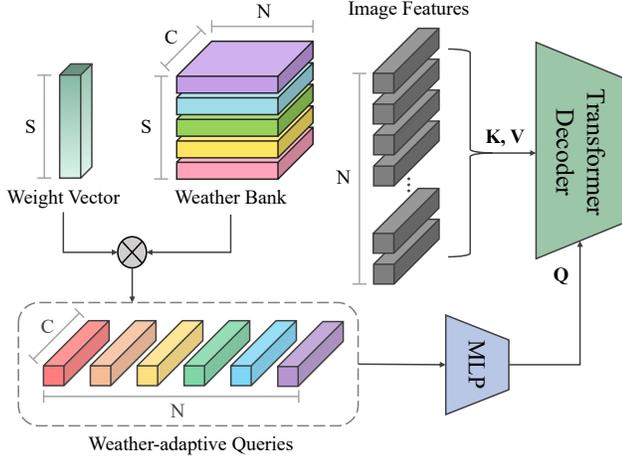


Figure 3. **Learning weather-adaptive queries.** A weight vector is combined with the prototypes in the weather bank via weighted sum to generate N patch-aware weather queries for the transformer decoder.

weather queries are constrained by the contrastive loss which can guarantee the learned queries to be weather-related.

Pipeline. The learning process for weather-adaptive queries is shown in the green regions of Figure 2 and Figure 3. First, the weather encoder predicts a weight vector for the weather bank based on the image features. This weight vector can construct the weather queries by the weighted summation of each prototype in the weather bank. To constrain these weather queries, we apply the contrastive loss to enforce the weather queries to learn weather information. A multilayer perceptron (MLP) is then adopted to conduct non-linear operations for the generated weather queries. Finally, these weather queries and image features are adopted in the transformer decoder to decode the weather-aware crowd counting features.

Weather Bank. The weather bank is a collection of multiple weather prototypes which can be leveraged to construct weather queries. The dimension of the weather bank is $S \times N \times C$ where S , N , and C denote the number of weather prototypes, the number of tokens to represent the weather prototype, and the number of channels, respectively. We set S , N , and C as 8, 48, and 512 in our method. These weather prototypes serve as reference templates for constructing weather queries, which capture weather-related information in input images. They are trainable parameters that are learned with the network simultaneously. Since they can be learned automatically, our network does not need to adopt the annotations of the weather type to optimize them.

Weight Vector. To construct input-dependent weather queries, we propose to compute a weight vector to repre-

sent the weather type of an input image. The weight vector is computed by the weather encoder based on the image features whose dimension is S . With the learned weather bank, the vector can represent the input weather type and construct input-dependent weather queries with dimension $N \times C$ for the contrastive loss.

Contrastive Loss. To learn the weather-adaptive queries, we propose the contrastive loss \mathcal{L}_{Con} :

$$\mathcal{L}_{Con} = -\log \left[\frac{\exp(\varphi(v, v^+)/\tau)}{\exp(\varphi(v, v^+)/\tau) + \sum_{r=1}^R \exp(\varphi(v, v_r^-)/\tau)} \right], \quad (1)$$

where v , v^+ , and v^- are the weather queries of the input image, those of the positive sample, and those of the negative samples, respectively. In addition, $\varphi(\cdot, \cdot)$ is the cosine similarity function, τ is the scale temperature, and R denotes the total number of negative samples.

The main idea of our method is that the weather queries of the input image should be similar to those of the positive sample while dissimilar to those of the negative samples. To this end, we adopt the random crop and the random flip operations to construct the positive sample. Since the positive sample is generated with spatial augmentation, the weather information should be the same as the input image, and thus the two weather queries should be similar. We use the rest of the images as our negative samples whose weather queries should be dissimilar. Although the negative samples may possibly have a similar weather type as the input image, according to [51], the contrastive loss contains the tolerance to potential positive samples. Thus, the images with the same weather types still have the smaller distance in the feature space. We show one example in Figure 6.

Compact Prototype Loss. To reduce the redundancy of learned prototypes in the weather bank, we propose the compact prototype loss \mathcal{L}_{CP} :

$$\mathcal{L}_{CP} = \sum_{i=1}^S \sum_{j \neq i} |\varphi(P^i, P^j)|, \quad (2)$$

where P^i and P^j present the i^{th} and j^{th} prototype in the weather bank. With this operation, the diversity and distinctiveness of prototypes in the weather bank are improved. The similarity between different prototypes is reduced, promoting better representation of various weather conditions for enhanced adaptability in crowd-counting tasks.

3.3. Overall Loss

The proposed AWCC-Net is optimized via multiple loss functions including the compact prototype loss, the contrastive loss, and the crowd counting loss. The first two loss functions have been defined in the previous sections. The crowd counting loss \mathcal{L}_{CC} aims to constrain the learning of crowd counting and we adopt the Bayesian loss [33, 29] for

the robustness and the better performance:

$$\mathcal{L}_{CC} = \left| \sum_k P_0(k) \cdot D_k \right| + \sum_{i=1}^L \left| 1 - \sum_k P_i(k) \cdot D_k \right|, \quad (3)$$

where i and L denote the index of the annotated point and the number of annotated points, respectively. In addition, D is the computed density map, $P_i(k)$ indicates the posterior of the occurrence of the i^{th} annotation given the position k , and $P_0(k)$ denotes the background likelihood at position k .

The overall loss function of the AWCC-Net is:

$$\mathcal{L}_{AWCC} = \mathcal{L}_{CC} + \lambda_1 \mathcal{L}_{CP} + \lambda_2 \mathcal{L}_{Con}, \quad (4)$$

where λ_1 and λ_2 are scaling factors.

4. Implementation Details

4.1. Datasets

We use the ShanghaiTech [69], UCF-QNRF [19], JHU-Crowd++ [45], and NWPU-CROWD [54] datasets to evaluate the performance of the proposed method against the state-of-the-art approaches.

JHU-Crowd++ [45]. It consists of 4,372 images and 1.51 million annotated points totally. This dataset divides 2,272 images for training, 500 for validation, and the remaining 1,600 images for testing. In the testing set, there are 191 images are under adverse weather and 1409 images under normal scenes.

ShanghaiTech A [69]. It contains 482 images and 244, 167 annotated points. 300 images are split for training and the rest of 182 images are for testing.

UCF-QNRF [19]. It includes 1,535 high-resolution images collected from the Web, with 1.25 million annotated points. There are 1,201 images in the training set and 334 images in the testing set. The UCF-QNRF dataset has a wide range of people count between 49 and 12,865.

NWPU-CROWD [54]. It consists of 5,109 images with 2.13 million annotated points. 3,109 images are divided into the training set and 500 images are in the validation set, and the remaining 1,500 images are for testing.

4.2. Evaluation Protocol

Similar to the existing approaches, we apply Mean Absolute Error (MAE) and Mean Squared Error (MSE) for performance evaluation:

$$MAE = \frac{1}{Q} \sum_{i=1}^Q |GT_i - N_i|, MSE = \sqrt{\frac{1}{Q} \sum_{i=1}^Q (GT_i - N_i)^2}, \quad (5)$$

where Q is the number of images, GT_i and N_i denote the ground truth and predicted crowd count of the i -th image.

Method	Clear		Adverse Weather		Average	
	MAE	MSE	MAE	MSE	MAE	MSE
SFCN [55]	71.4	225.3	122.8	606.3	77.5	297.6
BL [33]	66.2	200.6	140.1	675.7	75.0	299.9
BL [33]-U	65.3	208.4	134.0	645.6	73.5	296.6
BL [33]-UF	62.6	205.7	130.4	638.4	70.7	293.1
LSCCNN [40]	103.8	399.2	178.0	744.3	112.7	454.4
CG-DRCN-V [45]	74.7	253.4	138.6	654.0	82.3	328.0
CG-DRCN-R [45]	64.4	205.9	120.0	580.8	71.0	278.6
UOT [34]	53.1	148.2	114.9	610.7	60.5	252.7
GL [50]	54.2	159.8	115.9	602.1	61.6	256.5
GL [50]-U	64.4	207.0	127.2	617.3	71.9	288.5
GL [50]-UF	56.3	174.1	127.6	658.5	64.8	280.1
CLTR [26]	52.7	148.1	109.5	568.5	59.5	240.6
MAN [29]	46.5	137.9	105.3	478.4	<u>53.4</u>	<u>209.9</u>
MAN [29]-U	56.9	182.5	<u>100.7</u>	548.2	62.1	255.4
MAN [29]-UF	60.8	187.7	117.1	623.2	67.6	278.2
AWCC-Net	<u>47.6</u>	153.9	87.3	430.1	52.3	207.2

Table 1. **Quantitative comparison on the JHU-Crowd++ [45] dataset with existing methods.** We evaluate the performance in adverse weather scenes and clear scenes. The words with **boldface** indicate the best results, and those with underline indicate the second-best results.

4.3. Training Details

In this work, the learning rate is 10^{-5} , and the Adam optimizer is applied. We set the batch size to be 1, and randomly crop all input images to 512×512 in the training process. The proposed network is trained on an Nvidia Tesla V100 GPU and implemented using the PyTorch framework. Our model is based on a hybrid CNN-Transformer backbone which contains a VGG-19 model pre-trained on ImageNet, a transformer decoder, and a CNN-based regression layer. The scaling factors of the losses λ_1 and λ_2 are both set to 1. In addition, the total number of negative samples R is 64, and the scale temperature τ is 0.2.

In the spatial augmentation, we crop two patches of the training image. We apply random horizontal flipping on the first patch. Then, we crop a second patch overlapped with the first one with an overlap factor sampled from a uniform distribution. The first patch serves as the training input for crowd counting and the anchor for the contrastive loss. The second patch is the positive sample for the contrastive loss and is not directly involved in crowd counting.

5. Experimental Results

In this section, we present the evaluation results of the proposed method for crowd counting in both adverse weather scenes and normal scenes. More results are available in the supplementary material.

5.1. Performance Evaluation

We evaluate our method against the state-of-the-art approaches including SFCN [55], BL [33], LSCCNN [40], CG-DRCN [45], UOT [34], S3 [28], GL [50], ChfL [42],

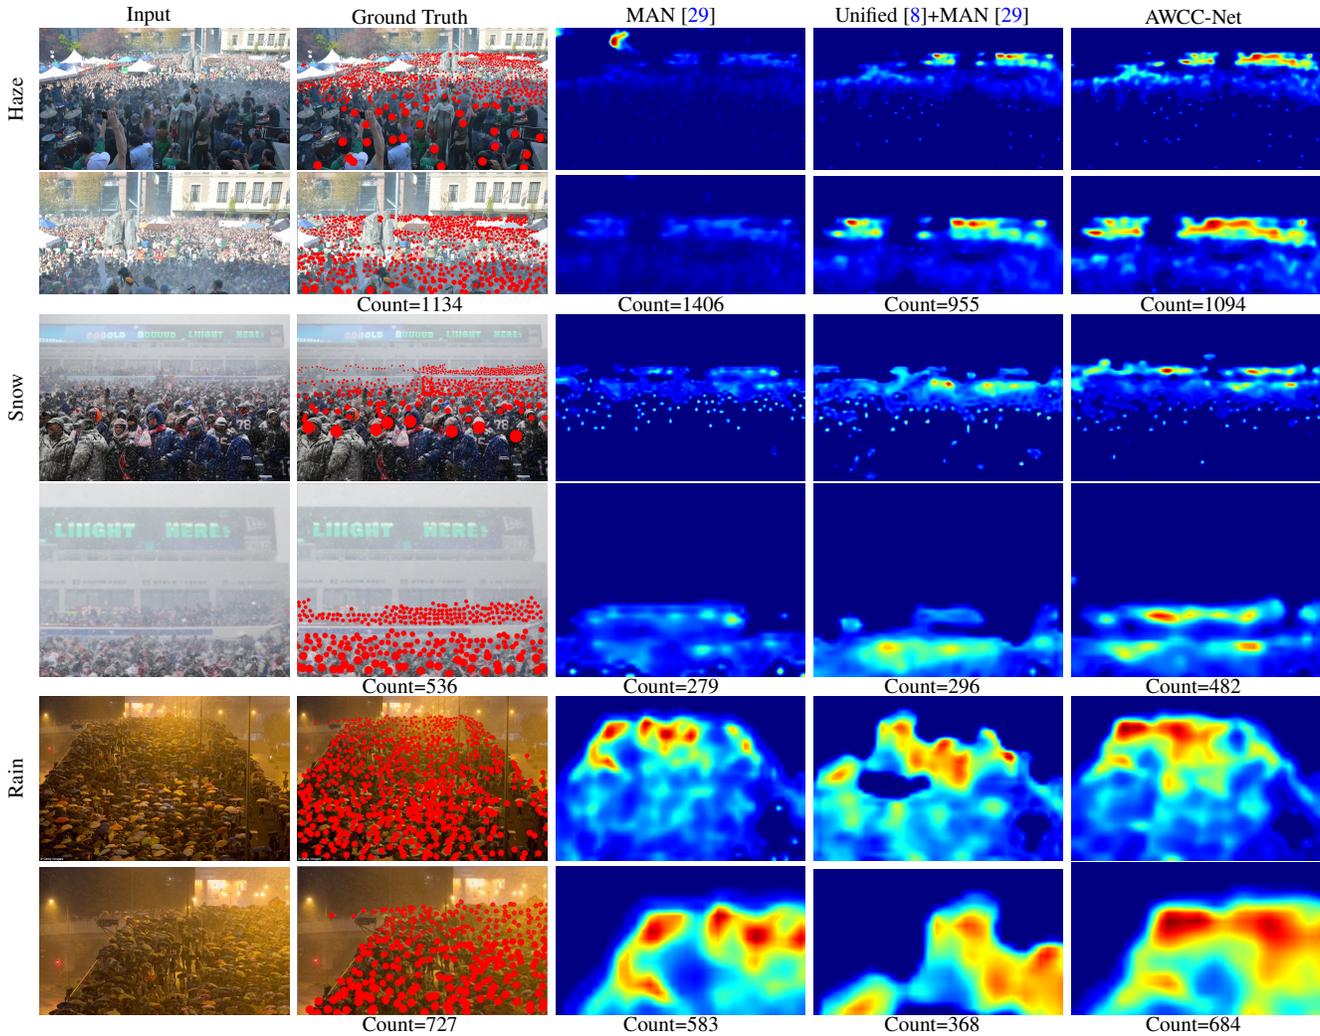


Figure 4. Comparison of density maps of the proposed method and other baselines in adverse weather (i.e., haze, snow, rain). The proposed method can predict more accurate density maps compared to the results estimated by other strategies.



Figure 5. Probe the weather type of an image of the ShanghaiTech A dataset from the JHU-Crowd++ dataset based on the weather-adaptive queries. The learned weather prototypes and the weight vector can well represent the weather type of the input image from the unseen dataset.

CLTR [26], MAN [29], and GauNet [9]. We present the evaluation results of our AWCC-Net on the JHU-Corwd++ dataset for adverse weather crowd counting since only this

dataset contains the annotations of bad weather. We also report the performance on the ShanghaiTech A, UCF-QNRF, and NWPU-CROWD datasets.

Evaluation on the JHU-Crowd++ dataset. We use two settings in the experiments:

1. **Vanilla crowd counting:** We apply vanilla crowd counting methods trained with the training set of the JHU-Crowd++ dataset.
2. **Two-stage strategy:** We use the Unified model [8] to restore image content under adverse weather and then apply vanilla crowd counting methods to estimate the counts. We denote this strategy with '-U'. Note the crowd counting models are not fine-tuned by restored images. For fair comparisons, we adopt these restored images to finetune the crowd counting models. We denote this strategy with '-UF'.

Table 1 shows that, although existing methods achieve

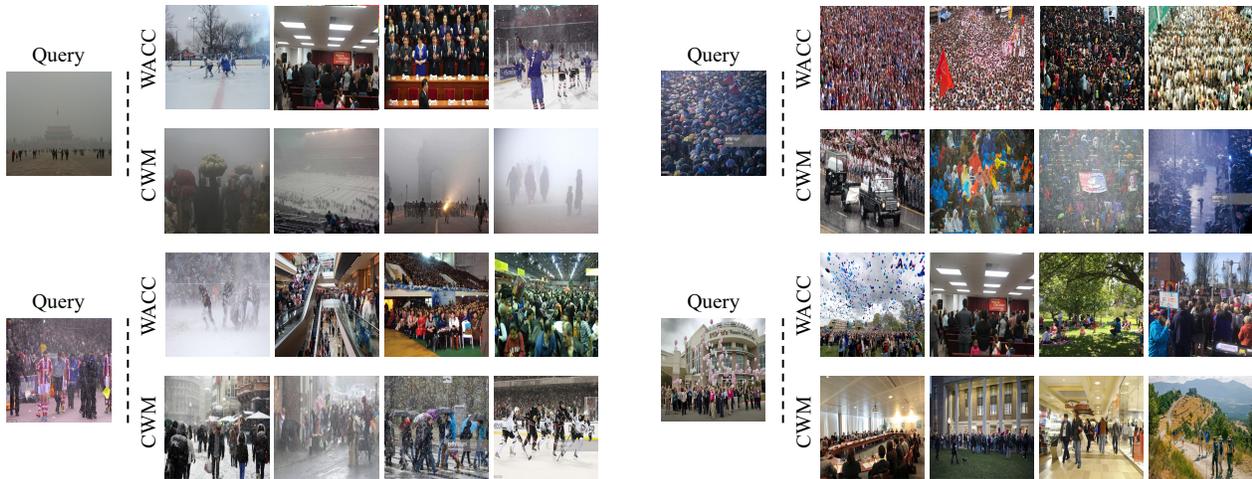


Figure 6. **Visual comparison of the proposed contrastive weather-adaptive queries.** The proposed contrastive weather-adaptive queries can achieve a better representation of the weather type. We term the proposed contrastive weather-adaptive module as 'CWM'.

state-of-the-art results in clear scenes, they do not perform well in adverse weather scenes. Second, the two-stage strategy may not help the improvement of the performance compared to the vanilla crowd counting approaches since the goal of the image restoration process is not the same as counting. This is because the image restoration methods are designed for restoring content rather than crowd counting, and thus the restored results may not suitable for crowd computation. Third, our method performs favorably against state-of-the-art approaches in adverse weather. For the clear scenes, our method can achieve comparable performance in MAE. In terms of average performance, our method can achieve the first place in both MAE and MSE. Overall, the proposed method is effective and robust when the input images are degraded by bad weather while it can retain the performance of clear scenes. In addition, we present the quantitative results across different weather types in Table 2 and demonstrate the density maps under different adverse weather predicted by the proposed method and other algorithms in Figure 4. The results show that the proposed method can compute the more accurate density distribution and counts of crowds under bad weather.

Evaluation on other datasets. We present the evaluation results of our algorithm on other datasets in Table 3. As these images do not contain annotations of weather types in the images, we can only compare them in the averaged performance. The results indicate that the proposed method performs favorably in the JHU-Crowd++ and UCF-QNRF datasets while comparable performance in the ShanghaiTech A and NWPU-CROWD datasets.

Generalization to other datasets. We analyze the generalization ability of our method on other datasets. Specifically, we train the AWCC-Net and other baselines on the JHU-Crowd++ dataset and directly test on the ShanghaiTech A,

Weather Method	Clear		Fog		Rain		Snow	
	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
MAN	46.5	137.9	85.2	183.4	112.4	375.3	117.4	666.6
MAN-U	56.9	182.5	86.3	231.5	<u>71.1</u>	<u>200.1</u>	131.0	816.5
AWCC-Net	<u>47.6</u>	<u>153.9</u>	77.2	162.5	67.9	179.2	107.7	641.2

Table 2. **Quantitative results across different weather types.**

and UCF-QNRF datasets. Table 4 shows that the proposed method achieves the best performance in terms of MAE and MSE. Moreover, we use an image in the ShanghaiTech A dataset and search for the most similar two images in the JHU-Crowd++ dataset. The two images whose weather queries from the training dataset are most similar to the weather queries of the input from the unseen dataset are presented in Figure 5. Our model can locate the images with the similar weather type in the training set based on the combination of the weight vector and the weather bank. Thus, the proposed weather-aware crowd counting mechanism and the weather-adaptive queries can benefit our network to be robust to the weather types in the unseen datasets.

Comparison with Other Baseline. We compare another simple solution with our method in Table 5 which is to assign higher weights of loss functions to adverse weather images during the training phase on JHU-Crowd++ dataset. This modified approach is based on MAN, which we refer to as "MAN-W", allocates weights of 10 and 1 to the adverse weather data and clear data, respectively. This weight allocation is based on the data distribution in the dataset. Comparative results demonstrate that our method outperforms others, especially in adverse weather and average scenarios.

5.2. Ablation Study

We evaluate the effectiveness of weather-aware crowd counting and contrastive weather-adaptive queries using the

Dataset Method	ShanghaiTechA		UCF-QNRF		JHU-Crowd++		NWPU-CROWD	
	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
SFCN [55]	64.8	107.5	102.0	171.4	77.5	297.6	105.7	424.1
BL [33]	62.8	101.8	88.7	154.8	75.0	299.9	105.4	454.2
LSCCNN [40]	66.5	101.8	120.5	218.2	112.7	454.4	-	-
CG-DRCN-VGG16 [45]	64.0	98.4	112.2	176.3	82.3	328.0	-	-
CG-DRCN-Res101 [45]	60.2	94.0	95.5	164.3	71.0	278.6	-	-
UOT [34]	58.1	95.9	83.3	142.3	60.5	252.7	87.8	387.5
S3 [28]	57.0	96.0	80.6	139.8	59.4	244.0	83.5	346.9
GL [50]	61.3	95.4	84.3	147.5	59.9	259.5	79.3	346.1
ChfL [42]	57.5	94.3	80.3	137.6	57.0	235.7	76.8	343.0
CLTR [26]	56.9	95.2	85.8	141.3	59.5	240.6	74.3	333.8
MAN [29]	56.8	<u>90.3</u>	<u>77.3</u>	<u>131.5</u>	<u>53.4</u>	<u>209.9</u>	76.5	323.0
GauNet [9]	54.8	89.1	81.6	153.7	58.2	245.1	-	-
AWCC-Net	<u>56.2</u>	91.3	76.4	130.5	52.3	207.2	<u>74.4</u>	<u>329.1</u>

Table 3. Quantitative comparison on the ShanghaiTech A [69], UCF-QNRF [19], JHU-Crowd++ [45], and NWPU-CROWD [54] datasets with existing methods. The words with **boldface** indicate the best results, and those with underline indicate the second-best results.

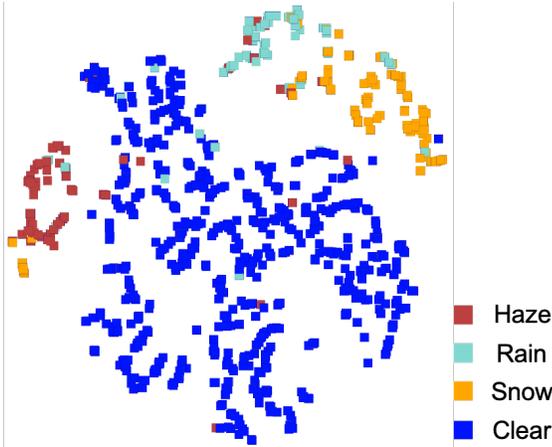


Figure 7. t-SNE visualization of the training set from JHU-Crowd++ based on weather queries.

Dataset Method	ShanghaiTechA		UCF-QNRF	
	MAE	MSE	MAE	MSE
BL [33]	102.6	186.1	164.7	297.0
GL [50]	103.7	188.8	155.0	283.9
MAN [29]	81.2	160.4	132.8	241.9
AWCC-Net	73.2	132.4	120.9	216.7

Table 4. Analysis on the generalization ability of the proposed AWCC-Net. Our method presents better generalization ability on other datasets.

JHU-Crowd++ dataset.

Effectiveness of Weather-aware Crowd Counting. We evaluate the proposed WACC mechanism in Table 6. With only the proposed WACC architecture (MAE of 60.1 and MSE of 229.5), moderate improvements over the baseline model (MAE of 60.7 and MSE of 230.4) are achieved. As

Weather Method	Clear		Adverse Weather		Average	
	MAE	MSE	MAE	MSE	MAE	MSE
MAN	46.5	137.9	105.3	478.4	<u>53.4</u>	209.9
MAN-U	56.9	182.5	<u>100.7</u>	548.2	62.1	255.4
MAN-W	49.8	177.4	103.4	<u>469.6</u>	56.2	232.5
AWCC-Net	<u>47.6</u>	<u>153.9</u>	87.3	430.1	52.3	207.2

Table 5. Comparison of model performance using different weighting strategies for adverse weather and clear data during training on JHU-Crowd++.

Module	Metric					
	WACC	IWQ	\mathcal{L}_{Con}	\mathcal{L}_{CP}	MAE	MSE
-	-	-	-	-	60.7	230.4
✓	-	-	-	-	60.1	229.5
✓	✓	-	-	-	59.3	226.1
✓	✓	✓	-	-	53.1	211.4
✓	✓	✓	✓	-	52.3	207.2

Table 6. Analysis on alternatives of AWCC-Net. Note that we term the input-dependent weather queries as 'IWQ'. We demonstrate that the contrastive weather-adaptive queries can improve the performance of crowd counting effectively.

Probe Dataset	JHU-Crowd++	
Query Dataset	QNRF	JHU-Crowd++
WACC	0.12	0.18
CWM	0.42	0.53

Table 7. Quantitative results of using adverse weather images from QNRF and JHU-Crowd++ validation sets, highlighting top-4 closest matches from the JHU-Crowd++ training set and their mean average precision (mAP) based on weather annotations.

stated in 3.1, it is due to the learning of weather-unrelated

Strategy	MAE	MSE
AWCC-Net-Label	56.1	220.3
AWCC-Net	52.3	207.2

Table 8. **Analysis on the necessity of using weather label in the AWCC-Net.** We demonstrate adopting weather labels may potentially degrade the performance of crowd counting.

information in the weather queries.

Effectiveness of Contrastive Weather-adaptive Module.

The contrastive weather-adaptive module facilitates the weather queries in WACC to be effective for crowd counting, as shown in Table 6. The proposed input-dependent weather queries and contrastive loss can improve the performance of WACC effectively since the learned weather queries are constrained to learn weather information effectively. Moreover, the compact prototype loss facilitates learning compact prototypes in the weather bank.

We validate the effectiveness of the proposed module visually in Figure 6. We compute the weather queries on images outside of the training set. Then, we randomly pick images of different weather types (i.e., snow, rain, and haze) and clear scenes as queries. Based on the chosen weather queries, we identify the images with the top-4 smallest distance from the remaining weather queries. Also, in Table 7, we conduct a quantitative evaluation to evaluate the effectiveness by using 30 and 64 images in adverse weather from the QNRF and JHU-Crowd++ validation sets as queries. These queries are used to probe the images in the training set of JHU-Crowd++ dataset, selecting the top-4 images with the smallest distance. We calculate the mean average precision (mAP) considering the adverse weather annotations. The results show that the proposed mechanism can represent the various weather types effectively while the results without the proposed mechanism may find the images with different weather types since the learned weather queries can not well represent weather information. Additionally, as shown in Figure 7, we provide a t-SNE visualization based on the weather queries from the training set in JHU-Crowd++. Our approach effectively discriminates among different weather types.

Necessity of Weather Annotations. We analyze the performance of adopting annotations of weather type and the proposed learning weather queries strategy in the AWCC-Net. We construct a baseline based on the proposed AWCC-Net and leverage the annotations of weather types as guidance. Specifically, we remove the weight vector in the contrastive weather-adaptive module and set four weather prototypes in the weather bank since the JHU-Crowd++ dataset only contains annotations of haze, rain, snow, and clear scenes. We manually adopt one of these prototypes as the queries in the transformer according to the weather label of the input image. The results are reported in Table 8 where the

aforementioned baseline is denoted as 'AWCC-Net-Label'. The results show that using annotations of weather type in our method may degrade the performance of crowd counting since human-defined weather labels may potentially be mislabeled and contain the weather ambiguity problem. In contrast, our network does not adopt the human label of the weather types and learns the weather prototypes automatically, which shows the better performance in adverse weather crowd counting problems.

6. Conclusion

In this paper, we propose the AWCC-Net model to address the crowd counting problem under adverse weather. We introduce the weather query mechanism to the crowd counting network, which enables our network to learn weather-aware feature extraction. To enforce the weather queries to learn the weather information effectively, we propose a module and contrastive loss to learn weather-adaptive queries for robust crowd counting. Moreover, the compact prototype loss is proposed to improve the model performance. Extensive experimental results show that the proposed method performs favorably against the state-of-the-art methods in both adverse weather and clear images.

7. Acknowledgement

We thank to National Center for High-performance Computing (NCHC) for providing computational and storage resources. This research was supported by the National Science and Technology Council, Taiwan under Grant 111-2221-E-002-136-MY3.

References

- [1] Shahira Abousamra, Minh Hoai, Dimitris Samaras, and Chao Chen. Localization in the crowd with topological constraints. In *AAAI*, 2021. 1, 2
- [2] Dana Berman, Shai Avidan, et al. Non-local image dehazing. In *CVPR*, 2016. 2
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 3
- [4] Antoni B Chan and Nuno Vasconcelos. Bayesian poisson regression for crowd counting. In *ICCV*, 2009. 2
- [5] Ke Chen, Chen Change Loy, Shaogang Gong, and Tony Xiang. Feature mining for localised crowd counting. In *BMVC*, 2012. 2
- [6] Wei-Ting Chen, Jian-Jiun Ding, and Sy-Yen Kuo. Pms-net: Robust haze removal based on patch map for single images. In *CVPR*, 2019. 2
- [7] Wei-Ting Chen, Hao-Yu Fang, Cheng-Lin Hsieh, Cheng-Che Tsai, I Chen, Jian-Jiun Ding, Sy-Yen Kuo, et al. All snow removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In *ICCV*, 2021. 2
- [8] Wei-Ting Chen, Zhi-Kai Huang, Cheng-Che Tsai, Hao-Hsiang Yang, Jian-Jiun Ding, and Sy-Yen Kuo. Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive regularization: Toward a unified model. In *CVPR*, 2022. 1, 3, 6
- [9] Zhi-Qi Cheng, Qi Dai, Hong Li, Jingkuan Song, Xiao Wu, and Alexander G Hauptmann. Rethinking spatial invariance of convolutional networks for object counting. In *CVPR*, 2022. 6, 8
- [10] Qili Deng, Ziling Huang, Chung-Chi Tsai, and Chia-Wen Lin. Hardgan: A haze-aware representation distillation gan for single image dehazing. In *ECCV*, 2020. 2
- [11] Yingjun Du, Jun Xu, Qiang Qiu, Xiantong Zhen, and Lei Zhang. Variational image deraining. In *WACV*, 2020. 2
- [12] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *CVPR*, 2017. 2
- [13] Junyu Gao, Tao Han, Qi Wang, and Yuan Yuan. Domain-adaptive crowd counting via inter-domain features segregation and gaussian-prior reconstruction. *arXiv preprint arXiv:1912.03677*, 2(5), 2019. 1, 2
- [14] Junyu Gao, Tao Han, Yuan Yuan, and Qi Wang. Learning independent instance maps for crowd localization. *arXiv preprint arXiv:2012.04164*, 2020. 2
- [15] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *TPAMI*, 2010. 2
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [17] Ming Hong, Yuan Xie, Cuihua Li, and Yanyun Qu. Distilling image dehazing with heterogeneous task imitation. In *CVPR*, 2020. 2
- [18] Haroon Idrees, Imran Saleemi, Cody Seibert, and Mubarak Shah. Multi-source multi-scale counting in extremely dense crowd images. In *CVPR*, 2013. 2
- [19] Haroon Idrees, Muhammad 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 *ECCV*, 2018. 1, 2, 5, 8
- [20] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, and Junjun Jiang. Multi-scale progressive fusion network for single image deraining. In *CVPR*, 2020. 2
- [21] Tai-Xiang Jiang, Ting-Zhu Huang, Xi-Le Zhao, Liang-Jian Deng, and Yao Wang. A novel tensor-based video rain streaks removal approach via utilizing discriminatively intrinsic priors. In *CVPR*, 2017. 2
- [22] Xiaoheng Jiang, Li Zhang, Mingliang Xu, Tianzhu Zhang, Pei Lv, Bing Zhou, Xin Yang, and Yanwei Pang. Attention scaling for crowd counting. In *CVPR*, 2020. 1, 2
- [23] Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *CVPR*, 2020. 2
- [24] Yuhong Li, Xiaofan Zhang, and Deming Chen. Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes. In *CVPR*, 2018. 1
- [25] Zhi Li, Juan Zhang, Zhijun Fang, Bo Huang, Xiaoyan Jiang, Yongbin Gao, and Jenq-Neng Hwang. Single image snow removal via composition generative adversarial networks. *IEEE Access*, 7:25016–25025, 2019. 2
- [26] Dingkang Liang, Wei Xu, and Xiang Bai. An end-to-end transformer model for crowd localization. In *ECCV*, 2022. 1, 5, 6, 8
- [27] Dingkang Liang, Wei Xu, Yingying Zhu, and Yu Zhou. Focal inverse distance transform maps for crowd localization. *TMM*, 2022. 1, 2
- [28] Hui Lin, Xiaopeng Hong, Zhiheng Ma, Xing Wei, Yunfeng Qiu, Yaowei Wang, and Yihong Gong. Direct measure matching for crowd counting. *arXiv preprint arXiv:2107.01558*, 2021. 5, 8
- [29] Hui Lin, Zhiheng Ma, Rongrong Ji, Yaowei Wang, and Xiaopeng Hong. Boosting crowd counting via multifaceted attention. In *CVPR*, 2022. 1, 4, 5, 6, 8
- [30] Jiang Liu, Chenqiang Gao, Deyu Meng, and Alexander G Hauptmann. Decidenet: Counting varying density crowds through attention guided detection and density estimation. In *CVPR*, 2018. 2
- [31] Weizhe Liu, Nikita Durasov, and Pascal Fua. Leveraging self-supervision for cross-domain crowd counting. In *CVPR*, 2022. 1
- [32] Yuting Liu, Miaoqing Shi, Qijun Zhao, and Xiaofang Wang. Point in, box out: Beyond counting persons in crowds. In *CVPR*, 2019. 2
- [33] Zhiheng Ma, Xing Wei, Xiaopeng Hong, and Yihong Gong. Bayesian loss for crowd count estimation with point supervision. In *ICCV*, 2019. 4, 5, 8
- [34] Zhiheng Ma, Xing Wei, Xiaopeng Hong, Hui Lin, Yunfeng Qiu, and Yihong Gong. Learning to count via unbalanced optimal transport. In *AAAI*, 2021. 5, 8

- [35] Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, Yu-Wing Tai, et al. Learning dual convolutional neural networks for low-level vision. In *CVPR*, 2018. 2
- [36] Soo-Chang Pei, Yu-Tai Tsai, and Chen-Yu Lee. Removing rain and snow in a single image using saturation and visibility features. In *ICMEW*, 2014. 2
- [37] Yifei Qian, Liangfei Zhang, Xiaopeng Hong, Carl R Donovan, Ognjen Arandjelovic, UK Fife, and PR Harbin. Segmentation assisted u-shaped multi-scale transformer for crowd counting. In *BMVC*, 2022. 1
- [38] Xu Qin, Zhilin Wang, Yuanhao Bai, Xiaodong Xie, and Huizhu Jia. Ffa-net: Feature fusion attention network for single image dehazing. In *AAAI*, 2020. 2
- [39] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*, 2015. 2
- [40] Deepak Babu Sam, Skand Vishwanath Peri, Mukuntha Narayanan Sundararaman, Amogh Kamath, and R Venkatesh Babu. Locate, size, and count: accurately resolving people in dense crowds via detection. *TPAMI*, 2020. 2, 5, 8
- [41] Yuanjie Shao, Lerenhan Li, Wenqi Ren, Changxin Gao, and Nong Sang. Domain adaptation for image dehazing. In *CVPR*, 2020. 2
- [42] Weibo Shu, Jia Wan, Kay Chen Tan, Sam Kwong, and Antoni B Chan. Crowd counting in the frequency domain. In *CVPR*, 2022. 5, 8
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2, 3
- [44] Vishwanath A Sindagi and Vishal M Patel. Generating high-quality crowd density maps using contextual pyramid cnns. In *ICCV*, 2017. 2
- [45] Vishwanath A Sindagi, Rajeev Yasarla, and Vishal M Patel. Jhu-crowd++: Large-scale crowd counting dataset and a benchmark method. *TPAMI*, 2020. 1, 5, 8
- [46] 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 *ICCV*, 2021. 2
- [47] Guolei Sun, Yun Liu, Thomas Probst, Danda Pani Paudel, Nikola Popovic, and Luc Van Gool. Boosting crowd counting with transformers. *arXiv preprint arXiv:2105.10926*, 2021. 1
- [48] Ye Tian, Xiangxiang Chu, and Hongpeng Wang. Cctrans: Simplifying and improving crowd counting with transformer. *arXiv preprint arXiv:2109.14483*, 2021. 1
- [49] Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *CVPR*, 2022. 1, 3
- [50] Jia Wan, Ziquan Liu, and Antoni B Chan. A generalized loss function for crowd counting and localization. In *CVPR*, 2021. 1, 5, 8
- [51] Feng Wang and Huaping Liu. Understanding the behaviour of contrastive loss. In *CVPR*, 2021. 4
- [52] Hong Wang, Qi Xie, Qian Zhao, and Deyu Meng. A model-driven deep neural network for single image rain removal. In *CVPR*, 2020. 2
- [53] Hong Wang, Zongsheng Yue, Qi Xie, Qian Zhao, Yefeng Zheng, and Deyu Meng. From rain generation to rain removal. In *CVPR*, 2021. 2
- [54] Qi Wang, Junyu Gao, Wei Lin, and Xuelong Li. Nwpu-crowd: A large-scale benchmark for crowd counting and localization. *TPAMI*, 2020. 1, 5, 8
- [55] Qi Wang, Junyu Gao, Wei Lin, and Yuan Yuan. Learning from synthetic data for crowd counting in the wild. *CVPR*, 2019. 5, 8
- [56] Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson WH Lau. Spatial attentive single-image deraining with a high quality real rain dataset. In *CVPR*, 2019. 2
- [57] Yi Wang, Junhui Hou, Xinyu Hou, and Lap-Pui Chau. A self-training approach for point-supervised object detection and counting in crowds. *TIP*, 2021. 1, 2
- [58] Yinglong Wang, Shuaicheng Liu, Chen Chen, and Bing Zeng. A hierarchical approach for rain or snow removing in a single color image. *TIP*, 2017. 2
- [59] Wei Wei, Lixuan Yi, Qi Xie, Qian Zhao, Deyu Meng, and Zongben Xu. Should we encode rain streaks in video as deterministic or stochastic? In *ICCV*, 2017. 2
- [60] Chenfeng Xu, Dingkan Liang, Yongchao Xu, Song Bai, Wei Zhan, Xiang Bai, and Masayoshi Tomizuka. Autoscale: learning to scale for crowd counting. *IJCV*, 2022. 2
- [61] Hao-Hsiang Yang, Chao-Han Huck Yang, and Yu-Chiang Frank Wang. Wavelet channel attention module with a fusion network for single image deraining. In *ICIP*, 2020. 2
- [62] Wenhan Yang, Jiaying Liu, Shuai Yang, and Zongming Guo. Scale-free single image deraining via visibility-enhanced recurrent wavelet learning. *TIP*, 2019. 2
- [63] Wenhan Yang, Robby T Tan, Jiashi Feng, Zongming Guo, Shuicheng Yan, and Jiaying Liu. Joint rain detection and removal from a single image with contextualized deep networks. *TPAMI*, 2019. 2
- [64] Yuntong Ye, Yi Chang, Hanyu Zhou, and Luxin Yan. Closing the loop: Joint rain generation and removal via disentangled image translation. In *CVPR*, 2021. 2
- [65] Shujian Yu, Yixiao Zhao, Yi Mou, Jinghui Wu, Lu Han, Xiaopeng Yang, and Baojun Zhao. Content-adaptive rain and snow removal algorithms for single image. In *ISNN*, 2014. 2
- [66] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021. 2
- [67] He Zhang and Vishal M Patel. Densely connected pyramid dehazing network. In *CVPR*, 2018. 2
- [68] Kaihao Zhang, Rongqing Li, Yanjiang Yu, Wenhan Luo, and Changsheng Li. Deep dense multi-scale network for snow removal using semantic and depth priors. *TIP*, 2021. 2
- [69] Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*, 2016. 1, 5, 8

- [70] Xianhui Zheng, Yinghao Liao, Wei Guo, Xueyang Fu, and Xinghao Ding. Single-image-based rain and snow removal using multi-guided filter. In *ICONIP*, 2013. [2](#)