

# JHU-CROWD++: Large-Scale Crowd Counting Dataset and A Benchmark Method

Vishwanath A. Sindagi, *Student Member, IEEE*, Rajeev Yasarla, *Student Member, IEEE*, and Vishal M. Patel, *Senior Member, IEEE*

**Abstract**—We introduce a new large scale unconstrained crowd counting dataset (JHU-CROWD++) that contains “4,372” images with “1.51 million” annotations. In comparison to existing datasets, the proposed dataset is collected under a variety of diverse scenarios and environmental conditions. Specifically, the dataset includes several images with weather-based degradations and illumination variations, making it a very challenging dataset. Additionally, the dataset consists of a rich set of annotations at both image-level and head-level. Several recent methods are evaluated and compared on this dataset. The dataset can be downloaded from <http://www.crowd-counting.com>.

Furthermore, we propose a novel crowd counting network that progressively generates crowd density maps via residual error estimation. The proposed method uses VGG16 as the backbone network and employs density map generated by the final layer as a coarse prediction to refine and generate finer density maps in a progressive fashion using residual learning. Additionally, the residual learning is guided by an uncertainty-based confidence weighting mechanism that permits the flow of only high-confidence residuals in the refinement path. The proposed Confidence Guided Deep Residual Counting Network (CG-DRCN) is evaluated on recent complex datasets, and it achieves significant improvements in errors.

**Index Terms**—crowd counting, dataset.

## 1 INTRODUCTION

WITH burgeoning population and rapid urbanization, crowd gatherings have become more prominent in the recent years. Consequently, computer vision-based crowd analytics and surveillance [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17] have received increased interest. Furthermore, algorithms developed for the purpose of crowd analytics have found applications in other fields such as agriculture monitoring [18], microscopic biology [19], urban planning and environmental survey [2, 20]. Current state-of-the-art counting networks achieve impressive error rates on a variety of datasets that contain numerous challenges. Their success can be broadly attributed to two major factors: (i) development and publication of challenging datasets [3, 4, 5, 21], and (ii) design of novel convolutional neural network (CNN) architectures specifically for improving count performance [4, 7, 22, 23, 24, 25, 26, 27]. In this paper, we consider both of the above factors in an attempt to further improve the crowd counting performance.

First, we identify the next set of challenges that require attention from the crowd counting research community and collect a large-scale dataset collected under a variety of conditions. Existing efforts like UCF\_CROWD\_50 [3], World Expo '10 [4] and ShanghaiTech [28] have progressively increased the complexity of the datasets in terms of average count per image, image diversity *etc.* While these datasets have enabled rapid progress in the counting task, they suffer

from shortcomings such as limited number of training samples, limited diversity in terms of environmental conditions, dataset bias in terms of positive samples, and limited set of annotations. Idrees *et al.*[21] proposed a new dataset called UCF-QNRF that alleviates some of these challenges. Most recently, Wang *et al.*[29] released a large-scale crowd counting dataset consisting of 5,109 images with 2.13 million annotations. Specifically, the images are collected under a variety of illumination conditions. Nevertheless, they do not specifically consider some of the challenges such as adverse environmental conditions, dataset bias and limited annotation data<sup>1</sup>.

To address these issues, we propose a new large-scale unconstrained dataset (JHU-CROWD++) with a total of 4,372 images (containing 1,515,005 head annotations) that are collected under a variety of conditions. Specific care is taken to include images captured under various weather-based degradations. Additionally, we include a set of distractor images that are similar to the crowd images that contain complex backgrounds which may be confused for crowd. Fig. 1 illustrates representative samples of the images in the JHU-CROWD++ dataset under various categories. Furthermore, the dataset also provides a much richer set of annotations at both image-level and head-level. These annotations include point-wise annotations, approximate sizes, blur-level, occlusion-level, weather-labels, *etc.* We also benchmark several representative counting networks, providing an overview of the state-of-the-art performance.

Next, we consider the design of network architecture for the task of counting. Design of novel networks specifically for the task of counting has improved the counting error by

• V. A. Sindagi, R. Yasarla and V. M. Patel are with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, 21218.  
E-mail: (vishwanathsindagi@jhu.edu, ryasarla19@jhu.edu, vpa-tel36@jhu.edu)

Manuscript received April 19, 2015; revised August 26, 2015.

1. Existing datasets provide only point-wise annotations.



Fig. 1. Representative samples of the images in the JHU-CROWD++ dataset. (a) Overall (b) Rain (c) Snow (d) Haze (e) Distractors.

large margins. Architectures have evolved from the simple ones like [4] which consisted of a set of convolutional and fully connected layers, to the most recent complex architectures like SA-Net [26] which consists of a set of scale aggregation modules. Typically, most existing works ([4, 5, 6, 7, 22, 23, 24, 25, 26, 26, 30, 31]) have designed their networks by laying a strong emphasis on addressing large variations of scale in crowd images. While this strategy of developing robustness towards scale changes has resulted in significant performance gains, it is nevertheless important

to exploit other properties like in [25, 32, 33] to further the improvements.

In a similar attempt, we exploit residual learning mechanism for the purpose of improving crowd counting. Specifically, we present a novel design based on the VGG16 network [34], and it employs residual learning to progressively generate better quality crowd density maps. This use of residual learning is inspired by its success in several other tasks like super-resolution [35, 36, 36, 37, 38]. Although this technique results in improvements in performance, it is

important to ensure that only highly confident residuals are used in order to ensure the effectiveness of residual learning. To address this issue, we draw inspiration from the success of uncertainty-based learning mechanism [39, 40, 41]. We propose an uncertainty-based confidence weighting module that captures high-confidence regions in the feature maps to focus on during the residual learning. The confidence weights ensure that only highly confident residuals get propagated to the output, thereby increasing the effectiveness of the residual learning mechanism. Furthermore, we exploit the additional image-level labels in the proposed dataset to extend the uncertainty-based confidence weighting module by conditioning it on the labels to improve the performance specifically in the adverse weather conditions.

To summarize, the following are our key contributions:

- We propose a new large-scale unconstrained crowd counting dataset that contains 4,372 images and 1.51 million annotations. The dataset specifically includes a number of images collected under adverse weather conditions. Furthermore, this is the first counting dataset that provides a rich set of annotations such as occlusion, blur, scale, and image-level labels, *etc.*
- We propose a crowd counting network that progressively incorporates residual mechanism to estimate high quality density maps. Furthermore, a set of uncertainty-based confidence weighting modules are introduced in the network to improve the efficacy of residual learning.

Note that this work is an extension of our ICCV 2019 work [42]. Compared to our earlier work, we attempt to improve both dataset and proposed method. These improvements are summarized below:

- **Dataset:** Specifically, we provide 3 key improvements as compared to the JHU-CROWD dataset [42]:
  - (i) *More number of images:* We increase the number of images in the dataset from 4,250 to 4,372.
  - (ii) *More number of annotations:* The new dataset contains 31% more annotations. (1.51 million v/s 1.15 million).
  - (iii) *Better scale annotations:* The earlier version of the dataset contains size indicators for each head in the image. In the new dataset, we provide better scale annotations which consist of approximate width and height of each head.
- **Method:** We provide the following improvements in the proposed method:
  - (i) *Class conditioning:* We extend the CGDRN method proposed in [42] to improve the counting performance in the adverse weather conditions. Specifically, we condition the uncertainty-guided residual estimation on the image level labels to incorporate weather-based information into the learning process.
  - (ii) *New backbone:* We demonstrate that the proposed uncertainty-based residual learning mechanism generalizes to other backbone networks like Res101 [43]. This results in further improvements on all the datasets.
- **Experiments:** We conduct the following new experiments:
  - (i) *Ablation study:* We conduct additional ablation studies where we evaluate the efficacy of number of branches and different network architecture.
  - (ii) *Benchmarking:* We benchmark recently published methods on the newly proposed JHU-CROWD++ dataset.

## 2 RELATED WORK

**Crowd Datasets.** Crowd counting datasets have evolved over time with respect to a number of factors such as size, crowd densities, image resolution, and diversity. UCSD [8] is among one of the early datasets proposed for counting and it contains 2000 video frames of low resolution with 49,885 annotations. The video frames are collected from a single frame and typically contain low density crowds. Zhang *et al.*[4] addressed the limitations of UCSD dataset by introducing the WorldExpo dataset that contains 108 videos with a total of 3,980 frames belonging to 5 different scenes. While the UCSD and WorldExpo datasets contain only low/low-medium densities, Idrees *et al.*[3] proposed the UCF\_CROWD\_50 dataset specifically for very high density crowd scenarios. However, the dataset consists of only 50 images rendering it impractical for training deep networks. Zhang *et al.*[5] introduced the ShanghaiTech dataset which has better diversity in terms of scenes and density levels as compared to earlier datasets. The dataset is split into two parts: Part A (containing high density crowd images) and Part B (containing low density crowd images). The entire dataset contains 1,198 images with 330,165 annotations. Recently, Idrees *et al.*[21] proposed a new large-scale crowd dataset containing 1,535 high density images with a total of 1.25 million annotations. Wang *et al.*[44] introduced a synthetic crowd counting dataset that is based on GTA V electronic game. The dataset consists of 15,212 crowd images under a diverse set of scenes. In addition, they proposed a SSIM based CycleGAN [45] for adapting the network trained on synthetic images to real world images. Most recently, Wang *et al.*[29] released a large-scale crowd counting dataset (NWPUCrowd) consisting of 5,109 images with 2.13 million annotations.

**Crowd Counting.** Traditional approaches for crowd counting from single images are based on hand-crafted representations and different regression techniques. Loy *et al.*[46] categorized these methods into (1) detection-based methods [47] (2) regression-based methods [3, 48, 49] and (3) density estimation-based methods [19, 50, 51]. Interested readers are referred to [1, 49] for a more comprehensive study of different crowd counting methods.

Recent advances in CNNs have been exploited for the task of crowd counting and these methods [4, 5, 6, 7, 7, 22, 23, 52, 53, 54, 55, 56] have demonstrated significant improvements over the traditional methods. A recent survey [57] categorizes these approaches based on the network property and the inference process. Walach *et al.*[22] used CNNs with layered boosting approach to learn a non-linear function between an image patch and count. Recent work [5, 23] addressed the scale issue using different architectures. Sam *et al.*[7] proposed a VGG16-based switching classifier that first identifies appropriate regressor based on the content of the input image patch. More recently, Sindagi *et al.*[6] proposed to incorporate global and local context from the input image into the density estimation network. In another approach, Cao *et al.*[26] proposed an encoder-decoder network with scale aggregation modules.

In contrast to these methods that emphasize on specifically addressing large-scale variations in head sizes, the

most recent methods ([30], [32], [33], [58], [25]) have focused on other properties of the problem. For instance, Babu *et al.* [30] proposed a mechanism to incrementally increase the network capacity conditioned on the dataset. Shen *et al.* [32] overcame the issue of blurred density maps by utilizing adversarial loss. In a more recent approach, Ranjan *et al.* [25] proposed a two-branch network to estimate density map in a cascaded manner. Shi *et al.* [33] employed deep negative correlation based learning for more generalizable features. Liu *et al.* [58] used unlabeled data for counting by proposing a new framework that involves learning to rank.

Recent approaches like [31, 59, 60, 61, 62, 63] have aimed at incorporating various forms of related information like attention [59], semantic priors [60], segmentation [61], inverse attention [62], and hierarchical attention [31] respectively into the network. Other techniques such as [64, 65, 66, 67, 68] leverage features from different layers of the network using different techniques like trellis style encoder decoder [64], explicitly considering perspective [65], context information [66], adaptive density map generation [68] and multiple views [67]. More recently, Sam *et al.* [69] introduced a detection framework for densely crowded scenarios where the network is trained using estimated bounding-boxes. Ma *et al.* [70] proposed a novel Bayesian loss function for training counting networks, which involves supervision on the count expectation at each annotated point. While most of the existing approaches are focused on counting in 2D plane, Zhang *et al.* [71] propose to solve the multi-view crowd counting task through 3D feature fusion with 3D scene-level density maps. For a comprehensive study on various crowd counting techniques, the reader is referred to detailed surveys like [57, 72].

### 3 JHU-CROWD++: LARGE-SCALE CROWD COUNTING DATASET

In this section, we first motivate the need for a new crowd counting dataset, followed by a detailed description of the various factors and conditions while collecting the dataset.

#### 3.1 Motivation and dataset details

As discussed earlier, existing datasets (such as UCF\_CROWD\_50 [3], World Expo '10 [4] and ShanghaiTech [28]) have enabled researchers to develop novel counting networks that are robust to several factors such as variations in scale, pose, view *etc.* Several recent methods have specifically addressed the large variations in scale by proposing different approaches such as multi-column networks [5], incorporating global and local context [6], scale aggregation network [26], *etc.* These methods are largely successful in addressing issues in the existing datasets, and there is pressing need to identify newer set of challenges that require attention from the crowd counting community.

In what follows, we describe the shortcomings of existing datasets and discuss the ways in which we overcome them:

(i) *Limited number of training samples:* Typically, crowd counting datasets have limited number of images available for training and testing. For example, ShanghaiTech dataset

[5] has only 1,198 images and this low number of images results in lower diversity of the training samples. Due to this issue, networks trained on this dataset will have reduced generalization capabilities. Although datasets like Mall [49], WorldExpo '10 [4] have higher number of images, it is important to note that these images are from a set of video sequences from surveillance cameras and hence, they have limited diversity in terms of background scenes and number of people. Most recently, Idrees *et al.* [21] addressed this issue by introducing a high-quality dataset (UCF-QNRF) that has images collected from various geographical locations under a variety of conditions and scenarios. Although it has a large set of diverse scenarios, the number of samples is still limited from the perspective of training deep neural networks.

To address this issue, we collect a new large scale unconstrained dataset with a total of 4,372 images that are collected under a variety of conditions. Such a large number of images results in increased diversity in terms of count, background regions, scenarios, *etc.* as compared to existing datasets. The images are collected from several sources on the internet using different keywords such as crowd, crowd+marathon, crowd+walking, crowd+India, *etc.* A summary of the keywords used for the search purpose is illustrated in Fig. 3.

(ii) *Absence of adverse conditions:* Typical application of crowd counting is video surveillance in outdoor scenarios which involve regular weather-based degradations such as haze, snow, rain *etc.* It is crucial that networks, deployed under such conditions, are robust.

To overcome this issue, specific care is taken during our dataset collection efforts to include images captured under various weather-based degradations such as rain, haze, snow, *etc.* (as shown in Fig. 1(b-d)). Table 2 summarizes images collected under adverse conditions.

(iii) *Dataset bias:* Existing datasets focus on collecting only images with crowd, due to which a deep network trained on such a dataset may end up learning bias in the dataset. Due to this error, the network will erroneously predict crowd even in scenes that do not contain crowd.

In order to address this, we include a set of distractor images that are similar to crowd images but contain very few people. These images can enable the network to avoid learning bias in the dataset. The total number of distractor images in the dataset is 106. Fig. 1(e) shows sample distractor images.

(iv) *Limited annotations:* Typically, crowd counting datasets provide point-wise annotations for every head/person in the image, *i.e.*, each image is provided with a list of  $x, y$  locations of the head centers. While these annotations enable the networks to learn the counting task, absence of more information such as occlusion level, head sizes, blur level *etc.* limits the learning ability of the networks. For instance, due to the presence of large variations in perspective, size of the head is crucial to determine the precise count. One of the reasons for these missing annotations is that crowd images typically contain several people and it is highly labor intensive to obtain detailed annotations such as size.



Fig. 2. Examples of head-level annotations: (a) Dots (b) Approximate sizes (c) Blur-level.

TABLE 1

Comparison of different datasets. P: Point-wise annotations for head locations, O: Occlusion level per head, B: Blur level per head, S: Size indicator per head, S<sup>†</sup>: Approximate size (w×h), I: Image level labels.

Dataset	Num of Images	Num of Annotations	Avg Count	Max Count	Avg H×W	Weather degradations	Distractors	Type of annotations
UCSD [8]	2,000	49,885	25	46	158×238	✗	✗	P
Mall [49]	2,000	62,325	-	53	320×240	✗	✗	P
UCF_CROWD_50 [3]	50	63,974	1,279	4,543	2101×2888	✗	✗	P
WorldExpo '10 [4]	3,980	199,923	50	253	576×720	✗	✗	P
ShanghaiTech [5]	1,198	330,165	275	3,139	598×868	✗	✗	P
UCF-QNRF [21]	1,535	1,251,642	815	12,865	2,013×2,902	✗	✗	P
NWPU-CROWD [29]	5,109	2,133,238	418	20,033	2,311×3,383	✗	✓	P
JHU-CROWD (ours)	4,250	1,114,785	262	7,286	900×1,450	✓	✓	P, O, B, S, I
JHU-CROWD++ (ours)	4,372	1,515,005	346	25,791	910×1,430	✓	✓	P, O, B, S <sup>†</sup> , I

TABLE 2

Summary of images collected under adverse conditions.

Degradation type	Rain	Snow	Fog/Haze	Total
No. of images	145	201	168	514
No. of annotations	40,328	47,347	48,821	136,496

To enable more effective learning, we collect a much richer set of annotations at both head-level/point-level and

image-level. These are described below:

- Head-level/point-level annotations include  $x, y$  locations of heads and corresponding occlusion level, blur level and size level. The total number of point-level annotations in the dataset are 1,515,005. Occlusion label has three levels:  $\{un\text{-occluded}, \text{partially occluded}, \text{fully occluded}\}$ . Blur level has two labels:  $\{blur, no\text{-blur}\}$ . In JHU-CROWD [42], each head is labeled with a size indicator. We improve over

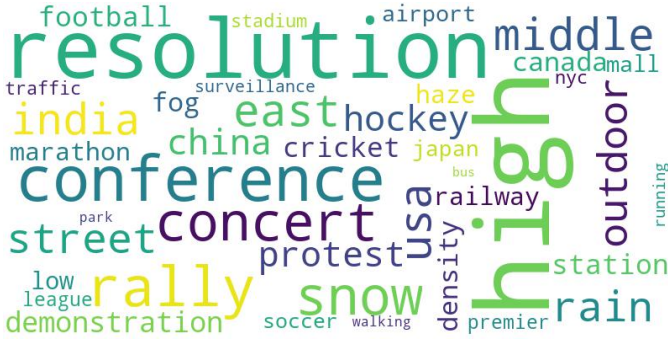


Fig. 3. Summary of keywords used to scrape the internet for images.

TABLE 3

Distribution of images under different densities.

Density	Low (0-50)	Med (51-500)	High (500+)	Total
No. of images	1,228	2,512	632	4,372

these size annotations by providing “approximate” size (width and height) for each head annotation. To obtain these, annotators were instructed to annotate bounding boxes for a set of neighbouring heads which have similar sizes. Note that these bounding boxes are only “approximate” and are not as accurate as the ones found in detection datasets. Fig. 2 illustrates sample annotations provided in our dataset.

- Image level annotations include scene-labels (such as *marathon, mall, railway station, stadium, etc.*) and the weather-labels (rain, snow and fog). Fig. 4 illustrates the distribution of scene-labels in the proposed dataset.

### 3.2 Data collection process

We used different sources like google images, bing images, flickr, etc for collecting the images. The keywords used for searching were carefully selected to ensure diversity in terms of density, weather, geographical locations, scene-type, resolution, events, physical structures etc. The keywords are summarized in Fig. 3 and Table 4. After the collection process, duplicate images were detected and eliminated from the dataset. The remaining images were annotated with the help of Amazon Mechanical Turk workers. The workers were explicitly instructed to annotate the center of head. Additionally, they were also instructed to label the occlusion type and blur-level. The annotated images were then evaluated manually verified to filter out incorrectly labeled images. Such images were sent again to the annotation and verification processes.

TABLE 4

List of keywords used for searching.

Factor	Keyword
Crowd density	low crowd, small crowd, large crowd, high density crowd
Geographical location	US, China, India, Iran, Iraq, Canada
Structures	airport, mall, railway station, street, park, stadium, traffic
Events	concert, protest, rally, festival, sports, cricket, football, soccer, hockey, premier league, conference, marathon
Weather	rain, snow, fog, haze, low-light

### 3.3 Summary and evaluation protocol

Fig. 1 illustrates representative samples of the images in the JHU-CROWD++ dataset under various categories. Table




Fig. 4. Distribution of image-level labels.

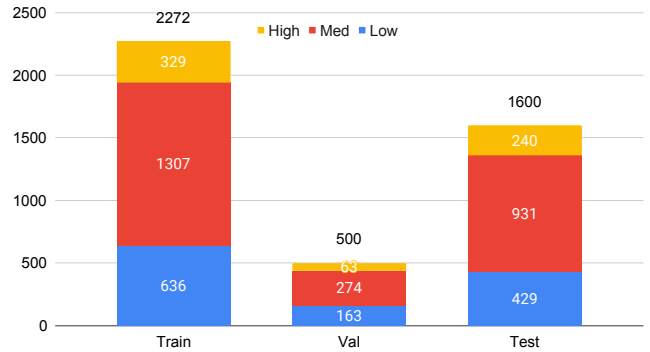


Fig. 5. Distribution of images of different density levels in train, val and test sets.




Fig. 6. Distribution of images of weather conditions in train, val and test sets.

1 summarizes the proposed dataset in comparison with the existing ones. It can be observed that the proposed dataset enjoys a host of properties such as a richer set of annotations, weather-based degradations and distractor images. With these properties, the proposed dataset will serve as a good complementary to other datasets such as UCF-QNRF and NWPU-CROWD. The dataset is randomly split into train, val and test sets, which contain 2722, 500 and 1600 images respectively.

Following the existing works, we perform evaluation using the standard MAE and MSE metrics. Furthermore, these metrics are calculated for the following sub-categories of images:

- (i) Low density: images containing count between 0 and 50,
- (ii) Med density: images with count between 51 and 500,

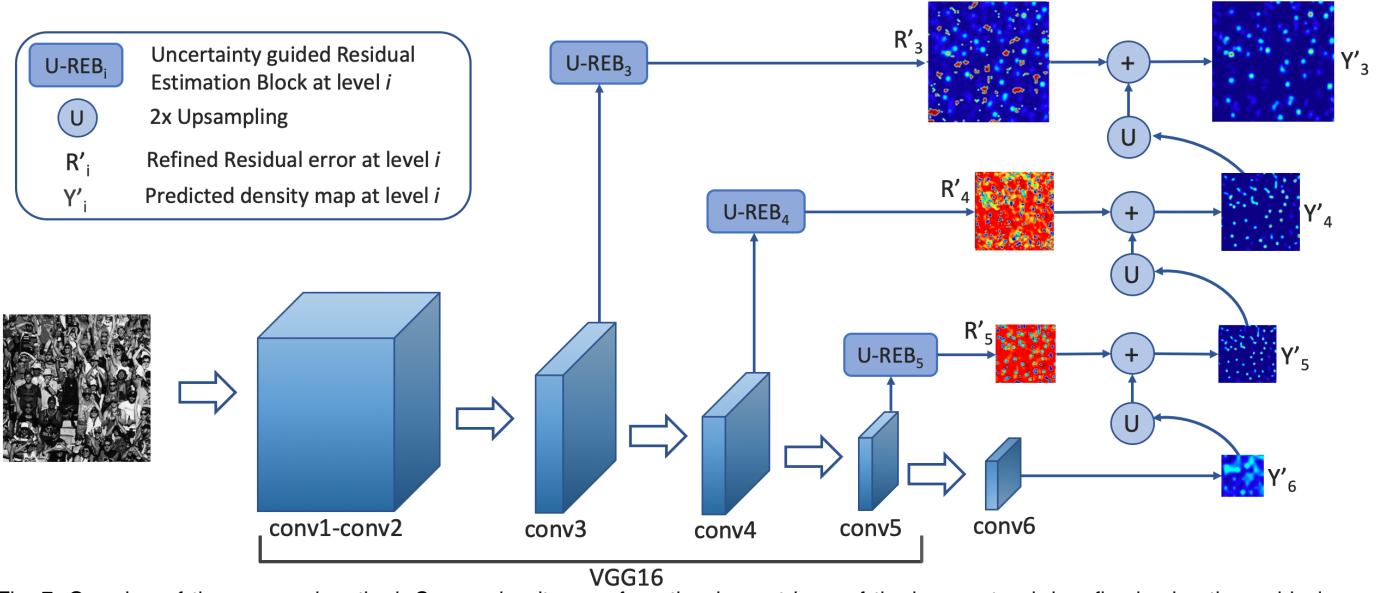


Fig. 7. Overview of the proposed method. Coarse density map from the deepest layer of the base network is refined using the residual map estimated by the shallower layer. The residual estimation is performed by  $U-REB_i$ . In the residual maps, red indicates negative values and cyan indicates positive value.

- (iii) High density: images with count more than 500 people,
- (iv) Weather degraded images, and
- (v) Overall.

The metrics under these sub-categories will provide a holistic understanding of the network performance.

Fig. 5 and Fig. 6 illustrate the distribution the number of images among the density and weather sub-categories respectively. Table 3 shows the distribution of images for different density-levels.

## 4 PROPOSED METHOD

In this section, we present the details of the proposed Confidence Guided Deep Residual Crowd Counting (CG-DRCN) along with the training and inference specifics. Fig. 7 shows the architecture of the proposed network.

### 4.1 Base network

Following recent approaches [6, 7, 26], we perform counting based on the density estimation framework. In this framework, the network is trained to estimate the density map ( $\hat{Y}$ ) from an input crowd image ( $X$ ). The target density map ( $Y$ ) for training the network is generated by imposing normalized 2D Gaussian at head locations provided by the dataset annotations:

$$Y(x) = \sum_{x_g \in S} \mathcal{N}(x - x_g, \sigma), \quad (1)$$

where,  $S$  is the set of all head locations ( $x_g$ ) in the input image and  $\sigma$  is scale parameter of 2D Gaussian kernel. Due to this formulation, the density map contains per-pixel density information of the scene, which when integrated results in the count of people in the image.

The proposed network consists of conv1~conv5 layers ( $C_1 - C_5$ ) of the VGG16 architecture as a part of the backbone, followed by a conv block ( $CB_6$ ) and a max-pooling layer with stride 2. First, the input image (of size

$W \times H$ ) is passed through  $C_1 - C_5$ ,  $CB_6$  and the max pooling layer to produce the corresponding density map ( $\hat{Y}_6$ ) of size  $\frac{W}{32} \times \frac{H}{32}$ .  $CB_6$  is defined by  $\{\text{conv}_{512,32,1}\text{-relu-conv}_{32,32,3}\text{-relu-conv}_{32,1,3}\}^2$ . Due to its low resolution, ( $\hat{Y}_6$ ) can be considered as a coarse estimation, and learning this will implicitly incorporate global context in the image due the large receptive field at the deepest layer in the network.

### 4.2 Residual learning

Although  $\hat{Y}_6$  provides a good estimate of the number of people in the image, the density map lacks several local details as shown in Fig. 9 (a). This is because deeper layers learn to capture abstract concepts and tend to lose low level details in the image. On the other hand, the shallower layers have relatively more detailed local information as compared to their deeper counterparts [73]. Based on this observation, we propose to refine the coarser density maps by employing shallower layers in a residual learning framework. This refinement mechanism is inspired in part by several leading work on super-resolution [35, 36, 37] that incorporate residual learning to learn finer details required to generate a high quality super-resolved image. Specifically, features from  $C_5$  are forwarded through a uncertainty guided residual estimation block ( $U-REB_5$  to generate a residual map  $\hat{R}_5$ , which is then added to an appropriately up-sampled version of  $\hat{Y}_6$  to produce the density map  $\hat{Y}_5$  of size  $\frac{W}{16} \times \frac{H}{16}$ , i.e.,

$$\hat{Y}_5 = \hat{R}_5 + up(\hat{Y}_6). \quad (2)$$

Here,  $up()$  denotes up-sampling by a factor of  $2 \times$  via bilinear interpolation. By enforcing  $U-REB_5$  to learn a residual map, the network focuses on the local errors emanating from the deeper layer, resulting in better learning of the offsets required to refined the coarser density map.  $U-REB$  is described in Section 4.3.

2.  $conv_{N_i, N_o, k}$  denotes conv layer (with  $N_i$  input channels,  $N_o$  output channels,  $k \times k$  filter size),  $relu$  denotes ReLU activation

The above refinement is further repeated to recursively generate finer density maps  $\hat{Y}_4$  and  $\hat{Y}_3$  using the feature maps from the shallower layers  $C_4$  and  $C_3$ , respectively. Specifically, the output of  $C_4$  and  $C_3$  are forwarded through  $U-REB_4$ ,  $U-REB_3$  to learn residual maps  $\hat{R}_4$  and  $\hat{R}_3$ , which are then added to the appropriately up-sampled versions of the coarser maps  $\hat{Y}_5$  and  $\hat{Y}_4$  to produce  $\hat{Y}_4$  and  $\hat{Y}_3$  respectively in that order. Specifically,  $\hat{Y}_4$  and  $\hat{Y}_3$  are obtained as follows:

$$\hat{Y}_4 = \hat{R}_4 + up(\hat{Y}_5), \quad \hat{Y}_3 = \hat{R}_3 + up(\hat{Y}_4) \quad (3)$$

### 4.3 Uncertainty guided residual learning (U-REB)

In this section, we provide a detailed description of the uncertainty guided residual estimation block ( $U-REB$ ) that is used to refine the residual estimation process. Specifically, features ( $F_i$ ) from the main branch are forwarded through a conv block ( $CB_i$ ) which estimates the residual map  $R_i$ . In order to improve the efficacy of the residual learning mechanism, we propose an uncertainty guided confidence estimation block ( $CEB$ ) that guides the refinement process. The task of conv blocks  $CB_i$  is to capture residual errors that can be incorporated into the coarser density maps to produce high quality density maps in the end. For this purpose, these conv blocks employ feature maps from shallower conv layers  $C_i$  from the main branch.




Fig. 8. Uncertainty-guided residual estimation block ( $U-REB$ ).

Since the conv layers in the main branch are primarily trained for estimating the coarsest density map, their features have high responses in regions where crowd is present, and hence, they may not necessarily produce effective residuals. In order to overcome this issue, we propose to gate the residuals that are not effective. This is achieved by using uncertainty estimation. Inspired by uncertainty estimation in CNNs [39, 40, 41, 74], we aim to model pixel-wise aleatoric uncertainty of the residuals estimated by  $CB_i$ . That is we, predict the pixel-wise confidence (inverse of the uncertainties) of the residuals which are then used to gate the residuals before being passed on to the subsequent outputs. This ensures that only highly confident residuals get propagated to the output. Note that  $CB_5, CB_4, CB_3$  are defined as follows:

$$\begin{aligned} CB_5: & \{\text{conv}_{512,32,1}\text{-relu-conv}_{32,32,3}\text{-relu-conv}_{32,1,3}\}^2. \\ CB_4: & \{\text{conv}_{512,32,1}\text{-relu-conv}_{32,32,3}\text{-relu-conv}_{32,1,3}\}^2. \\ CB_3: & \{\text{conv}_{256,32,1}\text{-relu-conv}_{32,32,3}\text{-relu-conv}_{32,1,3}\}^2. \end{aligned}$$

In terms of the overall architecture, we introduce a set of  $U-REBs$  as shown in Fig. 7. Each residual branch consists

of one such block. Fig. 8 illustrates the mechanism of the proposed  $U-REB$ .  $UREB_i$  takes the feature map  $F_i$  from the main branch and forwards them through a conv block  $CB_i$  to produce residual map ( $R_i$ ). This residual map is then concatenated with dimensionality reduced features<sup>3</sup> from the main branch and forwarded through confidence estimation block ( $CEB_i$ ). This block is defined by  $\{\text{conv}_{33,32,1}\text{-relu-conv}_{32,16,3}\text{-relu-conv}_{16,16,3}\text{-relu-conv}_{16,1,1}\}$  and it produces a confidence map  $CM_i$  which is then multiplied element-wise with the input to form the refined residual map:

$$\hat{R}_i = R_i \odot CM_i, \quad (4)$$

where  $\odot$  denotes element-wise multiplication.

In order to learn these confidence maps, the loss function  $L_f$  used to train the network is defined as follows,

$$L_f = L_d - \lambda_c L_c, \quad (5)$$

where,  $\lambda_c$  is a regularization constant,  $L_d$  is the pixel-wise regression loss to minimize the density map prediction error and is defined as:

$$L_d = \sum_{i \in \{3,4,5,6\}} \|(CM_i \odot Y_i) - (CM_i \odot \hat{Y}_i)\|_2, \quad (6)$$

where,  $\hat{Y}_i$  is the predicted density map,  $i$  indicates the index of the conv layer from which the predicted density map is taken,  $Y_i$  is the corresponding target.

$L_c$  is the confidence guiding loss, defined as,

$$L_c = \sum_{i \in \{3,4,5,6\}} \sum_{j=1}^H \sum_{k=1}^W \log(CM_i^{j,k}), \quad (7)$$

where,  $W \times H$  is the dimension of the confidence map  $CM_i$ . As it can be seen from Eq. (5), the loss  $L_f$  has two parts  $L_d$  and  $L_c$ . The first term minimizes the Euclidean distance between the prediction and target features, whereas  $L_c$  maximizes the confidence scores  $CM_i$  by making them closer to 1.

Note that optimizing  $L_c$  (Equation 7) forces the values of confidence map close to 1. This is true when the loss function is Equation 7 alone. However, our final loss function (Equation 5) contains both  $L_d$  and  $L_c$ . When we minimize Equation 5, if  $CM_i$  is driven to “one” due to loss  $L_c$ , the first loss in Equation 6 ( $L_d$ ) increases. Due to this, the network tries to reduce  $L_d$  by forcing the network to predict lower values of  $CM_i$  in addition to learning to predict accurate density maps. Hence, due to the presence of  $CM_i$  in both the parts of the loss function  $L_f$  (Equation 5), the network “does not” learn to trivially predict ones for  $CM_i$ . This is analogous to the use of aleatoric uncertainty prediction for improving the network performance described in [39].

Fig. 9 illustrates the output density maps ( $\hat{Y}_6, \hat{Y}_5, \hat{Y}_4, \hat{Y}_3$ ) generated by the proposed network for a sample crowd image. It can be observed that the density maps progressively improve in terms of fine details and the count value.

Fig. 10 illustrates the residual maps generated with and without the confidence gating. It can be clearly observed that the use of confidence scores aids in better feature learning.

3. We use a dimensionality reduction (DR) block which consists of  $1 \times 1$  conv layer to reduce the number of channels to 32.






Fig. 9. Density maps estimated by different layers of the proposed network. (a)  $\hat{Y}_6$  (b)  $\hat{Y}_5$  (c)  $\hat{Y}_4$  (d)  $\hat{Y}_3$  (e)  $Y$  (ground-truth). It can be observed that the output of the deepest layer ( $\hat{Y}_6$ ) looks very coarse, and it is refined in a progressive manner using the residual learned by  $U-REB_5$ ,  $U-REB_4$ ,  $U-REB_3$  to obtain the  $\hat{Y}_5, \hat{Y}_4, \hat{Y}_3$  respectively. Note that fine details and the total count in the density maps improve as we move from  $\hat{Y}_6$  to  $\hat{Y}_3$ .

#### 4.4 Class-conditioned Uncertainty guided residual learning (U-REBC)

In order to leverage additional information provided in the proposed JHU-CROWD++ dataset, we propose to condition the residual estimation based on the image-level labels (specifically, weather labels). That is, we augment the  $U-REB$  module with additional class conditioning ( $CC$ ) block as shown in Fig. 11. This block consists of a set of 2 conv relu-layers ( $\{\text{conv}_{32,32,3}\text{-relu-conv}_{32,4,3}\}^2$ ) followed by an average-pool layer and a soft-max layer. Note that the output of this block is 4 classes corresponding to *rain, fog, haze* and *normal*. The  $CC$  block is trained via cross-entropy error using labels available in the dataset. To condition the uncertainty estimation on the classes, the feature maps ( $F_i^c$ ) prior to the average-pool layer in  $CC$  are concatenated with the residual map  $R_i$  and the dimensionality reduced features from the main branch. These concatenated feature maps are then forwarded through the confidence estimation block  $CEB_i$  to predict the confidences as described earlier in Section 4.3.

For training the network, we modify the loss function in Eq. 5 as follows:

$$L_f = L_d - \lambda_c L_c + \lambda_w L_w, \quad (8)$$

where,  $L_w$  is the cross-entropy loss for the weather classification and  $\lambda_w$  is a weighting factor and we set it to 0.01. Note that the distribution of weather images is imbalanced. Hence, we weight the each class proportionately based on the number of samples in each category.

#### 4.5 Training and inference details

The training dataset is obtained by cropping patches from multiple random locations in each training image. The cropped patch-size is  $256 \times 256$ . For JHU-CROWD++, we use




Fig. 10. Residual maps. *Top row*: Without confidence gating. *Bottom row*: With confidence gating. (a)  $R_5$  (b)  $R_4$  (c)  $R_3$ . Red indicates negative values and cyan indicates positive values. The use of confidence gating improves the residual maps significantly, especially for the shallower layers.

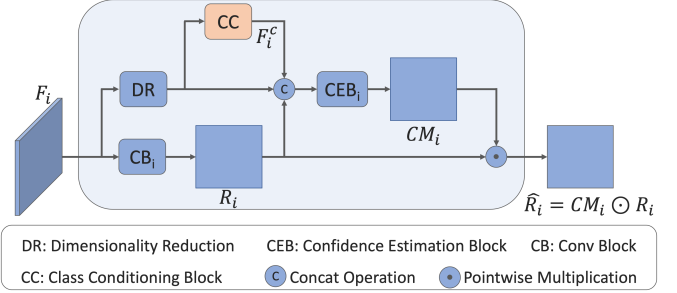


Fig. 11. Class-conditioned uncertainty-guided residual estimation block ( $U-REBC$ ).

the validation set for model selection and hyper-parameter tuning. For other datasets, we use 10% of the training images as validation set. We use the Adam optimizer to train the network. We use a learning rate of 0.00001 and a momentum of 0.9 with a batch-size of 24. Before cropping, we resize all the images such that the minimum dimension is 512 and maximum dimension is 2048 while maintaining the aspect ratio.

For inference, the density map  $\hat{Y}_3$  is considered as the final output. The count performance is measured using the standard error metrics: mean absolute error ( $MAE$ ) and mean squared error ( $MSE$ ). These metrics are defined as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^N |C(Y^i) - C(\hat{Y}_3^i)|, \quad (9)$$

$$MSE = \sqrt{\frac{1}{N} \sum_{i=1}^N |C(Y^i) - C(\hat{Y}_3^i)|^2},$$

where  $N$  is the number of test samples,  $Y^i$  is the ground-truth count and  $\hat{Y}_3^i$  is the estimated count corresponding to the  $i^{th}$  sample, and  $C(Y)$  denotes the sum of all the values in  $Y$ .

## 5 ABLATION STUDY

In this section, we discuss the results of different ablation studies conducted to analyze (i) the effect of different components in the proposed network, (ii) generalizability to other network architectures, (iii) the effect of different branches in the proposed architecture for residual estimation, and (iv) effect of reduced data during training. Due to the presence of various complexities such as high density crowds, large variations in scales, occlusion, etc. we choose to perform the ablation study on JHU-CROWD++ val set.

### 5.1 Residual Learning, Uncertainty and Class-conditioning

The ablation study consisted of evaluating the following configurations of the proposed method:

- (i) *Base network*: VGG16 network with an additional conv block ( $CB_6$ ) at the end.
- (ii) *Base network + R*: the base network with residual learning.
- (iii) *Base network + R + U-REB* ( $\lambda_c = 0$ ): the base network with residual learning guided by the confidence estimation blocks as discussed in Section 4.3. In this configuration, we aim to measure the performance due to the addition of the confidence estimation blocks without the uncertainty estimation mechanism by setting  $\lambda_c$  is set to 0.
- (iv) *Base network + R + U-REB* ( $\lambda_c = 1$ ): the base network with residual learning guided by the confidence estimation blocks as discussed in Section 4.3.
- (v) *Base network + R + U-REBC* ( $\lambda_w = 0$ ): the base network with residual learning guided by the class-conditioned confidence estimation blocks as discussed in Section 4.4. In this configuration, we aim to measure the performance due to the addition of conv block in the class conditioning  $CC$  module without image-level training by setting  $\lambda_w$  is set to 0.
- (vi) *Base network + R + U-REBC* ( $\lambda_w = 0.01$ ): the base network with residual learning guided by the class-conditioned confidence estimation blocks as discussed in Section 4.4.

The results of these experiments are shown in Table 5. It can be seen that there are considerable improvements in the performance due to the inclusion of residual learning into the network. The use of confidence-based weighting of the residuals results in further improvements, thus highlighting its significance in improving the efficacy of uncertainty-based residual learning.

TABLE 5  
Results of ablation study using “VGG16” base network on the JHU-CROWD++ dataset (val-set).

Method	MAE	MSE
Base network	81.1	300.5
Base network + R	77.5	290.6
Base network + R + UREB ( $\lambda_c = 0$ )	77.1	290.5
Base network + R + UREB ( $\lambda_c = 1$ )	74.1	275.5
Base network + R + UREB-C ( $\lambda_w = 0$ )	74.6	274.1
Base network + R + UREB-C ( $\lambda_w = 0.01$ )	67.9	262.1

Further, we perform additional ablation studies to demonstrate the effectiveness of conditioning the estimation based on the class labels. Specifically, we conduct the following three experiments (see Table 6):

- (i) *W/o conditioning*: This is the baseline experiment that does not involve any use of the weather labels,
- (ii) *With multi-task learning*: Here, we use the weather labels to train the  $CC$  block. However, the feature maps from the  $CC$  block “are not” concatenated with the residual map  $R_i$ . This experiment demonstrates the benefit of using the class labels in a naive manner.
- (iii) *With multi-task learning & conditioning*: This corresponds to the proposed method of conditioning the residual estimation on the class labels (as described in 4.4). It can be observed that conditioning results in better improvements

as compared to the naive use class labels. This leads to significant improvements in the overall error.

Fig. 12 visualizes the density map estimation results from the above experiments. It can be observed that class-conditioning improves the quality of density maps and prediction error.

TABLE 6  
Ablation results: “Class-conditioning” for weather-conditions study on the JHU-CROWD++ weather dataset (val-set).

Method	MAE	MSE
W/o conditioning	78.4	170.5
With multi-task learning	74.4	140.1
With multi-task learning & conditioning	63.6	116.6

### 5.2 Res101 backbone network

In order to demonstrate that the proposed uncertainty-guided residual learning mechanism is not network-dependent, we evaluate the method using a different base network: Res101 [43]. To employ the Res101 architecture as the base network: we (i) add the uncertainty-based residual estimation blocks  $U-REB_3$ ,  $U-REB_4$  and  $U-REB_5$  after layers 2, 3 and 4 in Res101 respectively, (ii) add conv6 layer after layer 5 with the input number of channels changed appropriately to match the number of output channels of layer 4 in Res101, and (iii) change the number of input channels in the conv blocks in  $U-REB_i$ 's to match the number of output channels of the respective blocks in the main branch of Res101. Furthermore, since the  $U-REB_3$  is added to a shallower layer, we weight the loss function corresponding to  $\hat{Y}_6$ . That is, we modify Eq. 6 as follows:

$$L_d = \sum_{i \in \{3,4,5,6\}} \lambda_i \|(CM_i \odot Y_i) - (CM_i \odot \hat{Y}_i)\|_2. \quad (10)$$

In the above equation, we set  $\lambda_3 = 0.1$  and  $\lambda_4 = \lambda_5 = \lambda_6 = 1$ .

Table 7 shows the results of the proposed network using Res101 backbone network. We make similar observations as in the case of VGG16 base network. That is, the use of residual learning results in better performance compared to the base network. Further, incorporating uncertainty-guided residual estimation and class conditioning results in further improvements. From this experiment, we can observe that the proposed method can generalize to other types of network architectures.

TABLE 7  
Results of ablation study using “Res101” base network on the JHU-CROWD++ dataset (val-set).

Method	MAE	MSE
Base network	72.1	280.5
Base network + R	68.5	270.9
Base network + R + UREB ( $\lambda_c = 0$ )	68.2	271.2
Base network + R + UREB ( $\lambda_c = 1$ )	62.5	258.1
Base network + R + UREB-C ( $\lambda_w = 0$ )	63.1	259.9
Base network + R + UREB-C ( $\lambda_w = 0.01$ )	57.6	244.4

### 5.3 Number of branches

Since the proposed method involves residual learning at multiple scales of the base network, we conduct a set of experiments to understand the effectiveness of using multiple



Fig. 12. Ablation results: “**Class-conditioning**” for weather-conditions study on the JHU-CROWD++ weather dataset (val-set). (a) Input (b) Without class-conditioning (c) With class-conditioning (d) Ground truth. Class-conditioning improves the quality of density maps and prediction error. Boxes marked in red highlight the regions where class-conditioning improves the quality.

TABLE 8

Results of ablation on the “**branches**” used for density estimation on the JHU-CROWD++ dataset (val-set).

Base network	VGG16		Res101	
Branch	MAE	MSE	MAE	MSE
$\hat{Y}_6$	81.1	300.5	72.1	280.5
$\hat{Y}_6 + \hat{Y}_5$	72.1	280.1	60.6	251.4
$\hat{Y}_6 + \hat{Y}_5 + \hat{Y}_4$	70.7	270.5	58.8	249.4
$\hat{Y}_6 + \hat{Y}_5 + \hat{Y}_4 + \hat{Y}_3$	67.9	262.1	57.6	244.4

scales. We evaluate for two backbone architectures: VGG16 and Res101. Specifically, we conduct experiments where we sequentially add the residual estimation blocks at conv5, conv4 and conv3 for VGG16 and at layer4, layer 3 and layer 2 for Res101. Table 8 shows the results of these experiments. It can be observed for both architectures that as we add more residual estimation blocks at different layers, the errors drops by considerable margins.

### 5.4 Effect of reduced training data

In order to demonstrate the need for a large dataset with high diversity, we conducted a set of experiments by varying the number of data samples. More specifically, we trained CG-DRCN-CC-Res101 with different percentages (5%, 25%, 50%, 75% and 100%) of the JHU-CROWD++ train dataset. Note that the evaluation is performed on the val set. The results of this experiment are provided in Table 9. As it can be observed, the error reduces as we increase the number of training data samples.

TABLE 9

Results with varying training dataset size. MAE reduces with increase in the number of samples used for training.

% of training data	5%	25%	50%	75%	100%
MAE	124.3	105.4	93.1	78.9	57.6

TABLE 10  
Results on JHU-CROWD++ dataset (“Val Set”). **RED** indicates best error and **BLUE** indicates second-best error.

Category	Method	Model	Low		Medium		High		Weather		Overall	
			MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
MCNN [5] (CVPR 16)	Custom	Custom	90.6	202.9	125.3	259.5	494.9	856.0	241.1	532.2	160.6	377.7
CMTL [24] (AVSS 17)	Custom	Custom	50.2	129.2	88.1	170.7	583.1	986.5	165.0	312.9	138.1	379.5
CSR-Net [75](CVPR 18)	VGG16	VGG16	22.2	40.0	49.0	99.5	<b>302.5</b>	<b>669.5</b>	83.0	168.7	<b>72.2</b>	<b>249.9</b>
SA-Net [26](ECCV 18)	VGG16	VGG16	13.6	26.8	50.4	78.0	397.8	749.2	<b>72.2</b>	<b>126.7</b>	82.1	272.6
CACC [66] (CVPR 19)	VGG16	VGG16	34.2	69.5	65.6	115.3	336.4	<b>619.7</b>	101.8	179.3	89.5	<b>239.3</b>
DSSI-Net [76] (ICCV 19)	VGG16	VGG16	50.3	85.9	82.4	164.5	436.6	814.0	155.7	314.8	116.6	317.4
MBTTBF [63] (ICCV 19)	VGG16	VGG16	23.3	48.5	53.2	119.9	<b>294.5</b>	674.5	88.2	200.8	73.8	256.8
LSC-CNN [69] (PAMI 20)	VGG16	VGG16	<b>6.8</b>	<b>10.1</b>	<b>39.2</b>	<b>64.1</b>	504.7	860.0	77.6	187.2	87.3	309.0
CG-DRCN-CC-VGG16 (ours)	VGG16	VGG16	<b>17.1</b>	<b>44.7</b>	<b>40.8</b>	<b>71.2</b>	317.4	719.8	<b>63.5</b>	<b>116.6</b>	<b>67.9</b>	262.1
SFCN [44] (CVPR 19)	ResNet-101	ResNet-101	11.8	19.8	<b>39.3</b>	<b>73.4</b>	297.3	679.4	<b>52.3</b>	<b>93.6</b>	62.9	247.5
BCC [70](ICCV 19)	VGG19	VGG19	<b>6.9</b>	<b>10.3</b>	39.7	85.2	<b>279.8</b>	<b>620.4</b>	58.9	124.7	<b>59.3</b>	<b>229.2</b>
CG-DRCN-CC-Res101 (ours)	ResNet-101	ResNet-101	<b>11.7</b>	<b>24.8</b>	<b>35.2</b>	<b>57.5</b>	<b>273.9</b>	<b>676.8</b>	<b>54.0</b>	<b>106.8</b>	<b>57.6</b>	<b>244.4</b>

TABLE 11  
Results on JHU-CROWD++ dataset (“Test Set”). **RED** indicates best error and **BLUE** indicates second-best error.

Category	Method	Model	Low		Medium		High		Weather		Overall	
			MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
MCNN [5] (CVPR 16)	Custom	Custom	97.1	192.3	121.4	191.3	618.6	1,166.7	330.6	852.1	188.9	483.4
CMTL [24] (AVSS 17)	Custom	Custom	58.5	136.4	81.7	144.7	635.3	1,225.3	261.6	816.0	157.8	490.4
CSR-Net [75] (CVPR 18)	VGG16	VGG16	27.1	64.9	43.9	71.2	<b>356.2</b>	<b>784.4</b>	141.4	640.1	85.9	<b>309.2</b>
SA-Net [26] (ECCV 18)	VGG16	VGG16	<b>17.3</b>	<b>37.9</b>	46.8	69.1	397.9	817.7	154.2	685.7	91.1	320.4
CACC [66] (CVPR 19)	VGG16	VGG16	37.6	78.8	56.4	86.2	384.2	789.0	155.4	<b>617.0</b>	100.1	314.0
DSSI-Net [76] (ICCV 19)	VGG16	VGG16	53.6	112.8	70.3	108.6	525.5	1,047.4	229.1	760.3	133.5	416.5
MBTTBF [63] (ICCV 19)	VGG16	VGG16	19.2	58.8	41.6	66.0	<b>352.2</b>	<b>760.4</b>	<b>138.7</b>	<b>631.6</b>	<b>81.8</b>	<b>299.1</b>
LSCCNN [69] (PAMI 20)	VGG16	VGG16	<b>10.6</b>	<b>31.8</b>	<b>34.9</b>	<b>55.6</b>	601.9	1,172.2	178.0	744.3	112.7	454.4
CG-DRCN-CC-VGG16 (ours)	VGG16	VGG16	19.5	58.7	<b>38.4</b>	<b>62.7</b>	367.3	837.5	<b>138.6</b>	654.0	<b>82.3</b>	328.0
SFCN [44] (CVPR 19)	ResNet-101	ResNet-101	16.5	55.7	38.1	59.8	<b>341.8</b>	<b>758.8</b>	<b>122.8</b>	<b>606.3</b>	77.5	<b>297.6</b>
BCC [70] (ICCV 19)	VGG19	VGG19	<b>10.1</b>	<b>32.7</b>	<b>34.2</b>	<b>54.5</b>	352.0	768.7	140.1	675.7	<b>75.0</b>	299.9
CG-DRCN-CC-Res101 (ours)	ResNet-101	ResNet-101	<b>14.0</b>	<b>42.8</b>	<b>35.0</b>	<b>53.7</b>	<b>314.7</b>	<b>712.3</b>	<b>120.0</b>	<b>580.8</b>	<b>71.0</b>	<b>278.6</b>

## 6 BENCHMARKING ON JHU-CROWD++ DATASET

In this section, we present results of benchmarking of several recent algorithms including the proposed method on the JHU-CROWD++ dataset. Specifically, we evaluate the following recent works: multi-column network (MCNN) [5], cascaded multi-task learning for crowd counting (CMTL) [24], CSR-Net [75], SA-Net [26], context-aware crowd counting (CACC) [66], spatial fully convolutional network (SFCN) [44], deep structured scale integration network (DSSI-Net) [76], multi-level bottom-top and top-bottom feature fusion [63], Bayesian loss for counting (BCC) [70] and locate-size-count-CNN (LSC-CNN) [69]. In addition, we also evaluate the proposed class-conditioned uncertainty-guided residual estimation method (CG-DRCN-CC) and demonstrate its effectiveness over the other methods.

All the networks are trained using the training set. We use the validation set for model selection. Table 10 and 11 show the results of the above experiments for various sub-categories of images. Based on these results we make the following observations:

- (i) The proposed method (CG-DRCN-CC) with Res101 base network achieves lowest overall MAE while obtaining comparable performance for validation set.
- (ii) The proposed method (CG-DRCN-CC) with Res101 base network achieves lowest overall MAE/MSE as compared to

all the other methods on the test set. In addition, it achieves best errors for the “high-density” and “weather” categories while obtaining comparable performance for the rest of the categories.

- (iii) The proposed method (CG-DRCN-CC) with VGG16 base network achieves comparable performance in all categories with respect to the other methods.

(iv) BCC [70] and LSC-CNN [69] achieve lowest errors in the “low-density” categories. These methods do not follow the traditional density-estimation based approach for supervising the networks. Instead they incorporate size information during the training through strategies like Bayesian-loss and bounding box-based supervision.

- (v) Res101-based methods tend to perform better compared to VGG16-based approaches in terms of overall error.

## 7 EVALUATION ON OTHER DATASETS

In this section, we evaluate the proposed method on other datasets like ShanghaiTech [5] and UCF-QNRF [21]. In addition, we compare the proposed method with several recent methods and demonstrate that our method is able to achieve comparable performance with respect to the state-of-the-art methods.



Fig. 13. Results of the proposed dataset on sample images from the JHU-CROWD++ dataset. (a) Input image (b) Ground-truth density map (c) Estimated density map.

### 7.1 ShanghaiTech Dataset [5]

This dataset contains 1,198 annotated images with a total of 330,165 people. This dataset consists of two parts: Part A with 482 images and Part B with 716 images. Both parts are further divided into training and test datasets with training set of Part A containing 300 images and that of Part B containing 400 images. Rest of the images are used as test set.

The proposed network is trained on the train splits using the same strategy as discussed in Section 4.5. Table 12 shows

the results of the proposed method on ShanghaiTech as compared with several recent approaches: CP-CNN[6], IG-CNN [30], D-ConvNet [33], Liu *et al.*[58], CSR-Net [75], ic-CNN [25], SA-Net[26], ACSCP [32] and Jian *et al.*[64], CA-Net [76], BCC [70], DSSI-Net [76], MBTTBF [63] and LSC-CNN [69]. It can be observed that the proposed method outperforms all existing methods on Part A of the dataset, while achieving comparable performance on Part B.

TABLE 12  
Results on “ShanghaiTech” dataset [5].

Method	Part-A		Part-B	
	MAE	MSE	MAE	MSE
CP-CNN [6]	73.6	106.4	20.1	30.1
IG-CNN [30]	72.5	118.2	13.6	21.1
Liu <i>et al.</i> [58]	73.6	112.0	13.7	21.4
D-ConvNet [33]	73.5	112.3	18.7	26.0
CSRNet [75]	68.2	115.0	10.6	16.0
ic-CNN [25]	69.8	117.3	10.7	16.0
SA-Net [26]	67.0	104.5	8.4	13.6
ACSCP [32]	75.7	102.7	17.2	27.4
Jian <i>et al.</i> [64]	64.2	109.1	8.2	12.8
CA-Net [66]	61.3	100.0	7.8	12.2
BCC [70]	62.8	117.0	8.1	12.7
DSSI-Net [76]	60.6	96.0	<b>6.8</b>	<b>10.3</b>
MBTTBF [63]	<b>60.2</b>	<b>94.1</b>	8.0	15.5
LSC-CNN [69]	66.5	101.8	7.7	12.7
CG-DRCN-VGG16 (ours)	64.0	98.4	8.5	14.4
CG-DRCN-Res101 (ours)	<b>60.2</b>	<b>94.0</b>	<b>7.5</b>	<b>12.1</b>

## 7.2 UCF-QNRF Dataset[21]

UCF-QNRF is a large crowd counting dataset with 1535 high-resolution images and 1.25 million head annotations. There are 1201 training images and 334 test images. It contains extremely congested scenes where the maximum count of an image can reach up to 12,865.

Table 13 shows results on the UCF-QNRF dataset. The proposed method is compared with the following recent methods: Idrees *et al.*[3], MCNN [5], CMTL [24], Switching-CNN [7], Idrees *et al.*[21], Jian *et al.*[64], CA-Net [76], BCC [70], DSSI-Net [76], MBTTBF [63] and LSC-CNN [69]. It can be observed that the proposed method achieves comparable performance with respect to the recent state-of-the-art methods.

TABLE 13  
Results on “UCF-QNRF” dataset [21].

Method	MAE	MSE
Idrees <i>et al.</i> [3]	315.0	508.0
Zhang <i>et al.</i> [4]	277.0	426.0
CMTL <i>et al.</i> [24]	252.0	514.0
Switching-CNN [7]	228.0	445.0
Idrees <i>et al.</i> [21]	132.0	191.0
Jian <i>et al.</i> [64]	113.0	188.0
CA-Net [66]	107.0	183.0
DSSI-Net [76]	99.1	159.2
MBTTBF [63]	97.5	165.2
BCC [70]	<b>88.7</b>	<b>154.8</b>
LSC-CNN [69]	120.5	218.2
CG-DRCN-VGG16 (ours)	112.2	176.3
CG-DRCN-Res101 (ours)	<b>95.5</b>	<b>164.3</b>

## 8 CONCLUSIONS

In this work, we introduce a new large scale unconstrained crowd counting dataset (JHU-CROWD++) consisting of 4,372 images with 1.51 million annotations. The new dataset is collected under a variety of conditions and includes images with weather-based degradations and other distractors. Additionally, the dataset provides a rich set of

annotations such as head locations, blur-level, occlusion-level, approximate bounding boxes and other image-level labels. In addition, we benchmark several recent state-of-the-art crowd counting techniques on the new dataset.

Furthermore, we present a novel crowd counting network that employs residual learning mechanism in a progressive fashion to estimate coarse to fine density maps. The efficacy of residual learning is further improved by introducing an uncertainty-based confidence weighting mechanism that is designed to enable the network to propagate only high-confident residuals to the output. Additionally, we incorporate class-conditioning mechanism to leverage the image-level labels in the new dataset for improving the performance in adverse weather conditions. The proposed method is evaluated on recent datasets and we demonstrate that it achieves comparable performance with respect to the state-of-the-art methods.

## ACKNOWLEDGMENTS

This work was supported by the NSF grant 1910141.

We would like to express our deep gratitude to everyone who contributed to the creation of this dataset including the members of the JHU-VIU lab and the numerous Amazon Mturk workers. We would like to specially thank Kumar Siddhanth, Poojan Oza, A. N. Sindagi, Jayadev S, Supriya S, Shruthi S and S. Sreevali for providing assistance in annotation and verification efforts.

## REFERENCES

- [1] T. Li, H. Chang, M. Wang, B. Ni, R. Hong, and S. Yan, “Crowded scene analysis: A survey,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 25, no. 3, pp. 367–386, 2015. **1, 3**
- [2] B. Zhan, D. N. Monekosso, P. Remagnino, S. A. Velastin, and L.-Q. Xu, “Crowd analysis: a survey,” *Machine Vision and Applications*, vol. 19, no. 5-6, pp. 345–357, 2008. **1**
- [3] H. Idrees, I. Saleemi, C. Seibert, and M. Shah, “Multi-source multi-scale counting in extremely dense crowd images,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2547–2554. **1, 3, 4, 5, 14**
- [4] C. Zhang, H. Li, X. Wang, and X. Yang, “Cross-scene crowd counting via deep convolutional neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 833–841. **1, 2, 3, 4, 5, 14**
- [5] Y. Zhang, D. Zhou, S. Chen, S. Gao, and Y. Ma, “Single-image crowd counting via multi-column convolutional neural network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 589–597. **1, 2, 3, 4, 5, 12, 13, 14**
- [6] V. A. Sindagi and V. M. Patel, “Generating high-quality crowd density maps using contextual pyramid cnns,” in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. **1, 2, 3, 4, 7, 13, 14**
- [7] D. B. Sam, S. Surya, and R. V. Babu, “Switching convolutional neural network for crowd counting,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. **1, 2, 3, 7, 14**

- [8] A. B. Chan, Z.-S. J. Liang, and N. Vasconcelos, "Privacy preserving crowd monitoring: Counting people without people models or tracking," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. IEEE, 2008, pp. 1–7. [1](#), [3](#), [5](#)
- [9] M. Rodriguez, I. Laptev, J. Sivic, and J.-Y. Audibert, "Density-aware person detection and tracking in crowds," in *2011 International Conference on Computer Vision*. IEEE, 2011, pp. 2423–2430. [1](#)
- [10] F. Zhu, X. Wang, and N. Yu, "Crowd tracking with dynamic evolution of group structures," in *European Conference on Computer Vision*. Springer, 2014, pp. 139–154. [1](#)
- [11] W. Li, V. Mahadevan, and N. Vasconcelos, "Anomaly detection and localization in crowded scenes," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 1, pp. 18–32, 2014. [1](#)
- [12] V. Mahadevan, W. Li, V. Bhalodia, and N. Vasconcelos, "Anomaly detection in crowded scenes." in *CVPR*, vol. 249, 2010, p. 250. [1](#)
- [13] M. Marsden, K. McGuinness, S. Little, C. E. Keogh, and N. E. O'Connor, "People, penguins and petri dishes: Adapting object counting models to new visual domains and object types without forgetting," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [1](#)
- [14] V. A. Sindagi and V. M. Patel, "Dafe-fd: Density aware feature enrichment for face detection," in *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2019, pp. 2185–2195. [1](#)
- [15] D. B. Sam, N. N. Sajjan, H. Maurya, and R. V. Babu, "Almost unsupervised learning for dense crowd counting," in *Thirty-Third AAAI Conference on Artificial Intelligence*, 2019. [1](#)
- [16] A. B. Chan and N. Vasconcelos, "Bayesian poisson regression for crowd counting," in *2009 IEEE 12th International Conference on Computer Vision*. IEEE, 2009, pp. 545–551. [1](#)
- [17] K. Chen, S. Gong, T. Xiang, and C. Change Loy, "Cumulative attribute space for age and crowd density estimation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 2467–2474. [1](#)
- [18] H. Lu, Z. Cao, Y. Xiao, B. Zhuang, and C. Shen, "Tasselnet: Counting maize tassels in the wild via local counts regression network," *Plant Methods*, vol. 13, no. 1, p. 79, 2017. [1](#)
- [19] V. Lempitsky and A. Zisserman, "Learning to count objects in images," in *Advances in Neural Information Processing Systems*, 2010, pp. 1324–1332. [1](#), [3](#)
- [20] G. French, M. Fisher, M. Mackiewicz, and C. Needle, "Convolutional neural networks for counting fish in fisheries surveillance video," in *British Machine Vision Conference Workshop*. BMVA Press, 2015. [1](#)
- [21] H. Idrees, M. Tayyab, K. Athrey, D. Zhang, S. Al-Maadeed, N. Rajpoot, and M. Shah, "Composition loss for counting, density map estimation and localization in dense crowds," in *European Conference on Computer Vision*. Springer, 2018, pp. 544–559. [1](#), [3](#), [4](#), [5](#), [12](#), [14](#)
- [22] E. Walach and L. Wolf, "Learning to count with cnn boosting," in *European Conference on Computer Vision*. Springer, 2016, pp. 660–676. [1](#), [2](#), [3](#)
- [23] D. Onoro-Rubio and R. J. López-Sastre, "Towards perspective-free object counting with deep learning," in *European Conference on Computer Vision*. Springer, 2016, pp. 615–629. [1](#), [2](#), [3](#)
- [24] V. A. Sindagi and V. M. Patel, "Cnn-based cascaded multi-task learning of high-level prior and density estimation for crowd counting," in *Advanced Video and Signal Based Surveillance (AVSS), 2017 IEEE International Conference on*. IEEE, 2017. [1](#), [2](#), [12](#), [14](#)
- [25] V. Ranjan, H. Le, and M. Hoai, "Iterative crowd counting," in *European Conference on Computer Vision*. Springer, 2018, pp. 278–293. [1](#), [2](#), [4](#), [13](#), [14](#)
- [26] X. Cao, Z. Wang, Y. Zhao, and F. Su, "Scale aggregation network for accurate and efficient crowd counting," in *European Conference on Computer Vision*. Springer, 2018, pp. 757–773. [1](#), [2](#), [3](#), [4](#), [7](#), [12](#), [13](#), [14](#)
- [27] D. B. Sam and R. V. Babu, "Top-down feedback for crowd counting convolutional neural network," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. [1](#)
- [28] C. Zhang, K. Kang, H. Li, X. Wang, R. Xie, and X. Yang, "Data-driven crowd understanding: A baseline for a large-scale crowd dataset," *IEEE Transactions on Multimedia*, vol. 18, no. 6, pp. 1048–1061, 2016. [1](#), [4](#)
- [29] Q. Wang, J. Gao, W. Lin, and X. Li, "Nwpu-crowd: A large-scale benchmark for crowd counting," *arXiv preprint arXiv:2001.03360*, 2020. [1](#), [3](#), [5](#)
- [30] D. Babu Sam, N. N. Sajjan, R. Venkatesh Babu, and M. Srinivasan, "Divide and grow: Capturing huge diversity in crowd images with incrementally growing cnn," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3618–3626. [2](#), [4](#), [13](#), [14](#)
- [31] V. A. Sindagi and V. M. Patel, "Ha-ccn: Hierarchical attention-based crowd counting network," *arXiv preprint arXiv:1907.10255*, 2019. [2](#), [4](#)
- [32] Z. Shen, Y. Xu, B. Ni, M. Wang, J. Hu, and X. Yang, "Crowd counting via adversarial cross-scale consistency pursuit," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#), [4](#), [13](#), [14](#)
- [33] Z. Shi, L. Zhang, Y. Liu, X. Cao, Y. Ye, M.-M. Cheng, and G. Zheng, "Crowd counting with deep negative correlation learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#), [4](#), [13](#), [14](#)
- [34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations*, 2015. [2](#)
- [35] Y. Tai, J. Yang, and X. Liu, "Image super-resolution via deep recursive residual network," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, no. 2, 2017, p. 5. [2](#), [7](#)
- [36] J. Kim, J. Kwon Lee, and K. Mu Lee, "Accurate image super-resolution using very deep convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1646–1654. [2](#), [7](#)
- [37] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-

- resolution,” in *The IEEE conference on computer vision and pattern recognition (CVPR) workshops*, vol. 1, no. 2, 2017, p. 4. [2, 7](#)
- [38] W. Ke, J. Chen, J. Jiao, G. Zhao, and Q. Ye, “Srn: Side-output residual network for object symmetry detection in the wild,” *arXiv preprint arXiv:1703.02243*, 2017. [2](#)
- [39] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” in *Advances in neural information processing systems*, 2017, pp. 5574–5584. [3, 8](#)
- [40] L. Zhu and N. Laptev, “Deep and confident prediction for time series at uber,” in *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 2017, pp. 103–110. [3, 8](#)
- [41] T. DeVries and G. W. Taylor, “Learning confidence for out-of-distribution detection in neural networks,” *arXiv preprint arXiv:1802.04865*, 2018. [3, 8](#)
- [42] V. A. Sindagi, R. Yasarla, and V. M. Patel, “Pushing the frontiers of unconstrained crowd counting: New dataset and benchmark method,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1221–1231. [3, 5](#)
- [43] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778. [3, 10](#)
- [44] Q. Wang, J. Gao, W. Lin, and Y. Yuan, “Learning from synthetic data for crowd counting in the wild,” *arXiv preprint arXiv:1903.03303*, 2019. [3, 12](#)
- [45] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2223–2232. [3](#)
- [46] C. C. Loy, K. Chen, S. Gong, and T. Xiang, “Crowd counting and profiling: Methodology and evaluation,” in *Modeling, Simulation and Visual Analysis of Crowds*. Springer, 2013, pp. 347–382. [3](#)
- [47] M. Li, Z. Zhang, K. Huang, and T. Tan, “Estimating the number of people in crowded scenes by mid based foreground segmentation and head-shoulder detection,” in *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*. IEEE, 2008, pp. 1–4. [3](#)
- [48] D. Ryan, S. Denman, C. Fookes, and S. Sridharan, “Crowd counting using multiple local features,” in *Digital Image Computing: Techniques and Applications, 2009. DICTA'09*. IEEE, 2009, pp. 81–88. [3](#)
- [49] K. Chen, C. C. Loy, S. Gong, and T. Xiang, “Feature mining for localised crowd counting,” in *European Conference on Computer Vision*, 2012. [3, 4, 5](#)
- [50] V.-Q. Pham, T. Kozakaya, O. Yamaguchi, and R. Okada, “Count forest: Co-voting uncertain number of targets using random forest for crowd density estimation,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 3253–3261. [3](#)
- [51] B. Xu and G. Qiu, “Crowd density estimation based on rich features and random projection forest,” in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2016, pp. 1–8. [3](#)
- [52] C. Wang, H. Zhang, L. Yang, S. Liu, and X. Cao, “Deep people counting in extremely dense crowds,” in *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 2015, pp. 1299–1302. [3](#)
- [53] C. Arteta, V. Lempitsky, and A. Zisserman, “Counting in the wild,” in *European Conference on Computer Vision*. Springer, 2016, pp. 483–498. [3](#)
- [54] L. Boominathan, S. S. Kruthiventi, and R. V. Babu, “Crowdnet: A deep convolutional network for dense crowd counting,” in *Proceedings of the 2016 ACM on Multimedia Conference*. ACM, 2016, pp. 640–644. [3](#)
- [55] Z. Wang, Z. Xiao, K. Xie, Q. Qiu, X. Zhen, and X. Cao, “In defense of single-column networks for crowd counting,” *arXiv preprint arXiv:1808.06133*, 2018. [3](#)
- [56] D. Oñoro-Rubio, R. J. López-Sastre, and M. Niepert, “Learning short-cut connections for object counting,” in *British Machine Vision Conference 2018, BMVC 2018, Northumbria University, Newcastle, UK, September 3-6, 2018*, 2018. [3](#)
- [57] V. A. Sindagi and V. M. Patel, “A survey of recent advances in cnn-based single image crowd counting and density estimation,” *Pattern Recognition Letters*, 2017. [3, 4](#)
- [58] X. Liu, J. van de Weijer, and A. D. Bagdanov, “Leveraging unlabeled data for crowd counting by learning to rank,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [4, 13, 14](#)
- [59] N. Liu, Y. Long, C. Zou, Q. Niu, L. Pan, and H. Wu, “Adcrowdnet: An attention-injective deformable convolutional network for crowd understanding,” *arXiv preprint arXiv:1811.11968*, 2018. [4](#)
- [60] J. Wan, W. Luo, B. Wu, A. B. Chan, and W. Liu, “Residual regression with semantic prior for crowd counting,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4036–4045. [4](#)
- [61] M. Zhao, J. Zhang, C. Zhang, and W. Zhang, “Leveraging heterogeneous auxiliary tasks to assist crowd counting,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12736–12745. [4](#)
- [62] V. A. Sindagi and V. M. Patel, “Inverse attention guided deep crowd counting network,” *arXiv preprint*, 2019. [4](#)
- [63] —, “Multi-level bottom-top and top-bottom feature fusion for crowd counting,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1002–1012. [4, 12, 13, 14](#)
- [64] X. Jiang, Z. Xiao, B. Zhang, X. Zhen, X. Cao, D. Dörmann, and L. Shao, “Crowd counting and density estimation by trellis encoder-decoder network,” *arXiv preprint arXiv:1903.00853*, 2019. [4, 13, 14](#)
- [65] M. Shi, Z. Yang, C. Xu, and Q. Chen, “Revisiting perspective information for efficient crowd counting,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7279–7288. [4](#)
- [66] W. Liu, M. Salzmann, and P. Fua, “Context-aware crowd counting,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5099–5108. [4, 12, 14](#)
- [67] Q. Zhang and A. B. Chan, “Wide-area crowd counting via ground-plane density maps and multi-view fusion cnns,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8297–8306. [4](#)



- [68] J. Wan and A. Chan, "Adaptive density map generation for crowd counting," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1130–1139. 4
- [69] D. B. Sam, S. V. Peri, A. Kamath, R. V. Babu *et al.*, "Locate, size and count: Accurately resolving people in dense crowds via detection," *arXiv preprint arXiv:1906.07538*, 2019. 4, 12, 13, 14
- [70] Z. Ma, X. Wei, X. Hong, and Y. Gong, "Bayesian loss for crowd count estimation with point supervision," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 6142–6151. 4, 12, 13, 14
- [71] Q. Zhang and A. B. Chan, "3d crowd counting via multi-view fusion with 3d gaussian kernels," *arXiv preprint arXiv:2003.08162*. 4
- [72] G. Gao, J. Gao, Q. Liu, Q. Wang, and Y. Wang, "Cnn-based density estimation and crowd counting: A survey," *arXiv preprint arXiv:2003.12783*, 2020. 4
- [73] R. Ranjan, V. M. Patel, and R. Chellappa, "Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 1, pp. 121–135, Jan 2019. 7
- [74] R. Yasarla and V. M. Patel, "Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 8
- [75] Y. Li, X. Zhang, and D. Chen, "Csrnet: Dilated convolutional neural networks for understanding the highly congested scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1091–1100. 12, 13, 14
- [76] L. Liu, Z. Qiu, G. Li, S. Liu, W. Ouyang, and L. Lin, "Crowd counting with deep structured scale integration network," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1774–1783. 12, 13, 14



**Rajeev Yasarla** [S'18] is a PhD student in the Dept. Of Electrical & Computer Engineering at The Johns Hopkins University. Prior to joining Johns Hopkins, he worked for Qualcomm India. He graduated from IIT-Madras with a Master's degree in Communication Engineering. His research interests lie in the image restoration and object detection tasks. Currently, he is working on deraining, deblurring and dehazing tasks.



**Vishal M. Patel** [SM'15] is an Assistant Professor in the Department of Electrical and Computer Engineering (ECE) at Johns Hopkins University. Prior to joining Hopkins, he was an A. Walter Tyson Assistant Professor in the Department of ECE at Rutgers University and a member of the research faculty at the University of Maryland Institute for Advanced Computer Studies (UMIACS). His current research interests include signal processing, computer vision, and pattern recognition with applications in biometrics and imaging. He has received a number of awards including the 2016 ONR Young Investigator Award, the 2016 Jimmy Lin Award for Invention, A. Walter Tyson Assistant Professorship Award, Best Paper Award at IEEE AVSS 2017, Best Paper Award at IEEE BTAS 2015, Honorable Mention Paper Award at IAPR ICB 2018, two Best Student Paper Awards at IAPR ICPR 2018, and Best Poster Awards at BTAS 2015 and 2016. He is an Associate Editor of the IEEE Signal Processing Magazine, IEEE Biometrics Compendium, Pattern Recognition Journal, and serves on the Information Forensics and Security Technical Committee of the IEEE Signal Processing Society. He is serving as the Vice President (Conferences) of the IEEE Biometrics Council. He is a member of Eta Kappa Nu, Pi Mu Epsilon, and Phi Beta Kappa.



**Vishwanath A Sindagi** [S'16] is a PhD student in the Dept. Of Electrical & Computer Engineering at The Johns Hopkins University. Prior to joining Johns Hopkins, he worked for Samsung R&D Institute-Bangalore. He graduated from IIT-Bangalore with a Master's degree in Information Technology. His research interests include deep learning based crowd analytics, object detection, applications of generative modeling, domain adaptation and low-level vision.