

# Spatiotemporal Dilated Convolution with Uncertain Matching for Video-based Crowd Estimation

Yu-Jen Ma, Hong-Han Shuai, *Member, IEEE*, and Wen-Huang Cheng, *Senior Member, IEEE*

**Abstract**—In this paper, we propose a novel SpatioTemporal convolutional Dense Network (STDNet) to address the video-based crowd counting problem, which contains the decomposition of 3D convolution and the 3D spatiotemporal dilated dense convolution to alleviate the rapid growth of the model size caused by the Conv3D layer. Moreover, since the dilated convolution extracts the multiscale features, we combine the dilated convolution with the channel attention block to enhance the feature representations. Due to the error that occurs from the difficulty of labeling crowds, especially for videos, imprecise or standard-inconsistent labels may lead to poor convergence for the model. To address this issue, we further propose a new patch-wise regression loss (PRL) to improve the original pixel-wise loss. Experimental results on three video-based benchmarks, i.e., the UCSD, Mall and WorldExpo’10 datasets, show that STDNet outperforms both image- and video-based state-of-the-art methods. The source codes are released at <https://github.com/STDNet/STDNet>.

**Index Terms**—Crowd counting, density map regression, spatiotemporal modeling, dilated convolution, patch-wise regression loss.

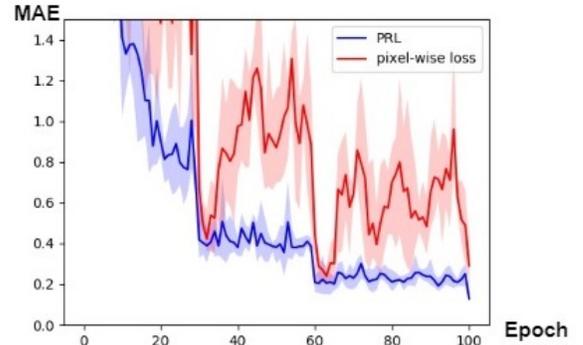
## I. INTRODUCTION

CROWD counting plays an important role in computer vision since it facilitates a variety of fundamental applications, e.g., automatic driving technologies [1]–[3], video surveillance [4], [5], safety management [6]–[8], and tracking [9]–[11]. The goal is to estimate the accurate number of people in a particular region from a video or an image. The early studies of crowd counting were based on detection frameworks with well-trained classifiers [12]–[14]. However, crowd counting is still a challenging task due to multiple issues, such as scale variations, perspective distortion, people occlusion, and light illumination.

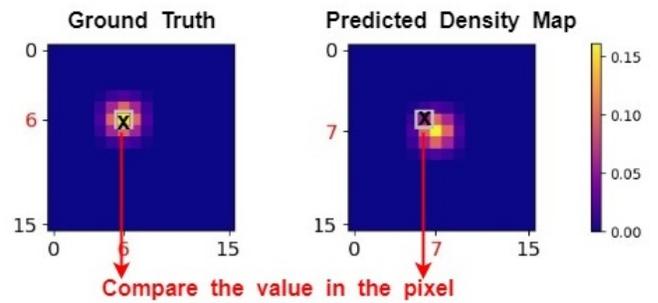
In recent years, benefiting from the powerful capability of feature extraction by deep convolutional neural networks (CNNs) [15]–[18], many CNN-based methods [19]–[30] have been developed to improve the performance on crowd counting benchmarks. However, most existing approaches, even if video datasets are available, still exploit image-based methods for crowd counting while ignoring the rich information in the temporal domain. In contrast, a recent line of research attempts to leverage the temporal correlation to improve the model [30]–[33]. For example, [31] proposed the ConvLSTM to incorporate the convolutional operation into the LSTM cell

Y.-J. Ma and H.-H. Shuai are with the Department of Electrical and Computer Engineering, National Chiao Tung University, Hsinchu, Taiwan. E-mail: {ninechi.eed08g,hshuai}@nctu.edu.tw.

W.-H. Cheng is with the Institute of Electronics, National Chiao Tung University, Hsinchu, 300 Taiwan, and the Artificial Intelligence and Data Science Program, National Chung Hsing University, Taichung, 400 Taiwan. E-mail: whcheng@nctu.edu.tw.



(a) Comparisons of training curve on the UCSD dataset



(b) Illustration of the pixel-wise regression loss

Fig. 1. An illustrative example of the problem of pixel-wise regression loss. (a) shows that the pixel-wise loss has a larger oscillation scale than the proposed Patch-wise Regression Loss. The reason may be attributed to (b), i.e., the pixel-wise loss is sensitive to such pixel-level difference between two Gaussian blobs.

to capture the temporal information, while E3D [30] integrated the 3D convolutional layer into the model to jointly consider the temporal and spatial dimensions. Moreover, [32] proposed the Locality-constrained Spatial Transformer (LSTN) module to improve the density map of the current frame by considering the information of its previous frame. TACCN [33] proposed three complementary attention mechanisms to achieve the more robust estimation.

Despite their comparable results, such models incorporated with temporal information still suffer from three issues. First, to capture the temporal information, the model sizes of 3D CNNs and ConvLSTM are much larger than those of the general 2D CNNs. When a very deep 3D CNN is used to extract discriminative features from scratch results, the computational cost and memory demand are expensive, which limits the feasibility of real-world applications (e.g., the running time or the hardware requirements). Moreover, ConvLSTM

updates the model by backward propagation through time (BPTT), which may result in inefficient training.<sup>1</sup> Second, the temporal features are encoded at a single temporal scale, while the different regions may require features at different temporal scales. As such, incorporating unnecessary temporal information may deteriorate the performance, which leads to a nonnegligible gap between state-of-the-art image-based methods and existing spatiotemporal methods (e.g., [30]–[33]). Third, it is time consuming and challenging for humans to annotate large-scale crowds, i.e., annotating dots to the centers of heads, especially for videos. As such, the annotations in videos may not be precise. To deal with imprecise dot annotation, traditional methods obtain a ground truth density map by convolving the annotation dot map with a Gaussian kernel with fixed width and minimize the pixel-wise regression loss between the ground truth density map and predicted density map. However, as shown in Fig. 1(b), even though two Gaussian blobs are close, the pixel-wise loss is still large, which leads to the unstable training since the ground truth (dot annotation) may not be precise. In contrast, as illustrated in Fig. 1(a), the proposed PRL stabilizes the training process since it does not enforce a strict alignment of the ground truth and the predicted density map. Moreover, the training curve for the pixel-wise regression loss shows a non-negligible oscillation, which results in a slower convergence rate of the model and may be sensitive to the hyperparameters (e.g., learning rate).

To address these issues, in this paper, we propose a novel SpatioTemporal convolutional Dense network (STDNet) that efficiently extracts spatial and temporal information for crowd counting. Specifically, following P3D-ResNet [34], we decouple the Conv3D layer into a Conv2D layer on the spatial dimension and a Conv1D layer on the temporal dimension to address the first issue. Moreover, to address the second issue, we extend the dilated dense convolution to the temporal domain by making the 2D Dense Dilated Convolution Block (DDCB) [20] to be the 3D DDCB, which extracts multiscale temporal features. Moreover, since the multiscale features are extracted to different channels, a general channel attention block is also proposed to attend important channels for different regions. To address the third issue, we propose a patch-wise regression loss (PRL) that calculates the difference between the ground truth density map and predicted density map patch-by-patch.<sup>2</sup> To demonstrate the improvements of the proposed STDNet, we conduct experiments on the three benchmark datasets UCSD [35], Mall [36], and WorldExpo'10 [37], [38], which show that STDNet outperforms the state-of-the-art methods by 10.5%, 3.3%, and 2% in terms of the mean absolute error, respectively.

The main contributions of this paper can be summarized as follows:

- We propose the SpatioTemporal convolutional Dense Net-

work (STDNet), which alleviates the issues encountered in general 3D CNNs by extending the dilated dense convolution to the temporal dimension with channel attention. The number of parameters of STDNet is only slightly greater than that of image-based methods. To the best of our knowledge, this is the first work that exploits dilated 3D convolution and dense networks in the crowd counting task to deal with video-based datasets.

- To handle the imprecise annotation of crowds, we propose the patch-wise regression loss (PRL) to mitigate the issues caused by the original pixel-wise regression loss, which stabilizes the training and is compatible with both image- and video-based approaches.
- Experiments demonstrate the effectiveness of each proposed component, which reveals that STDNet outperforms state-of-the-art models in terms of mean absolute error (MAE) and mean square error (MSE) on the three benchmark datasets. Moreover, the results also validate that the proposed PRL improves the commonly-used pixel-wise regression loss.

## II. RELATED WORKS

We introduce recent crowd counting works, which can be further categorized into A) detection-based methods, B) regression-based methods and C) density map-based methods. Moreover, we discuss some works that proposed the specialized loss function for the crowd counting task in D) the loss function for crowd counting.

### A. Detection-based Methods

To address the crowd counting task, detection-based methods aim to first detect and locate the head, parts of the body, or the full body for each person by using handcrafted and low-level features [12], [14], [39]–[44] and then accumulate the number of corresponding annotations. However, detection is too challenging for extremely crowded scenes due to occlusion and complex patterns, leading to a rapid performance drop. Although there have recently been many advanced deep CNN-based detection methods, e.g., Faster R-CNN [45], SSD [46], YOLO [47], DecideNet [22], and PSDDN [24], the performance gap between the detection-based methods and the density-map-based methods is still large, especially for densely crowded scenes.

### B. Regression-based Methods

Due to the failure of detection-based methods in extremely congested scenes, there are some studies aiming to directly estimate the total number of people by using machine learning methods, e.g., Gaussian processes [35], Bayesian inference [48], and neural networks [49]. However, such methods only consider the relations between the single scalar and the global features of input images while ignoring the informative indication of the local patterns. Moreover, even if the count number is accurate, it provides less information than the density map or detection bounding box since the single scalar cannot infer the location of people.

<sup>1</sup>For example, CSRNet [21] takes 35 seconds to train one epoch, whereas ConvLSTM [31] spends 530 seconds training one epoch under the same experimental setting.

<sup>2</sup>It is worth noting that the proposed PRL is general and compatible with both image- and video-based methods. We show the improvement of using PRL with other methods in the experiments.

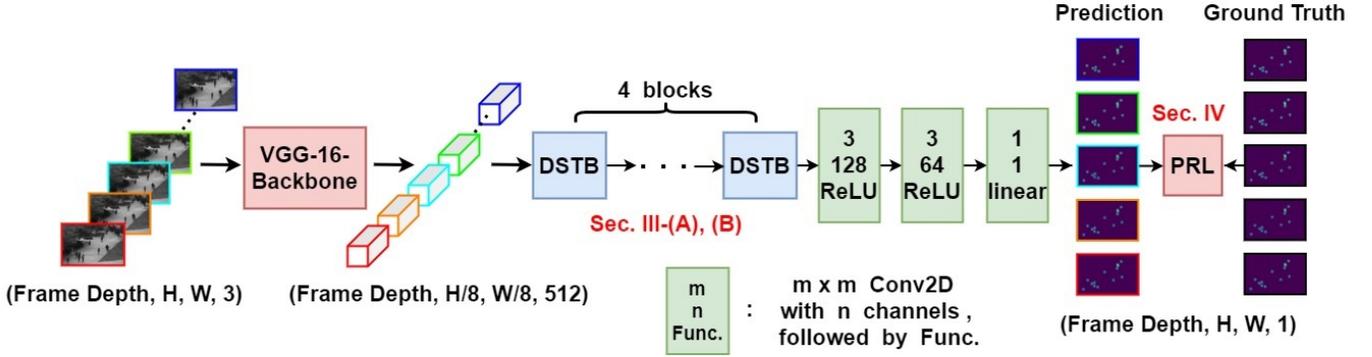


Fig. 2. The overall network architecture of our proposed STDNet.

### C. Density-map-based Methods

The density map, different from the count number, is the label to indicate the existence of humans geometrically, which provides more informative guidance to the training process. Lempisky *et al.* [50] approximate the counting task to the convex quadratic problem by using the proposed maximum excess over subarray (MESA) loss that optimizes the pixel-level subregion with a high discrepancy to the ground truth. Recently, with the advancement of the deep CNN, a recent line of research has studied the crowd counting problem by CNN-based methods [19]–[29], [51]–[54]. For example, Zhang *et al.* [29] proposed a multicolumn architecture to extract the multiscale features for solving the variant scale problem. Sam *et al.* [26] proposed classifying each patch in an image to determine which CNN branch is proper to estimate the number of people. Li *et al.* [21] further proposed CSRNet by using the 2D dilated convolutional layer to enlarge the receptive fields without significantly increasing the parameters. Dai *et al.* [20] proposed the dense dilated convolution block to consider continuously varied scale features. Sindagi *et al.* [54] proposed two attention modules that enhance low-level features by infusing spatial segmentation information (*e.g.*, [55], [56]) and improved the channel-wise information in the higher level layers. Despite their promising results, most of them put much effort into devising the network with the spatial mechanism without considering the rich information in the temporal domain of video datasets. Few studies [30]–[33] consider the temporal correlation between consecutive frames. However, the multiscale temporal information is not considered, as well as the imprecise annotations.

### D. Loss function for Crowd Counting

Since the learning target (*i.e.*, the density map) is 2D matrix-form, the intuitive idea is to regress each entry value between the predicted density map and the ground truth, *i.e.*, the pixel-wise regression loss. Although the loss is simple and effective for crowd counting, there are still some studies [57]–[59] that proposed a specialized loss function to avoid the lack of the original loss. For example, Cheng *et al.* [57] proposed the maximum excess over pixels (MEP) loss to find the pixel-level subregions that need to be updated to avoid the oversmooth gradients from the pixel-wise regression loss. Liu *et al.* [58]

proposed the multitask framework with unlabeled data, using the nature of crowd counting that the number of people in the subimage is definitely less than or equal to that of the corresponding original image, to train the model in a self-supervised manner. Ma *et al.* [59] did not directly regard the density map as the learning target; instead, they used the dot annotations and corresponding Gaussian distribution as the likelihood probability, proposing the Bayesian loss to address the training difficulty from complex and irregular situations (*e.g.*, congestion and occlusion) for crowd counting. Although these works effectively deal with some lack of the original loss, most of them still assume that the given label is accurate. However, it is inconsistent with reality since pixel-level dot annotations are very easily standard-inconsistent and imprecise, and such a problem may lead to inefficient training.

## III. MODEL ARCHITECTURE

To capture the multiscale spatiotemporal information for crowd counting without significantly increasing the number of parameters, we propose STDNet, which learns to generate a density map from  $T$  consecutive frames. Fig. 2 illustrates the model architecture of STDNet. Specifically, the first ten layers in VGG-16 [17] are used as the backbone to extract features.<sup>3</sup> Afterward, we stack four dense spatiotemporal blocks (DSTBs) without any pooling layers in the DSTBs to extract spatiotemporal information while retaining the size of the feature maps. Two Conv2D layers are then stacked after the DSTBs and followed by a  $1 \times 1$  Conv2D layer with a 1 channel to generate the density map. Since there are three pooling layers in the backbone, we simply use bilinear interpolation with a scale of 8 to upsample the output of the model to match the size of the ground truth. Finally, the predicted density map and the ground truth are fed into the patch-wise regression loss with a gradient-descent-based optimizer to train the parameters. In the following, we first present the proposed DSTB and channel attention block in Sec. III-A and Sec. III-B, respectively. Afterward, the definitions of the ground truth of the density map, geometry-adaptive kernels, and patch-wise regression loss (PRL) are detailed in Sec. IV.

<sup>3</sup>We use the VGG-16 as the backbone for a fair comparison since many baselines use the VGG-16 as the backbone (*e.g.*, CSRNet [21], PACNN [52], PaDNet [28] and CAN [23]).

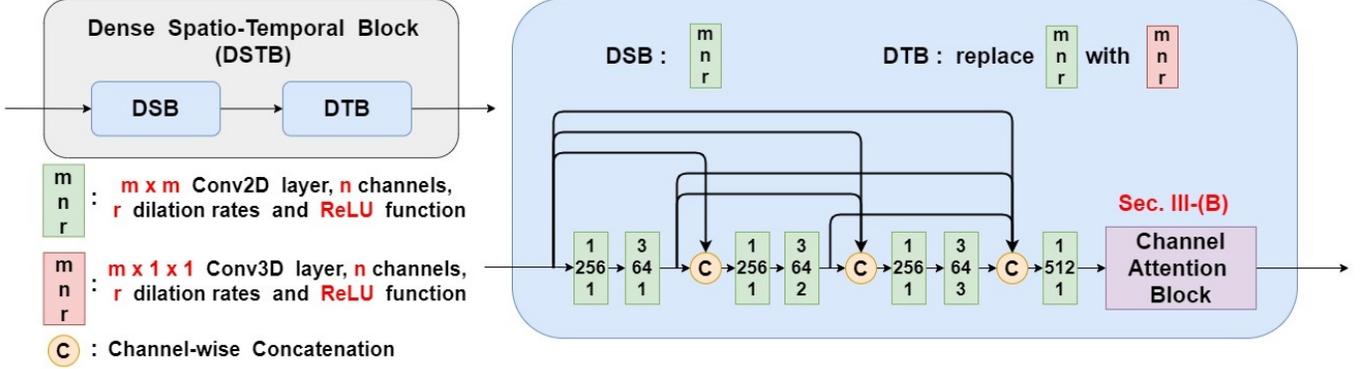


Fig. 3. The illustration of the proposed DSTB. It is worth noting that the inner structure of the DTB is similar to DSB, except for the Conv. layer. DSB utilizes Conv2D layer while DTB utilizes the Conv3D layer.

### A. Dense Spatiotemporal Block

1) *Decomposition of 3D CNNs*: To jointly capture the spatial and temporal information without directly using 3D CNNs, inspired by P3D-ResNet [34], we decompose a regular 3D CNN to a 2D CNN on the space dimension and a 1D CNN on the temporal dimension. Specifically, given a 3D convolutional filter with the size of  $n_t \times n_{s1} \times n_{s2}$ , where  $n_{s1}$  and  $n_{s2}$  denote the dimension of the spatial domain and  $n_t$  denotes the dimension of the temporal domain, we can decompose the 3D convolutional filter to a 2D convolutional filter with the size of  $1 \times n_{s1} \times n_{s2}$ , which are shared by the spatial domain at different time slots and a 1D convolutional filter with the size of  $n_t \times 1 \times 1$ . The decomposition reduces the number of parameters and takes the temporal features into account simultaneously. For example, when  $n_t = n_{s1} = n_{s2} = n$ , the number of parameters is reduced from  $O(n^3)$  to  $O(n^2)$  by using the decomposition method.

2) *Dense Spatiotemporal Block*: The architecture of the dense spatiotemporal block (DSTB) is shown in Fig. 3. The DSTB can be decomposed into the dense spatial block (DSB) and the dense temporal block (DTB). The inner structure of the DTB is similar to DSB, except for the size of the kernel, *i.e.*,  $3 \times 3$  for DSB, and  $3 \times 1 \times 1$  for DTB. The key idea of the dense spatiotemporal block fully exploits the dilated convolution with a dense reception field to capture the spatiotemporal information. Specifically, dilated convolution is regarded as a general and powerful operation for increasing the reception field without any pooling layer or any downsampling layer. As such, we can increase the reception field while retaining the size of the kernel to prevent the model from increasing the number of parameters. Generally, the output of a 3D dilated convolution with a kernel size of  $(2L+1) \times (2M+1) \times (2N+1)$  at position  $(s, i, j)$ , denoted by  $y(s, i, j)$ , can be calculated as:

$$\sum_{l=-L}^L \sum_{m=-M}^M \sum_{n=-N}^N (w(l, m, n) \cdot x(s + rl, i + rm, j + rn)), \quad (1)$$

where  $y(\cdot, \cdot, \cdot)$  is the output of the 3D dilated convolutional layer and  $x(\cdot, \cdot, \cdot)$  and  $w(\cdot, \cdot, \cdot)$  are the input and the convolutional kernel, respectively. The scalar  $r$  represents the dilation

rate. If  $r$  is set to 1, then Equation (1) can be viewed as a regular 3D convolution layer. The dilation rate controls the reception field in both the temporal and spatial domains, which is helpful for extracting multiscale features.

Inspired by DSNet [20], we use the same settings for the number of channels and the dilation rates to address the problem of scale variations. Dense channelwise concatenation is also used in the DSTB to present the multiscale features on the channel dimension. Moreover, the  $1 \times 1$  and  $1 \times 1 \times 1$  Conv. layers are used in front of each dilated Conv. layer to prevent the feature maps from increasing the channel size. It is worth noting that this is the first work combining two mutually beneficial concepts together, *i.e.*, the dense dilated structure provides the representative information on the spatial and temporal domains, while the decomposition trick effectively reduces the number of parameters required to construct a continuous reception field in a dense structure. Moreover, the combination of the decomposition method and the dilated operation can further mitigate the growth of the number of parameters compared with regular 3D convolution.

### B. Channel Attention Block

Different from DSNet [20], we use the general channel attention block inspired by [30], [60] to enhance the feature representations for each channel since different regions require features of different scales to predict the correct density map. Therefore, after DSTB extracts the multiscale features by using the different dilation rates and concatenates them on the channel dimension, we further exploit the channel attention block to enhance the extracted features. The general channel attention block can be simply split into the global average pooling layer (GAP) and the fully connected layer. The channel attention blocks in the DSB and the DTB are similar. However, the operation of the GAP layer in the DTB additionally considers the temporal dimension of the feature maps since one of the effects in the DTB is to leverage the temporal information. Specifically, let  $C_t$  and  $C_s$  denote the number of channels derived from DTB and DSB, respectively. For the input feature maps in the GAP layer, the importance

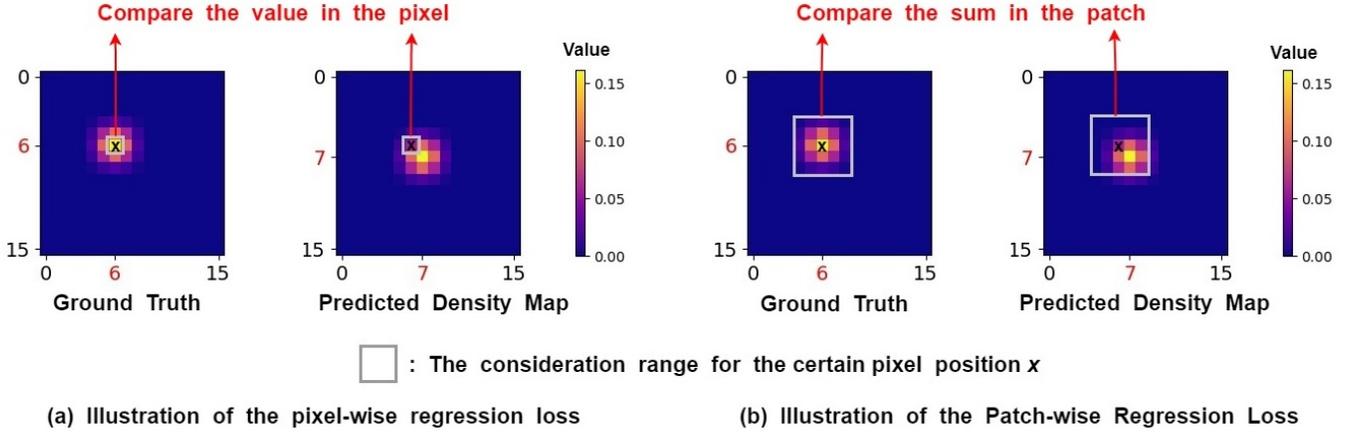


Fig. 4. An illustrative example of (a) pixel-wise regression loss and (b) Patch-wise Regression Loss (PRL). The main difference is the considered range for the certain pixel position  $\mathbf{x}$ : the former is to consider the value only in the pixel for (a), and the latter is to consider the sum of the value in the patch for (b).

vector  $\vec{a}_t = [a_t^{(1)}, \dots, a_t^{(c)}, \dots, a_t^{(C_t)}]$  for the channels in DTB can be calculated by:

$$a_t^{(c)} = \frac{1}{D_t \times H_t \times W_t} \sum_{k=1}^{D_t} \sum_{i=1}^{H_t} \sum_{j=1}^{W_t} x_t^{(c)}(k, i, j), \quad (2)$$

where  $x_t^{(c)} \in \mathbb{R}^{D_t \times H_t \times W_t}$  represents the  $c$ -th channel feature maps derived by DTB. Similarly, the importance vector  $\vec{a}_s = [a_s^{(1)}, \dots, a_s^{(c)}, \dots, a_s^{(C_s)}]$  for the channels in DSB is calculated by:

$$a_s^{(c)} = \frac{1}{H_s \times W_s} \sum_{i=1}^{H_s} \sum_{j=1}^{W_s} x_s^{(c)}(i, j), \quad (3)$$

where  $x_s^{(c)} \in \mathbb{R}^{H_s \times W_s}$  represents the feature maps derived by DSB. Following  $\vec{a}_t$ , we use 2 fully-connected layers to derive the attention vector  $\vec{\alpha}_t$ , i.e.,

$$\vec{\alpha}_t = \text{Sigmoid}(W_2 \times \text{ReLU}(W_1 \times \vec{a}_t)), \quad (4)$$

where  $\vec{\alpha}_t \in \mathbb{R}^C$  is the output vector.  $W_1$  and  $W_2$  represent the weight matrices of the two fully connected layers.  $\vec{\alpha}_s$  can be derived in a similar way. Finally, we use  $\vec{\alpha}_t$  and  $\vec{\alpha}_s$  to perform channelwise multiplication on the inputs. Since the range of the sigmoid function is defined on  $[0, 1]$ , the channel attention block can adaptively adjust the weights for each channel according to the inputs. Different from DSNet [20], which leverages dense spatial features, the proposed channel attention block effectively selects the required feature scales for different image regions to address the scale-variant problem. Moreover, the 3D decomposition effectively reduces the required parameters for extracting the spatiotemporal information so that the attention block only has to attend either different temporal scales or different spatial scales.

#### IV. PATCH-WISE REGRESSION LOSS

For the objective function, we use the density map as the ground truth and treat crowd counting as a regression problem. The generation of the density map follows previous work [50]. Given the  $i$ -th ground truth dot map  $D_i$ , where each dot

represents the location of a person and is marked as 1 on  $D_i$ , we use a 2D Gaussian kernel to blur  $D_i$ . The density map of  $D_i$ , denoted as  $Y_i^{GT}$ , can be defined as:

$$Y_i^{GT}(\mathbf{x}) = \sum_{\{\mathbf{x}' | D(\mathbf{x}')=1\}} \mathcal{N}(\mathbf{x} | \mathbf{x}', \sigma^2 I), \quad (5)$$

where  $\mathbf{x} = (x_1, x_2)$  is the pixel position in the density map, and  $\mathcal{N}(\mathbf{x} | \mu, \Sigma)$  is a multivariate Gaussian with mean  $\mu$  and covariance  $\Sigma$ . To address the perspective distortion, MCNN [29] proposed using geometry-adaptive kernels for generating the density map. The definition of geometry-adaptive kernels is similar to Equation (5), except for the standard deviation of the 2D Gaussian kernel, which is formulated as:

$$Y_i^{GT}(\mathbf{x}) = \sum_{\{\mathbf{x}' | D(\mathbf{x}')=1\}} \mathcal{N}(\mathbf{x} | \mathbf{x}', \sigma_{ij}^2 I), \quad (6)$$

where  $\sigma_{ij}$  denotes the standard deviation of the 2D Gaussian kernel for the  $j$ -th person in  $D_i$ .  $\sigma_{ij}$  is defined as  $\beta \times \bar{d}_{ij}$ , where  $\beta$  is a hyperparameter, and  $\bar{d}_{ij}$  represents the average distance of  $k$  nearest point annotations for the  $j$ -th person in the  $i$ -th image.<sup>4</sup>

Equipped with the density map, most previous works use the pixel-wise difference between the prediction and the ground truth as the objective loss, which is defined by:

$$\mathcal{L}(\theta) = \frac{1}{2N_b} \sum_{i=1}^{N_b} \|f(X_i; \theta) - Y_i^{GT}\|_2^2, \quad (7)$$

where  $X_i$  represents the  $i$ -th image,  $\theta$  denotes the parameters of the model,  $f(\cdot; \theta)$  is the model output, and  $N_b$  denotes the total number of images in a batch.<sup>5</sup>

However, imprecise (or inconsistent) annotations are very common in densely crowded scenes, especially in videos. Using a pixel-level loss with imprecise annotation may lead to poor performance. For example, if the position of the predicted density peak and that of the ground truth density peak are

<sup>4</sup>For a fair comparison, we set  $\beta$  and  $k$  as 0.3 and 3, respectively, which follows the setting of MCNN [29].

<sup>5</sup>The factor  $\frac{1}{2}$  is commonly used for gradient-based methods with the  $\ell_2$ -norm.

only one pixel away, which is caused by annotation error, it still induces nonnegligible loss. Therefore, since the semantic patterns for one person are usually filled with several pixels, we propose a patch-wise loss for uncertain matching. Fig. 4(a) illustrates the original pixel-wise regression loss. Assume that the mean of the Gaussian distribution in the ground truth is (6, 6), while the predicted distribution is located at (7, 7), to human beings, it is accurate enough to understand the total number of people in the particular region. However, it still induces a considerable error in the pixel-wise loss, which may be harmful to the training process and lead to unstable convergence.

To address the problem, we propose a new loss for crowd counting, i.e., patch-wise regression loss (PRL). The concept of the PRL is illustrated in Fig. 4(b). Given a pixel position  $\mathbf{x}$  in the density map, instead of comparing the value between the prediction and the ground truth, we consider the sum of the density in the patch centered at  $\mathbf{x}$ . As such, even if the pixel-level error is produced in the training process, it does not lead to a large loss and generates a dominant gradient to update parameters. Moreover, we can use different sizes of patches to consider different levels of precision, e.g.,  $1 \times 1$  (pixel-wise),  $3 \times 3$ , and  $5 \times 5$ . In addition, because the distances to the center pixel  $\mathbf{x}$  are different, the importance for each pixel in the patch should be different. Therefore, we use the 2D Gaussian kernel again to imply the weight in the patch. Let  $\mathcal{L}_p^{(z)}$  and  $n_p$  denote the loss of patch size  $(2z - 1) \times (2z - 1)$  and the number of different patch sizes, respectively. The patch-wise regression loss, denoted as  $\mathcal{L}_p$ , is summarized over  $n_p$  different patch losses as follows.

$$\mathcal{L}_p(\theta) = \sum_{z=1}^{n_p} \lambda_z \cdot \mathcal{L}_p^{(z)}(\theta), \quad (8)$$

$$\mathcal{L}_p^{(z)}(\theta) = \frac{1}{N_b} \sum_{i=1}^{N_b} \left\| G_\sigma^{(z)} * f(X_i; \theta) - G_\sigma^{(z)} * Y_i^{GT} \right\|_1, \quad (9)$$

where  $*$  denotes the convolution operator, and  $G_\sigma^{(z)}$  represents the 2D Gaussian kernel with a kernel size of  $(2z-1) \times (2z-1)$ , e.g.,  $G_\sigma^{(2)}$  is the 2D Gaussian kernel with a kernel size of  $3 \times 3$ . It is worth noting that we use the  $\ell_1$ -norm as our objective loss in Equation (9) since the  $\ell_2$ -norm loss is sensitive to the outliers from the wrong labeling, which is common to the crowd counting benchmarks.

## V. EXPERIMENTS

To evaluate the proposed STDNet, we conduct experiments on three public video-based datasets, including UCSD [35], Mall [36], and WorldExpo'10 [37], [38] datasets. In the following, we first introduce the evaluation metrics in Sec. V-A and present the implementation details in Sec. V-B. Afterward, the quantitative results on the three datasets with the comparison to the state-of-the-art methods are shown in Sec. V-C, followed by the qualitative results in Sec. V-D. In Sec. V-E, we provide the ablation study on UCSD [35] to demonstrate the improvement of each proposed component. Finally, we discuss the inference speed, compatibility of the proposed PRL, and visualization of channel attention in Sec. V-F.

### A. Evaluation Metrics

The most common evaluation metrics in crowd counting are the mean absolute error (MAE) and the mean square error (MSE), which are defined as follows.

$$MAE = \frac{1}{N_t} \sum_{i=1}^{N_t} |P_i - P_i^{GT}|, \quad (10)$$

$$MSE = \sqrt{\frac{1}{N_t} \sum_{i=1}^{N_t} (P_i - P_i^{GT})^2}, \quad (11)$$

where  $N_t$  is the total number of images in the testing data,  $P_i$  and  $P_i^{GT}$  represent the prediction and the ground truth of the crowd count of the  $i$ -th image, which are obtained by integrating the predicted density map and the ground truth density map, respectively.

### B. Implementation Details

Similar to CSRNet [21], we use the pretrained weight of the backbone layer (i.e., the first ten layers from VGG-16) to accelerate the training process instead of training from scratch. Adam is utilized as the optimizer [61] to minimize our proposed  $\mathcal{L}_p$ , while the learning rate is set to 0.0001 initially and divided by 2 per 30 epochs. Moreover, the ground truth generation settings follow CSRNet [21]. That is, a fixed kernel with  $\sigma = 3$  is set on the UCSD [35] and the WorldExpo'10 [37], [38], while the deometry-adaptive kernels are used on the Mall [36]. Additionally, we use the horizontal flip to perform the data augmentation. By using cross-validation, the number of patch sizes  $n_p$  is set to 3, and the weight  $\lambda_z$  for each  $L^{(z)}(\cdot)$  is set to 1, 15 and 3 for  $z = 1, 2$  and 3. The standard deviation of the 2D Gaussian kernel defined in Equation (9) is set to 1. In addition, the frame depths  $T$  mentioned in Sec. III on the three datasets (UCSD, Mall, and WorldExpo'10) are set to 10, 8, and 5, respectively, due to the different frame rates of the three datasets. The implementation code will be released on: <https://github.com/STDNet/STDNet>.

TABLE I  
COMPARISONS OF DIFFERENT METHODS ON UCSD DATASET.

	Method	MAE ↓	MSE ↓
Image-based	RPNet [62]	1.32	1.23
	CSRNet [21]	1.16	1.47
	SPANet+SANet [57]	1.00	1.28
	ADCrowdNet [63]	0.98	1.25
	PACNN [52]	0.89	1.18
	PaDNet [28]	0.85	1.06
Video-based	Bi-ConvLSTM [31]	1.13	1.43
	LSTN [32]	1.07	1.39
	E3D [30]	0.93	1.17
	STDNet (ours)	<b>0.76</b>	<b>1.01</b>

### C. Quantitative Results

**UCSD dataset.** The UCSD dataset [35] is composed of 2,000 gray-level frames by a fixed camera in the same scene. According to the settings in CSRNet [21] and MCNN [29], the 601-1,400th frames are selected as the training data,

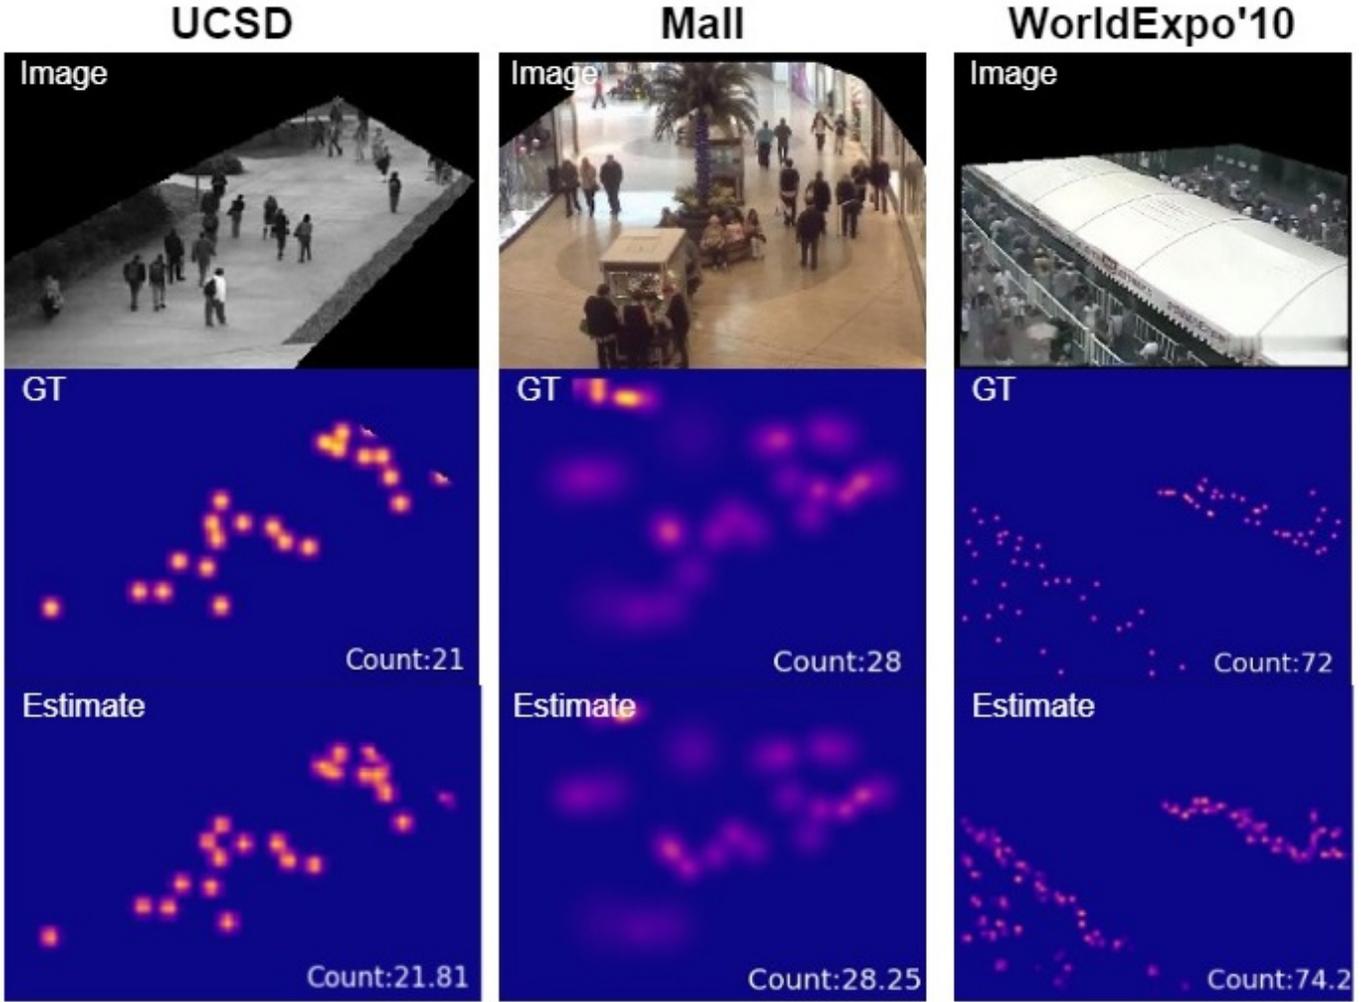


Fig. 5. Examples of the predicted results by STDNet on the three datasets, i.e., UCSD, Mall, and WorldExpo'10. The first row shows the input image, the second row represents the ground truth density map, and the last row is the predicted density map.

and the resting 1,200 frames are used as the testing set. Table I compares the proposed STDNet with the state-of-the-art image- and video-based methods [21], [28], [30], [31], [52], [57], [62], [63]. The results show that the proposed STDNet outperforms all of the previous methods in terms of MAE and MSE by at least 10.5% and 4.7% (compared to PaDNet [28]), respectively. Additionally, the proposed STDNet also outperforms video-based methods, i.e., ConvLSTM [31], LSTN [32], and E3D [30], since STDNet effectively extracts the multiscale temporal information by DSTB and further enhances the feature by channel attention.

It is worth noting that the improvement on the UCSD dataset is significant for the following two main reasons. First, the frame rate per second of the UCSD dataset is 10 fps, which is suitable for our work to extract multiscale temporal features since there are high correlations between consecutive frames. For example, if there is a person continuously appearing from the first frame to the last frame, the proposed method can extract the short-term information (e.g., from the 1st frame to the 2nd frame) and the long-term information (e.g., from the 1st frame to the 10th frame). As such, the estimations of both

the position and count are improved since the position of the person in each frame should present the monotonic direction, and the count in each frame should remain the same. Second, since the moving speed of every person is usually different, if we can extract multiscale temporal information, all people with different moving speeds can be simultaneously considered to predict better results.

**Mall dataset.** The Mall dataset [36] is also a public video-based dataset, which consists of 2,000 color images in the shopping mall with a resolution of  $480 \times 640$ . For a fair comparison with previous works, the first 800 frames are used as the training set, and the remaining part is used as the testing data.

Table II compares STDNet with previous works [22], [30]–[33], [36], [64]–[67]. Similar to the results on the UCSD dataset, STDNet outperforms the image-based methods by at least 3.3% in terms of MAE (compared with DecideNet [22]). In addition, STDNet also outperforms the four video-based methods [30]–[33] by at least 10.3% in terms of MAE. The results also show that if video-based methods do not leverage the temporal features effectively, the performance is not better

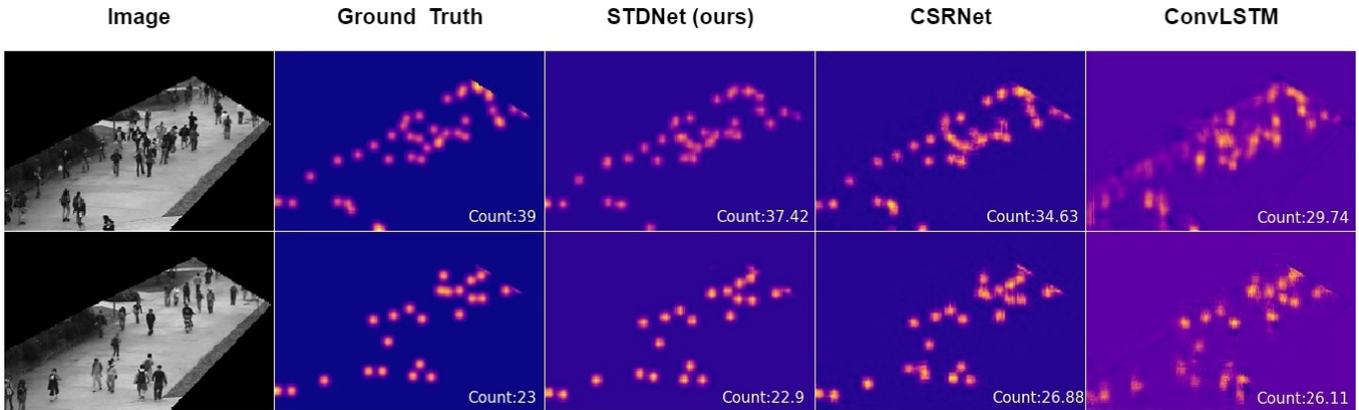


Fig. 6. Examples on the UCSD dataset with the comparison to the image- and video-based methods. The first row presents the failure case to all the three models, while the second row presents the success case for the proposed STDNet.

than that of the image-based methods.

Note that the frame rate per second of the Mall dataset is 2 fps. Therefore, the improvement benefiting from the temporal information may be smaller. However, we find that the scene in the Mall dataset is more complex than that in the UCSD dataset, *i.e.*, the higher perspective distortion and the occluded situation. Therefore, it is inclined to cause incorrect or imprecise annotations. In this case, the proposed PRL can address this issue since it can match the value between the predicted density map and the ground truth in a patch-by-patch manner, which results in more effective training guidance and outperforms state-of-the-art methods.

TABLE II  
COMPARISONS OF DIFFERENT METHODS ON MALL DATASET.

	Method	MAE ↓	MSE ↓
Image-based	Chen <i>et al.</i> [36]	3.15	15.7
	CNN-Boosting [66]	2.01	-
	MoC-CNN [64]	2.75	13.4
	HSRNet [67]	1.80	2.28
	DRSAN [65]	1.72	2.10
	DecideNet [22]	1.52	1.90
Video-based	Bi-ConvLSTM [31]	2.10	7.6
	LSTN [32]	2.00	2.50
	TACNN [33]	1.94	2.41
	E3D [30]	1.64	2.13
	STDNet (ours)	<b>1.47</b>	<b>1.88</b>

**WorldExpo’10 dataset.** WorldExpo’10 [37], [38] contains 3,980 frames made up of 1,130 videos in 108 scenes with a resolution of  $576 \times 720$ . For a fair comparison, 3,380 frames from 103 different scenes are used for training, and the remaining 600 frames from 5 different scenes are used for testing. Table III compares the proposed STDNet with existing methods [23], [25], [30], [31], [51]–[53], [62], [68], [69]. The MAE is reported for each scene, and the overall performance is averaged by the results of the different 5 scenes. From Table. III, the proposed STDNet performs the best in Scene 1, Scene 4, Scene 5 and the average, while ECAN [23] and PGCNet [53] perform the best in Scene 2 and Scene 3, respectively.

We further investigate Scenes 2 and 3 to find the possible reasons. First, the frame rates per second of the three datasets

are different. The frame rate per second of the WorldExpo’10 dataset is  $\frac{1}{30}$  only, which is much smaller than that of the UCSD and Mall datasets (10 fps and 2 fps, respectively). Therefore, the frame rate is too low to capture the correlation between the consecutive frames on the WorldExpo’10 dataset, leading to relatively small improvements compared with the state-of-the-art methods. Moreover, we observe that all the video-based approaches [30], [31] perform worse in Scene 2 and Scene 3 than in others. In fact, Scene 2 and Scene 3 show more obvious temporal patterns (crowded for watching the performance on the stage and regularly entering the venue in batches, respectively). However, previous frames provide out-of-date information (30 seconds ago) that cannot precisely match the current frame for a better prediction. For example, temporal information shows that people are entering the venue, but a person in the previous frame cannot be found in the current frame after 30 seconds. Since video-based approaches aim to extract the temporal features, the improvement is small if the previous frames provide useless or uncorrelated assistance. It is still worth noting that the average performance of the proposed STDNet is still better than all of the baselines.

TABLE III  
COMPARISONS OF DIFFERENT METHODS ON WORLDEXPO’10 DATASET.

Method	S1	S2	S3	S4	S5	Avg.
ic-CNN [25]	17.0	12.3	9.2	8.1	4.7	10.3
RpNet [62]	2.4	10.2	9.7	11.5	3.8	8.2
SDANet [69]	2.0	14.3	12.5	9.5	2.5	8.1
PGCNet [53]	2.5	12.7	<b>8.4</b>	13.7	3.2	8.1
TEDnet [51]	2.3	10.1	11.3	13.8	2.6	8.0
LSC-CNN [68]	2.9	11.3	9.4	12.3	4.3	8.0
PACNN [52]	2.3	12.5	9.1	11.2	3.8	7.8
ECAN [23]	2.4	<b>9.4</b>	8.8	11.2	4.0	7.2
Bi-ConvLSTM [31]	6.8	14.5	14.9	13.5	3.1	10.6
E3D [30]	2.8	12.5	12.9	10.2	3.2	8.32
STDNet (ours)	<b>1.83</b>	12.78	10.3	<b>7.88</b>	<b>2.5</b>	<b>7.05</b>

In addition, Scene 4 is a crowded scene with high perspective distortion, and STDNet still performs the best. It shows that even if the frame rate per second is dramatically distributed on different video-based datasets, our STDNet still achieves state-of-the-art performance due to the proposed

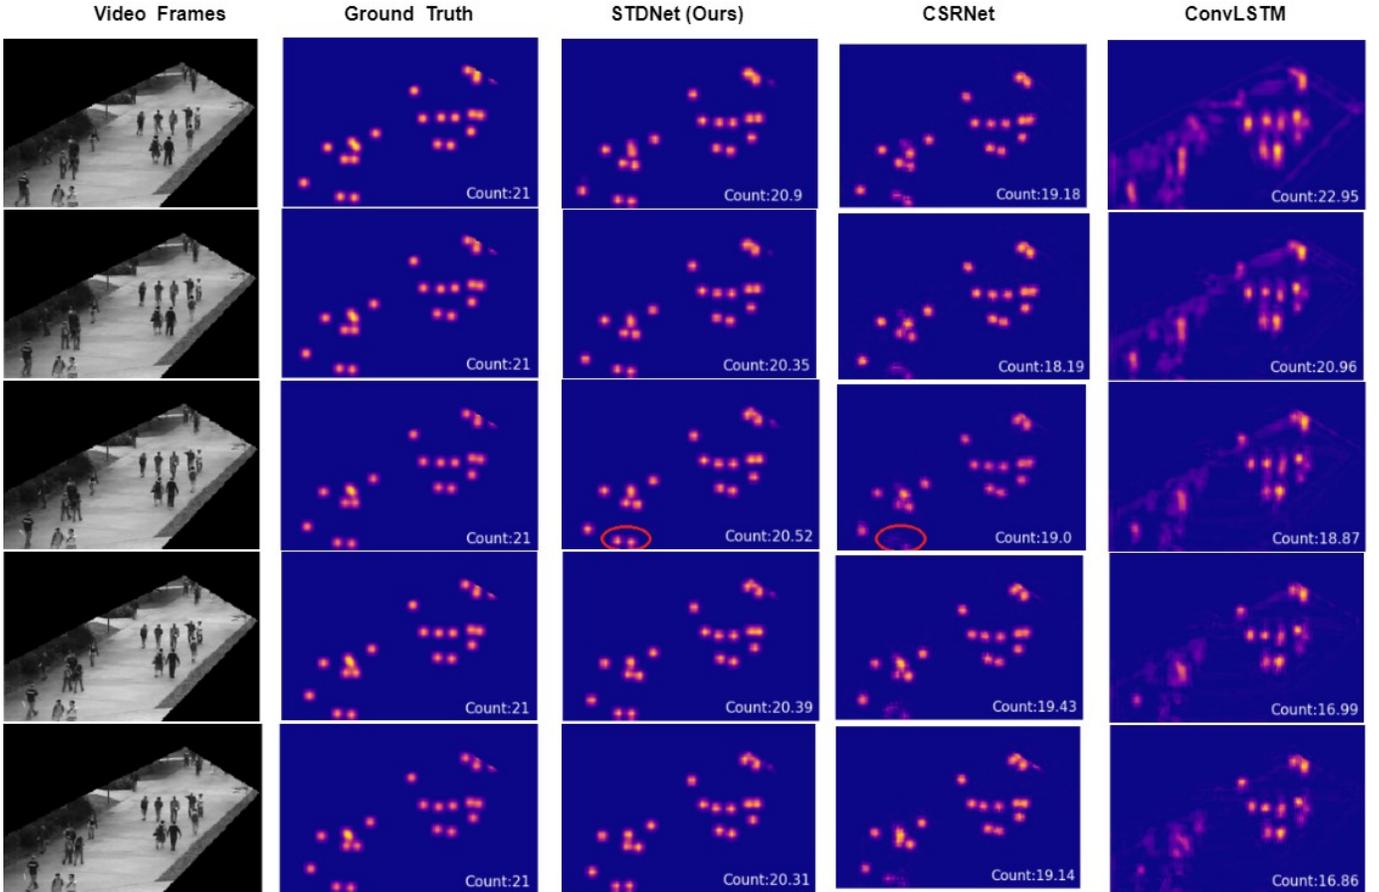


Fig. 7. Visual comparison of the image- and video-based methods on consecutive frames of UCSD dataset.

TABLE IV  
LEAVE-ONE-OUT ABLATION ANALYSIS WITH DIFFERENT CONFIGURATIONS ON THE UCSD DATASET.

Method	Configurations				MAE↓	MSE↓
	CA	DS	DT	PRL		
Ours					1.25	1.63
		✓	✓	✓	0.80	1.06
	✓		✓	✓	0.84	1.07
	✓	✓		✓	0.83	1.09
	✓	✓	✓		0.85	1.11
	✓	✓	✓	✓	<b>0.76</b>	<b>1.01</b>

DSTB and the channel attention block for extracting and enhancing the different-scale temporal features.

#### D. Quantitative Results

As shown in Fig. 5, the proposed STDNet can not only calculate the crowd count accurately but also precisely estimate the crowd density map. In addition, we also provide some qualitative results compared with the image-based method [21] and video-based method [31] in Fig. 6 (failure case in the first row and success case in the second row). By the success case in Fig. 6, it shows that even if the crowd is sparse with a simple scene, a nonnegligible error still occurs if methods do not carefully extract the information for the prediction. In addition, since our STDNet can adaptively aggregate spatial

and temporal features with short-term and long-term information, even in the failure case of Fig. 6 with severe occlusion, it still achieves better estimation results compared with other methods.

Fig. 7 visually compares the results of ground truth, STDNet, CSRNet, and ConvLSTM for consecutive frames. The results show that the proposed STDNet better predicts the density maps than CSRNet and ConvLSTM. For example, as shown inside the red circle in the third row and fourth column, CSRNet does not correctly recognize the people, even though it correctly recognize the people in the first frame. This is because CSRNet does not leverage the temporal correlation. Moreover, the results of ConvLSTM are blurry in consecutive frames since ConvLSTM only considers single-scale temporal information while different people may have different walking speeds. In contrast, the proposed STDNet better estimates the density map than the other two methods by extracting the multiscale temporal information and mitigating the inconsistent label issue with PRL.

#### E. Ablation Study

To further investigate the improvement of each proposed component, including the channel attention block (CA), the dense dilated convolution on the spatial (DS) and temporal (DT) dimensions, and the patch-wise regression loss (PRL), we perform a leave-one-out ablation analysis on the UCSD

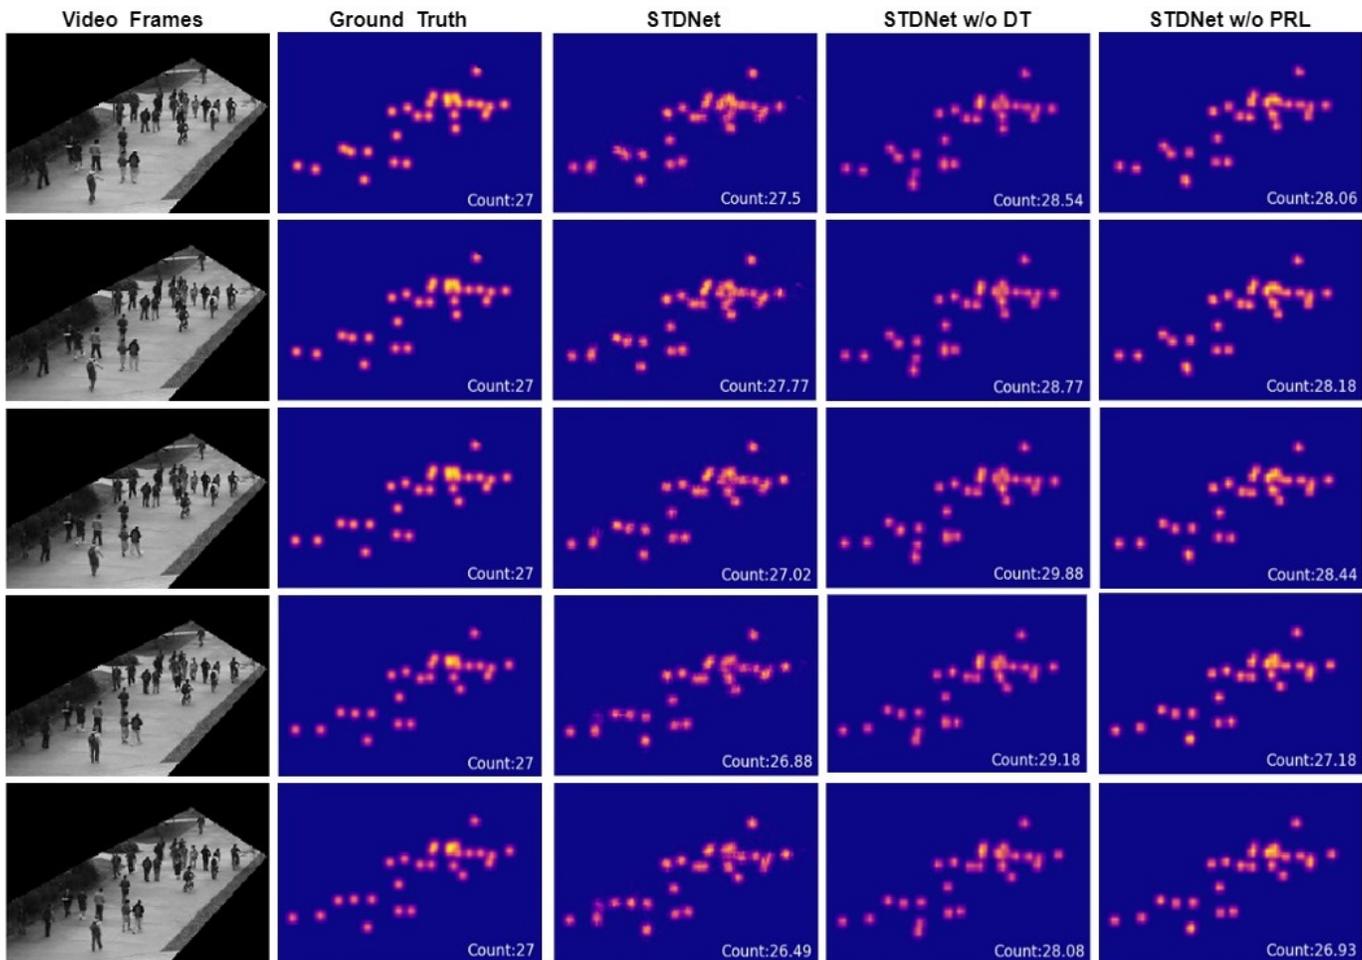


Fig. 8. Visual comparison for the ablation studies on consecutive frames of UCSD dataset.

dataset. Table IV shows the MAE and MSE with 6 different configurations. When DS or DT is removed, we use the regular convolutional layer in all the DSBs or DTBs instead of the dilated convolutional layer. Moreover, when PRL is removed, we use the original pixel-wise loss as the objective function.

The results show that MAE and MSE are 1.25 and 1.63, respectively, if we remove all the proposed components (the first row of results), while MAE and MSE are improved by 39.2% and 38.0%, respectively, if we integrate the four proposed methods into our model (the last row). Moreover, any removal of the proposed module increases the error, which indicates that all four proposed modules are important for video-based crowd counting.

Fig. 8 visualizes the qualitative results of the ablation study, i.e., STDNet without the dilated convolution in temporal domain (DT) or without the proposed PRL. The results show that the temporal dilated convolution and the PRL significantly improves the model. This is because the dilated convolution on the temporal domain extracts multiscale temporal features, which improves the model for some complex situations. For example, if an occlusion occur in the current frame, STDNet can estimate the density map more precisely by considering the multiscale features of previous frames. Moreover, the proposed PRL alleviates the issue of imprecise labels to better train the

model.

## F. Discussion

1) *Model Size and Processing Speed:* As shown in Table V, we compare STDNet with CSRNet [21] and ConvLSTM [31] in terms of training time, number of parameters, inference speed, GFLOPS, and MAE on the UCSD dataset. The number of parameters in STDNet is 18.14 million only, and the training time per epoch is 50 seconds since we use the decomposition method and the dilated operation. In addition, the inference time of STDNet is 34 fps. Compared with the image-based CSRNet, our proposed STDNet considers temporal information but only increases the number of parameters by 11.6%. In contrast, ConvLSTM requires 150.6% more parameters<sup>6</sup> to deal with temporal information. On the other hand, although the GFLOPS of ConvLSTM is smaller than that of STDNet, the inference speed of ConvLSTM is slower than that of STDNet since the non-linear dependencies between sequence elements of LSTM prevent parallelizing inference over sequence length.

<sup>6</sup>E3D is similar to ConvLSTM but cannot be compared here due to the lack of the design details, e.g., the number of channels in each Conv3D layer.

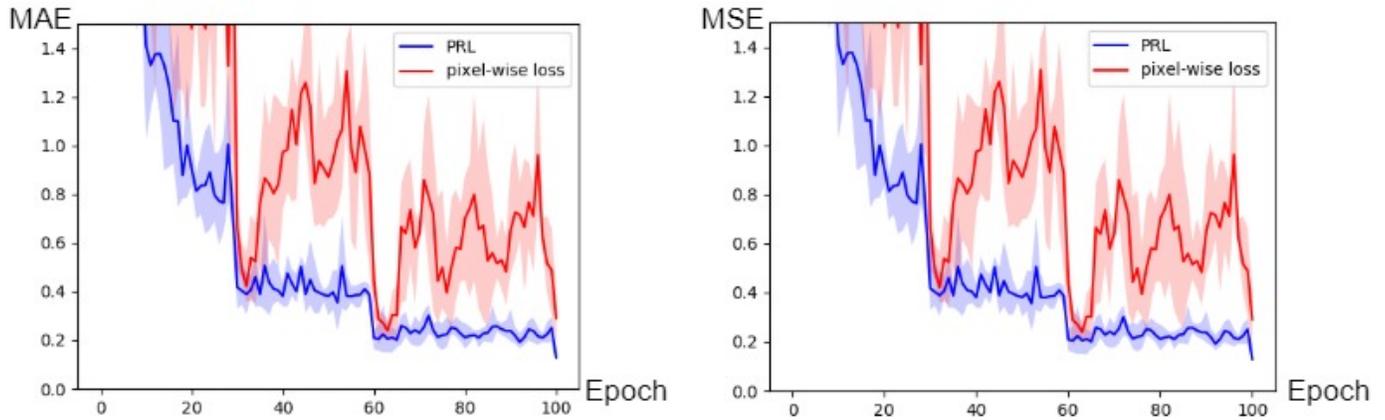


Fig. 9. The learning curve of the proposed PRL and the original pixel-wise loss in the training stage in terms of MAE and MSE. The red curve represents the pixel-wise regression loss, while the blue one is the Patch-wise Regression Loss (PRL).

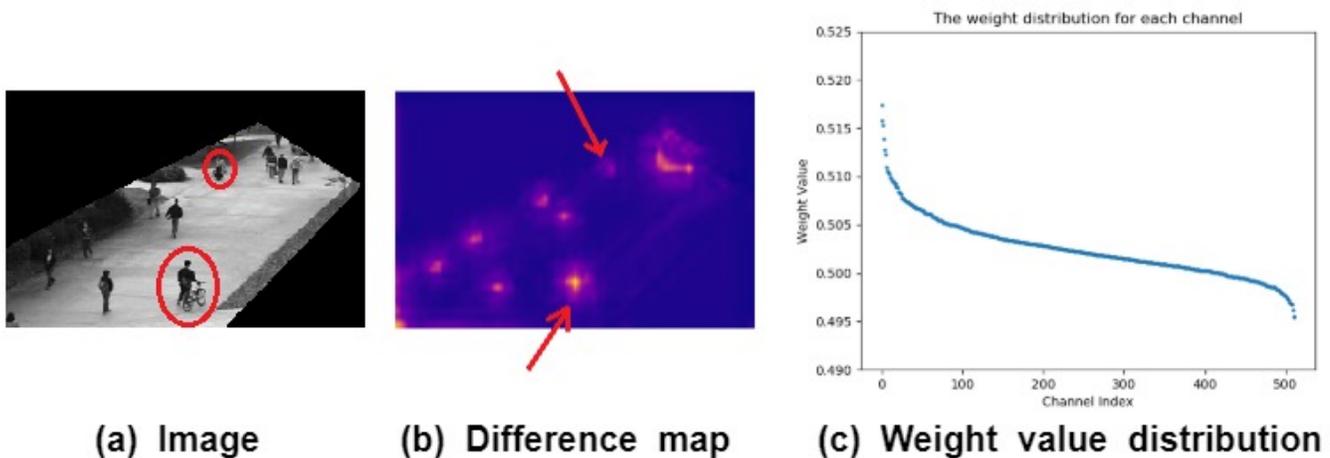


Fig. 10. Visualization of the proposed Channel Attention Block in the Dense Spatial Block (Spatial part in the proposed DSTB)

2) *Generalization and Analysis of PRL*: To demonstrate that the proposed patch-wise regression loss (PRL) is general and compatible with any existing crowd counting methods, including both image- and video-based methods, we compare the results using the original pixel-wise loss and PRL for CSRNet [21], CAN [23], and ConvLSTM [31]. As shown in Table VI, using the proposed PRL can further improve the results for both image- and video-based methods, including the state-of-the-art ones (*e.g.*, CAN [23]). Therefore, the proposed PRL is better than the original pixel-wise loss for crowd counting with imprecise annotations.

Moreover, to further analyze the PRL, Fig. 9 shows the training curve in terms of MAE and MSE for the STDNet with PRL and pixel-wise loss. It is worth noting that we repeat the experiments several times to eliminate the randomness, and depict the max, mean and min of the curve. The results show that 1) the proposed PRL (the blue curve) achieves better convergence (solid line) than the original pixel-wise loss (the red curve) for both MAE and MSE, and 2) the oscillation scale (shaded region) of the training curves for the PRL is smaller than that of the pixel-wise loss. This is because it is difficult for the model to precisely estimate the density map in a pixel-level

TABLE V  
COMPARISONS OF THE TRAINING SPEED, NUMBER OF PARAMETERS, INFERENCE TIME, GFLOPs AND MAE ON UCSD DATASET UNDER THE SAME TRAINING SETTINGS. NOTE THAT THE NOTATION “T” MEANS THE METHOD CONSIDERS THE TEMPORAL INFORMATION.

Method	T	Training time (sec./epoch)	# parameters	Inference speed (fps)	GFLOPs	MAE
CSRNet [21]		35	16.26M	48	2.71	1.16
ConvLSTM [31]	✓	530	40.61M	13	3.27	1.30
STDNet (ours)	✓	50	18.14M	34	4.23	0.76

manner and thus induce a large pixel-level difference between the ground truth density map and the predicted density. As such, pixel-wise loss induces a large loss that dominates the gradients and hinders the learning process. In contrast, Fig. 9 shows that our proposed PRL can mitigate the oscillation of the curve and achieve better convergence since we compare the sum of the density value patch-by-patch, avoiding unnecessary guidance from pixel-wise loss; in other words, in the crowd counting task, it is more proper to predict the number of people in a patch-level manner for any method.

TABLE VI

COMPARISONS OF THE PROPOSED PATCH-WISE REGRESSION LOSS (PRL) AND THE ORIGINAL PIXEL-WISE LOSS FOR THREE METHODS ON UCSD DATASET.

Method	Image- or Video-based	pixel-wise loss		PRL	
		MAE	MSE	MAE	MSE
CSRNet [21]	Image	1.16	1.47	<b>1.06</b>	<b>1.35</b>
CAN [23]	Image	1.08	1.32	<b>1.02</b>	<b>1.25</b>
ConvLSTM [31]	Video	1.30	1.79	<b>1.23</b>	<b>1.72</b>
STDNet (ours)	Video	0.85	1.11	<b>0.76</b>	<b>1.01</b>

3) *Visualization of Channel Attention Block*: To investigate the effect of channel attention, we visualize the proposed channel attention block. Figs. 10(a)-(c) represent the input image, the difference between the input and output of the channel attention block, and the weight value distribution for each channel. It is worth noting that the difference map in Fig. 10 is derived by 1) adding up the input and output feature maps over the channel dimension and 2) taking the absolute value of the difference between the summation of input feature maps and that of output feature maps. Moreover, in Fig. 10(c), we sort the weight values for each channel in descending order for better visualization. In Fig. 10(a), we can observe that the scales of human patterns are different (*e.g.*, the two red circles in the image); therefore, by using the proposed channel attention block, it can adaptively highlight the features at different scales for the different regions (*e.g.*, the regions annotated by the two red arrows in Fig. 10(b)) to help the model estimate the density map. Fig. 10(c) also shows that the channel attention block indeed differentiates the importance of different channels.

## VI. CONCLUSION

In this paper, we propose a novel spatiotemporal architecture named STDNet for the crowd counting task. STDNet effectively utilizes video data by the proposed spatiotemporal dilated convolution and channel attention block and alleviates the issues generated from the regular 3D CNN or ConvLSTM. Moreover, to handle uncertain or imprecise annotations, PRL is proposed to improve the model. Experimental results on three public datasets show that our model can achieve competitive results compared with other state-of-the-art methods, and the ablation study also demonstrates the effect of each proposed component. In the future, we plan to improve PRL by further exploiting the spatiotemporal information for video-based crowd estimation. Moreover, since the proposed STDNet is compatible with other image-based methods, we will investigate the performance of replacing the spatial convolutional layers with other state-of-the-art image-based methods.

## VII. ACKNOWLEDGEMENT

This work was supported in part by the Ministry of Science and Technology of Taiwan under Grants MOST-109-2221-E-009-114-MY3, MOST-109-2634-F-009-018, MOST-109-2221-E-001-015, MOST-108-2218-E-002-055, MOST-109-2223-E-009-002-MY3, MOST-109-2218-E-009-025 and MOST-109-2218-E-002-015. We are grateful to the National Center for High-performance Computing for computer time and facilities.

## REFERENCES

- [1] A. Basalamah, "Automatic update of crowd and traffic data using device monitoring," Jul 2016. US Patent 9,401,086.
- [2] Z.-X. Xia, W.-C. Lai, L.-W. Tsao, L.-F. Hsu, C.-C. Hu Yu, H.-H. Shuai, and W.-H. Cheng, "Human-like traffic scene understanding system: A survey," *IEEE Industrial Electronics Magazine*, 2020.
- [3] W.-C. Lai, Z.-X. Xia, H.-S. Lin, L.-F. Hsu, H.-H. Shuai, I.-H. Jhuo, and W.-H. Cheng, "Trajectory prediction in heterogeneous environment via attended ecology embedding," in *Proceedings of the ACM International Conference on Multimedia*, 2020.
- [4] S. Saxena, F. Brémond, M. Thonnat, and R. Ma, "Crowd behavior recognition for video surveillance," in *International Conference on Advanced Concepts for Intelligent Vision Systems*, pp. 970–981, Springer, Oct. 2008.
- [5] Z. Zhang, M. Wang, and X. Geng, "Crowd counting in public video surveillance by label distribution learning," *Neurocomputing*, vol. 166, pp. 151–163, Oct. 2015.
- [6] N. Ihaddadene and C. Djeraba, "Real-time crowd motion analysis," in *19th International Conference on Pattern Recognition*, pp. 1–4, IEEE, Dec. 2008.
- [7] M. Xu, C. Li, P. Lv, N. Lin, R. Hou, and B. Zhou, "An efficient method of crowd aggregation computation in public areas," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 28, pp. 2814–2825, Jul. 2017.
- [8] D.-Y. Hsieh, H.-H. Lu, R.-F. Jheng, H.-C. Chen, H.-H. Shuai, and W.-H. Cheng, "Learning to navigate from synthetic data for friction-adaptive autonomous driving," in *Proceedings of the AAAI Workshop on Artificial Intelligence of Things*, 2020.
- [9] W. Ren, D. Kang, Y. Tang, and A. B. Chan, "Fusing crowd density maps and visual object trackers for people tracking in crowd scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5353–5362, Jun. 2018.
- [10] C. H. Sio, Y.-J. Ma, H.-H. Shuai, J.-C. Chen, and W.-H. Cheng, "S2siamfc: Self-supervised fully convolutional siamese network for visual tracking," in *Proceedings of the ACM International Conference on Multimedia*, 2020.
- [11] C. H. Sio, H.-H. Shuai, and W.-H. Cheng, "Multiple fisheye camera tracking via real-time feature clustering," in *Proceedings of the ACM International Conference on Multimedia in Asia*, 2019.
- [12] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, pp. 743–761, Apr. 2011.
- [13] V. B. Subburaman, A. Descamps, and C. Carincotte, "Counting people in the crowd using a generic head detector," in *2012 IEEE Ninth International Conference on Advanced Video and Signal-Based Surveillance*, pp. 470–475, Sep. 2012.
- [14] P. Viola, M. J. Jones, and D. Snow, "Detecting pedestrians using patterns of motion and appearance," *International Journal of Computer Vision*, vol. 63, pp. 153–161, Feb. 2005.
- [15] 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*, pp. 770–778, Jun. 2016.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, pp. 1097–1105, Dec. 2012.
- [17] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proceedings of the IEEE International Conference on Learning Representations*, May 2015.
- [18] I.-C. Shen and W.-H. Cheng, "Gestalt rule feature points," *IEEE Transactions on Multimedia*, vol. 17, pp. 526–537, April 2015.
- [19] X. Cao, Z. Wang, Y. Zhao, and F. Su, "Scale aggregation network for accurate and efficient crowd counting," in *Proceedings of the European Conference on Computer Vision*, pp. 734–750, Sep. 2018.
- [20] F. Dai, H. Liu, Y. Ma, J. Cao, Q. Zhao, and Y. Zhang, "Dense scale network for crowd counting," *arXiv preprint arXiv:1906.09707*, Jun. 2019.
- [21] 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*, pp. 1091–1100, Jun. 2018.
- [22] J. Liu, C. Gao, D. Meng, and A. G. Hauptmann, "Decidenet: Counting varying density crowds through attention guided detection and density estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5197–5206, Jun. 2018.

- [23] W. Liu, M. Salzmann, and P. Fua, "Context-aware crowd counting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5099–5108, Jun. 2019.
- [24] Y. Liu, M. Shi, Q. Zhao, and X. Wang, "Point in, box out: Beyond counting persons in crowds," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6469–6478, Jun. 2019.
- [25] V. Ranjan, H. Le, and M. Hoai, "Iterative crowd counting," in *Proceedings of the European Conference on Computer Vision*, pp. 270–285, Sep. 2018.
- [26] D. B. Sam, S. Surya, and R. V. Babu, "Switching convolutional neural network for crowd counting," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4031–4039, IEEE, Jul. 2017.
- [27] V. A. Sindagi and V. M. Patel, "Generating high-quality crowd density maps using contextual pyramid cnns," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1861–1870, Oct. 2017.
- [28] Y. Tian, Y. Lei, J. Zhang, and J. Z. Wang, "Padnet: Pan-density crowd counting," *IEEE Transactions on Image Processing*, vol. 29, pp. 2714–2727, Nov. 2019.
- [29] 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*, pp. 589–597, Jun. 2016.
- [30] Z. Zou, H. Shao, X. Qu, W. Wei, and P. Zhou, "Enhanced 3d convolutional networks for crowd counting," in *British Machine Vision Conference (BMVC)*, Sep. 2019.
- [31] F. Xiong, X. Shi, and D.-Y. Yeung, "Spatiotemporal modeling for crowd counting in videos," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5151–5159, Oct. 2017.
- [32] Y. Fang, B. Zhan, W. Cai, S. Gao, and B. Hu, "Locality-constrained spatial transformer network for video crowd counting," in *IEEE International Conference on Multimedia and Expo*, pp. 814–819, IEEE, 2019.
- [33] Q. Wu, C. Zhang, X. Kong, M. Zhao, and Y. Chen, "Triple attention for robust video crowd counting," in *IEEE International Conference on Image Processing*, pp. 1966–1970, IEEE, 2020.
- [34] Z. Qiu, T. Yao, and T. Mei, "Learning spatio-temporal representation with pseudo-3d residual networks," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5533–5541, Oct. 2017.
- [35] A. B. Chan, Z.-S. J. Liang, and N. Vasconcelos, "Privacy preserving crowd monitoring: Counting people without people models or tracking," in *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–7, IEEE, Jun. 2008.
- [36] K. Chen, C. C. Loy, S. Gong, and T. Xiang, "Feature mining for localised crowd counting," in *British Machine Vision Conference (BMVC)*, vol. 1, p. 3, Sep. 2012.
- [37] 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, pp. 1048–1061, Mar. 2016.
- [38] 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*, pp. 833–841, Jun. 2015.
- [39] G. J. Brostow and R. Cipolla, "Unsupervised bayesian detection of independent motion in crowds," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 594–601, Jun. 2006.
- [40] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 886–893, IEEE, Jun. 2005.
- [41] 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*, pp. 2547–2554, Jun. 2013.
- [42] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2, pp. 1150–1157, Sep. 1999.
- [43] V. Rabaud and S. Belongie, "Counting crowded moving objects," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 705–711, Jun. 2006.
- [44] J. Sanchez-Riera, K.-L. Hua, Y.-S. Hsiao, T. Lim, S. C. Hidayati, and W.-H. Cheng, "A comparative study of data fusion for rgb-d based visual recognition," *Pattern Recognition Letters*, vol. 73, pp. 1–6, 2016.
- [45] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems*, pp. 91–99, Dec. 2015.
- [46] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *Proceedings of the European Conference on Computer Vision*, pp. 21–37, Oct. 2016.
- [47] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 779–788, Jun. 2016.
- [48] A. B. Chan and N. Vasconcelos, "Bayesian poisson regression for crowd counting," in *2009 IEEE 12th International Conference on Computer Vision*, pp. 545–551, IEEE, Sep. 2009.
- [49] S.-Y. Cho, T. W. Chow, and C.-T. Leung, "A neural-based crowd estimation by hybrid global learning algorithm," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 29, pp. 535–541, Aug. 1999.
- [50] V. Lempitsky and A. Zisserman, "Learning to count objects in images," in *Neural Information Processing Systems*, Dec. 2010.
- [51] X. Jiang, Z. Xiao, B. Zhang, X. Zhen, X. Cao, D. Doermann, and L. Shao, "Crowd counting and density estimation by trellis encoder-decoder networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6133–6142, Jun. 2019.
- [52] 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*, pp. 7279–7288, Jun. 2019.
- [53] Z. Yan, Y. Yuan, W. Zuo, X. Tan, Y. Wang, S. Wen, and E. Ding, "Perspective-guided convolution networks for crowd counting," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 952–961, Oct. 2019.
- [54] V. A. Sindagi and V. M. Patel, "Ha-ccn: Hierarchical attention-based crowd counting network," *IEEE Transactions on Image Processing*, vol. 29, pp. 323–335, 2020.
- [55] Y.-X. Lin, D. S. Tan, W.-H. Cheng, Y.-Y. Chen, and K.-L. Hua, "Spatially-aware domain adaptation for semantic segmentation of urban scenes," in *IEEE International Conference on Image Processing*, 2019.
- [56] Y.-X. Lin, D. S. Tan, W.-H. Cheng, and K.-L. Hua, "Adapting semantic segmentation of urban scenes via mask-aware gated discriminator," in *IEEE International Conference on Multimedia & Expo*, 2019.
- [57] Z.-Q. Cheng, J.-X. Li, Q. Dai, X. Wu, and A. G. Hauptmann, "Learning spatial awareness to improve crowd counting," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 6152–6161, Oct. 2019.
- [58] X. Liu, J. Van De Weijer, and A. D. Bagdanov, "Leveraging unlabeled data for crowd counting by learning to rank," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7661–7669, Jun. 2018.
- [59] 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*, pp. 6142–6151, Oct. 2019.
- [60] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7132–7141, 2018.
- [61] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proceedings of the IEEE International Conference on Learning Representations*, May 2015.
- [62] Y. Yang, G. Li, Z. Wu, L. Su, Q. Huang, and N. Sebe, "Reverse perspective network for perspective-aware object counting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4374–4383, Jun. 2020.
- [63] N. Liu, Y. Long, C. Zou, Q. Niu, L. Pan, and H. Wu, "Adcrowdnet: An attention-injective deformable convolutional network for crowd understanding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3225–3234, Jun. 2019.
- [64] S. Kumagai, K. Hotta, and T. Kurita, "Mixture of counting cnns," *Machine Vision and Applications*, vol. 29, pp. 1119–1126, July 2018.
- [65] L. Liu, H. Wang, G. Li, W. Ouyang, and L. Lin, "Crowd counting using deep recurrent spatial-aware network," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, Jul. 2018.
- [66] E. Walach and L. Wolf, "Learning to count with cnn boosting," in *Proceedings of the European Conference on Computer Vision*, pp. 660–676, Springer, Oct. 2016.
- [67] Z. Zou, Y. Liu, S. Xu, W. Wei, S. Wen, and P. Zhou, "Crowd counting via hierarchical scale recalibration network," in *European Conference on Artificial Intelligence (ECAI)*, Aug. 2020.
- [68] D. B. Sam, S. V. Peri, M. N. Sundararaman, A. Kamath, and V. B. Radhakrishnan, "Locate, size and count: Accurately resolving people in

dense crowds via detection,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

- [69] Y. Miao, Z. Lin, G. Ding, and J. Han, “Shallow feature based dense attention network for crowd counting,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, pp. 11765–11772, Feb. 2020.



**Yu-Jen Ma** received the B.S. degrees from the Department of Electrical Engineering, National Sun Yat-Sen University (NSYSU), Kaohsiung, Taiwan, R.O.C., in 2019. His research interests include artificial intelligence, deep learning, and computer vision.



**Hong-Han Shuai** received the B.S. degree from the Department of Electrical Engineering, National Taiwan University (NTU), Taipei, Taiwan, R.O.C., in 2007, the M.S. degree in computer science from NTU in 2009, and the Ph.D. degree from Graduate Institute of Communication Engineering, NTU, in 2015. He is now an assistant professor in NCTU. His research interests are in the area of multimedia processing, machine learning, social network analysis, and data mining. His works have appeared in top-tier conferences such as MM, CVPR, AAAI,

KDD, WWW, ICDM, CIKM and VLDB, and top-tier journals such as TKDE, TMM and JIOT. Moreover, he has served as the PC member for international conferences including MM, AAAI, IJCAI, WWW, and the invited reviewer for journals including TKDE, TMM, JVCJ and JIOT.



**Wen-Huang Cheng** is Professor with the Institute of Electronics, National Chiao Tung University (NCTU), Hsinchu, Taiwan. He is also Jointly Appointed Professor with the Artificial Intelligence and Data Science Program, National Chung Hsing University (NCHU), Taichung, Taiwan. Before joining NCTU, he led the Multimedia Computing Research Group at the Research Center for Information Technology Innovation (CITI), Academia Sinica, Taipei, Taiwan, from 2010 to 2018. His current research interests include multimedia, artificial intelligence,

computer vision, and machine learning. He has actively participated in international events and played important leading roles in prestigious journals and conferences and professional organizations, like Associate Editor for *IEEE Transactions on Multimedia*, General co-chair for *IEEE ICME* (2022) and *ACM ICMR* (2021), Chair-Elect for *IEEE MSA* technical committee, governing board member for *IAPR*. He has received numerous research and service awards, including the 2018 *MSRA Collaborative Research Award*, the 2017 *Ta-Yu Wu Memorial Award* from Taiwan’s Ministry of Science and Technology (the highest national research honor for young Taiwanese researchers under age 42), the 2017 *Significant Research Achievements of Academia Sinica*, the *Top 10% Paper Award* from the 2015 *IEEE MMSP*, and the *K. T. Li Young Researcher Award* from the *ACM Taipei/Taiwan Chapter* in 2014. He is *IET Fellow* and *ACM Distinguished Member*.