

Learning to Count in the Crowd from Limited Labeled Data

Vishwanath A. Sindagi¹, Rajeev Yasarla¹, Deepak Sam Babu², R. Venkatesh Babu², and Vishal M. Patel¹

¹ Johns Hopkins University, Baltimore MD 21218, USA

² Indian Institute of Science, Bangalore 560012, India

{vishwanathsindagi,ryasar11,vpatel36}@jhu.edu

{deepaksam,venky}@iisc.ac.in

Abstract. Recent crowd counting approaches have achieved excellent performance. However, they are essentially based on fully supervised paradigm and require large number of annotated samples. Obtaining annotations is an expensive and labour-intensive process. In this work, we focus on reducing the annotation efforts by learning to count in the crowd from limited number of labeled samples while leveraging a large pool of unlabeled data. Specifically, we propose a Gaussian Process-based iterative learning mechanism that involves estimation of pseudo-ground truth for the unlabeled data, which is then used as supervision for training the network. The proposed method is shown to be effective under the reduced data (semi-supervised) settings for several datasets like ShanghaiTech, UCF-QNRF, WorldExpo, UCSD, *etc.* Furthermore, we demonstrate that the proposed method can be leveraged to enable the network in learning to count from synthetic dataset while being able to generalize better to real-world datasets (synthetic-to-real transfer).

Keywords: Crowd counting, semi-supervised learning, pseudo-labeling, domain adaptation, synthetic to real transfer

1 Introduction

Due to its significance in several applications (like video surveillance [12,50,44], public safety monitoring [58], microscopic cell counting [15], environmental studies [23], *etc.*), crowd counting has attracted a lot of interest from the deep learning research community. Several convolutional neural network (CNN) based approaches have been developed that address various issues in counting like scale variations, occlusion, background clutter [17,59,18,36,42,22,39,37,3,28,19,43,33,34,2], *etc.* While these methods have achieved excellent improvements in terms of the overall error rate, they follow a fully-supervised paradigm and require several labeled data samples. There is a wide variety of scenes and crowded scenarios that these networks need to handle to in the real world. Due to a distribution gap between the training and testing environments, these networks have limited generalization abilities and hence, procuring annotations becomes especially important.

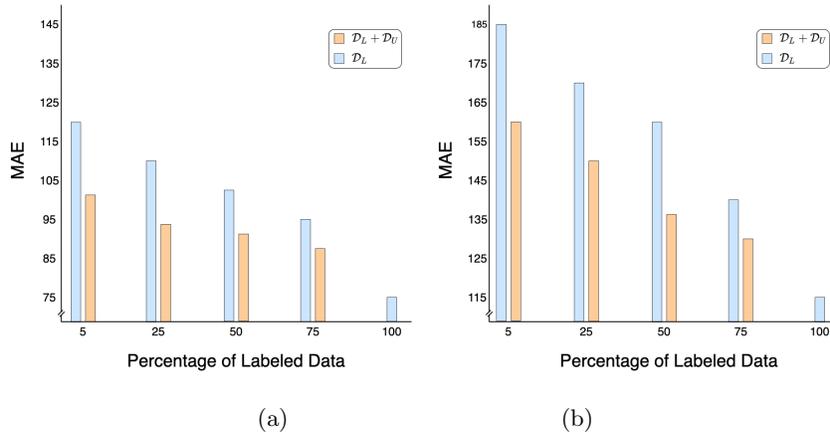


Fig. 1. Results of semi-supervised learning experiments. (a) ShanghaiTech A (b) UCF-QNRF. For both datasets, the error increases with reduction in the %-age of labeled data. By leveraging the unlabeled dataset using the proposed GP-based framework, we are able to reduce the error considerably. Note that \mathcal{D}_L and \mathcal{D}_U indicate labeled and unlabeled dataset, respectively.

However, annotating data for crowd counting typically involves obtaining point-wise annotations at head locations, and this is a labour intensive and expensive process. Hence, it is infeasible to procure annotations for all possible scenarios. Considering this, it is crucial to reduce the annotation efforts, especially for crowd counting methods which get deployed in a wide variety of scenarios.

With the exception of a few works [6,22,?], reducing annotation efforts while maintaining good performance is relatively less explored for the task of crowd counting. Hence, in this work, we focus on learning to count using limited labeled data while leveraging unlabeled data to improve the performance. Specifically, we propose a Gaussian Process (GP) based iterative learning framework where we augment the existing networks with capabilities to leverage unlabeled data, thereby resulting in overall improvement in the performance. Inspired by [57], the proposed framework follows a pseudo-labeling approach, where we estimate the pseudo-ground truth (pseudo-GT) for the unlabeled data, which is then used to supervise the network. The network is trained iteratively on labeled and unlabeled data. In the labeled stage, the network weights are updated by minimizing the L_2 error between predictions and the ground-truth (GT) for the labeled data. In addition, we save the latent space vectors of the labeled data along with the ground-truths. In the unlabeled stage, we first model the relationship between the latent space vectors of the labeled images along with the corresponding ground-truth and unlabeled latent space vectors jointly using GP. Next, we estimate the pseudo-GT for the unlabeled inputs using the GP modeled earlier. This pseudo-GT is then used to supervise the network for the unlabeled data. Minimizing the error between the unlabeled data predictions and the pseudo-GT results in improved performance. Fig. 1 illustrates the effectiveness

of the proposed GP-based framework in exploiting unlabeled data on two datasets (ShanghaiTech-A [61] and UCF-QNRF[10]) in the reduced data setting. It can be observed that the proposed method is able to leverage unlabeled data effectively resulting in lower error across various settings.

The proposed method is evaluated on different datasets like ShanghaiTech [61], UCF-QNRF [10], WorldExpo [59], UCSD [4], *etc.* in the reduced data settings. In addition to obtaining lower error as compared to the existing methods [22], the performance drop due to less data is improved by a considerable margin. Furthermore, the proposed method is effective for learning to count from synthetic data as well. More specifically, we use labeled synthetic crowd counting dataset (GCC [55]) and unlabeled real-world datasets (ShanghaiTech [61], UCF-QNRF [10], WorldExpo [59], UCSD [5]) in our framework, and show that it is able to generalize better to real-world datasets as compared to recent domain adaptive crowd counting approaches [55]. To summarize, the following are our contributions:

- We propose a GP-based framework to effectively exploit unlabeled data during the training process, resulting in improved overall performance. The proposed method consists of iteratively training over labeled and unlabeled data. For the unlabeled data, we estimate the pseudo-GT using the GP modeled during labeled phase.
- We demonstrate that the proposed framework is effective in semi-supervised and synthetic-to-real transfer settings. Through various ablation studies, we show that the proposed method is generalizable to different network architectures and various reduced data settings.

2 Related Work

Crowd Counting. Traditional approaches in crowd counting ([16,31,7,9,15,27,56]) typically involved feature extraction techniques and training regression algorithms. Recently, CNN-based approaches like [54,59,36,1,51,26,61,36,42] have surpassed the traditional approaches by a large margin in terms of the overall error rate. Most of these methods focus on addressing the issue of large variations in scales. Approaches like [61,36,42] focus on improving the receptive field. Different from these, approaches like [28,41,47,32] focus on effective ways of fusing multi-scale information from deep networks. In addition to scale variation, recent approaches have addressed other issues in crowd counting like improving the quality of predicted density maps using adversarial regularization [42,?], use of deep negative correlation-based learning for obtaining more generalizable features, and scale-based feature aggregation [3]. Most recently, several methods have employed additional information like segmentation and semantic priors [62,53], attention [20,45,46], perspective [38], context information [21], multiple-views [60] and multi-scale features [11], adaptive density maps [52] into the network. In other efforts, researchers have made important contributions by creating large-scale datasets for counting like UCF-QNRF [10], GCC [55] and JHU-CROWD [48,49].

For a more detailed discussion on these methods, the reader is referred to recent comprehensive surveys [43,8].

Learning from limited data. Recent research in crowd counting has been largely focused on improving the counting performance in the fully-supervised paradigm. Very few works like [6,?,55] have made efforts on minimizing annotation efforts. Loy *et al.*[6] proposed a semi-supervised regression framework that exploit underlying geometric structures of crowd patterns to assimilate the count estimation of two nearby crowd pattern points in the manifold. However, this approach is specifically designed for video-based crowd counting.

Recently, Liu *et al.*[22] proposed to leverage additional unlabeled data for counting by introducing a learning to rank framework. They assume that any sub-image of a crowded scene image is guaranteed to contain the same number or fewer persons than the super-image. They employ pairwise ranking hinge loss to enforce this ranking constraint for unlabeled data in addition to the L_2 error to train the network. In our experiments we observed that this constraint is almost always satisfied, and it provides relatively less supervision over unlabeled data.

Babu *et al.*[35] focus on a different approach, where they train 99.9% of their parameters from unlabeled data using a novel unsupervised learning framework based on winner-takes-all (WTA) strategy. However, they still train the remaining set of parameters using labeled data.

Wang *et al.*[55] take a totally different approach to minimize annotation efforts by creating a new synthetic crowd counting dataset (GCC). Additionally, they propose a Cycle-GAN based domain adaptive approach for generalizing the network trained on synthetic dataset to real-world dataset. However, there is a large gap in terms of the style and also the crowd count between the synthetic and real-world scenarios. Domain adaptive approaches have limited abilities in handling such scenarios. In order to obtain successful adaptation, the authors in [55] manually select the samples from the synthetic dataset that are closer to the real-world scenario in terms of crowd count for training the network. This selection is possible when one has information about the count from the real-world datasets, which violates the assumption of lack of unlabeled data in the target domain for unsupervised domain adaptation.

Considering the drawbacks of existing approaches, we propose a new GP-based iterative training framework to exploit unlabeled data.

3 Preliminaries

In this section, we briefly review the concepts (crowd counting, semi-supervised learning and Gaussian Process) that are used in this work.

Crowd counting. Following recent works [59,61], we employ the approach of density estimation technique. That is, an input crowd image is forwarded through the network, and the network outputs a density map. This density map indicates the per-pixel count of people in the image. The count in the image is obtained by integrating over the density map. For training the network using labeled data,

the ground-truth density maps are obtained by imposing 2D Gaussians at head location x_g using $D(x) = \sum_{x_g \in S} \mathcal{N}(x - x_g, \sigma)$. Here, σ is the Gaussian kernel’s scale and S is the list of all locations of people.

Problem formulation. We are given a set of labeled dataset of input-GT pairs $(\{x, y\} \in \mathcal{D}_{\mathcal{L}})$ and a set of unlabeled input data samples $x \in \mathcal{D}_{\mathcal{U}}$. The objective is to fit a mapping-function $f(x|\phi)$ (with parameters defined by ϕ) that accurately estimates target label y for unobserved samples. Note that this definition applies to both semi-supervised setting and synthetic-to-real transfer setting. In the case of synthetic-to-real transfer, the synthetic dataset is labeled and hence, can be used as the labeled dataset ($\mathcal{D}_{\mathcal{L}}$). Similarly, the real-world dataset is unlabeled and can be used as the unlabeled dataset ($\mathcal{D}_{\mathcal{U}}$).

In order to learn the parameters, both labeled and unlabeled datasets are exploited. Typically, loss functions such as L_1 , L_2 or cross entropy error are used for labeled data. For exploiting unlabeled data $\mathcal{D}_{\mathcal{U}}$, existing approaches augment $f(x|\phi)$ with information like shape of the data manifold [25] via different techniques such as enforcing consistent regularization [13], virtual adversarial training [24] or pseudo-labeling [14]. In this work, we employ pseudo-labeling based approach where we estimate pseudo-GT for unlabeled data, and then use them for supervising the network using traditional supervised loss functions.

Gaussian process. A Gaussian process (GP) $f(v)$ is an infinite collection of random variables, any finite subset of which have a joint Gaussian distribution. A GP is fully specified by its mean function ($m(v)$) and covariance function $K(v, v')$. These are defined below:

$$m(v) = \mathbb{E}[f(v)], \quad (1)$$

$$K(v, v') = \mathbb{E}[(f(v) - m(v))(f(v') - m(v'))], \quad (2)$$

where $v, v' \in \mathcal{V}$ denote the possible inputs that index the GP. The covariance matrix is computed from the covariance function K which expresses the notion of smoothness of the underlying function. GP can then be formulated as follows:

$$f(v) \sim \mathcal{GP}(m(v), K(v, v') + \sigma_{\epsilon}^2 I), \quad (3)$$

where I is identity matrix and σ_{ϵ}^2 is the variance of the additive noise. Any collection of function values is then jointly Gaussian as follows

$$f(V) = [f(v_1), \dots, f(v_n)]^T \sim \mathcal{N}(\mu, K(V, V') + \sigma_{\epsilon}^2 I), \quad (4)$$

with mean vector and covariance matrix defined by the GP as mentioned earlier. To make predictions at unlabeled points, one can compute a Gaussian posterior distribution in closed form by conditioning on the observed data. For more details, we refer the reader to [29].

4 GP-based iterative learning

Fig. 2 gives an overview of the proposed method. The network is constructed using an encoder $f_e(x, \phi_e)$ and a decoder $f_d(z, \phi_d)$, that are parameterized by ϕ_e and ϕ_d , respectively. The proposed framework is agnostic to the encoder network, and we show in the experiments section that it generalizes well to architectures such as VGG16 [40], ResNet-50 and ResNet-101 [30]. The decoder consists of a set of 2 conv-relu layers (see supplementary material for more details). Typically, an input crowd image x is forwarded through the encoder network to obtain the corresponding latent space vector z . This vector is then forwarded through the decoder network to obtain the crowd density output y , *i.e.*, $y = f_d(f_e(x, \phi_e), \phi_d)$.

We are given a training dataset, $\mathcal{D} = \mathcal{D}_{\mathcal{L}} \cup \mathcal{D}_{\mathcal{U}}$, where $\mathcal{D}_{\mathcal{L}} = \{x_l^i, y_l^i\}_{i=1}^{N_l}$ is a labeled dataset containing N_l training samples and $\mathcal{D}_{\mathcal{U}} = \{x_u^i\}_{i=1}^{N_u}$ is an unlabeled dataset containing N_u training samples. The proposed framework effectively leverages both the datasets by iterating the training process over labeled $\mathcal{D}_{\mathcal{L}}$ and unlabeled datasets $\mathcal{D}_{\mathcal{U}}$. More specifically, the training process consists of two stages: (i) Labeled training stage: In this stage, we employ supervised loss function \mathcal{L}_s to learn the network parameters using labeled dataset, and (ii) Unlabeled training stage: We generate pseudo GTs for the unlabeled data points using the GP formulation, which is then used for supervising the network on the unlabeled dataset. In what follows, we describe these stages in detail.

4.1 Labeled stage

Since the labeled dataset $\mathcal{D}_{\mathcal{L}}$ comes with annotations, we employ L_2 error between the predictions and the GTs as supervision loss for training the network. This loss objective is defined as follows:

$$\mathcal{L}_s = \mathcal{L}_2 = \|y_l^{pred} - y_l\|_2, \quad (5)$$

where $y_l^{pred} = g(z_l, \phi_d)$ is the predicted output, y_l is the ground-truth, $z = h(x, \phi_e)$ is the intermediate latent space vector. Note that, the subscript l in the above quantities indicate that these are defined for labeled data.

Along with performing supervision on the labeled data, we additionally save feature vectors z_l^i 's from the intermediate latent space in a matrix F_{z_l} . Specifically, $F_{z_l} = \{z_l^i\}_{i=1}^{N_l}$. This matrix is used for computing the pseudo-GTs for unlabeled data at a later stage. The dimension of F_{z_l} matrix is $N_l \times M$. Here, M is the dimension of the latent space vector z_l . In our case, the latent space vector dimension is $64 \times 32 \times 32$ (see supplementary material for more details), which is reshaped to $1 \times 65,536$. Hence, $M = 65,536$.

4.2 Unlabeled stage

Since the unlabeled data $\mathcal{D}_{\mathcal{U}}$ does not come with any GT annotations, we estimate pseudo-GTs which are then used as supervision for training the network on unlabeled data. For this purpose, we model the relationship between the latent

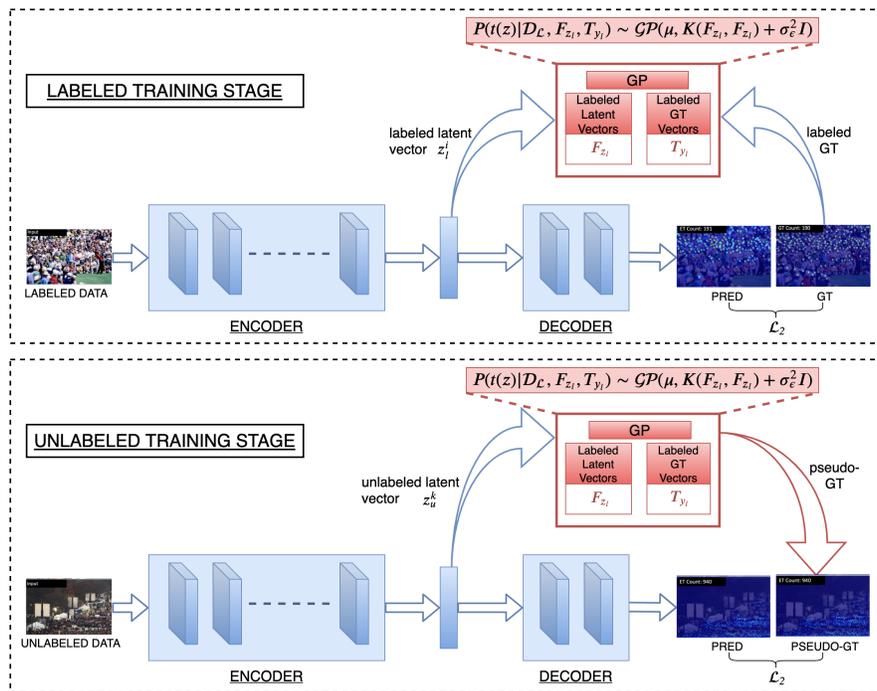


Fig. 2. Illustration of the proposed framework. Training is performed iteratively over labeled and unlabeled data. For labeled data, we minimize the L_2 error between the predictions and GT. For unlabeled data, we minimize the L_2 error between the predictions and pseudo-GT.

space vectors of the labeled images F_{z_l} along with the corresponding GT T_{y_l} and unlabeled latent space vectors z_u^{pred} jointly using GP.

Estimation of pseudo-GT: As discussed earlier, the training process iterates over labeled $\mathcal{D}_{\mathcal{L}}$ and unlabeled data $\mathcal{D}_{\mathcal{U}}$. After the labeled stage, the labeled latent space vectors F_{z_l} and their corresponding GT density maps T_{y_l} are used to model the function t which maps the relationship between the latent vectors and the output density maps as, $y = t(z)$. Using GP, we model this function $t(\cdot)$ as an infinite collection of functions of which any finite subset is jointly Gaussian. More specifically, we jointly model the distribution of the function values $t(\cdot)$ of the latent space vectors of the labeled and the unlabeled samples using GP as follows:

$$P(t(z)|\mathcal{D}_{\mathcal{L}}, F_{z_l}, T_{y_l}) \sim \mathcal{GP}(\mu, K(F_{z_l}, F_{z_l}) + \sigma_{\epsilon}^2 I), \quad (6)$$

where μ is the function value computed using GP, σ_{ϵ}^2 is set equal to 1, and K is the kernel function. Based on this, the conditional joint distribution for the latent space vector z_u^k of the k^{th} unlabeled sample x_u^k can be expressed as the

following Gaussian distribution:

$$P(t(z_u^k)|\mathcal{D}_{\mathcal{L}}, F_{z_l}, T_{z_l}) = \mathcal{N}(\mu_u^k, \Sigma_u^k), \quad (7)$$

where

$$\mu_u^k = K(z_u^k, F_{z_l})[K(F_{z_l}, F_{z_l}) + \sigma_\epsilon^2 I]^{-1} T_{y_l}, \quad (8)$$

$$\Sigma_u^k = K(z_u^k, z_u^k) - K(z_u^k, F_{z_l})[K(F_{z_l}, F_{z_l}) + \sigma_\epsilon^2 I]^{-1} K(F_{z_l}, z_u^k) + \sigma_\epsilon^2 \quad (9)$$

where σ_ϵ^2 is set equal to 1 and K is a kernel function with the following definition:

$$K(Z, Z)_{k,i} = \kappa(z_u^k, z_l^i) = \frac{\langle z_u^k, z_l^i \rangle}{|z_u^k| \cdot |z_l^i|}. \quad (10)$$

Considering the large dimensionality of the latent space vector, $K(F_{z_l}, F_{z_l})$ can grow quickly in size especially if the number of labeled data samples N_l is high. In such cases, the computational and memory requirements become prohibitively high. Additionally, all the latent vectors may not be necessarily effective since these vectors correspond to different regions of images in terms of content and size/density of the crowd. In order to overcome these issues, we use only those labeled vectors that are similar to the unlabeled latent vector. Specifically, we consider only N_n nearest labeled vectors corresponding to an unlabeled vector. That is, we replace F_{z_l} by $F_{z_l, n}$ in Eq. (7)-(9). Here $F_{z_l, n} = \{z_l^j : z_l^j \in \text{nearest}(z_u^k, F_{z_l}, N_n)\}$, and $T_{y_l, n} = \{y_l^j : z_l^j \in \text{nearest}(z_u^k, F_{z_l}, N_n)\}$ with $\text{nearest}(p, Q, N_n)$ being a function that finds top N_n nearest neighbors of p in Q .

The pseudo-GT for unlabeled data sample is given by the mean predicted in Eq. (8), *i.e.*, $y_{u, pseudo}^k = \mu_u^k$. The L_2 distance between the predictions $y_{u, pred}^k = g(z_u^k, \phi_e)$ and the pseudo-GT $y_{u, pseudo}^k$ is used as supervision for updating the parameters of the encoder $f_e(\cdot, \phi_e)$ and the decoder $f_d(\cdot, \phi_d)$.

Furthermore, the pseudo-GT estimated using Eq. (8) may not be necessarily perfect. Errors in pseudo-GT will limit the performance of the network. To overcome this, we explicitly exploit the variance modeled by the GP. Specifically, we minimize the predictive variance by considering Eq. (9) in the loss function. As discussed earlier, using all the latent space vectors of labeled data may not be necessarily effective. Hence, we minimize the variance $\Sigma_{u, n}^k$ computed between z_u^k and the N_n nearest neighbors in the latent space vectors using GP. Thus, the loss function during the unlabeled stage is defined as:

$$\mathcal{L}_{un} = \frac{1}{|\Sigma_{u, n}^k|} \|y_{u, pred}^k - y_{u, pseudo}^k\|_2 + \log \Sigma_{u, n}^k, \quad (11)$$

where $y_{u, pred}^k$ is the crowd density map prediction obtained by forwarding an unlabeled input image x_u^k through the network, $y_{u, pseudo}^k = \mu_u^k$ is the pseudo-GT (see Eq. (8)), and $\Sigma_{u, n}^k$ is the predictive variance obtained by replacing F_{z_l} in Eq. (9) with $F_{z_l, n}$. Note that the prediction error (the first term) is scaled by loss by inverse of the variance. This ensures that the loss from uncertain pseudo-gts are down-weighted and hence, only accurate pseudo-gts are used for training.

4.3 Final objective function

We combine the supervised loss Eq. (5) and unsupervised loss Eq. (11) to obtain the final objective function as follows:

$$\mathcal{L}_f = \mathcal{L}_s + \lambda_{un}\mathcal{L}_{un}, \quad (12)$$

where λ_{un} is a hyper-parameter that weighs the unsupervised loss.

5 Experiments and results

In this section, we discuss the details of the various experiments conducted to demonstrate the effectiveness of the proposed method. Since the proposed method is able to leverage unlabeled data to improve the overall performance, we performed evaluation in two settings: (i) *Semi-supervised settings*: In this setting, we varied the percentage of labeled samples from 5% to 75%. We first show that with the base network, there is performance drop due to the reduced data. Later, we show that the proposed method is able to recover a major percentage of the performance drop. (ii) *Synthetic-to-real transfer settings*: In this setting, the goal is to train on synthetic dataset (labeled), while adapting to real-world dataset. Unlabeled images from the real-world are available during training. In both settings, the proposed method is able to achieve better results as compared to recent methods. Details of the datasets are provided in the supplementary material.

5.1 Semi-supervised settings

In this section, we conduct experiments in the semi-supervised settings by reducing the amount of labeled data available during training. The rest of the samples in the dataset are considered as unlabeled samples wherever applicable. In the following sub-sections, we present comparison of the proposed method in the 5% setting with other recent methods. For comparison, we used 4 datasets: ShanghaiTech (SH-A/B)[61], UCF-QNRF [10], WorldExpo [59] and UCSD [4]. This is followed by a detailed ablation study involving different architectures and various percentages of labeled data used during training. For ablation, we chose ShanghaiTech-A and UCF-QNRF datasets since they contain a wide diversity of scenes and large variation in count and scales.

Implementation details. We train the network using Adam optimizer with a learning rate of $10e - 5$ and a momentum of 0.9 on an NVIDIA Titan Xp GPU. We use batch size of 24. During training, random crops of size 256×256 are used. During inference, the entire image is forwarded through the network. For evaluation, we use mean absolute error (MAE) and mean squared error (MSE) metrics, which are defined as: $MAE = \frac{1}{N} \sum_{i=1}^N |y_i - y'_i|$ and $MSE = \sqrt{\frac{1}{N} \sum_{i=1}^N |y_i - y'_i|^2}$, respectively. Here, N is the total number of test images, y_i is the ground-truth/target count of people in the image and y'_i is the predicted count of people

Table 1. Comparison of results in SSL settings. Reducing labeled data to 5% results in performance drop by a big margin as compared to 100% data. ResNet-50 was used as the encoder network for all the methods. RL: Ranking-Loss. GP: Gaussian-Process. AG: Average Gain %³.

Method	\mathcal{D}_L	\mathcal{D}_U	SH-A			SH-B			UCF-QNRF			WExpo		UCSD		
			MAE	MSE	AG	MAE	MSE	AG	MAE	MSE	AG	MAE	AG	MAE	MSE	AG
ResNet-50 (Oracle)	100%	-	76	126	-	8.4	14.5	-	114	195	-	10.1	-	1.7	2.1	-
ResNet-50 (\mathcal{D}_L -only)	5%	-	118	211	-	21.2	34.2	-	186	295	-	14.2	-	2.2	2.8	-
ResNet-50+RL	5%	95%	115	208	2.0	20.1	32.9	4.0	182	291	1.7	14.0	0.01	2.2	2.8	0
ResNet-50+GP(Ours)	5%	95%	102	172	16	15.7	27.9	22	160	275	10	12.8	10	2.0	2.4	12

in to the i^{th} image. We set aside 10% of the training set for the purpose of validation. The hyper-parameter λ_{un} was chosen based on the validation performance. More details are provided in the supplementary.

Comparison with recent approaches. Here, we compare the effectiveness of the proposed method with a recent method by Liu *et al.*[22] on 4 different datasets. In order to get a better understanding of the overall improvements, we also provide the results of the base network with (i) 100% labeled data supervision that is the oracle performance, and (ii) 5% labeled data supervision.

For all the methods (except oracle), we limited the labeled data used during training to 5% of the training dataset. Rest of the samples were used as unlabeled samples. We used ResNet-50 as the encoder network. The results of the experiments are shown in Table 1. For all the experiments that we conducted, we report the average of the results for 5 trials. The standard deviations are reported in the supplementary. We make the following observations for all the datasets: (i) Compared to using the entire dataset, reducing the labeled data during training (to 5%) leads to significant increase in error. (ii) The proposed GP-based framework is able to reduce the performance drop by a large margin. Further, the proposed method achieves an average gain (AG)³ of anywhere between 10%-22% over the \mathcal{D}_L -only baseline across all datasets. (iii) The proposed method is able to leverage the unlabeled data more effectively as compared to Liu *et al.*[22]. This is because the authors in [22] using a ranking loss on the unlabeled data which is based on the assumption that sub-image of a crowded scene is guaranteed to contain the same or fewer number of people compared to the entire image. We observed that this constraint is satisfied naturally for most of the unlabeled images, and hence it provides less supervision (see supplementary material for a detailed analysis).

Ablation of labeled data percentage. We conducted an ablation study where we varied the percentage of labeled data used during the training process. More specifically, we used 4 different settings: 5%, 25%, 50% and 75%. The remaining data were used as unlabeled samples. We used ResNet-50 as the network encoder for all the settings. This ablation study was conducted on 2 datasets:

³ $AG = \frac{G_{mae} + G_{mse}}{2}$, $G_{mae} = \frac{mae(\mathcal{D}_U + \mathcal{D}_L) - mae(\mathcal{D}_L)}{mae(\mathcal{D}_L)}$, $G_{mse} = \frac{mse(\mathcal{D}_U + \mathcal{D}_L) - mse(\mathcal{D}_L)}{mse(\mathcal{D}_L)}$

Table 2. Results of ablation study with different %-ages of labeled data. The proposed method achieves significant gains across different percentages of labeled data. We used ResNet-50 as the encoder network for all the experiments. AG: Average Gain %³.

\mathcal{D}_L %	SH-A					UCF-QNRF				
	No-GP (\mathcal{D}_L -only)		GP ($\mathcal{D}_L + \mathcal{D}_U$)		AG %	No-GP (\mathcal{D}_L -only)		GP ($\mathcal{D}_L + \mathcal{D}_U$)		AG %
	MAE	MSE	MAE	MSE		MAE	MSE	MAE	MSE	
5	118	211	102	172	16	186	295	160	275	10
25	110	160	91	149	12	178	252	147	226	14
50	102	149	89	148	6.1	158	250	136	218	13
75	93	146	88	139	4.7	139	240	129	210	9.8
100	76	126	-	-	-	114	195	-	-	-

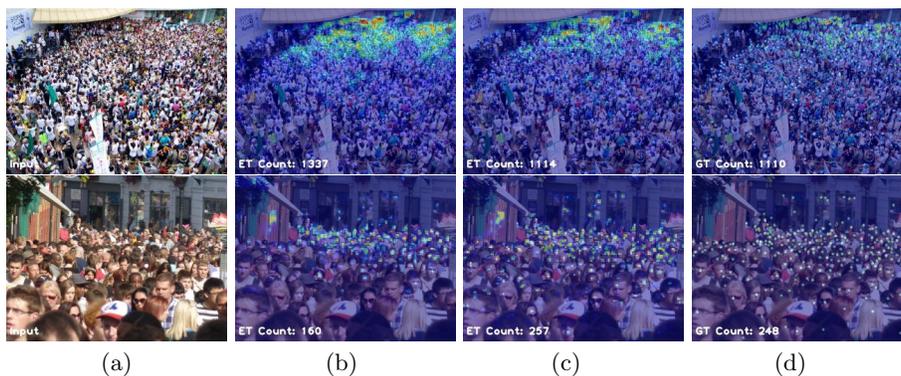


Fig. 3. Results of SSL experiments on the ShanghaiTech-A [61] dataset using the 5% labeled data setting. (a): Input. (b) No-GP (c) Proposed Method (d) Ground-truth.

ShanghaiTech-A (SH-A) and UCF-QNRF. The results of this ablation study are shown in Table 2. It can be observed for both datasets that as the percentage of labeled data is reduced, the performance of the baseline network drops significantly. However, the proposed GP-based framework is able to leverage unlabeled data in all the cases to reduce this performance drop by a considerable margin. Fig. 3 show sample qualitative results on ShanghaiTech-A dataset for the semi-supervised protocol with 5% labeled data setting. It can be observed that the proposed method is able to predict the density maps more accurately as compared to the baseline method that does not consider unlabeled data.

Architecture ablation. We conducted an ablation study where we evaluated the proposed method using different architectures. More specifically, we used different networks like ResNet-50, ResNet-101 and VGG16 as encoder network. The ablation was performed on 2 datasets: ShanghaiTech-A (SH-A) and UCF-QNRF. For all the experiments, we used 5% of the training dataset as labeled dataset, and the rest were used as unlabeled samples. The results of this experiment are shown in Table 3. Based on these results, we make the following observations: (i) Since networks like VGG16 and ResNet-101 have higher number of parameters, they tend to overfit more in the reduced-data setting as compared to ResNet-50. (ii)

Table 3. Results of ablation study with different networks. The proposed method is able to exploit unlabeled data irrespective of different architectures. We used 5% of the training data as labeled set, and the rest as unlabeled samples. AG: Average Gain %³.

Net	$\mathcal{D}_L\%$	SH-A					UCF-QNRF				
		No-GP(\mathcal{D}_L -only)		GP($\mathcal{D}_L + \mathcal{D}_U$)		AG	No-GP (\mathcal{D}_L -only)		GP ($\mathcal{D}_L + \mathcal{D}_U$)		AG
		MAE	MSE	MAE	MSE	%	MAE	MSE	MAE	MSE	%
ResNet-50	100	76	126	-	-	-	114	195	-	-	-
	5	118	211	102	172	16	186	295	160	275	10
ResNet-101	100	76	117	-	-	-	116	197	-	-	-
	5	131	200	110	162	18	196	324	174	288	11
VGG16	100	74	118	-	-	-	120	197	-	-	-
	5	121	205	112	163	14	188	316	175	291	7.4

The proposed GP-based method obtains consistent gains by leveraging unlabeled dataset across different architectures.

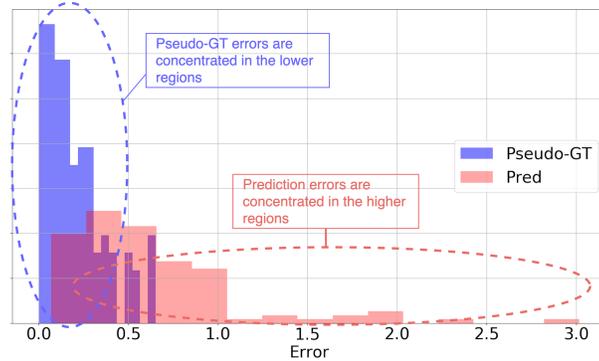


Fig. 4. Histogram for pseudo-GT errors (err_{pseudo}^u) and prediction errors (err_{pred}^u) on unlabeled data during training. Note that pseudo-GT errors are concentrated on the lower end, implying that they are more closer to the ground truth as compared to the predictions. Hence, pseudo-GTs provide meaningful supervision.

Pseudo-GT Analysis. In order to gain a deeper understanding about the effectiveness of the proposed approach, we plot the histogram of normalized errors with respect to the predictions y_{pred}^u of the network and the pseudo-GT y_{pseudo}^u for the unlabeled data during the training process. Specifically, we plot histograms of err_{pred}^u and err_{pseudo}^u , where $err_{pred}^u = \frac{|y_{pred}^u - y_{gt}^u|}{y_{gt}^u}$ and $err_{pseudo}^u = \frac{|y_{pseudo}^u - y_{gt}^u|}{y_{gt}^u}$. Here, y_{gt}^u is the actual GT corresponding to the unlabeled data sample. The plot is shown in Fig. 4. It can be observed that the pseudo-GT errors are concentrated in the lower end of the error region as compared to the prediction errors. This implies that the pseudo-GTs are more closer to the GTs than the predictions.

Table 4. Comparison of results in synthetic-to-real transfer settings. We train the network on synthetic crowd counting dataset (GCC), and leverage the training set of real-world datasets without any labels. We used the same network as described in [55].

Method	SH-A		SH-B		UCF-QNRF		UCF-CC-50		WExpo
	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
No Adapt	160	217	22.8	30.6	276	459	487	689	42.8
Cycle GAN [63]	143	204	24.4	39.7	257	401	405	548	32.4
SE Cycle GAN [55]	123	193	19.9	28.3	230	384	373	529	26.3
Proposed Method	121	181	12.8	19.2	210	351	355	505	20.4

Hence, the pseudo-GTs obtained using the proposed method are able to provide good quality supervision on the unlabeled data.

5.2 Synthetic-to-Real transfer setting

Recently, Wang *et al.*[55] proposed a synthetic crowd counting dataset (GCC) that consists of 15,212 images with a total of 7,625,843 annotations. The primary purpose of this dataset is to reduce the annotation efforts by training the networks on the synthetic dataset, thereby eliminating the need for labeling. However, due to a gap between the synthetic and real-world data distributions, the networks trained on synthetic dataset perform poorly on real-world images. In order to overcome this issue, the authors in [55] proposed a Cycle-GAN based domain adaptive approach that additionally enforces SSIM consistency. More specifically, they first learn to translate from synthetic crowd images to real-world images using SSIM-based Cycle-GAN. This transfers the style in the synthetic image to more real-world style. The translated synthetic images are then used to train a counting network (SFCN) that is based on ResNet-101 architecture.

While this approach improves the error over the baseline methods, its performance is essentially limited in the case of large distribution gap between real and synthetic images. Moreover, the authors in [55] perform a manual selection of synthetic samples for training the network. This selection ensures that only samples that are closer to the real-world images in terms of the count are used for training. Such a selection is not feasible in the case of unsupervised domain adaptation where we have no access to labels in the target dataset.

The proposed GP-based framework overcomes these drawbacks easily and can be extended to the synthetic-to-real transfer setting as well. We consider the synthetic data as labeled training set and real-world training set as unlabeled dataset, and train the network to leverage the unlabeled dataset. The results of this experiment are reported in Table 4. We used the same network (SFCN) and training process as described in [55]. As it can be observed, the proposed method achieves considerable improvements compared to the recent approach. Since we estimate the pseudo-GT for unlabeled real-world images and use it as supervision directly, the distribution gap that the network needs to handle is much lesser. This results in better performance compared to the domain adaptive approach [55]. Unlike [55], we train the network on the unlabeled data and hence, we do not



Fig. 5. Results of Synthetic-to-Real transfer experiments on ShanghaiTech-A dataset. (a): Input. (b) No Adapt (c) Proposed Method (d) Ground-truth.

need to perform any synthetic sample selection. Fig. 5 show sample qualitative results on the ShanghaiTech-A dataset for the synthetic-to-real transfer protocol. The proposed method is able to predict the density maps more accurately as compared to the baseline.

6 Conclusions

In this work, we focused on learning to count in the crowd from limited labeled data. Specifically, we proposed a GP-based iterative learning framework that involves estimation of pseudo-GT for unlabeled data using Gaussian Processes, which is then used as supervision for training the network. Through various experiments, we show that the proposed method can be effectively used in a variety of scenarios that involve unlabeled data like learning with less data or synthetic to real-world transfer. In addition, we conducted detailed ablation studies to demonstrate that the proposed method generalizes well to different network architectures and is able to achieve consistent gains for different amounts of labeled data.

Acknowledgement

This work was supported by the NSF grant 1910141.

References

1. Arteta, C., Lempitsky, V., Zisserman, A.: Counting in the wild. In: European Conference on Computer Vision. pp. 483–498. Springer (2016)
2. Babu Sam, D., Sajjan, N.N., Venkatesh Babu, R., Srinivasan, M.: 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. pp. 3618–3626 (2018)
3. Cao, X., Wang, Z., Zhao, Y., Su, F.: Scale aggregation network for accurate and efficient crowd counting. In: European Conference on Computer Vision. pp. 757–773. Springer (2018)
4. Chan, A.B., Liang, Z.S.J., Vasconcelos, N.: Privacy preserving crowd monitoring: Counting people without people models or tracking. In: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on. pp. 1–7. IEEE (2008)
5. Chan, A.B., Vasconcelos, N.: Bayesian poisson regression for crowd counting. In: 2009 IEEE 12th International Conference on Computer Vision. pp. 545–551. IEEE (2009)
6. Change Loy, C., Gong, S., Xiang, T.: From semi-supervised to transfer counting of crowds. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2256–2263 (2013)
7. Chen, K., Loy, C.C., Gong, S., Xiang, T.: Feature mining for localised crowd counting. In: European Conference on Computer Vision (2012)
8. Gao, G., Gao, J., Liu, Q., Wang, Q., Wang, Y.: Cnn-based density estimation and crowd counting: A survey. arXiv preprint arXiv:2003.12783 (2020)
9. Idrees, H., Saleemi, I., Seibert, C., Shah, M.: 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 (2013)
10. Idrees, H., Tayyab, M., Athrey, K., Zhang, D., Al-Maadeed, S., Rajpoot, N., Shah, M.: Composition loss for counting, density map estimation and localization in dense crowds. In: European Conference on Computer Vision. pp. 544–559. Springer (2018)
11. Jiang, X., Xiao, Z., Zhang, B., Zhen, X., Cao, X., Doermann, D., Shao, L.: Crowd counting and density estimation by trellis encoder-decoder network. arXiv preprint arXiv:1903.00853 (2019)
12. Kang, D., Ma, Z., Chan, A.B.: Beyond counting: Comparisons of density maps for crowd analysis tasks-counting, detection, and tracking. arXiv preprint arXiv:1705.10118 (2017)
13. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. arXiv preprint arXiv:1610.02242 (2016)
14. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks
15. Lempitsky, V., Zisserman, A.: Learning to count objects in images. In: Advances in Neural Information Processing Systems. pp. 1324–1332 (2010)
16. Li, M., Zhang, Z., Huang, K., Tan, T.: 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. pp. 1–4. IEEE (2008)
17. Li, T., Chang, H., Wang, M., Ni, B., Hong, R., Yan, S.: Crowded scene analysis: A survey. IEEE Transactions on Circuits and Systems for Video Technology **25**(3), 367–386 (2015)

18. Li, W., Mahadevan, V., Vasconcelos, N.: Anomaly detection and localization in crowded scenes. *IEEE transactions on pattern analysis and machine intelligence* **36**(1), 18–32 (2014)
19. Li, Y., Zhang, X., Chen, D.: 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 (2018)
20. Liu, N., Long, Y., Zou, C., Niu, Q., Pan, L., Wu, H.: Adcrowdnet: An attention-injective deformable convolutional network for crowd understanding. *arXiv preprint arXiv:1811.11968* (2018)
21. Liu, W., Salzmann, M., Fua, P.: Context-aware crowd counting. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5099–5108 (2019)
22. Liu, X., van de Weijer, J., Bagdanov, A.D.: Leveraging unlabeled data for crowd counting by learning to rank. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
23. Lu, H., Cao, Z., Xiao, Y., Zhuang, B., Shen, C.: Tasselnet: Counting maize tassels in the wild via local counts regression network. *Plant Methods* **13**(1), 79 (2017)
24. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)
25. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. In: *Advances in Neural Information Processing Systems*. pp. 3235–3246 (2018)
26. Onoro-Rubio, D., López-Sastre, R.J.: Towards perspective-free object counting with deep learning. In: *European Conference on Computer Vision*. pp. 615–629. Springer (2016)
27. Pham, V.Q., Kozakaya, T., Yamaguchi, O., Okada, R.: 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*. pp. 3253–3261 (2015)
28. Ranjan, V., Le, H., Hoai, M.: Iterative crowd counting. In: *European Conference on Computer Vision*. pp. 278–293. Springer (2018)
29. Rasmussen, C.E.: Gaussian processes in machine learning. In: *Summer School on Machine Learning*. pp. 63–71. Springer (2003)
30. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: *Advances in neural information processing systems*. pp. 91–99 (2015)
31. Ryan, D., Denman, S., Fookes, C., Sridharan, S.: Crowd counting using multiple local features. In: *Digital Image Computing: Techniques and Applications, 2009. DICTA'09*. pp. 81–88. IEEE (2009)
32. Sam, D.B., Babu, R.V.: Top-down feedback for crowd counting convolutional neural network. In: *Thirty-second AAAI conference on artificial intelligence* (2018)
33. Sam, D.B., Peri, S.V., Kamath, A., Babu, R.V., et al.: Locate, size and count: Accurately resolving people in dense crowds via detection. *arXiv preprint arXiv:1906.07538* (2019)
34. Sam, D.B., Peri, S.V., Mukuntha, N., Babu, R.V.: Going beyond the regression paradigm with accurate dot prediction for dense crowds. In: *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 2853–2861. IEEE (2020)

35. Sam, D.B., Sajjan, N.N., Maurya, H., Babu, R.V.: Almost unsupervised learning for dense crowd counting. In: Thirty-Third AAAI Conference on Artificial Intelligence (2019)
36. Sam, D.B., Surya, S., Babu, R.V.: Switching convolutional neural network for crowd counting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2017)
37. Shen, Z., Xu, Y., Ni, B., Wang, M., Hu, J., Yang, X.: Crowd counting via adversarial cross-scale consistency pursuit. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
38. Shi, M., Yang, Z., Xu, C., Chen, Q.: Revisiting perspective information for efficient crowd counting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7279–7288 (2019)
39. Shi, Z., Zhang, L., Liu, Y., Cao, X., Ye, Y., Cheng, M.M., Zheng, G.: Crowd counting with deep negative correlation learning. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
40. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (2015)
41. Sindagi, V.A., Patel, V.M.: 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)
42. Sindagi, V.A., Patel, V.M.: Generating high-quality crowd density maps using contextual pyramid cnns. In: The IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
43. Sindagi, V.A., Patel, V.M.: A survey of recent advances in cnn-based single image crowd counting and density estimation. *Pattern Recognition Letters* (2017)
44. Sindagi, V.A., Patel, V.M.: Dafe-fd: Density aware feature enrichment for face detection. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 2185–2195. IEEE (2019)
45. Sindagi, V.A., Patel, V.M.: Ha-cnn: Hierarchical attention-based crowd counting network. arXiv preprint arXiv:1907.10255 (2019)
46. Sindagi, V.A., Patel, V.M.: Inverse attention guided deep crowd counting network. arXiv preprint (2019)
47. Sindagi, V.A., Patel, V.M.: Multi-level bottom-top and top-bottom feature fusion for crowd counting. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1002–1012 (2019)
48. Sindagi, V.A., Yasarla, R., Patel, V.M.: Pushing the frontiers of unconstrained crowd counting: New dataset and benchmark method. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1221–1231 (2019)
49. Sindagi, V.A., Yasarla, R., Patel, V.M.: Jhu-crowd++: Large-scale crowd counting dataset and a benchmark method. arXiv preprint arXiv:2004.03597 (2020)
50. Toropov, E., Gui, L., Zhang, S., Kottur, S., Moura, J.M.: Traffic flow from a low frame rate city camera. In: Image Processing (ICIP), 2015 IEEE International Conference on. pp. 3802–3806. IEEE (2015)
51. Walach, E., Wolf, L.: Learning to count with cnn boosting. In: European Conference on Computer Vision. pp. 660–676. Springer (2016)
52. Wan, J., Chan, A.: Adaptive density map generation for crowd counting. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1130–1139 (2019)
53. Wan, J., Luo, W., Wu, B., Chan, A.B., Liu, W.: Residual regression with semantic prior for crowd counting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4036–4045 (2019)

54. Wang, C., Zhang, H., Yang, L., Liu, S., Cao, X.: Deep people counting in extremely dense crowds. In: Proceedings of the 23rd ACM international conference on Multimedia. pp. 1299–1302. ACM (2015)
55. Wang, Q., Gao, J., Lin, W., Yuan, Y.: Learning from synthetic data for crowd counting in the wild. arXiv preprint arXiv:1903.03303 (2019)
56. Xu, B., Qiu, G.: Crowd density estimation based on rich features and random projection forest. In: 2016 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1–8. IEEE (2016)
57. Yasarla, R., Sindagi, V.A., Patel, V.M.: Syn2real transfer learning for image de-raining using gaussian processes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2726–2736 (2020)
58. Zhan, B., Monekosso, D.N., Remagnino, P., Velastin, S.A., Xu, L.Q.: Crowd analysis: a survey. *Machine Vision and Applications* **19**(5-6), 345–357 (2008)
59. Zhang, C., Li, H., Wang, X., Yang, X.: Cross-scene crowd counting via deep convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 833–841 (2015)
60. Zhang, Q., Chan, A.B.: 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. pp. 8297–8306 (2019)
61. Zhang, Y., Zhou, D., Chen, S., Gao, S., Ma, Y.: 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 (2016)
62. Zhao, M., Zhang, J., Zhang, C., Zhang, W.: Leveraging heterogeneous auxiliary tasks to assist crowd counting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12736–12745 (2019)
63. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)