

CNN-based Cascaded Multi-task Learning of High-level Prior and Density Estimation for Crowd Counting

Vishwanath A. Sindagi Vishal M. Patel

Department of Electrical and Computer Engineering, Rutgers University
94 Brett Road, Piscataway, NJ, 08854, USA

vishwanath.sindagi@rutgers.edu, vishal.m.patel@rutgers.edu

Abstract

Estimating crowd count in densely crowded scenes is an extremely challenging task due to non-uniform scale variations. In this paper, we propose a novel end-to-end cascaded network of CNNs to jointly learn crowd count classification and density map estimation. Classifying crowd count into various groups is tantamount to coarsely estimating the total count in the image thereby incorporating a high-level prior into the density estimation network. This enables the layers in the network to learn globally relevant discriminative features which aid in estimating highly refined density maps with lower count error. The joint training is performed in an end-to-end fashion. Extensive experiments on highly challenging publicly available datasets show that the proposed method achieves lower count error and better quality density maps as compared to the recent state-of-the-art methods.

1. Introduction

Crowd analysis has gained a lot of interest in recent years due to its variety of applications such as video surveillance, public safety design and traffic monitoring. Researchers have attempted to address various aspects of analyzing crowded scenes such as counting [3], [4], [23], [10], density estimation [12], [32], [31], [18], [29], [2], segmentation [11], behavior analysis [24], tracking [21], scene understanding [25] and anomaly detection [19]. In this paper, we specifically focus on the joint task of estimating crowd count and density map from a single image.

One of the many challenges faced by researchers working on crowd counting is the issue of large variations in scale and appearance of the objects that occurs due to severe perspective distortion of the scene. Many methods


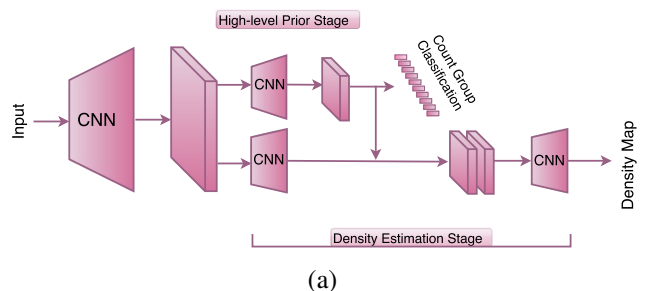


Fig. 1: Proposed method and results. (a) Cascaded architecture for learning high-level prior and density estimation. (b) Input image (from the ShanghaiTech dataset [32]). (c) Ground truth density map. (d) Density map generated by the proposed method.

have been developed that incorporate scale information into the learning process using different methods. Some of the early methods relied on multi-source and hand-crafted representations and catered only to low density crowded scenes [10]. These methods are rendered ineffective in high density crowds and the results are far from optimal. Inspired by the success of Convolutional Neural Networks (CNNs) for various computer vision tasks, many CNN-based methods have been developed to address the problem of crowd counting [2], [1], [31]. Considering scale issue as a limiting factor to achieve better accuracies, certain CNN-based methods specifically cater to the issue of scale changes via multi-column or multi-resolution network [32], [17], [23]. Though these methods demonstrated robustness to scale changes, they are still restricted to the

scales that are used during training and hence are limited in their capacity to learn well-generalized models.

The aim of this work is to learn models that cater to a wide variety of density levels present in the dataset by incorporating a high-level prior into the network. The high-level prior learns to classify the count into various groups whose class labels are based on the number of people present in the image. By exploiting count labels, the high-level prior is able to estimate coarse count of people in the entire image irrespective of scale variations thereby enabling the network to learn more discriminative global features. The high-level prior is jointly learned along with density map estimation using a cascade of CNN networks as shown in Fig. 1 (a). The two tasks (crowd count classification and density estimation) share an initial set of convolutional layers which is followed by two parallel set of networks that learn high-dimensional feature maps relevant to high-level prior and density estimation, respectively. The global features learned by the high-level prior are concatenated with the feature maps obtained from the second set of convolutional layers and further processed by a set of fractionally strided convolutional layers to produce high resolution density maps. Results of the proposed method on a sample input image are shown in Fig. 1 (c)-(d).

2. Related work

Traditional approaches for crowd counting from single images relied on hand-crafted representations to extract low level features. These features were then mapped to count or density map using various regression techniques. Loy *et al.* [14] categorized existing methods into (1) detection-based methods (2) regression-based methods and (3) density estimation-based methods.

Detection-based methods typically employ sliding window-based detection algorithms to count the number of object instances in an image [26]. These methods are adversely affected by the presence of high density crowd and background clutter. To overcome these issues, researchers attempted to count by regression where they learn a mapping between features extracted from local image patches to their counts [22], [6]. Using a similar approach, Idrees *et al.* [10] fused count from multiple sources. The authors also introduced an annotated dataset (UCF_CC_50) of 50 images containing 64000 humans.

Detection and regression methods ignore key spatial information present in the images as they regress on the global count. Hence, in order to incorporate spatial information present in the images, Lempitsky *et al.* [12] introduced a new approach of learning a linear mapping between local patch features and corresponding object density maps. Instead of a linear mapping, Pham *et al.* in [18] proposed to learn a non-linear function using a

random forest framework. Wang and Zou [29] computed the relationship between image patches and their density maps in two distinct feature spaces. Recently, Xu and Qiu [30] proposed to use much richer and extensive set of features for crowd density estimation. A more comprehensive survey of different crowd counting methods can be found in [6], [13].

More recently, due to the success of CNNs in various computer vision tasks, several CNN-based approaches have been developed for crowd counting [28], [31], [15], [16]. Walach *et al.* [27] used CNNs with layered training approach. In contrast to the existing patch-based estimation methods, Shang *et al.* [23] proposed an end-to-end estimation method using CNNs by simultaneously learning local and global count on the whole sized input images. Observing that the existing approaches cater to a single scale due to their fixed receptive fields, Zhang *et al.* [32] proposed a multi-column architecture to extract features at different scales. In addition, they also introduced a large scale annotated dataset (ShanghaiTech dataset). Onoro-Rubio and López-Sastre in [17] addressed the scale issue by proposing a scale aware counting model called Hydra CNN. Boominathan *et al.* in [2] proposed to tackle the issue of scale variation using a combination of shallow and deep networks along with an extensive data augmentation by sampling patches from multi-scale image representations.

Zhang *et al.* [32] and Onoro *et al.* [17] demonstrated that designing networks that are robust to scale variations is crucial for achieving better performance as compared to other CNN-based approaches. However, these methods rely on architectures that cater to selected set of scales thereby limiting their abilities to learn more generalized models. Additionally, the recent approaches individually regress either on crowd count or density map. Among the approaches that estimate density maps, the presence of pooling layers in the existing approaches reduce the resolution of the output density map prohibiting one to regress on full resolution density maps. This results in the loss of crucial details especially in images containing large variation in scales. Considering these drawbacks, we present a novel end-to-end cascaded CNN network that jointly learns a high-level global prior and density estimation. The high-level prior enables the network to learn globally relevant and discriminative features that aid in estimating density maps from images with large variations in scale and appearance.

3. Proposed method

Inspired by the success of cascaded convolutional networks for related multiple tasks [5], [8], [20], we propose to learn two related sub-tasks: crowd count classification (which we call as high-level prior) and density map estima-




Fig. 2: Overview of the proposed cascaded architecture for jointly learning high-level prior and density estimation.

tion in a cascaded fashion as shown in Fig. 2. The network takes an image of arbitrary size, and outputs crowd density map. The cascaded network has two stages corresponding to the two sub-tasks, with the first stage learning high-level prior and the second stage performing density map estimation. Both stages share a set of convolutional features. The first stage consists of a set of convolutional layers and spatial pyramid pooling to handle arbitrarily sized images followed by a set of fully connected layers. The second stage consists of a set of convolutional layers followed by fractionally-strided convolutional layers for upsampling the previous layer’s output to account for the loss of details due to earlier pooling layers. Two different set of loss layers are used at the end of the two stages, however, the loss of the second layer is dependent on the output of the earlier stage. The following sub-sections discuss the details of all the components of the proposed network.

3.1.Shared convolutional layers

The initial shared network consists of 2 convolutional layers with a Parametric Rectified Linear Unit (PReLU) activation function after every layer. The first convolutional layer has 16 feature maps with a filter size of 9×9 and the second convolutional layer has 32 feature maps with a filter size of 7×7 . The feature maps generated by this shallow network are shared by the two stages: high-level prior stage and density estimation stage.

3.2.High-level prior stage

Classifying the crowd into several groups is an easier problem as compared to directly performing classification or regression for the whole count range which requires a larger amount of training data. Hence, we quantize the crowd count into ten groups and learn a crowd count group classifier which also performs the task of incorporating high-level prior into the network. The high-level prior stage takes feature maps from the previous shared convolutional layers. This stage consists of 4 convolutional layers with a PReLU activation function after every layer. The first two

layers are followed by max pooling layers with a stride of 2. At the end, the high-level prior stage consists of three fully connected (FC) layers with a PReLU activation function after every layer. The first FC layer consists of 512 neurons whereas the second FC layer consists of 256 neurons. The final layer consists of a set of 10 neurons followed by a sigmoid layer, indicating the count class of the input image. To enable the use of arbitrarily sized images for training, Spatial Pyramid Pooling (SPP) [9] is employed as it eliminates the fixed size constraint of deep networks which contain fully connected layers. The SPP layer is inserted after the last convolutional layer. The SPP layer aggregates features from the convolutional layers to produce fixed size outputs and can be fed to the fully connected layers. Cross-entropy error is used as the loss layer for this stage.

3.3.Density estimation

The feature maps obtained from the shared layers are processed by an another CNN network that consists of 4 convolutional layers with a PReLU activation function after every layer. The first two layers are followed by max pooling layers with a stride of 2, due to which the output of CNN layers is downsampled by a factor of 4. The first convolutional layer has 20 feature maps with a filter size of 7×7 , the second convolutional layer has 40 feature maps with a filter size of 5×5 , the third layer has 20 feature maps with a filter size of 5×5 and the fourth layer has 10 feature maps with a filter size of 5×5 . The output of this network is combined with that of the last convolutional layer of high-level prior stage using a set of 2 convolutional and 2 fractionally strided convolutional layers. The first two convolutional layers have a filter size of 3×3 with 24 and 32 feature maps, respectively. These layers are followed by 2 sets of fractionally strided convolutional layers with 16 and 18 feature maps, respectively. In addition to integrating high-level prior from an earlier stage, the fractionally strided convolutions learn to upsample the feature maps to the original input size thereby restoring the details lost due to earlier max-pooling layers. The use of these layers results in upsampling of the CNN output by a factor of 4, thus enabling us to regress on full resolution density maps. Standard pixel-wise Euclidean loss is used as the loss layer for this stage. Note that this loss depends on intermediate output of the earlier cascade, thereby enforcing a causal relationship between count classification and density estimation.

3.4.Objective function

The cross-entropy loss function for the high-level prior stage is defined as follows:

$$L_c = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M [(y^i = j) F_c(X_i, \Theta)], \quad (1)$$

where N is number of training samples, Θ is a set of network parameters, X_i is the i^{th} training sample, $F_c(X_i, \Theta)$ is the classification output, y^i is the ground truth class and M is the total number of classes.

The loss function for the density estimation stage is defined as:

$$L_d = \frac{1}{N} \sum_{i=1}^N \|F_d(X_i, C_i, \Theta) - D_i\|_2, \quad (2)$$

where $F_d(X_i, C_i, \Theta)$ is the estimated density map, D_i is the ground truth density map, and C_i are the feature maps obtained from the last convolutional layer of the high-level prior stage.

The entire cascaded network is trained using the following unified loss function:

$$L = \lambda L_c + L_d, \quad (3)$$

where λ is a weighting factor.

This loss function is unlike traditional multi-task learning, because the loss term of the last stage depends on the output of the earlier one.

3.5. Training and implementation details

In this section, details of the training procedure are discussed. To create the training dataset, patches of size $1/4^{th}$ the size of original image are cropped from 100 random locations. Other augmentation techniques like horizontal flipping and noise addition are used to create another 200 patches. The random cropping and augmentation resulted in a total of 300 patches per image in the training dataset. Note that the cropping is used only as a data augmentation technique and the resulting patches are of arbitrary sizes.

Several sophisticated methods are proposed in the literature for calculating the ground truth density map [31], [32]. We use a simple method in order to ensure that the improvements achieved are due to the proposed method and are not dependent on the sophisticated methods for calculating the ground truth density maps. Ground truth density map D_i corresponding to the i^{th} training patch is calculated by summing a 2D Gaussian kernel centered at every person's location x_g as defined below:

$$D_i(x) = \sum_{x_g \in S} \mathcal{N}(x - x_g, \sigma), \quad (4)$$

where σ is the scale parameter of the 2D Gaussian kernel and S is the set of all points at which people are located.

The training and evaluation was performed on NVIDIA GTX TITAN-X GPU using Torch framework [7]. λ was set to 0.0001 in (3). Adam optimization with a learning rate of 0.00001 and momentum of 0.9 was used to train the model. Additionally, for the classification (high-level

prior) stage, to account for the imbalanced datasets, the losses for each class were weighted based on the number of samples available for that particular class. The training took approximately 6 hours.

4. Experimental results

In this section, we present the experimental details and evaluation results on two publicly available datasets: ShanghaiTech [32] and UCF_CROWD_50 [10]. For the purpose of evaluation, the standard metrics used by many existing methods for crowd counting were used. These metrics are defined as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - y'_i|, \quad MSE = \sqrt{\frac{1}{N} \sum_{i=1}^N |y_i - y'_i|^2},$$

where MAE is mean absolute error, MSE is mean squared error, N is number of test samples, y_i is ground truth count and y'_i is estimated count corresponding to the i^{th} sample.

4.1. ShanghaiTech dataset

The ShanghaiTech dataset was introduced by Zhang *et al.* [32] and it contains 1198 annotated images with a total of 330,165 people. This dataset consists of two parts: Part A with 482 images and Part B with 716 images. Both parts are further divided into training and test datasets with training set of Part A containing 300 images and that of Part B containing 400 images. Rest of the images are used as test set. The results of the proposed method are compared with two recent approaches: Zhang *et al.* [31] and MCNN by Zhang *et al.* [32] (Table I). The authors in [31] proposed a switchable learning function where they learned their network by alternatively training on two objective functions: crowd count and density estimation. In the other approach by Zhang *et al.* in [32], the authors proposed a multi-column convolutional network (MCNN) to address scale issues and a sophisticated ground truth density map generation technique. It can be observed from Table I, that the proposed method is able to achieve significant improvements without the use of multi-column networks or sophisticated ground truth map generation. Furthermore, to demonstrate the improvements obtained by incorporating high-level prior via cascaded architecture, we evaluated our network without the high-level prior stage (Single stage CNN) on ShanghaiTech dataset. It can be observed from Table I, that the cascaded learning of count classification and density estimation reduces the count error by a large margin as compared to the single stage CNN.

Fig. 3 illustrates the density map results obtained using the proposed method as compared to Zhang *et al.* [32] and single stage CNN. It can be observed that in addition to achieving lower count error, the proposed method results in




Fig. 3: Density estimation results using proposed method on ShanghaiTech dataset. (a) Input (b) Ground truth (c) Output.

higher quality density maps due to the use of fractionally strided convolutional layers.

TABLE I: Comparison results: Estimation errors on the ShanghaiTech dataset. The proposed method achieves lower error compared to existing approaches involving multi column CNNs and sophisticated density maps.

Method	Part A		Part B	
	MAE	MSE	MAE	MSE
Zhang <i>et al.</i> [31]	181.8	277.7	32.0	49.8
MCNN [32]	110.2	173.2	26.4	41.3
Single stage CNN	130.4	190.9	29.3	40.5
Proposed method	101.3	152.4	20.0	31.1

4.2. UCF_CC_50 dataset

The UCF_CC_50 is an extremely challenging dataset introduced by Idrees *et al.* [10]. The dataset contains 50 annotated images of different resolutions and aspect ratios crawled from the internet. There is a large variation in densities across images. Following the standard protocol discussed in [10], a 5-fold cross-validation was performed for evaluating the proposed method. The results are compared with five recent approaches: Idrees *et al.* [10], Zhang *et al.* [31], MCNN [32], Onoro *et al.* [17] and Walach *et al.* [27]. The authors in [10] proposed to combine information from multiple sources such as head detections, Fourier analysis and texture features (SIFT). Onoro *et al.* in [17] proposed a scale aware CNN to learn a multi-scale non-linear regression model using a pyramid of image patches extracted at multiple scales. Walach *et al.* [27] proposed a layered approach of learning CNNs for crowd counting by iteratively adding CNNs where every new CNN is trained on residual error of the previous layer. It can be observed

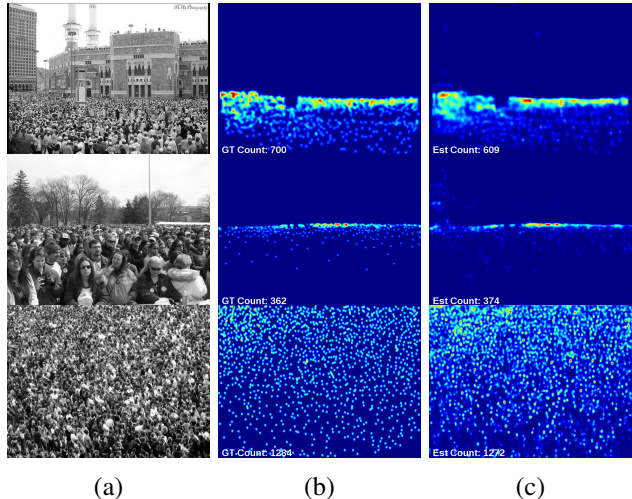


Fig. 4: Density estimation results using proposed method on UCF_CC_50 dataset. (a) Input (b) Ground truth (c) Output.

from Table II that our network achieves the lowest MAE and comparable MSE score. Density maps obtained using the proposed method on sample images from UCF_CC_50 dataset are shown in Fig. 4.

TABLE II: Comparison results: Estimation errors on the UCF_CC_50 dataset.

Method	MAE	MSE
Idrees <i>et al.</i> [10]	419.5	541.6
Zhang <i>et al.</i> [31]	467.0	498.5
MCNN [32]	377.6	509.1
Onoro <i>et al.</i> [17]	465.7	371.8
Walach <i>et al.</i> [27]	364.4	341.4
Proposed method	322.8	397.9

5. Conclusions

In this paper, we presented a multi-task cascaded CNN network for jointly learning crowd count classification and density map estimation. By learning to classify the crowd count into various groups, we are able to incorporate a high-level prior into the network which enables it to learn globally relevant discriminative features thereby accounting for large count variations in the dataset. Additionally, we employed fractionally strided convolutional layers at the end so as to account for the loss of details due to max-pooling layers in the earlier stages there by allowing us to regress on full resolution density maps. The entire cascade was trained in an end-to-end fashion. Extensive experiments performed on challenging datasets and comparison with recent state-of-the-art approaches demonstrated the significant improvements achieved by the proposed method.

Acknowledgement

This work was supported by US Office of Naval Research (ONR) Grant YIP N00014-16-1-3134.

References

- [1] A. Bansal and K. Venkatesh. People counting in high density crowds from still images. *arXiv preprint arXiv:1507.08445*, 2015.
- [2] L. Boominathan, S. S. Kruthiventi, and R. V. Babu. Crowdnet: A deep convolutional network for dense crowd counting. In *Proceedings of the 2016 ACM on Multimedia Conference*, pages 640–644. ACM, 2016.
- [3] A. B. Chan, Z.-S. J. Liang, and N. Vasconcelos. Privacy preserving crowd monitoring: Counting people without people models or tracking. In *IEEE CVPR*, pages 1–7. IEEE, 2008.
- [4] A. B. Chan and N. Vasconcelos. Counting people with low-level features and bayesian regression. *IEEE Transactions on Image Processing*, 21(4):2160–2177, 2012.
- [5] J.-C. Chen, A. Kumar, R. Ranjan, V. M. Patel, A. Alavi, and R. Chellappa. A cascaded convolutional neural network for age estimation of unconstrained faces. In *International Conference on BTAS*, pages 1–8. IEEE, 2016.
- [6] K. Chen, C. C. Loy, S. Gong, and T. Xiang. Feature mining for localised crowd counting. In *ECCV*, 2012.
- [7] R. Collobert, K. Kavukcuoglu, and C. Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, number EPFL-CONF-192376, 2011.
- [8] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *IEEE CVPR*, pages 3150–3158, 2016.
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*, pages 346–361. Springer, 2014.
- [10] H. Idrees, I. Saleemi, C. Seibert, and M. Shah. Multi-source multi-scale counting in extremely dense crowd images. In *IEEE CVPR*, pages 2547–2554, 2013.
- [11] K. Kang and X. Wang. Fully convolutional neural networks for crowd segmentation. *arXiv preprint arXiv:1411.4464*, 2014.
- [12] V. Lempitsky and A. Zisserman. Learning to count objects in images. In *NIPS*, pages 1324–1332, 2010.
- [13] T. Li, H. Chang, M. Wang, B. Ni, R. Hong, and S. Yan. Crowded scene analysis: A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(3):367–386, 2015.
- [14] C. C. Loy, K. Chen, S. Gong, and T. Xiang. Crowd counting and profiling: Methodology and evaluation. In *Modeling, Simulation and Visual Analysis of Crowds*, pages 347–382. Springer, 2013.
- [15] M. Marsden, K. McGuinness, S. Little, and N. O’Connor. Fully convolutional crowd counting on highly congested scenes. In *The International Conference on Computer Vision Theory and Applications*, 2017.
- [16] M. Marsden, K. McGuinness, S. Little, and N. E. O’Connor. Resnetcrowd: A residual deep learning architecture for crowd counting, violent behaviour detection and crowd density level classification. *arXiv preprint arXiv:1705.10698*, 2017.
- [17] D. Onoro-Rubio and R. J. López-Sastre. Towards perspective-free object counting with deep learning. In *ECCV*, pages 615–629. Springer, 2016.
- [18] V.-Q. Pham, T. Kozakaya, O. Yamaguchi, and R. Okada. Count forest: Co-voting uncertain number of targets using random forest for crowd density estimation. In *Proceedings of the IEEE ICCV*, pages 3253–3261, 2015.
- [19] H. Rabiee, J. Haddadnia, H. Mousavi, M. Kalantarzadeh, M. Nabi, and V. Murino. Novel dataset for fine-grained abnormal behavior understanding in crowd. In *IEEE International Conference on AVSS*, pages 95–101. IEEE, 2016.
- [20] R. Ranjan, V. Patel, and R. Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE transactions on PAMI*, 2016.
- [21] M. Rodriguez, I. Laptev, J. Sivic, and J.-Y. Audibert. Density-aware person detection and tracking in crowds. In *IEEE ICCV*, pages 2423–2430. IEEE, 2011.
- [22] D. Ryan, S. Denman, C. Fookes, and S. Sridharan. Crowd counting using multiple local features. In *Digital Image Computing: Techniques and Applications, 2009. DICTA’09.*, pages 81–88. IEEE, 2009.
- [23] C. Shang, H. Ai, and B. Bai. End-to-end crowd counting via joint learning local and global count. In *IEEE ICIP*, pages 1215–1219. IEEE, 2016.
- [24] J. Shao, C. Change Loy, and X. Wang. Scene-independent group profiling in crowd. In *IEEE CVPR*, pages 2219–2226, 2014.
- [25] J. Shao, K. Kang, C. C. Loy, and X. Wang. Deeply learned attributes for crowded scene understanding. In *Proceedings of the IEEE CVPR*, pages 4657–4666. IEEE, 2015.
- [26] I. S. Topkaya, H. Erdogan, and F. Porikli. Counting people by clustering person detector outputs. In *IEEE International Conference on AVSS*, pages 313–318. IEEE, 2014.
- [27] E. Walach and L. Wolf. Learning to count with cnn boosting. In *ECCV*, pages 660–676. Springer, 2016.
- [28] C. Wang, H. Zhang, L. Yang, S. Liu, and X. Cao. Deep people counting in extremely dense crowds. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 1299–1302. ACM, 2015.
- [29] Y. Wang and Y. Zou. Fast visual object counting via example-based density estimation. In *IEEE ICIP*, pages 3653–3657. IEEE, 2016.
- [30] B. Xu and G. Qiu. Crowd density estimation based on rich features and random projection forest. In *2016 IEEE WACV*, pages 1–8. IEEE, 2016.
- [31] C. Zhang, H. Li, X. Wang, and X. Yang. Cross-scene crowd counting via deep convolutional neural networks. In *Proceedings of the IEEE CVPR*, pages 833–841, 2015.
- [32] 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 CVPR*, pages 589–597, 2016.